Performance Analysis of an Optical Switch Using Importance Sampling

Phil Whiting Carl Nuzman Xiao Zhang

Bell Laboratories
Lucent Technologies
Murray Hill, NJ 07974
{pawhiting,cnuzman}@lucent.com

Div. of Applied Math.
Brown University
Providence, RI 02912
jimzhang@cfm.brown.edu

Abstract

Packet loss is examined in an optical switch which uses a shared pool of wavelength converters to reduce contention. A key component of packet loss arises from wavelength converter pool exhaustion. Since loss probablilities of the order of 10^{-10} are desired, conventional simulation methods are impractical. However, a recent large deviations analysis enables efficient and accurate estimates of this loss component via importance sampling. The analysis identifies a family of extremal trajectories, which give the least cost paths (in exponent) to pool exhaustion. Process level changes of measure based on these extremals then drive the importance sampling simulation. These changes of measure can be independent of simulation state, or may be adapted as the simulation progresses. Our results confirm an approximate bound which was previously derived and also agree with refined Bahadur-Rao approximations to the loss probability.

1 Introduction

Wavelength division multiplexed (WDM) transmission on optical fibre has enabled the link capacities of voice and data networks to grow quickly and economically. So far, switching of data streams at network nodes has been performed electronically, with banks of processors scaled and speeded up to keep pace with the new bandwidths. In recent years, technologies have been developed which allow switching to be performed in the optical domain, without electronic processing. An even more radical change being explored is to perform this switching at a packet or burst level, enabling all-optical packet neworks. All-optical switching has the potential for large cost and power savings, if the associated technical challenges can be met [1].

A basic all-optical switch can direct signals from each incoming fibre to each outgoing fibre, but cannot change a signal's modulating wavelength. Wavelength blocking occurs when multiple signals using the same wavelength on different fibres require the same output fibre. This problem can be eliminated by adding a bank of wavelength converters to the output ports of the switch. A potentially less expensive approach is to use a relatively small, shared pool of wavelength converters (see e. g. [2] for connection-based and [3] for packet-based networks).

This paper studies the blocking performance of a switch with shared wavelength converters, under a synchronous packet traffic model. It extends [4, 5], which used large

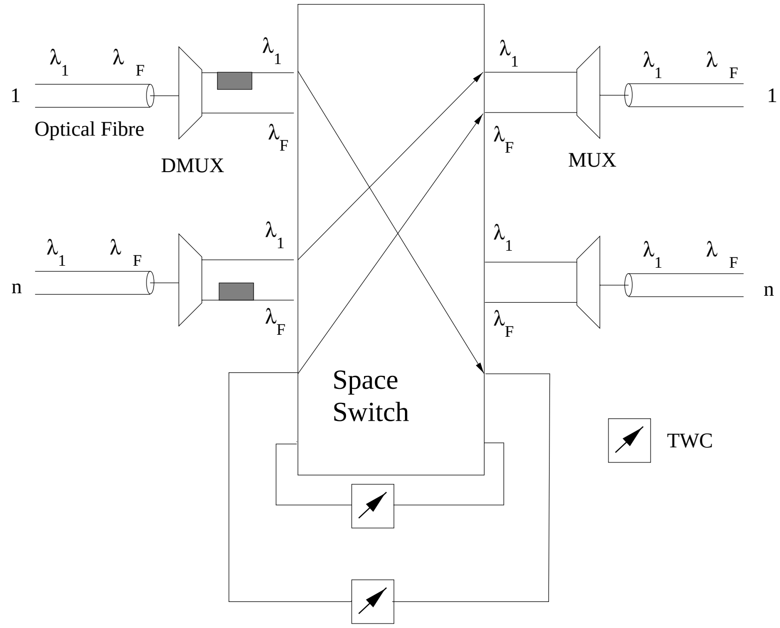

Figure 1: Optical switch with shared wavelength converters.

deviations analysis of the classical occupancy problem to estimate the probability of packet loss for a given number of shared wavelength converters. This work shows two complementary ways of computing the wavelength blocking more precisely. An analytic approach combines combinatorial expressions for the classical occupancy problem with the Bahadur-Rao approximation for the sum of independent variables. A simulation based approach uses importance sampling driven by the previous large deviations analysis. The importance sampling method is motivated in part by ongoing work (not reported here) on generalizations of the switch problem in which the switch may choose to route packets to any one of a set of output fibres.

