A Large Deviations Analysis of the Multiplexing of Periodic Traffic Sources

Richard Weber

Columbia, Applied Probability Day

5 May 2000

Overflow in a multiplexer

The figure below shows a 2×2 switch, where output links are served at rate C cells per second.

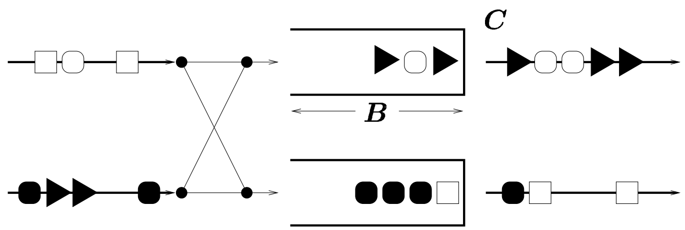

To say how many virtual circuits can use this output link, subject to a given Quality of Service constraint, requires an estimate of the probability that the size of the queue, Q, exceeds the buffer of size B.

When $Q = B$ there is cell loss and cell loss probability (CLP) should be small, say 10^{-8}.

$\mathbb{P}(Q = B)$ should be small.

Two possible source models

A cell scale source model

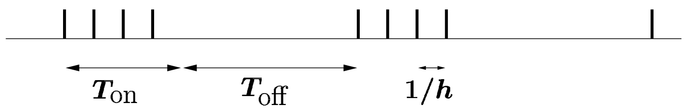

The source has on and off phases.

In the on phase it delivers one cell to the buffer every $1/h$ seconds.

h cells per second is the 'peak rate'.

A fluid source model

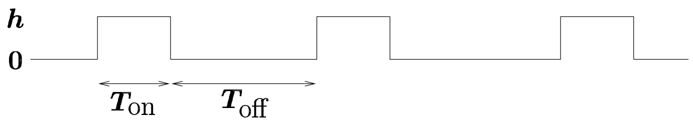

The source has on and off phases.

In the on phase it delivers fluid to the buffer a rate of h units per seconds.

h cells per second is the 'peak rate'.

Large deviations

We shall estimate $P(Q = B)$ (and CLP) using large deviations.

- The performance of many systems is limited by events which have a small probability of occurring, but which have severe consequences when they occur.
- LD deals with rare events, and is asymptotic in nature.
- It can be viewed as a refinement of the law of large numbers.
- It is useful when simulation or numerical techniques become increasingly difficult as a parameter tends to its limit.
- It has many applications:

  queueing and communications models,

  information theory,

  simulation techniques,

  parameter estimation,

  hypothesis testing, $\ldots$

Overflow in a $M/M/1/B$ queue

Consider a $M/M/1/B$ queue, with finite buffer, here being shared by traffic sources with combined Poisson arrival rate λ.

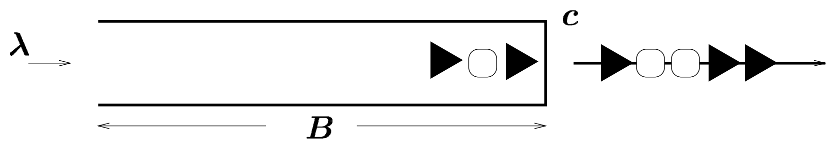

We know

$$\mathbb{P}(Q = B) = \left[\frac{1-(\lambda/c)}{1-(\lambda/c)^{B+1}}\right](\lambda/c)^B.$$

Hence

$$\mathbb{P}(Q = B) \sim e^{-B\log(c/\lambda)} \quad \text{for large } B,$$

where $\sim$ means

$$\lim_{B\to\infty}\frac{1}{B}\log\mathbb{P}(Q = B) = -\log(c/\lambda).$$

This result is a typical large deviations result.

The 'large N' asymptotic: identical sources

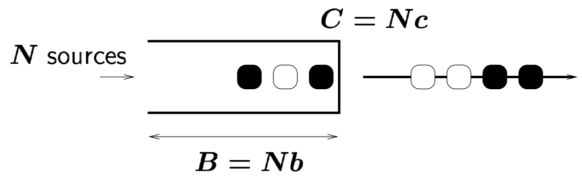

Suppose buffer and bandwidth scale with N.

Sources share a common output buffer of size $B = Nb$ and an output link with bandwidth $C = Nc$.

Denote the log moment generating function

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

where $X[0,t]$ is the load produced by a single source in $[0,t]$.

Sources are independent and have stationary increments.

Then [Courcoubetis and Weber, 1996]

$$\mathbb{P}\left(Q^N > Nb\right) \approx \text{CLP} \sim e^{-NI},$$

where

$$I = \inf_{t>0}\sup_{s>0}\{s(b+ct) - \varphi(s,t)\}.$$

LD for i.i.d. random variables

Suppose $X_1, X_2, \ldots$ are i.i.d. random variables. Then

$$\mathbb{P}(X_1 + \cdots + X_N > aN) \sim e^{-NI(a)}$$

where $I(a) = \inf_{s>0} \left\{ sa - \log \mathbb{E} e^{sX_1} \right\}$.

E.g., if $X_1, X_2, \ldots$ are $B(1,p)$,

$$\begin{aligned} I(a) &= \inf_{s>0} \{ sa - (q + pe^s) \} \\ &= a \log \left(\frac{a}{p} \right) + (1-a) \log \left(\frac{1-a}{1-p} \right) \end{aligned}$$

An upper bound is easy. For all $s > 0$,

$$\begin{aligned} &\mathbb{P}(X_1 + \cdots + X_N > aN) \\ &\qquad \le \mathbb{E} \left[1_{\{X_1 + \cdots + X_N > aN\}} e^{s(X_1 + \cdots + X_N - aN)} \right] \\ &\qquad \le \mathbb{E} \left[e^{s(X_1 + \cdots + X_N - aN)} \right] \\ &\qquad = \exp \left(-N \left\{ sa - \log E e^{sX_1} \right\} \right) \end{aligned}$$

Hence

$$\begin{aligned} &\mathbb{P}(X_1 + \cdots + X_N > aN) \\ &\qquad \le \exp \left(-N \inf_{s>0} \left\{ sa - \log E e^{sX_1} \right\} \right) \end{aligned}$$

Intuition for the large N aymptotic

$$\mathbb{P} \left(Q^N > Nb \right) \sim e^{-NI},$$

where

$$I = \inf_{t>0} \sup_{s>0} \{ s(b + ct) - \varphi(s,t) \} .$$

Intuition:

$$\mathbb{P} \left(Q^N > Nb \right) = \sup_{t>0} \mathbb{P} \left(\sum_{j=1}^{N} X_j[-t, 0] > Nb + Nct \right)$$

The optimizing t, say t^*, can be interpreted as the typical time over which overflow occurs.

