EURO XIV (14th European Conference on Operations Research, Jerusalem, July 3–6, 1995)

ESTIMATION OF OVERFLOW PROBABILITIES FOR STATE-DEPENDENT SERVICE OF TRAFFIC STREAMS WITH DEDICATED BUFFERS

Costas Courcoubetis, ICS Forth and University of Crete

Richard Weber, University of Cambridge

Richard Weber, 5 July 1995

Abstract

Large deviation asymptotics can be used to estimate the rate with which cell loss occurs because of overflow at the buffer of an ATM switch. These asymptotics may be appropriate when either the buffer is large or when a large number of traffic sources are multiplexed through the switch. In some cases these estimates lead to a natural definition of effective bandwidths for the sources.

As a step towards generalising and applying these ideas to networks in which different qualities of service are to be guaranteed for different sources, we consider M traffic sources, each with its own buffer of size B, which are served by a a single deterministic server of bandwidth c. The server implements a state-dependent service discipline (for example, to share its effort fairly amongst the buffers).

We show the frequency of buffer overflow, Φ, has an asymptotic of $\log\Phi = -I^*B + g(B)$, where $\lim_{B\to\infty} g(B)/B = 0$, and where I^* can be computed as the solution to an optimal control problem posed in terms of rate functions I_i, $i = 1, \ldots, M$ for the M traffic sources.

ATM/BISDN networks

- Data is transmitted in cells (53 bytes).
- Traffic sources are heterogeneous:

  voice, video, file transfer, email, etc.
- Traffic sources have different quality of service requirements.
- Traffic sources are bursty.
- Cells from a single call follow a 'virtual circuit' (VC).

Here we have four VCs and three 2×2 switches.

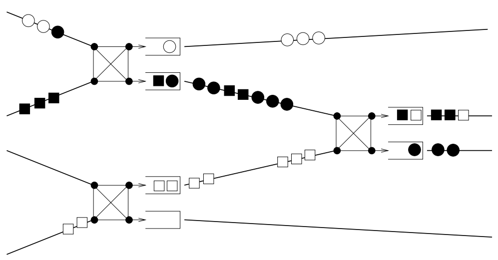

Important issues for ATM

- Quality of Service

  Cell loss (due to buffer overflow).

  This should be very small.

  Cell Delay.
- Call acceptance control (CAC)
- Call routing
- How best to use resources:

  Buffers, bandwidth, alternative routes,

  Statistical multiplexing,

  Signalling,

  Flow control.
- Charging and accounting

A discrete time model

- Discrete time, with epochs $k = 1, 2, \dots$.
- M independent sources.
- Source i produces U_{it} cells in epoch t.
- $U_{i1}, U_{i2}, \dots$ can be seen as a dependent sequence of random variables.
- Sources share a single buffer of size B.
- Buffer is served at the rate of c cells per epoch.
- $X_{t+1} = \max\left\{X_t + \sum_{i=1}^{M} U_{i,t+1} - c, 0\right\}$.

A 2×2 switch with 4 virtual circuits

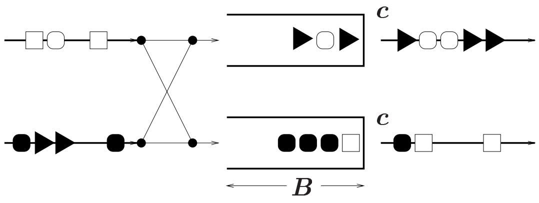

The problem of estimating buffer overflow frequency

We concentrate on a single buffer and the overflow frequency of this buffer.

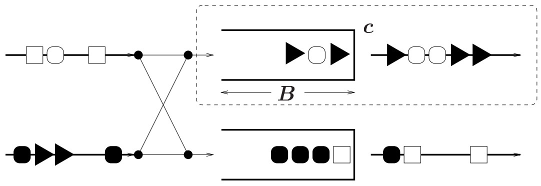

In order to know how many virtual circuits may be allowed to use this output link, for a given Quality of Service constraint, we need to estimate the probability of buffer overflow.

