Bit dropping in ATM: A large deviations analysis

Adam Shwartz
Electrical Engineering
Technion—Israel Institute of Technology
Israel

Alan Weiss
AT&T Bell Laboratories
Murray Hill, NJ

October 1995

First Workshop on ATM Traffic Management, Paris, 6-8 December 1995.

1 Introduction.

The AMS (Anick-Mitra-Sondhi) model is the basic model of high-speed packet communication. The model consists of sources, a buffer, and a communications link with a given capacity. The size of packets is assumed to be small relative to the typical number of packets in a message, so that packets can be modeled as fluid. Traffic sources are modeled as Markov chains, switching between "off" (no transmission) and "on" states. Different "on" states correspond to different service requirements, which are manifest as different rates at which fluid is produced. The sources feed a buffer, which is employed to overcome bursts in service requirements. Service is assumed deterministic with constant rate: that is, fluid is removed from the buffer at a constant rate.

Several control schemes to prevent buffer overflow have been proposed. We show that open loop control (that is, control that does not use information about the buffer size) may be ineffective. This is due to the fact that limiting the number of sources that are simultaneously "on" is effective only if the limit is rather severe, in which case the buffer is not utilized, and performance suffers.

In some audio and video applications, it is possible to drop certain bits from packets, giving poorer quality transmission but not dropping any packets; this is a more graceful degradation of performance. We investigate the effect of various control schemes, using such mechanisms. Thus, the control schemes under investigation never deny access to a source (that is, they do not prevent any source from becoming active). Rather, they lower the quality of service when the requested rate is too high. Specifically, the control schemes we investigate include

- drop bits as long as a "critical level" for the buffer is exceeded,

- start dropping bits above some critical buffer level, and continue until the buffer drops below a lower level,
- drop bits from packets that are already stored in the buffer, once a "critical level" is exceeded.

We use tools from the theory of large deviations, a mathematical theory that deals with probabilities of rare events (such as an overflow in a stable system). Using this theory we can obtain the (leading term of the) probabilities of overflow, as well as the behavior of the sources that lead to overflow. This allows us to quantify the effectiveness of various control schemes, and to "tune" them so that they are indeed effective.

The purpose of this paper is to illustrate the ideas behind this technique and to demonstrate its applicability. For more details on the general theory, on applications to the performance of other models and further details on the model under investigation here the reader is referred to [SW]. A sample of recent references to other work using the increasingly popular tool—large deviations—is also available in [SW].

2 The model

In order to provide a Markovian model of this system, we model each source as a Markov chain with a finite state space. Associated with state i is an activity level a_i: this is the rate at which a source at state i pumps fluid (= traffic) into the buffer.

Naturally, sources of different type (e.g. voice, video, interactive work, etc.) have different states, different rates of transitions and different activity levels. We assume that there is a fixed, finite number of types (although we will eventually let the number of sources increase without bound). Now let $\{1, 2, \ldots, d\}$ be an enumeration of all states of all types. We can then consider any of the Markov chains, of any type, as evolving on the state space $\{1, 2, \ldots, d\}$: depending on the initial state, all states corresponding to other types will never be reached. Thus we arrive at the model where all Markov chains share the same state space.

Since now the sources are indistinguishable except for the initial state, we do not need to know which state each one is in. In order to know the total flow produced by the sources, and the probabilistic evolution of this flow, we only need to know the number of sources in each state at each point in time. So, define

$$x(t) \stackrel{\triangle}{=} (\text{number of sources in state } 1, \ldots, \text{number of sources in state } d). \tag{1}$$

The vector $x(t)$ is again a Markov process, and a little reflection will convince the reader that it contains all the information we need about the behavior and the statistics of the demand process (the total flow into the buffer).

Let us now scale the system, to reflect a situation where the number of users n is very large, and buffer capacity and link capacity are proportionately large. To simplify the notation, let us assume that $d = 2$, so that there is only one "on" state and one "off" state. Define

$$z_n(t) \triangleq \frac{1}{n} \cdot \text{ number of "on" sources} \tag{2}$$

$$b_n(t) \triangleq \frac{1}{n} \cdot \text{ buffer contents.} \tag{3}$$

Then z_n is a Markov process, and it contains all relevant information about the flow into the buffer. The dynamics of the buffer are given by

$$\frac{d}{dt} b_n(t) = \begin{cases} z_n(t) - C & \text{if } z_n(t) > C \text{ or } b_n(t) > 0 \\ 0 & \text{if } z_n(t) \leq C \text{ and } b_n(t) = 0 \end{cases} \tag{4}$$

where C is the capacity of the link, or the rate at which fluid is removed from the buffer. Note that the implied scaling is linear: as the number of sources increases, so does capacity, and capacity per source remains fixed. We also assume that buffer size per source remains fixed at the level B.