The optimizing s, say s^*, reflects the multiplexing gain. If s^* is near 0 then there is a substantial multiplexing gain and the 'effective bandwidth' of a source is close to its mean. If s^* is large then there is little multiplexing gain and the effective bandwidth of a source is close to its peak rate.

The 'large N' asymptotic: heterogeneous sources

Suppose there are $n_j = N\rho_j$ sources of type j, $\sum_{j \in J} \rho_j = 1$.

All sources are independent and have stationary increments.

Let $X^j[0,t]$ be the total load produced by a source of type j in an interval $[0,t]$.

Then

$$\lim_{N\to\infty} \frac{1}{N} \log \mathbb{P}(Q^N \geq Nb) = -NI(c,b)$$

where

$$I(c,b) = \inf_{t>0} \sup_{s>0} \left\{ s(b+tc) - \sum_{j\in J} \rho_j \log \mathbb{E} e^{sX^j[0,t]} \right\}$$

Some questions

$$\mathbb{P}\left(Q^N > Nb\right) \sim e^{-NI},$$

where

$$I = \inf_{t>0} \sup_{s>0} \{ s(b+ct) - \varphi(s,t) \} .$$

1. Does t^* increase continuously with b?
2. Does s^* decrease continuously with b?
3. Does a fluid model give a good approximation to I, or must we model cell scale effects?

A periodic source — cell scale model

Consider a periodic source, that is on for 0.20 of the time, i.e., $m = 4/20$. An on burst delivers 4 cells, and is followed by an off phase lasting four times as long.

A periodic source

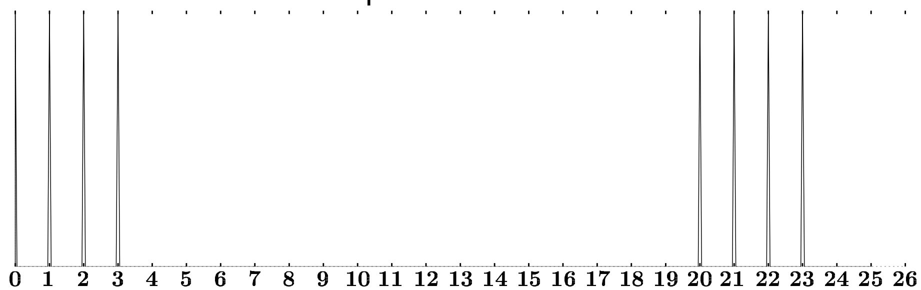

Suppose N such sources are input to a multiplexer that can serve $0.22Nt$ cells in time t. The beginnings of their on phases are independent and uniformly distributed.

How large a buffer guarantees no overflow?

What is the cell loss probability when the buffer is smaller?

The largest buffer that can be filled

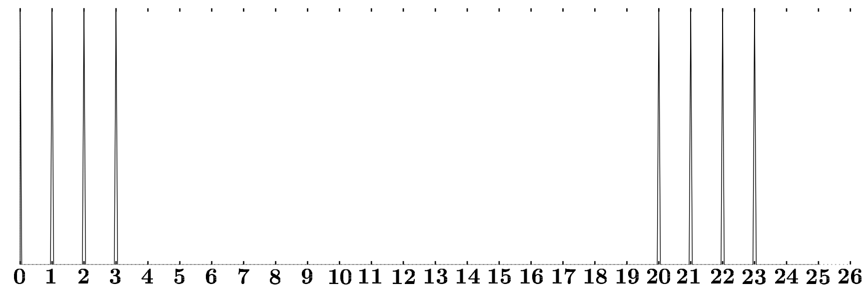

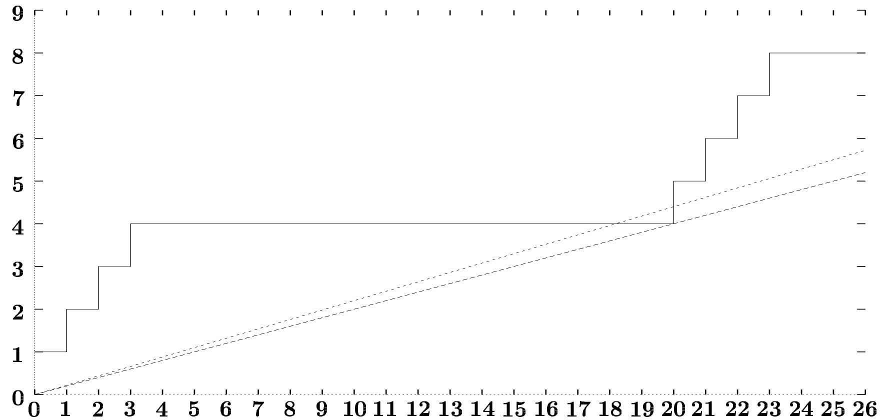

The largest buffer that can be filled is filled when all sources are synchronized to produce the maximal difference between inflow and outflow. This occurs when

$$B = (4 - 3c)N = (4 - 3(.22))N = 3.34N$$

The fluid source model

Sources are sometimes modelled as fluids. E.g., periodic fluid sources:

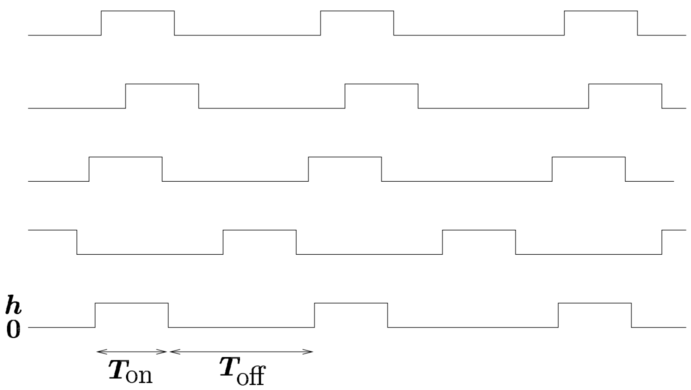

So the proportion of time that a source is on is

$$p = \frac{T_{on}}{T_{on} + T_{off}}$$

The mean rate of a source is $m = ph$.