$P(X_t \geq B)$ should be small.

The overflow probability in a $M/M/1/B$ queue

For a single server $M/M/1/B$ queue, for example (with finite buffer), being shared here by two VCs,

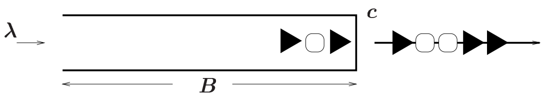

we know

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

Hence

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

This is typical.

Chernoff's theorem

Theorem 1 *Suppose* $U_1, U_2, \dots$ *is a sequence of i.i.d. random variables. Define the* logarithmic moment generating function

$$\varphi(\theta) = \log E[\exp(\theta U_1)]$$

and the rate function

$$I(u) = \sup_{\theta}[\theta u - \varphi(\theta)].$$

Then

$$\lim_{m\to\infty} \frac{1}{m} \log P\left(\frac{1}{m}\sum_{t=1}^{m} U_t \in [a,b]\right) = -\inf_{u\in[a,b]} I(u).$$

Informally, we write

$$P\left(\frac{1}{m}\sum_{t=1}^{m} U_t \in [a,b]\right) \sim e^{-m\inf_{u\in[a,b]} I(u)}$$

Note. $I(m) = 0$, where $m = EU_i$.

A typical rate function

Suppose $U_i = 0, 1$ with probabilities q, p. Then

$$\varphi(\theta) = \log(q + pe^{\theta}),$$

and

$$I(u) = \begin{cases} u\log\left(\frac{u}{p}\right) + (1-u)\log\left(\frac{1-u}{1-p}\right), & 0 \le u \le 1 \\ \infty, & \text{otherwise.} \end{cases}$$

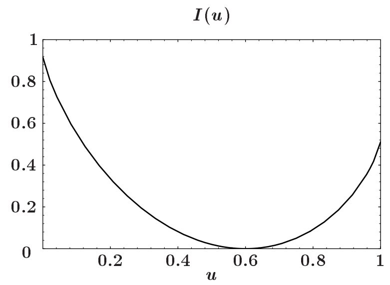

Here $p = 0.6$.

- $I(u)$ is convex.
- $I'(u) \to \infty$ as $u \to$ boundary of the set where $I(u)$ is finite.
- $I(m) = 0$, where $m = EU$.

Elements of large deviation theory

We have seen:

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

$$P\left(\frac{1}{m}\sum_{t=1}^{m} U_t \in [a,b]\right) \sim e^{-m\inf_{u\in[a,b]} I(u)}$$

This is typical. The general conclusions are:

- The frequency of occurence of rare events depends in an exponential manner on some parameters of the problem. E.g., B, m.
- If a rare events occurs then it occurs in the most likely way. E.g., $\inf_{u\in[a,b]}$.
- Rare events occur as a Poisson process.

Gärtner-Ellis's theorem

The Gärtner-Ellis theorem is similar to Chernoff's theorem but applies to a sequence of vector-valued, dependent random variables.

Theorem 2 *Suppose* $U_1, U_2, \ldots \in \mathbb{R}^M$ *is a sequence of random vectors and that the* asymptotic logarithmic moment generating function

$$\varphi(\theta) = \lim_{m\to\infty} \frac{1}{m} \log E\left[\exp\left\langle \theta, \sum_{t=1}^{m} U_t \right\rangle\right]$$

exists for all $\theta \in \mathbb{R}^M$*. Define the* rate function

$$I(u) = \sup_{\theta}[\theta u - \varphi(\theta)].$$

Then for any set $G \subset \mathbb{R}^M$

$$\underline{\lim}_{m\to\infty} \frac{1}{m} \log P\left(\frac{1}{m}\sum_{t=1}^{m} U_t \in G\right) \ge -\inf_{u\in G^o} I(u),$$

$$\overline{\lim}_{m\to\infty} \frac{1}{m} \log P\left(\frac{1}{m}\sum_{t=1}^{m} U_t \in G\right) \le -\inf_{u\in \bar{G}} I(u),$$

where G^o *and* $\bar{G}$ *are respectively the interior and closure of* G*.*

Note that in many cases the two infimums are equal and so the limit exists.