Section 2 introduces the switch and traffic model, and the next section briefly describes the analytic approach. Section 4 reviews the large deviations analysis of [4] and describes the new importance sampling approach. Finally, some numerical results are presented and discussed.

2 Switch Model

The switch, depicted schematically in Figure 1, is connected to n neighboring nodes, with n input fibres and n output fibres carrying packets from and to the neighbors. Each fibre contains F wavelength channels, and the switch has W tunable wavelength converters (TWC) in a shareable pool. Packets arriving on wavelength channel f may be connected directly to channel f on any output fibre, or may be connected to *any* channel on any output fibre after passing through a wavelength converter. During each time slot, $a_f(\nu)$ packets of wavelength f arrive from node ν. The $\{a_f(\nu)\}$ are modeled as independent and identical Bernoulli random variables with success probability a, where a is referred to as the traffic intensity. In the present model, each packet must go out on a particular fibre, chosen at random independently of all other packets. (In generalizations under investigation, each packet has a *set* of acceptable output fibres.) We denote by $z_f(\nu)$ the number of packets of wavelength f with intended destination ν. The number of packets of a wavelength f, $\sum_\nu a_f(\nu) = \sum_\nu z_f(\nu)$, is a binomial $\text{Bin}(n,a)$ random variable. The fixed-wavelength arrival processes $\{z_f(\cdot)\}$ are identically and independently distributed, and each is equivalent to an urn occupancy problem. Specifically, $z_f(\cdot)$ could be obtained by throwing a $\text{Bin}(n,a)$ number of balls (the packets) at random into n urns (representing the output fibres), and counting the number of balls in each bin.

2.1 Blocking Analysis

When the number of packets needing a switch resource exceeds the capacity of that resource, packets need to be dropped. Packets are lost in two ways.

- **Capacity Blocking** The number of packets requiring destination ν is $\sum_f z_f(\nu)$. Only F channels lead to a given destination, hence $B_c(\nu) = \left(\left(\sum_{f=1}^F z_f(\nu)\right) - F\right)^+$ packets must be discarded due to capacity constraints at destination ν, and $B_c = \sum_\nu B_c(\nu)$ are lost across the switch.

- **Wavelength Blocking** Ignoring capacity blocking for a moment, the number of packets of a given f, ν pair that would require wavelength conversion is $(z_f(\nu) - 1)^+$. The total number of packets headed for destination ν which would require conversion would be $\tilde{w}(\nu) = \sum_f (z_f(\nu) - 1)^+$. Note that by Jensens inequality, $\tilde{w}(\nu) \geq B_c(\nu)$. Thus in the capacity blocking step, we can *always* choose to drop packets which would overflow anyway in the wavelength blocking step, and every packet dropped in the capacity blocking step reduces the blocking in the wavelength blocking step. Since $\tilde{w} = \sum_\nu \tilde{w}(\nu)$ is the number of converters that would be needed, ignoring capacity blocking, the number actually needed is $w = \tilde{w} - B_c$. Since W converters are available, $(w - W)^+$ packets are lost due to wavelength blocking.

The minimum number of blocked packets overall is

$$B = B_c + (w - W)^+ = \max\{\ B_c\ ,\ (\tilde{w} - W)^+\ \}$$

The packet loss probability $P_b = \mathbf{E}[B]/(anF)$ is defined to be the expected number of blocked packets divided by the expected number of arriving packets.