Of course we require

$$m = ph < c < h \quad \text{and} \quad \rho = m/c < 1.$$

A periodic source — fluid model

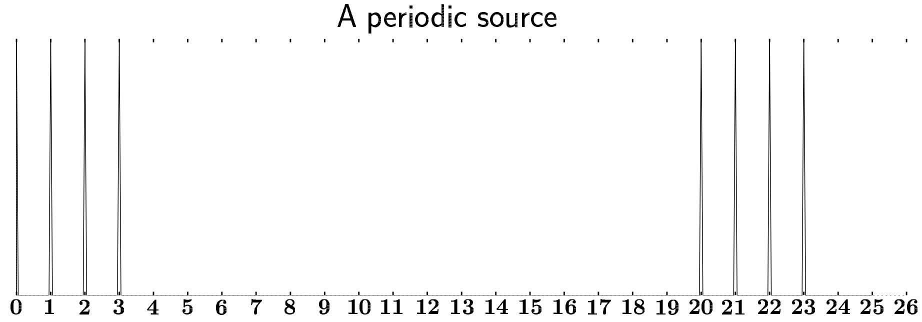

The fluid model of the source treats workload as a continuous fluid.

During a burst, fluid arrives at a rate of 1.

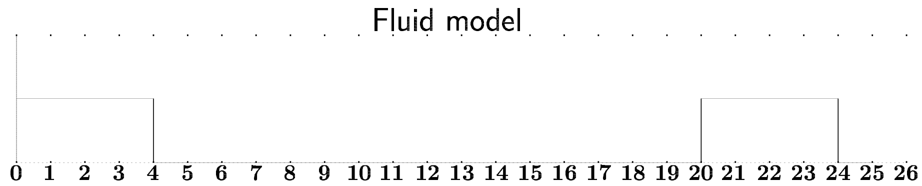

Maximal cumulative inflow

Here we show the maximal cumulative inflow for the cell scale source, the fluid source, the mean inflow $mt = 0.20t$ and multiplexer outflow $ct = 0.22t$.

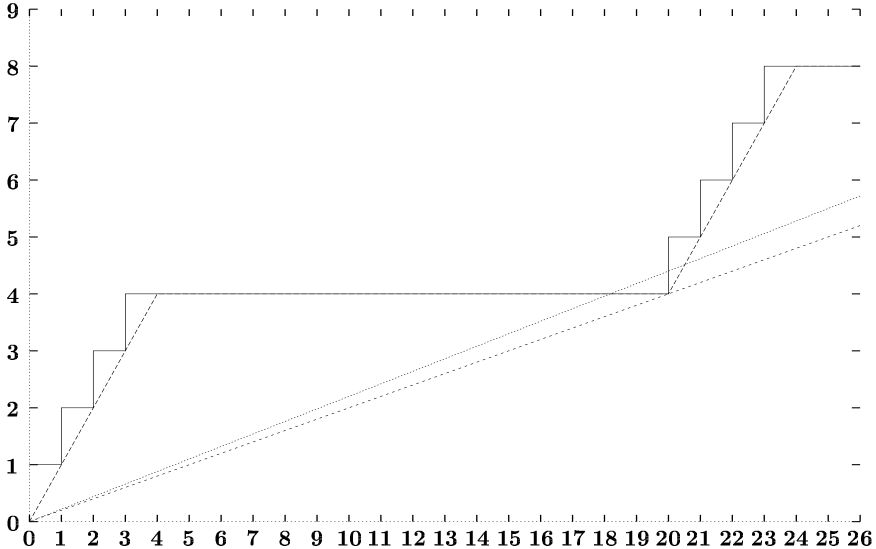

The largest buffer that can be filled by N fluid sources has size

$$B = 4N - 4cN = 3.12N$$

Note that this is smaller than $3.34N$ (the size of buffer that could be filled by N sources modelled at the cell scale.

A fluid source model estimate of CLP in the bufferless case

If there is no buffer, then overflow occurs as soon as the number of sources in the on phase exceeds c/h. Hence

$$\mathbb{P}\left(Q^N > 0\right) \sim e^{-NI_f(c/h)}$$

where

$$\begin{aligned} I_f(a) &= \sup_{s>0} \{sa - \log(q + pe^s)\} \\ &= a \log\left(\frac{a}{p}\right) + (1-a)\log\left(\frac{1-a}{1-p}\right) \end{aligned}$$

Note that

$$\mathbb{P}\left(Q^N > 0\right) \to 0 \text{ as } N \to \infty.$$

A cell scale estimate of the CLP in bufferless case

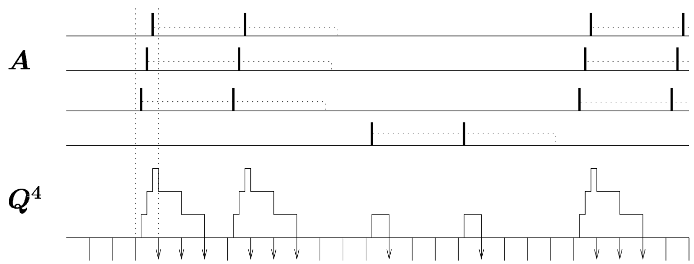

The CLP is the probability that more than one cell arrives in an interval of length $1/C$.

The number of cells produced by one source in an interval of length $1/C$ is distributed $B(1, m/C)$.