The logarithmic moment generating function

- For source i we define

$$\varphi_{im}(\theta) = \frac{1}{m} \log E \exp\left(\theta \sum_{k=1}^{m} U_{ik}\right).$$

- Suppose the asymptotic logarithmic moment generating function exists for all θ,

$$\varphi_i(\theta) = \lim_{m\to\infty} \varphi_{im}(\theta).$$

- Suppose the conditions of the Gärtner-Ellis theorem are satisfied. Then with

$$I_i(u) = \sup_{\theta}[\theta u - \varphi_i(\theta)]$$

we have,

$$\lim_{m\to\infty} \frac{1}{m} \log P\left(\frac{1}{m}\sum_{t=1}^{m} U_{it} \in [a,b]\right) = -\inf_{u\in[a,b]} I_i(u).$$

The large buffer asymptotic

Under the above assumptions, Kesidis, Walrand and Chang (1993) show that

$$-H(c) \equiv \lim_{B\to\infty} \frac{1}{B} \log P(X \ge B)$$
$$= -\frac{1}{B} \inf_{\alpha} \frac{B}{\alpha} I(c+\alpha)$$

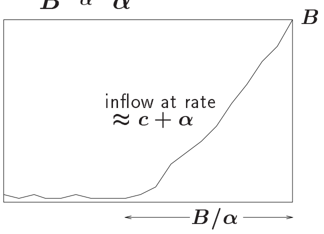

$$= -\inf_{\alpha} \frac{1}{\alpha} \sup_{\theta} \left\{ (c+\alpha)\theta - \sum_{i=1}^{M} \varphi_i(\theta) \right\}$$
$$= -\sup \left\{ \theta : \sum_{i=1}^{M} \varphi_i(\theta)/\theta \le c \right\}.$$

So $P(X \ge B) \sim e^{-H(c)B} \le e^{-\delta B}$, provided

$$\sum_{i=1}^{M} \varphi_i(\delta)/\delta \le c.$$

This motivates identifying $\varphi_i(\delta)/\delta$ as the effective bandwidth for source i.

The many sources asymptotic

- Suppose there are M identical sources.
- The buffer and bandwidth scale with M, so that the sources share a common output buffer of size $B = Mb$ and a bandwidth Mc.
- Courcoubetis and Weber (1995) have shown that,

$$-J(c,b) \equiv \lim_{M\to\infty} \frac{1}{M} \log P(X \ge Mb)$$
$$= -\inf_{m} \sup_{\theta} \left\{ (nc+b)\theta - m \sum_{i=1}^{M} \varphi_{im}(\theta) \right\}.$$

- This formula expresses the effect of statistical multiplexing over sources, whereas the large buffer asymptotic evaluates statistical multiplexing that occurs over time.
- It does not lead to a simple notion of effective bandwidths as with the large buffer asymptotic.
- However, $P(X \ge B) \approx e^{-J(c,b)M}$ usually provides a more accurate estimate of overflow probabilities than does $P(X \ge B) \approx e^{-H(Mc)B}$.

Comparison of the asymptotics

The following data is based on calculations of a Markov modulated fluid model of voice sources in a channel that is 66% utilized.

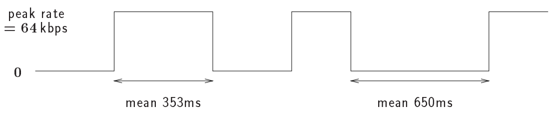

peak bandwidth per source = 64 kbps
mean bandwidth per source = 22.48 kbps
bandwidth per source = $c = 33.72$ kbps
buffer per source = b ranges from 0 to 100 bits