2.2 Upper Bound and Lower Bounds

In [4], the number of blocked packets was upper bounded by the expression

$$B' = B_c + (\tilde{w} - W)^+ > B.$$

This bound is tight when B_c is much smaller than $(\tilde{w} - W)^+$ (the converter-limited regime) or conversely when B_c dominates $(\tilde{w} - W)^+$ (the capacity limited regime). The approximation is loosest when the two components are approximately equal, in which casen the error is bounded by a factor of two. For packet loss probability, we have the upper bound

$$\mathbf{P_b} \leq \frac{\mathbf{E}[B_c] + \mathbf{E}[(\tilde{w} - W)^+)]}{anF}$$

and the lower bound

$$\mathbf{P_b} \geq \frac{\max\{\ \mathbf{E}[B_C]\ ,\ \mathbf{E}[(\tilde{w} - W)^+]\}}{anF}.$$

These bounds are convenient because the two expectation terms can be computed separately. The expected capacity blocking $\mathbf{E}[B_c]$ can be computed from the binomial distribution of the traffic, either directly or via a Poisson approximation. The next two sections below focus on computing the wavelength blocking term

$$\mu(W) = \mathbf{E}\left[(\tilde{w} - W)^+\right]$$

using Bahadur-Rao approximation and importance sampling. As Figure 5 in Section 5 illustrates, the upper and lower bounds agree closely except in the transition between capacity-dominated and converter-dominated packet loss.

3 Analytic Computation of Wavelength Blocking

Denote by

$$\tilde{w}_f = \sum_{\nu=1}^{n} (z_f(\nu) - 1)^+$$

the packets of a given wavelength f which are lost due to insufficient wavelength conversion (ignoring capacity blocking). The $\tilde{w}_f$ are i. i. d. random variables, whose distribution may be computed using the inclusion-exclusion principle [6, Chapter 4].

The total number of blocked packets $\tilde{w}$ is the sum of F i. i. d. random variables, and for large F, excursions above the mean of the form $\mathbf{P}\{\tilde{w} \geq \alpha\}$ can be estimated using large deviations estimates of the form $\mathbf{P}\{\tilde{w} \geq \alpha\} \approx e^{-FJ(\alpha/F)}$. Denoting by $\tilde{w}_f^{(F)}$ the blocked packets of a given wavelength in an F wavelength system, the exponent function is given by

$$J(\alpha) = \sup_t \; t\alpha - \lim_{F\to\infty} \frac{1}{F} \log \mathbf{E}\left[e^{t\tilde{w}_f^{(F)}}\right]. \tag{1}$$

The Bahadur-Rao approximation refines this approach with correction terms, leading to very accurate results even for relatively small values of F [7]. The approximation is given by

$$\mathbf{P}\{\tilde{w} \geq \alpha\} \approx \frac{e^{-FJ(\alpha/F)}}{\sqrt{2\pi F\sigma^2}\,\{1 - e^{-t^*_\alpha}\}}$$

where t^*_α is the maximizing argument in (1), and where σ^2 is the corresponding twisted variance. The desired expectation can finally be expressed

$$\mathbf{E}\left[(\tilde{w} - W)^+\right] = \sum_{k=W+1}^{(n-1)F} \mathbf{P}\{\tilde{w} \geq k\},$$

where in practice only the first few terms of the sum are significant.

4 Importance Sampling of Wavelength Blocking

We are interested in computing $\mu(W) = \mathbf{E}\left[(\tilde{w} - W)^+\right]$ in a regime in which packets loss is a rare event, for example when 1 in 10^{10} packets are lost. Straightforward Monte Carlo simulation is out of the question in this regime, because billions of realizations without loss must be generated for each lost packet. Importance sampling offers an efficient simulation technique [8, 9], which in this case provides accurate answers using only a few thousand realizations.