We assume that the system is stable by design, so that large buffers rarely occur. More formally, let λ be the transition rate from state "off" to state "on," and let μ be the transition rate from state "on" to state "off." Assume further that the rate of fluid production by a source in the "on" state is 1 (this amounts to a time-scaling). Then stability means that link capacity exceeds mean demand:

$$\frac{\lambda}{\lambda + \mu} < C. \tag{5}$$

so that buffer overflows should be rare. To avoid trivialities, assume $C < 1$.

3 Large Deviations

The theory of large deviations tells us what the probabilities of rare events are, and moreover, how rare events occur. More precisely, given a deterministic path r, we have (under technical conditions which we shall not specify here)

$$\mathbb{P}\left(\mathrm{z_n(t) \text{ near } r(t)}\right) = \exp\left(-nI_0^T(r) + o(n)\right) \tag{6}$$

$$I_0^T(r) \triangleq \int_0^T \ell(r(t), r'(t))\, dt. \tag{7}$$

The functional I_0^T is the rate function. The function ℓ is the local rate function which, in our case, can be written explicitly. (In more general models, we can only write an expression for ℓ as the solution to a maximization problem.) Thus, large deviations gives us the exponent, or leading term for the probabilities.

The events we shall be interested in can all be described by sets of paths. We shall calculate the probability that the buffer starts filling: this is just the probability that the process z_n reaches a level C. The event that, starting with an empty buffer at $t = 0$, the buffer overflows by time T is more complicated, but can also be described by a set of paths:

$$\left\{ r : \int_s^t (r(t) - C)\, dt = B \text{ for some } s < t \leq T, \int_s^u (r(t) - C)\, dt \geq 0 \text{ for all } u \leq t. \right\} \tag{8}$$

Here s is the last time before t that the buffer is empty. The probability of an event is roughly the probability of the most likely path in the definition of the event, so that

$$\mathbb{P}_n\,(\text{event}) \approx \exp\left[-n \inf_{r \in \text{event}} I_0^T(r) \right]. \tag{9}$$

Moreover, if an event occurs, it does near r^*—the minimizer in (9).

So, large deviations tells us that in order to compute probabilities, and to find how events occur, we need to solve a variational problem, and to find the minimum and the minimizer of the functional I.

These are all results over a finite time interval. The Friedlin-Wentzell theory provides information about steady state probabilities as follows. The most likely behavior for the scaled process is for the proportion of active sources to stay near $q \stackrel{\triangle}{=} \frac{\lambda}{\lambda+\mu}$. We can show that (under some technical conditions) the probability, in steady state, that the proportion of active sources exceeds C is approximately $\exp(-nI^*(C))$, where

$$I^*(C) \stackrel{\triangle}{=} \inf\{I_0^T(r) : r(0) = q,\ r(T) = C,\ T > 0\}. \tag{10}$$

This is a "free time" variational problem.

4 Overflow without control

The event of buffer overflow must happen in the following way. First, the number of active sources must increase so that $z_n > C$. Then, the buffer should stay non empty, and accumulate to level B. By the principle of optimality, each of these two parts should be an optimal solution of its respective variational problem. So, we need to solve

$$\inf \left\{ I_0^T(r) : r(0) = q,\ r(T) = C,\ T > 0 \right\} \tag{11}$$

so that the buffer starts filling, as well as the second variational problem:

$$\inf\left\{ I_0^T(r) : r(0) = C,\ \int_0^T (r(t) - C)\, dt = B \right\}. \tag{12}$$

Using soft arguments (no calculations) it is not hard to see that in the second problem there is no need to worry about paths that empty the buffer: they will be sub-optimal.

The optimal path for the first problem is an exponential: this as well as an explicit expression for the optimal cost follows from more general results for birth death processes. The optimal path for the second problem is a hyperbolic cosine: it is concave and symmetric. The calculation is rather involved, especially its justification: see [SW]. We can obtain explicit asymptotics as the buffer increases or decreases. We have

$$\mathbb{P}(\text{buffer overflow}) \approx \mathrm{e}^{-\mathrm{n}\mathrm{I}^*(\mathrm{B})+\mathrm{o}(\mathrm{n})} \tag{13}$$

$$I^*(B) = c_2\sqrt{B} + O(B) \qquad \text{as } B \to 0 \tag{14}$$

$$I^*(B) = c_3 B + c_4 + O(1) \qquad \text{as } B \to \infty. \tag{15}$$

Moreover, the way z_n behaves when overflows occur is to increase to a level above C, but strictly below some value $D < 1$, regardless of the size of the buffer. We conclude that limiting the number of active sources to stay below some D', with $D < D' < 1$ has no effect on the probability of buffer overflow! We need to act on the basis of buffer contents. Note that limiting the number of active sources to stay below some $D'' < D$ will certainly have an effect, but will also reduce the effectiveness of the buffer as a means to average short bursts in traffic.