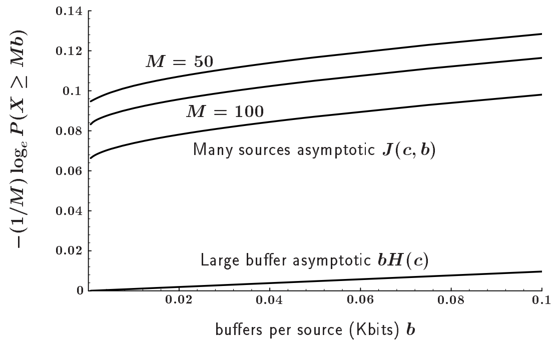

A switch with dedicated output buffers

A single output buffer per output link:

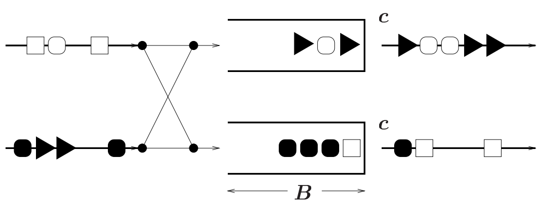

Dedicated output buffers per virtual circuit:

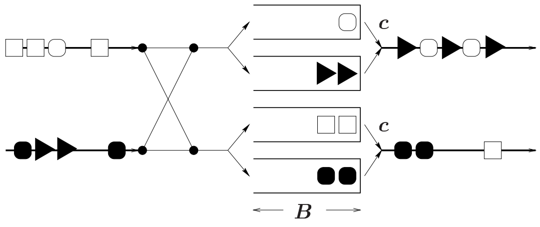

State dependent service disciplines

Now consider M VCs with separate buffers. The state is

$$X_t = (X_{1t}, \dots, X_{Mt}) \in \mathbb{R}^M.$$

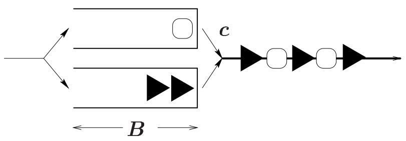

Some possible service disciplines are:

1. **Serve longest queue.** The entire bandwidth c is allocated to serving the longest queue, or shared equally amongst several equal longest queues.
2. **Weighted round robin.** A nonempty queue i is ensured a bandwidth of at least $\phi_i c$, where $\sum_i \phi_i = 1$. Any surplus bandwith that is available because some queues are empty is allocated to the nonempty queues, proportionally to their weights.
3. **Queue length weighted service.** Queue i is served at rate
$$\frac{\phi_i X_{it}}{\sum_j \phi_j X_{jt}} c.$$
4. **Threshold based disciplines.** Divide the bandwidth equally amongst buffers, but if one is more than 90% full, then it gets all the bandwidth.

Assumptions

We let the scaled buffer process be

$$X_t^B = \frac{1}{B} X_t$$

and take a state-dependent service allocation that *depends on the scaled queue lengths* X_t^B,

$$X_{t+1} = X_t + U_{t+1} - c(X_t^B).$$

Throughout the following the norm applied to $x \in \mathbb{R}^M$ is $|x| = \max_i |x_i|$. We make the following assumptions.

Assumptions

A1. $\left|U_{t+1} - c\left(X_t^B\right)\right| < K$.

A2 $c(\cdot)$ is Lipschitz continuous with constant L. That is, $|c(x) - c(x')| \le L|x - x'|$.

Assumptions

A3. Let $\cdots$ denote any past history of the process. Then $\forall \Delta > 0$,

$$\varphi(\theta) = \lim_{B\to\infty} \frac{1}{\Delta B} \log E_x \left[\exp \left\langle \theta, \sum_{t=\lambda B}^{\lambda B+\Delta} U_t \right\rangle \middle| \cdots \right]$$

exists and satisfies Gärtner-Ellis conditions, so that with the definition

$$I(u) = \sup_\theta [< \theta, u > - \varphi(\theta)]$$

and $G \subset \mathbb{R}^M$,

$$\overline{\lim_{B\to\infty}} \frac{1}{\Delta B} \log P \left(\frac{1}{\Delta B} \sum_{t=\lambda B}^{\lambda B+\Delta} U_t \in G \middle| \cdots \right) \le - \inf_{u \in \bar{G}} I(u),$$

$$\underline{\lim_{B\to\infty}} \frac{1}{\Delta B} \log P \left(\frac{1}{\Delta B} \sum_{t=\lambda B}^{\lambda B+\Delta} U_t \in G \middle| \cdots \right) \ge - \inf_{u \in G^o} I(u),$$

where G^o and $\bar{G}$ are respectively the interior and closure of G.