Denote the probability measure governing $z_f(\nu)$ in our model by P, and let Q represent some other measure (referred to as the *twisted measure*). After generating N realization of the process under the twisted measure and computing the associated $\{\tilde{w}(i)\}_{i=1}^N$, we form the unbiased estimate

$$\hat{\mu}(W) = \frac{1}{N}\sum_{i=1}^{N}(\tilde{w}(i) - W)^+ \frac{P(i)}{Q(i)} \tag{2}$$

where $P(i)$ and $Q(i)$ are the probabilities of the i-th realization under the two measures. The crux of the problem is to choose a measure Q which is easy to compute for any particular realization, and for which the estimate $\hat{\mu}(W)$ has small variance. For analytical as well as intuitive reasons, good twist measures typically give emphasis to realizations which are the most likely paths to the rare events in question. In a number of examples, most likely paths determined by large deviations analysis have been successfully used for importance sampling (see e. g. [10]).

4.1 Review of Large Deviations Analysis

To set the stage for our choice of twisted measure Q, we briefly review some large deviations results from [4]. As mentioned earlier, the process $z_f(\nu)$ is equivalent to F independent urn occupancy subproblems, where a random number of balls are thrown into n urns in each subproblem. Fixing attention on a single wavelength f, note that if r_f balls are thrown into n bins, occupying m_f bins, then $\tilde{w}_f = r_f - m_f$ balls enter previously occupied cells. The total number of wavelength converters required is $\tilde{w} = \sum_f \tilde{w}_f$. The large deviations analysis give asymptotically sharp (in n) estimates of

$$\frac{1}{n}\log \mathrm{P}(\tilde{w}/n >= W/n)\,.$$

when this is a rare event. Conditioned on the occurence of this rare event, the analysis shows that for large n it is likely that an unusually large number of balls r_f were thrown in each subproblem, that the balls piled up into a small number of bins m_f in each subproblem, and that degree of unusual piling up is roughly equal among the subproblems. The conditional excess traffic is characterized by a twist parameter $\beta > a$, and the excess piling up of the packets is described by a twist parameter $\rho > 1$, such that $r_f/n \approx \beta$ and $m_f/n \approx \left(1 - e^{-\rho\beta}\right)/\rho$.

The analysis also gave sample path results for the pattern in which the balls enter the urns when m is unusually small. Let $\gamma_1(x)$ represent the fraction of occupied urns after xn balls have been thrown into n urns. Conditioned on the terminal condition $\gamma_1(\beta) = m/n < 1 - e^{-\beta}$, it is likely that

$$\gamma_1(x) \approx \Gamma_1(x) := \frac{1}{\rho}\left(1 - e^{-\rho x}\right)$$

where $\rho > 1$ is (uniquely) chosen so that $\Gamma_1(\beta) = m/n$.

4.2 Importance Sampling procedure

The above analysis suggests the following procedure for importance sampling with twist Q. In the i-th of N iterations, a random variable r_f^i and sequence $\gamma_{1,f}^i(\cdot)$ is generated from a distribution Q for each wavelength f, and the number of packets requiring conversion

$\tilde{w}^i = \sum_{f=1}^{F} r_f^i - \gamma_{1,f}^i(\beta)$ is computed. The likelihood ratio $L(\mathbf{r}^i, \boldsymbol{\gamma_1}^i) = P(\mathbf{r}^i, \boldsymbol{\gamma_1}^i)/Q(\mathbf{r}^i, \boldsymbol{\gamma_1}^i)$ is also computed, and finally one forms the estimate (2). Under the twisted distribution Q, the number of arriving packets r_f is chosen from a binomial distribution with increased traffic intensity $\beta > a$. The realizations of $\gamma_{1,f}$ are generated conditioned on particular realizations of r_f, as described below.

The occupancy values are generated using a Markov chain tracking the process state $\gamma_1(x)$ as balls are thrown in one at a time. The chain begins with all urns empty, $\gamma_1(0) = 0$, and each ball thrown increments "time" by $1/n$. The probability of entering an unoccupied (resp. occupied) urn is denoted π_0 (resp. π_1). Under the true distribution P, the probability of hitting a set of urns is proportional to the size of the set and does not otherwise depend on time. That is, $\pi_1^P(\gamma_1, x) = \gamma_1$. In the conditional minimum-cost trajectory $\Gamma_1(x)$ arising in the large deviations analysis, the rate of increase $\dot{\Gamma}_1(x)$ of the number of occupied urns is the instantaneous, twisted probability of throwing the next ball into an empty urn. Hence, if the twisted probability is emphasize the large deviations behavior $\gamma_1(x) \approx \Gamma_1(x)$, it should satisfy

$$\pi_1^Q\left(\frac{1}{\rho}\left(1 - e^{-\rho x}\right), x\right) = 1 - \dot{\Gamma}_1(x) = 1 - e^{-\rho x}.$$