The number of cells produced by N sources in an interval of length $1/C$ is distributed $B(N, m/Nc)$. Thus $\rho = m/c$,

$$\mathbb{P}(\text{more than one cell produced in an interval of length } 1/C)$$
$$= \mathbb{P}\left(Q^N > 0\right) \to 1 - (1+\rho)e^{-\rho} \text{ as } N \to \infty.$$

Cell scale effects

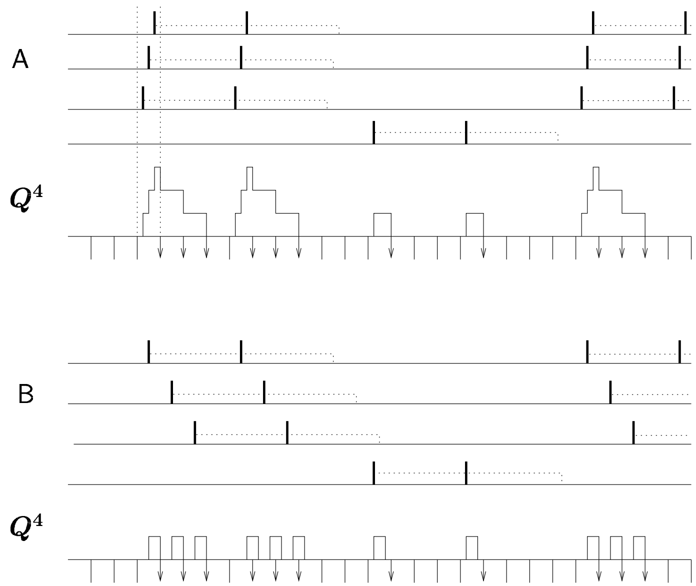

In case A a buffer of $B = 3$ cells will be needed to avoid cell loss. This is despite the fact that $C > Nh$. The overflow occurs because of an unfortuitous phasing of the sources.

In case B the phasing is fortuitous. Less buffering is required.

How cell scale and fluid models differ for small buffers

Recall $m = 0.20$, $c = 0.22$.