The large buffer scaling

First scale with respect to buffer size:

$$X_t^B = \frac{1}{B} X_t, \quad t = 0, \dots, \tau B$$

Then scale with respect to time, so that for $0 \le t \le \tau$,

$$x^B(t) = (1 - tB + [tB]) X_{[tB]}^B + (tB - [tB]) X_{[tB]+1}^B.$$

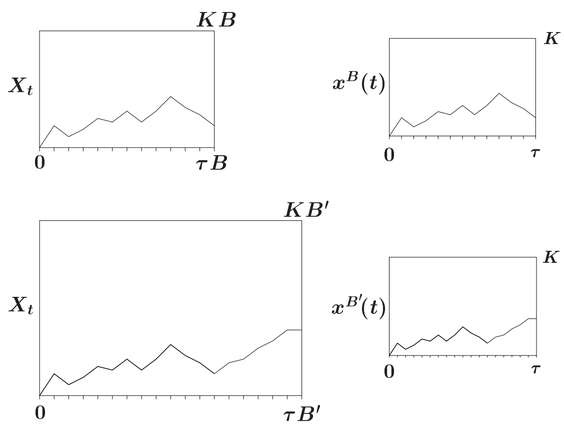

Note that $x^B(\cdot)$ is piecewise linear.

Upper and Lower Bounds

Theorem 3 *Suppose $A \subset C_{1,x}[0,\tau]$ and A1–A3 hold. Then*

$$\overline{\lim_{B\to\infty}} \frac{1}{B} \log P\left(x^B(\cdot) \in A\right) \leq - \inf_{\substack{x(\cdot),u(\cdot)\\ \dot{x}=u-c(x)\\ x(\cdot)\in\bar{A}}} \int_{t=0}^{\tau} I(u(t))\, dt,$$

$$\underline{\lim_{B\to\infty}} \frac{1}{B} \log P\left(x^B(\cdot) \in A\right) \geq - \inf_{\substack{x(\cdot),u(\cdot)\\ \dot{x}=u-c(x)\\ x(\cdot)\in A^o}} \int_{t=0}^{\tau} I(u(t))\, dt.$$

where $\bar{A}$ and A^o are the closure and interior of A respectively.

Note. The right hand side of the lower bound is $-\infty$ if A^o is empty.

Proof of the upper bound

Suppose $B\tau = nm$. Divide time into n blocks of m epochs. Divide the vertical axis into N increments of width $K\tau B/N$. Any path $x^B(\cdot) \in A$ must pass through certain 'gates' $C_{j_0}, \dots, C_{j_n}$.

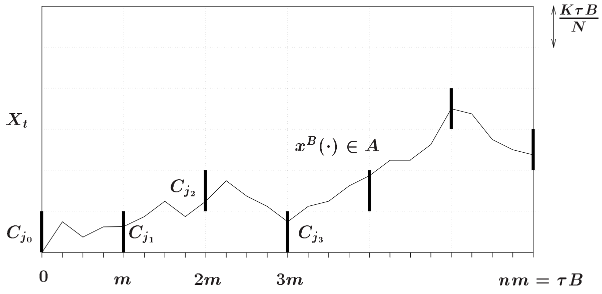

Hence writing $\Phi^B(A) = P(x^B(\cdot) \in A)$,

$$\Phi^B(A) \leq \sum_{j_0,\dots,j_n} \prod_{i=1}^{n} P\left(X_{im} \in C_{j_i} \middle| X_0, U_1, \dots, U_{(i-1)m}\right)$$

where the sum is over possible $j_0, \dots, j_n$.