There are a number of choices of π_1^Q which satisfy this constraint. An appealingly simple choice which preserves the state-dependence of P is is $\pi_1^Q(\gamma_1, x) = \rho\gamma_1$. To avoid arbitrarily large values of π_0^P/π_0^Q as $\gamma_1 \to 1/\rho$, in practice we used

$$\pi_1^Q(\gamma_1, x) = \begin{cases} \rho\gamma_1 & \gamma_1 < \omega_1 \\ 1 - (1 - \gamma_1)\frac{1-\rho\omega_1}{1-\omega_1} & \gamma_1 > \omega_1 \end{cases}, \quad (3)$$

where $\omega_1 = \beta - W/(nF) = (1 - e^{-\rho\beta})/\rho$ is the targeted fraction of occupied urns in the large deviations analysis. Another option is to ignore state dependence altogether for $0 < \gamma_1 < 1$, as

$$\pi_1^Q(\gamma_1, x) = \begin{cases} 0 & \gamma_1 = 0 \\ 1 - e^{-\rho x} & 0 < \gamma_1 < 1 \\ 1 & \gamma_1 = 1 \end{cases}. \quad (4)$$

Although the state-based twist (3) appears to be more natural than the time-based twist in this case, (4) turns out to be easier to generalize to the multi-dimensional state variables that occur when packets may be routed within sets of output fibres. Both twisted measures make it likely that approximately $n\omega_1$ bins are eventually occupied.

Twisted sequences $\{\gamma_1(i/n)\}_{i=0}^{r_\nu}$ are generated one ball at a time, using π_1^Q. Let $u(i) \in \{0, 1\}$ indicate whether or not the $(i+1)$-st ball entered an occupied urn. Then the complete likelihood ratio for each traffic realization is

$$\begin{aligned} L(\mathbf{r}, \boldsymbol{\gamma_1}) &= \prod_{f=1}^{F} \frac{P[r_f]}{Q[r_f]} \frac{P[\gamma_{1,f} \,|\, r_f]}{Q[\gamma_{1,f} \,|\, r_f]} \\ &= \prod_{f=1}^{F} \left(\frac{a}{\beta}\right)^{r_f} \left(\frac{1-a}{1-\beta}\right)^{n-r_f} \left[\prod_{i=0}^{r_f-1} \frac{\pi_{u(i)}^P(\gamma_1(i/n), i/n)}{\pi_{u(i)}^Q(\gamma_1(i/n), i/n)}\right]. \end{aligned}$$

5 Results

The importance sampling technique and Bahadur-Rao analysis proved to be extremely accurate, and the computational complexity of the simulation was largely independent of the blocking level being estimated. Figure 2 depicts the expected number of lost packets per timeslot $\hat{\mu}(W)$ as a function of the traffic intensity a for a switch with 25 fibres, 20 wavelengths per fibre, and a pool 30 wavelength converters. In the importance sampling simulations, 5000 simulation runs were used to calculate each data point. The simulation and analysis agree closely in the range of probabilities depicted. The plot also shows for comparison an estimate computed directly from large deviations as in [4]. The large deviations analysis is accurate to within a couple of orders of magnitude, but does not provide the precision of the other two techniques. The need for this extra precision depends on the intended application and on the accuracy of the model.