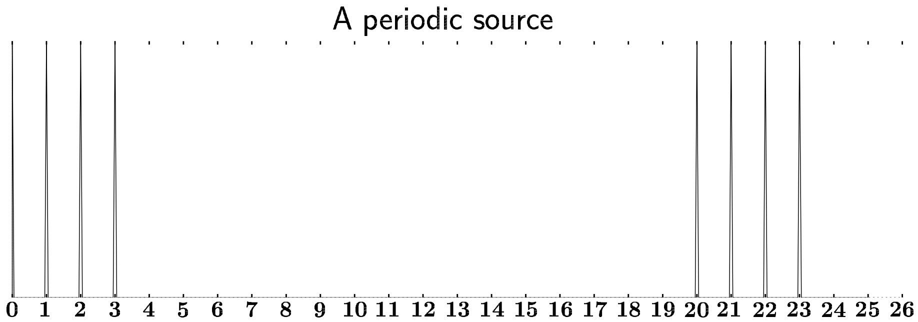

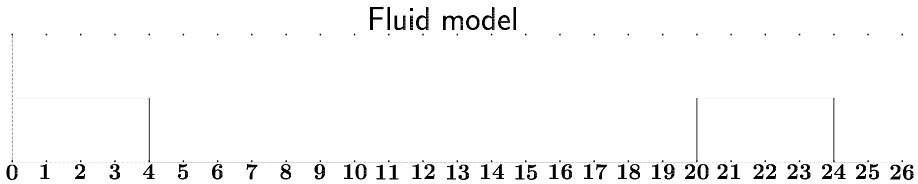

I against b

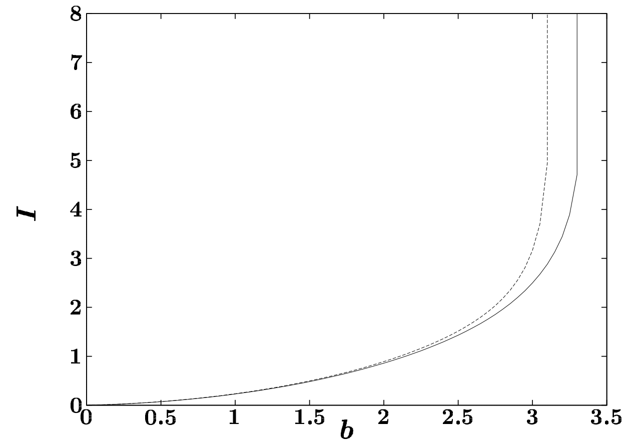

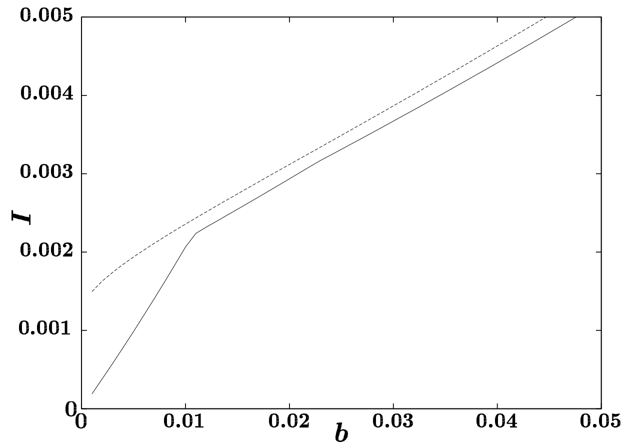

$\sup_s\{s(b+ct)-\varphi(s,t)\}$ against t

Higher curves correspond to higher values of b,

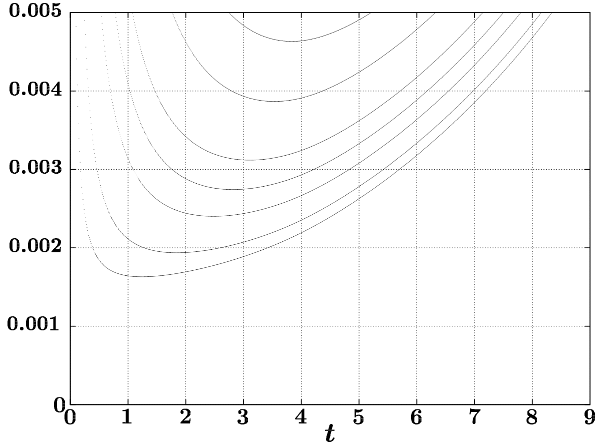

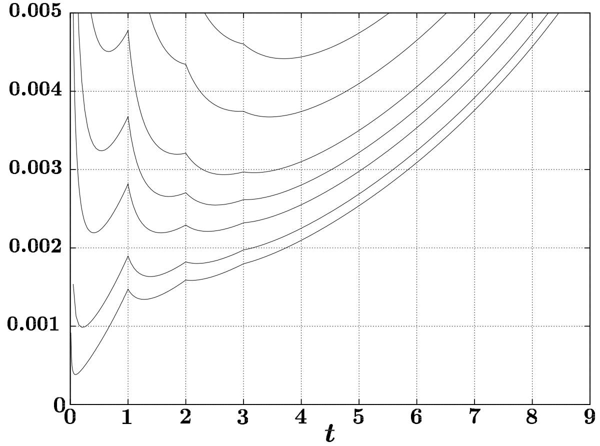

$b = 0.002, 0.005, 0.010565, 0.15, 0.2, 0.3, 0.4$.

Optimizing t^* does not increase with b

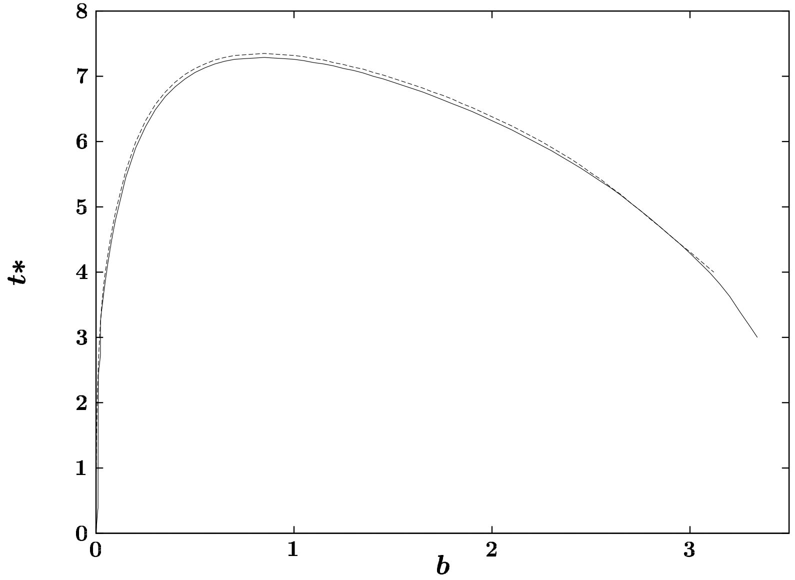

Optimizing t^* is not continuous in b

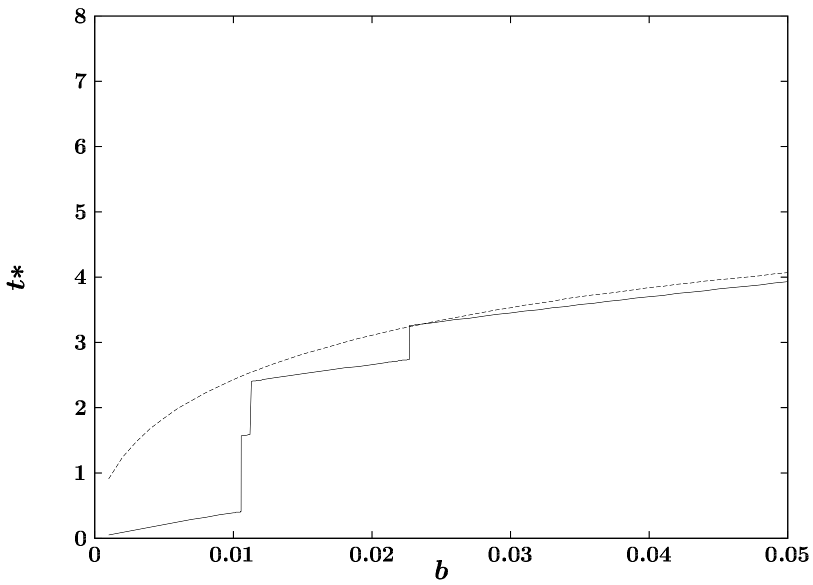

Small buffers

Recall

$$I = \inf_{t>0} \sup_{s>0} \{s(b+ct) - \varphi(s,t)\}\,.$$

Here t is the typical time over which the buffer content increases from 0 to B just prior to cell loss.

If t is large compared to $1/h$, but small compared to T_{on} and T_{off}, then we might expect a good approximation to be obtained from the bufferless fluid model:

$$I \approx I_f(c/h).$$

But it may be a poor approximation. Example: periodic sources, $N = 1990$, $C = 622$ Mbps, $h = 1$ Mbps, $m/h = .25$, $\rho = .8$, $B = 30$. Then

$$e^{-NI_f(c/h)} \sim 10^{-8.6},$$

but

$$e^{-NI} \sim 10^{-5.96}.$$

Constant bit rate sources

Suppose we have N sources, each producing a single cell every $1/h$ seconds. Let $\rho = h/c < 1$.

Theorem.

$$I = I_p(\rho, b)$$

where

$$\begin{aligned} I_p(\rho, b) &= \inf_{0<t<1/c} \sup_{s>0} \{s(b+ct) - \log[(1-ht) + hte^s]\} \\ &= \inf_{0<t<1} \sup_{s>0} \{s(b+t) - \log[(1-\rho t) + \rho te^s]\} \\ &= \inf_{0<t<1} \left\{ (b+t)\log\left[\frac{b+t}{\rho t}\right] \right. \\ &\qquad \left. +(1-b-t)\log\left[\frac{1-b-t}{1-\rho t}\right] \right\} \end{aligned}$$

The optimizing t is small and the overflow is completely due to unfortuitous relative phasing of the sources.