For n, N large, the path cannot only move between $C_{j_{i-1}}$ and C_{j_i} if $\frac{1}{m}\sum_{t=(i-1)m+1}^{im} U_t$ lies in some small set, say

$$\frac{1}{m} \sum_{t=(i-1)m+1}^{im} U_t \in S_{j_{i-1}}^{j_i}, \quad i = 1, \dots, n.$$

Conclusion of the proof of the upper bound

So we have

$$\Phi^B(A) \leq \sum_{j_0,\dots,j_n} \prod_{i=1}^{n} P\left(\frac{1}{m}\sum_{t=(i-1)m+1}^{im} U_t \in S_{j_{i-1}}^{j_i} \middle| \cdots\right)$$

Applying Assumption A3, as $B \to \infty$ via $m \to \infty$

$$\overline{\lim_{B\to\infty}} \frac{1}{B} \log P\left(\frac{1}{m}\sum_{t=(i-1)m+1}^{im} U_t \in S_{j_{i-1}}^{j_i} \middle| \cdots\right) \leq - \inf_{u_i \in S_{j_{i-1}}^{j_i}} \frac{\tau}{n} I(u).$$

Hence

$$\overline{\lim_{B\to\infty}} \frac{1}{B} \log \Phi^B(A) \leq - \inf_{\substack{j_0,\dots,j_n\\ u_1,\dots,u_n\\ u_i\in S_{j_{i-1}}^{j_i}}} \frac{\tau}{n} \sum_{i=1}^{n} I(u) \approx - \inf_{\substack{x(\cdot),u(\cdot)\\ \dot{x}=u-c(x)\\ x(\cdot)\in\bar{A}}} \int_0^{\tau} I(u)\, dt.$$

Proof of the lower bound

Again $B\tau = nm$ and divide time into n blocks of m epochs. Let $\tau_i = i\tau/n$. Pick any $\bar{x} \in A^o$.

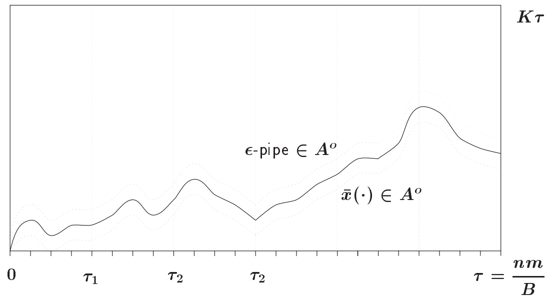

$$\frac{1}{m} \sum_{t=(i-1)m+1}^{im} U_t \approx \bar{u}(\tau_i), \ \forall i$$

$$\Rightarrow x^B(\cdot) \in \epsilon\text{-pipe} \subset A^o.$$

So $\underline{\lim_{B\to\infty}} \frac{1}{B} \log \Phi^B(A) \geq \underline{\lim_{B\to\infty}} \frac{1}{B} \log P(x^B(\cdot) \in \epsilon\text{-pipe})$

$$\geq \underline{\lim_{B\to\infty}} \sum_i \frac{1}{B} \log P\left(\frac{1}{m}\sum_{t=(i-1)m+1}^{im} U_t \text{ near } \bar{u}(\tau_i) \middle| \cdots\right)$$

$$\approx -\int_0^{\tau} I(\bar{u})\, dt - \eta(\epsilon).$$

Hence $\underline{\lim_{B\to\infty}} \frac{1}{B} \log \Phi^B(A) \geq - \inf_{\substack{x(\cdot),u(\cdot)\\ \dot{x}=u-c(x)\\ x(\cdot)\in A^o}} \int_0^{\tau} I(u)\, dt.$

The control problem

Theorem 4 *Suppose A1–A3 hold. Then*

$$\lim_{B\to\infty}\frac{1}{B}\log P(X_{1t}>B) = -\inf_{\substack{\tau,x(\cdot),u(\cdot)\\ \dot{x}=u-c(x)\\ x(0)=(0,0),\,x_1(\tau)=1}} \int_{t=0}^{\tau} I(u(t))\,dt.$$

Proof. Define A as the set of paths for which $x(0)=0$ and $x_1(\tau)\geq 1$. Show that as $\tau\to\infty$ the upper and lower bounds are the same, and independent of the starting state $x(0)=0$.