In (naive or importance sampled) simulation, it is desirable for the individual summands in (2) to have a reasonably small coeficient of variation c, defined as the ratio of the standard deviation s to the mean μ. The coefficient of variation C of the final estimate $\hat{\mu}(W)$ decreases with the number of independent samples N as $C = c/\sqrt{N}$, so that the number of samples needed to meet a desired level of accuracy is proportional to c^2. An importance sampling estimator is said to be asymptotically optimal if c is bounded by a constant, regardless of how small μ becomes. Figure 3 depicts the importance sampled estimate of blocking probability as a function of traffic intensity a. The estimates used $N = 5000$ simulated realizations for each data point, for a switch with $n = 25$ fibres, $F = 20$ wavelengths, and $W = 50$ converters. The estimate computed using (3) is labeled "state-based", while the estimate from (4) is labeled "time-based", and the two agree closely. The standard deviation s, computed empirically from the 5000 samples, is shown to track the probability being estimated within a roughly constant factor in both cases. For the same switch scenario, Figure 4 directly illustrates the (empirically estimated) variation c as a function of traffic level for the two importance sampling techniques. The state based method has slightly lower variance than the time based method, and for both methods the coefficient of variation increases only slightly for small blocking probabilities. In naive simulation (Q=P), each summand in (2) is a non-negative, integer-valued random variable with mean μ. When $\mu < 1$, the variance is lower bounded by the variance $\mu(1-\mu)$ of a Bernoulli random variable, so that $c \geq \sqrt{(1/\mu)-1}$ as $\mu \to 0$. This bound is depicted in Figure 4 for comparison with the importance sampling.

To compute overall packet loss probabilities, the wavelength blocking terms computed via the methods of Section 3 or 4 are combined with the capacity blocking term $\mathbf{E}\left[\mathbf{P_b}\right]/(anF)$. Overall packet loss for a switch with 25 fibres, 20 wavelengths, and various pool sizes is depicted in Figure 5. The upper and lower bounds coincide except at the transition point between capacity dominated and wavelength dominated blocking. Given a desired loss threshold such as $\mathbf{P_b} \leq 10^{-8}$, the capacity blocking curve can be used to determine a maximum acceptable traffic intensity, in this case $a \leq 0.2$. The wavelength blocking analysis can then be used to determine a minimum acceptable converter pool size, in this case roughly $W \geq 30$.

6 Discussion

Large deviations analysis is a useful technique for computing packet loss on switches of the type studied here. As shown in previous work, the large deviations exponent itself provides a reasonable approximation. Sample-path level analysis gives intuition into how rare events occur and forms the basis of twisted measures for efficient importance sampling. Ongoing work is extending the importance sampling approach to generalizations of the switch, for which the associated models are general overflow problems, rather than

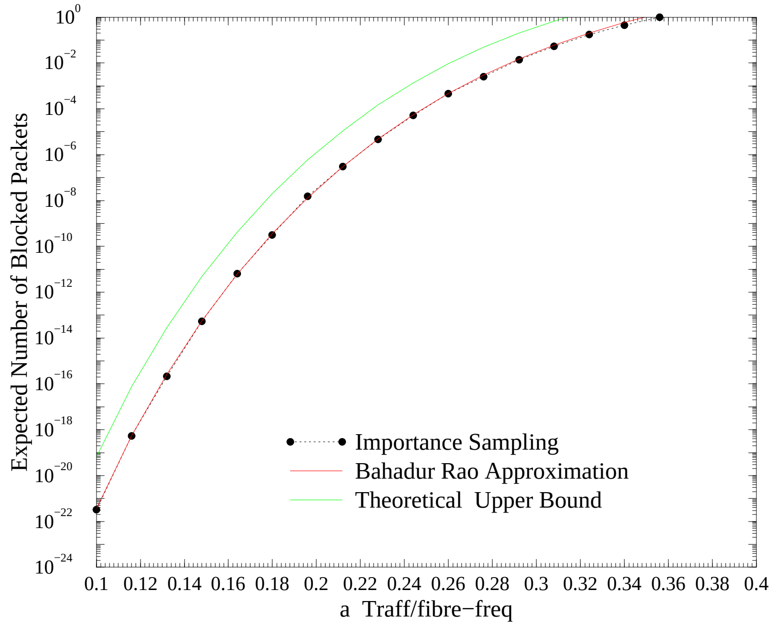

Figure 2: Average number of blocked packet per timeslot as a function of traffic intensity, computed by three methods: Bahadur-Rao approximation, importance sampling, and large deviations exponent. The switch has $n = 25$ fibres, $F = 20$ wavelengths, and $W = 30$ TWCs.