Two regimes

Suppose b is small, so that during the time that the buffer fills most sources remain fully on or fully off. Let us make the approximation that each source is fully on or fully off with probabilities p and q respectively. Then $I \simeq I'$ where

$$I' = \inf_{t>0} \sup_{s>0} \left[s(b+ct) - \log\left(q + p\mathbb{E}e^{sX[0,t]}\right) \right] \quad (1)$$

and $X[0,t]$ denotes the number of cells produced by a source that remains on continuously.

Theorem. *There exists $b^* \leq c/h < 1$, such that*

(a) for $b > b^*$,

$$I' = I_f(c/h)$$

and (1) is extremised by $t \to \infty$;

(b) for $b \leq b^*$,

$$I' = \inf_{0<a\leq c/h} [aI_p(ah/c, b/a) + I_f(a)]$$

and (1) is extremised by a t less than $1/h$.

Interpretation, $b > b^*$

$$I' = I_f(c/h)$$

The typical t over which the buffer fills is very large.

In this regime there is no cell scale effect.

This happens trivially for $b > 1$.

In fact, the cell scale effect vanishes for rather small $b < 1$.

E.g., $c/h = 0.3$, $p = 0.25$, $\rho = 0.8333$,
$b^* = 0.01667$.

Interpretation, $b \le b^*$

$$I' = \inf_{0<a\le c/h} [aI_p(ah/c, b/a) + I_f(a)]$$

The typical t over which the buffer fills is less than $1/h$.

In this regime there is a cell scale effect and an effect due to the number of sources deviating above the mean number pN.

The probability that aN sources are on is about $e^{-NI_f(a)}$.

These can be viewed as aN constant bit rate sources, with buffer per source of b/a and bandwidth per source of c/a. Hence the effective ρ is $\rho = h/(c/a)$ and

$$\mathbb{P}\left(Q^N > Nb\right) \sim e^{-aNI_p(ah/c,b/a)} \times e^{-NI_f(a)},$$

where a is chosen to maximize this probability.

Small buffers, changing scale effects

As b increases the contribution due to an unfortuitous number of sources being on increases. ($p = 0.25$, $c/h = 0.5$.) Note $a^* > p$.

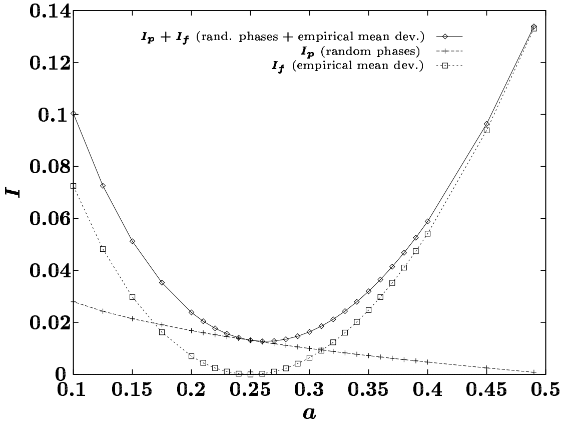

$b = .01$ cells

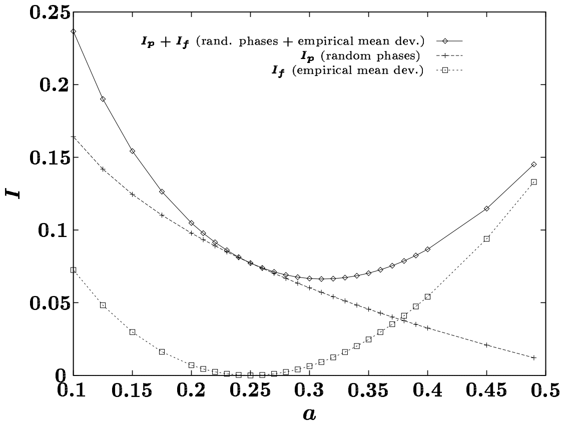

$b = .05$ cells

Small buffers, changing scale effects

As b increases the contribution due to an unfortuitous number of sources being on increases. ($p = 0.25$, $c/h = 0.5$.) Note $a^* > p$.