This is an optimal control problem that can be solved using Pontryagin's maximum principle. We choose u to maximize $H(x,u,\lambda)$, where

$$H = \lambda^T(u-c) - I(u),$$

$$\dot{\lambda} = -\frac{\partial}{\partial x_i}\lambda^T c,$$

$$\dot{x} = u - c.$$

Its solution for two buffers and fair shares

Consider the case in which there are just two buffers and the service discipline is 'weighted fair-shares'. That is, buffer i is guaranteed at least $\phi_i c$ when it is not empty. But if a buffer is empty the surplus bandwidth can be allocated to serving the other buffer. Then assuming $m_i = EU_{it} \leq \phi_i c$,

Theorem 5

$$\lim_{B\to\infty}\frac{1}{B}\log P(X_{1t}>B)$$

$$= -\min_{\tau>0}\min_{\substack{u_1+u_2=c+1/\tau,\\ u_2\leq\phi_2 c}} \tau[I_1(u_1)+I_2(u_2)]$$

$$= -\min_{\substack{u_1+u_2>c,\\ u_2\leq\phi_2 c}} \frac{I_1(u_1)+I_2(u_2)}{u_1-(c-u_2)}.$$

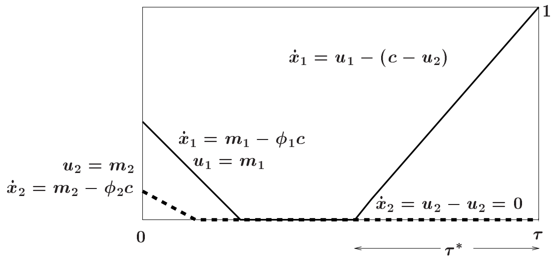

Character of the solution

Suppose we consider the most likely way that buffer 1 can fill over time $[0,\tau B]$, when $X_0=(0,0)$. For a given τ, the optimal choice of u_1, u_2 can take one of two forms:

1. $I_2'(\phi_2 c) \leq I_1'(\phi_1 c+\tau)$: In this case $u_1 = \phi_1 c + \tau$, $u_2 = \phi_2 c$.
2. $I_2'(\phi_2 c) > I_1'(\phi_1 c+\tau)$: In this case $u_1 = c+\tau-u_2$, $u_2 < \phi_2 c$, where these are chosen so that $I_2'(u_2) = I_1'(u_1)$.

   In this case it is as if the two sources shared a common buffer of size B.

Costs and risk-sensitivity

Suppose one wishes to choose a scheduling policy π to minimize over time $[0, B\tau]$ the cost function,

$$\Gamma_B = \inf_{\pi} E\left[\sum_{k=0}^{B\tau}\gamma\left(X_k^B, U_k/B, k/B\right)\right],$$

Following an idea of Whittle (1990), a 'risk-sensitive' form of this would have cost

$$\Gamma_B^\alpha = \inf_{\pi} E\left[\exp\left(\alpha\sum_{k=0}^{B\tau}\gamma\left(X_k^B, U_k/B, k/B\right)\right)\right].$$

The idea is that $\alpha>0$ makes one sensitive not only to the expected cost but also the variance of the cost. Taking $\alpha<0$ corresponds to 'risk-seeking' behaviour. By Varadhan's Lemma the large B asymptotic for Γ_B is computed by multiplying the cost of a path by its probability and finding the combination of greatest cost, so

$$\lim_{B\to\infty}\frac{1}{B}\log\Gamma_B^\alpha = \inf_{\pi}\sup_{\substack{x(\cdot),u(\cdot)\\ \dot{x}=u-c(x)}}\int_{t=0}^{\tau}\alpha\gamma(x,u,t) - I(u)\,dt.$$

where the infimum over π is understood as an infimum over a set of admissible $c(\cdot)$, e.g., a class of threshold policies. This may lead to interesting control problems and scheduling rules.