5 Control schemes

The control schemes we employ drop bits under certain conditions. The question is when to start, and when to stop. When bits are being dropped, the rate of fluid production by any source is reduced by a fixed fraction, say $\rho < 1$. Therefore, we expect to be solving the same variational problems, but with different parameters, over different intervals of time. We can again invoke the principle of optimality to split the variational problem into standard sub-problems.

Dropping bits (least significant) or packets (tagged).

1. Trim above watermark ($\rho < 1,\ B > B_1$)

$$\frac{d}{dt} b_n(t) = \begin{cases} z_n(t) - C & \text{if } z_n(t) > C \text{ and } 0 < b_n(t) < B_1 \\ \rho z_n(t) - C & \text{if } b_n(t) > B_1 \\ 0 & \text{if } z_n(t) \le C \text{ and } b_n(t) = 0. \end{cases} \tag{16}$$

By the principle of optimality, we need to find the cheapest way that makes $z_n(t) = C$, then makes $b_n(t)$ exceed B_1, and find the resulting z_n. Now make the buffer fill to level B with trimmed input rate $\rho z_n(t)$. So again, the first part of the path is exponential, which the second is a hyperbolic cosine, although with different parameters.

2. Trim all packets above watermark

Here, once the watermark is exceeded, we trim incoming as well as buffered bits (packets) until the buffer empties!

Dynamics: in each busy period if the watermark is exceeded, then we trim packets in the buffer by ρ, and

$$\frac{d}{dt} b_n(t) = \begin{cases} z_n(t) - C & \text{if } b_n(t) > 0 \text{ and } b_n(s) < B_1, \\ & \text{all } s < t \text{ in current busy period} \\ \rho z_n(t) - C & \text{if } b_n(t) > 0 \text{ and } b_n(s) > B_1, \\ & \text{some } s < t \text{ in current busy period} \\ 0 & \text{if } z_n(t) \leq C \text{ and } b_n(t) = 0 \end{cases} \tag{17}$$

We need an additional coordinate to make this Markov.

Analysis is in segments, as in the previous case, with an additional complication that the buffer content is no longer continuous due to the trimming of the buffer contents.

3. Two watermarks.

Start trimming if higher mark exceeded. Stop once below lower watermark. The dynamics are

$$\frac{d}{dt} b_n(t) = \begin{cases} z_n(t) - C & \text{if } b_n(t) > 0 \text{ and } b_n(s) < B_2, \\ & \text{all } s < t \text{ in current busy period} \\ z_n(t) - C & \text{if } 0 < b_n(t) < B_1 \\ \rho z_n(t) - C & \text{if } b_n(t) > B_1 \text{ and } b_n(s) > B_2, \\ & \text{some } s < t \text{ in current busy period} \\ 0 & \text{if } z_n(t) \leq C \text{ and } b_n(t) = 0 \end{cases} \tag{18}$$

We need an additional coordinate to make this Markov: trim status.

Analysis is in segments, as in the previous cases, with an additional complication that the buffer process now depends on an additional variable, describing whether trimming is in effect. Again, the segments are all either exponentials or hyperbolic cosines, depending on whether we are solving a crossing problem, or a problem of filling a buffer.

6 Sample calculation

Solving the variational problem for the level crossing problem: the first variational problem (q steady state) to be solved is to find the infimum and the minimizer, among all absolutely continuous paths r, in

$$\inf\left\{I_0^T(r) : r(0) = q,\ r(T) = C,\ T > 0\right\}.$$

This is in the form of "the simplest variational problem." The solution is characterized by the Euler necessary conditions

$$\ell(r, r') - r'\ell(r, r') = 0$$

where the right hand side is zero since this is a free time problem. In this case we obtain after some calculations

$$r'(t) = \mu(r(t)) - \lambda(r(t)),$$

a linear equation (since the rate transitions occur is linear in the number of sources). Moreover,

$$I_c^* = \int_q^C \log\left(\frac{\mu(r)}{\lambda(r)}\right)\, dr$$

from general results for birth-death processes. So, for this we have explicit solutions.

The second variational problem involves filling the buffer:

$$\inf\left\{I_0^T(r) : r(0) = C,\ \int_0^T (r(t) - C)\, dt = B\right\}$$

Introduce the buffer size as a constraint, and use Lagrange Multipliers to reduce this to a standard variational problem. Here our solution required numerical techniques, but we obtain explicit asymptotics.

Acknowledgment. Research of the first author was supported in part by the fund for promotion of research at the Technion.

Reference

[SW] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis: queues, communications and computing*, Chapman and Hall Stochastic Modelling Series, London, 1995.