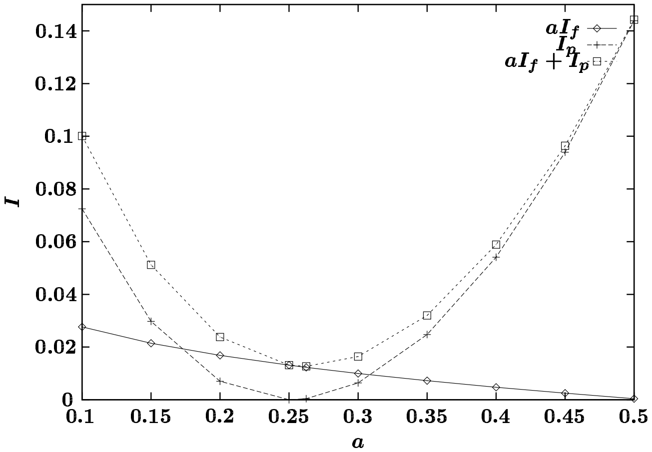

$b = .01$ cells

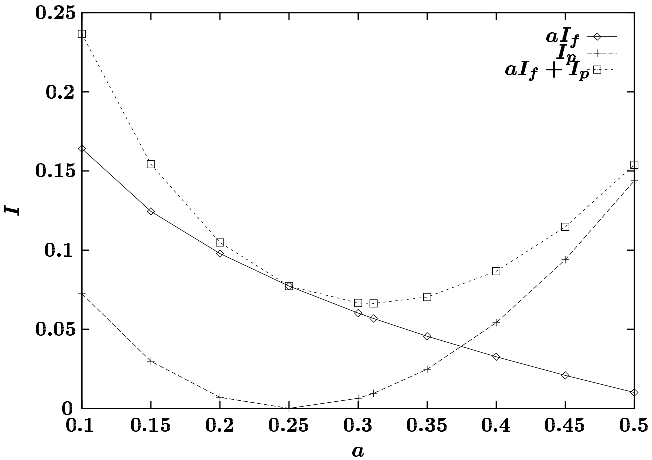

$b = .05$ cells

Boundary between regimes

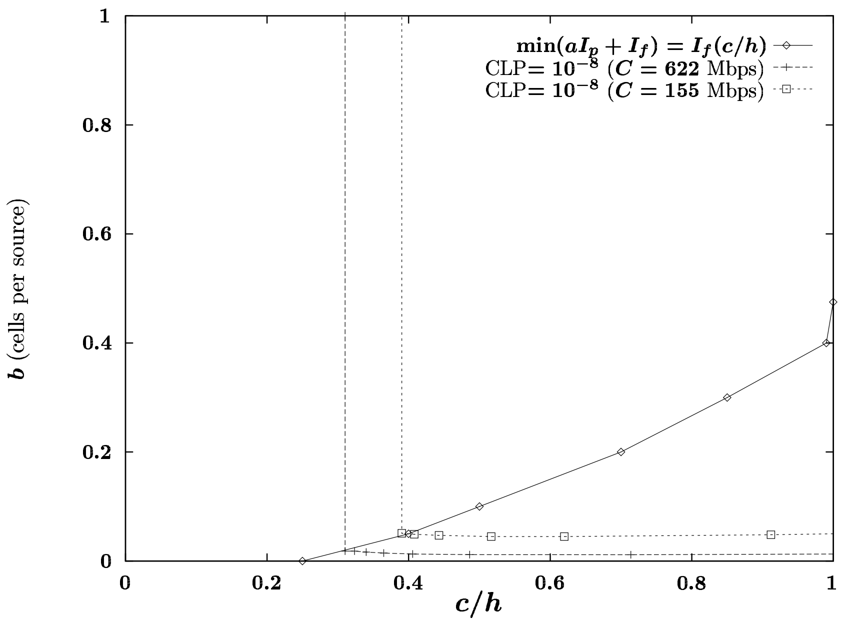

Here $p = 0.25$. The solid line shows b^*.

The dotted lines show the b necessary to achieve a CLP of 10^{-8} for various values of c/h.

Boundary between regimes

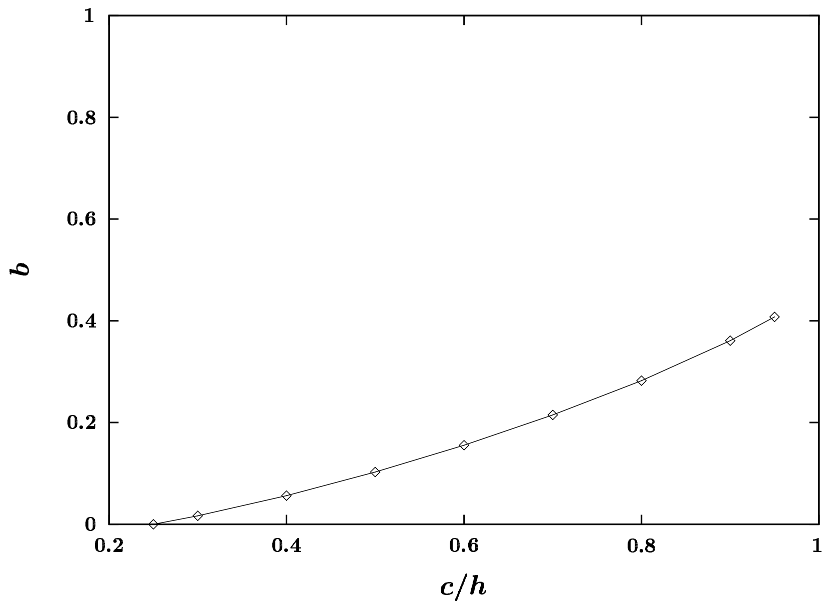

Here $p = 0.25$. The solid line shows b^*.

A CLP heuristic for small buffers

For CLP of practical interest b^* turns out to be very small.
E.g., $h = 1$ Mbps, $m/h = 0.25$, CLP$=10^{-8}$,
$b^* = 0.018$ cells for $C = 622$ Mbps, and
$b^* = 0.050$ cells for $C = 155$ Mbps.

For b small and ρ near 1,
[Fiche, Lorcher, Veyland and Oger, 1994] give

$$I_p(\rho, b) \approx 2b^2 + b(1 - \rho - \log \rho)\,.$$

When b^* is small, $a^* \approx p = m/h$. So $I_f(a^*) \approx 0$, and we have

$$\begin{aligned} I &\approx pI_p(ph/c, b/p) \\ &\approx 2b^2/p + b(1 - ph/c - \log(ph/c)) \\ &\approx b(1 - \rho - \log \rho) \end{aligned}$$

The 'heuristic'

$$\mathbb{P}\left(Q^N > Nb\right) = e^{-Nb(1-\rho-\log \rho)}$$

does a good job of estimating CLP for very small buffers.

Comparison of 'heuristic' and 'exact I'

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

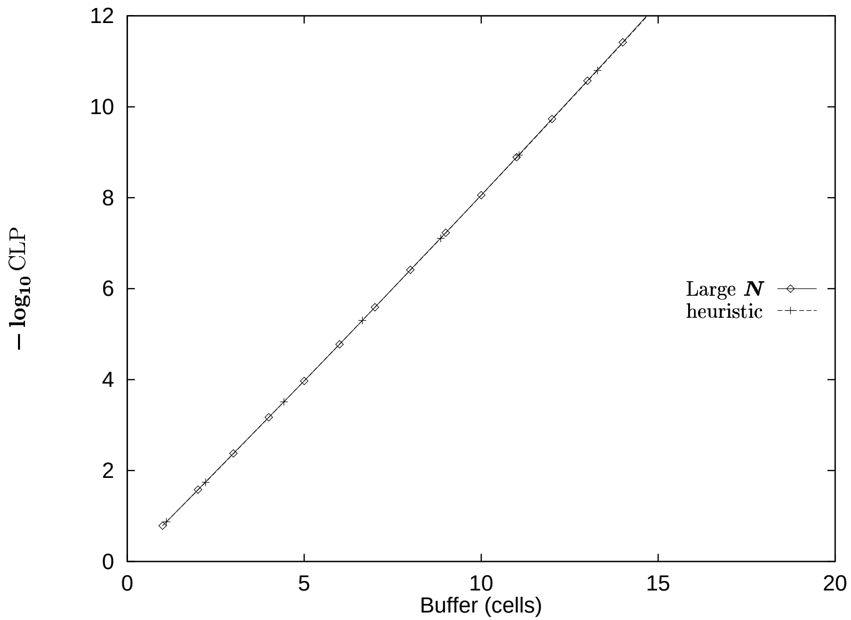

$c/h = 0.7$ ($\rho = 0.36$)

Comparison of 'heuristic' and 'exact I'

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

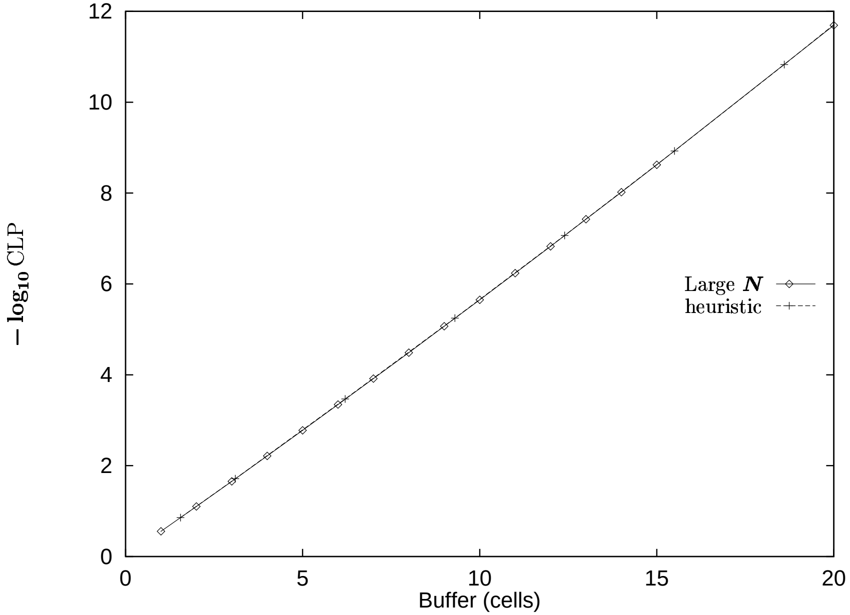

$c/h = 0.5$ ($\rho = 0.5$)

Comparison of 'heuristic' and 'exact I'

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

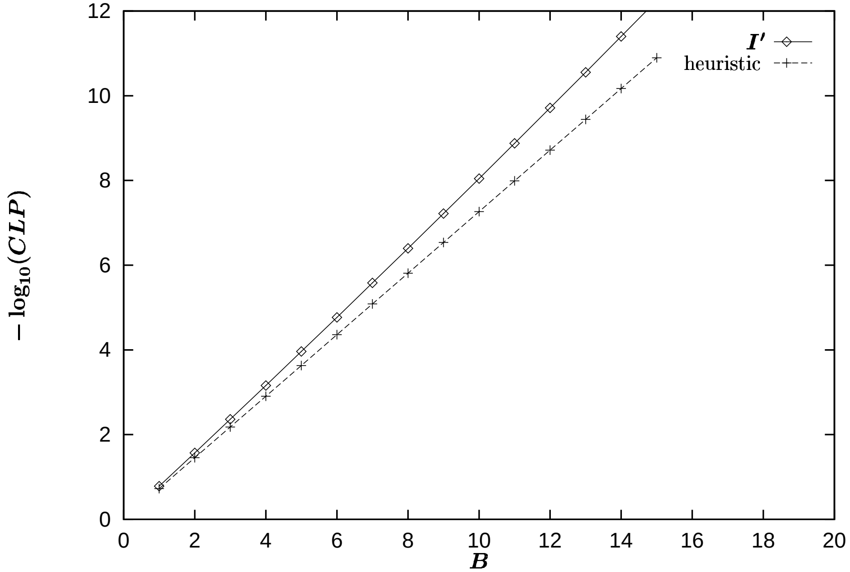

$c/h = 0.7$ ($\rho = 0.36$)

Comparison of 'heuristic' and 'exact I'

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

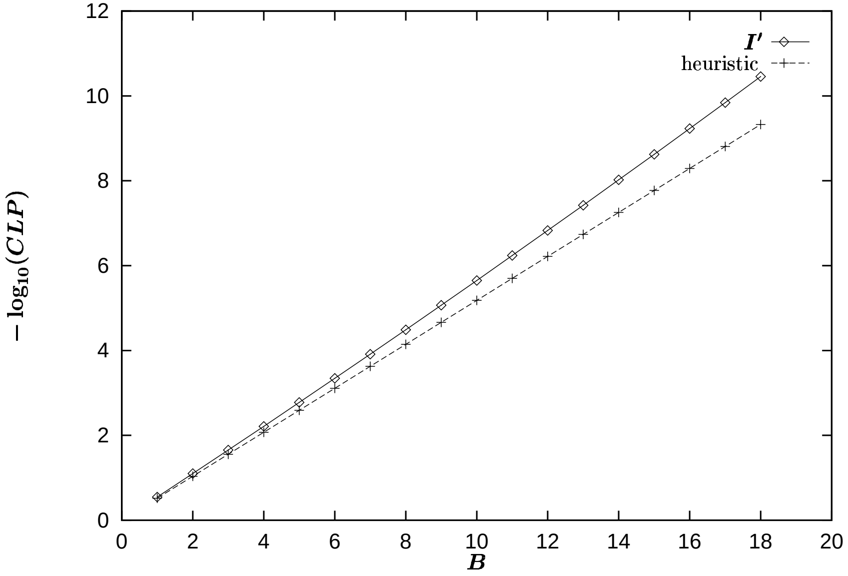

$c/h = 0.5$ ($\rho = 0.5$)

Comparison of 'fluid model/bufferless' and 'exact I'

	$\rho = 0.4, N = 248$		$\rho = 0.6, N = 372$		$\rho = 0.8, N = 496$	
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

The bufferless, fluid model estimate is good if the buffer is small, but not too small, — say about 100 cells —. This is typical of switches in practice.