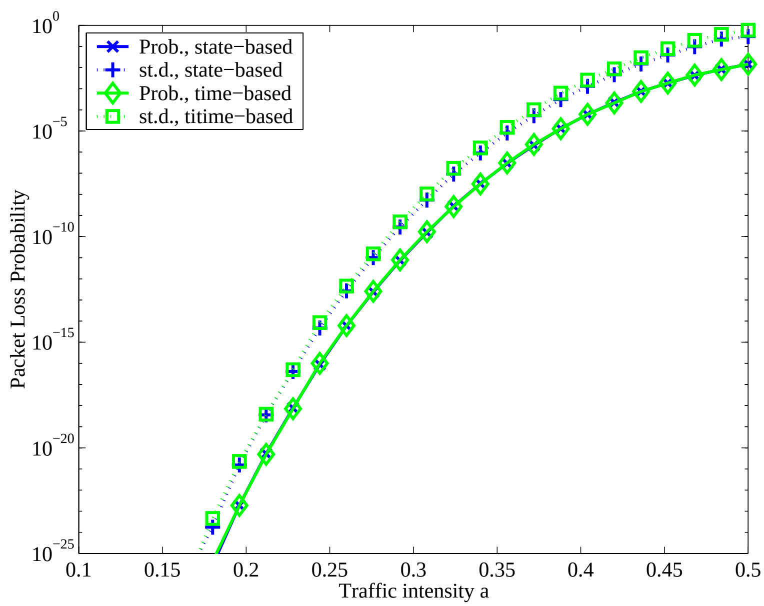

Figure 3: Packet loss probability due to wavelength blocking as a function of traffic intensity, computed by two importance sampling methods. Also shown is the standard deviation s of the estimate summands from (2). The switch has $n = 25$ fibres, $F = 20$ wavelengths, and $W = 50$ TWCs.

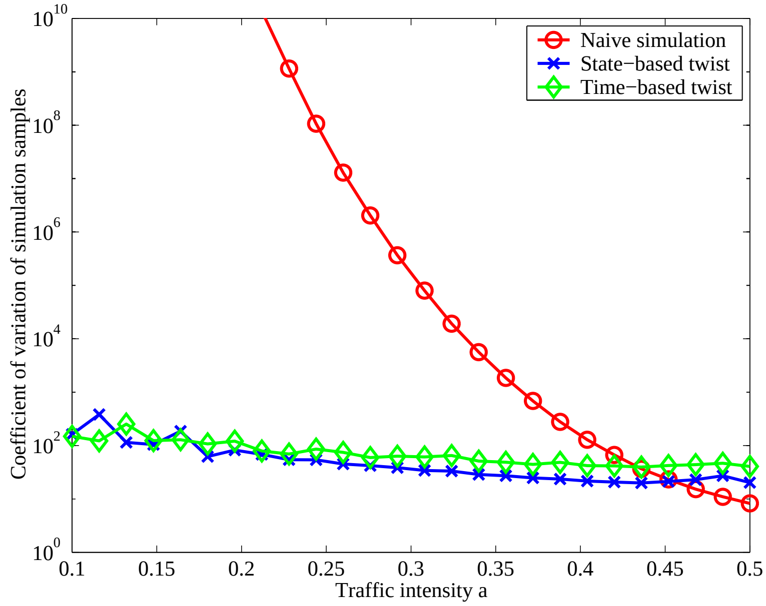

Figure 4: Coefficient of variation of estimate summands, as a function of traffic intensity, for naive simulation and two importance sampling methods. The switch has $n = 25$ fibres, $F = 20$ wavelengths, and $W = 50$ TWCs.

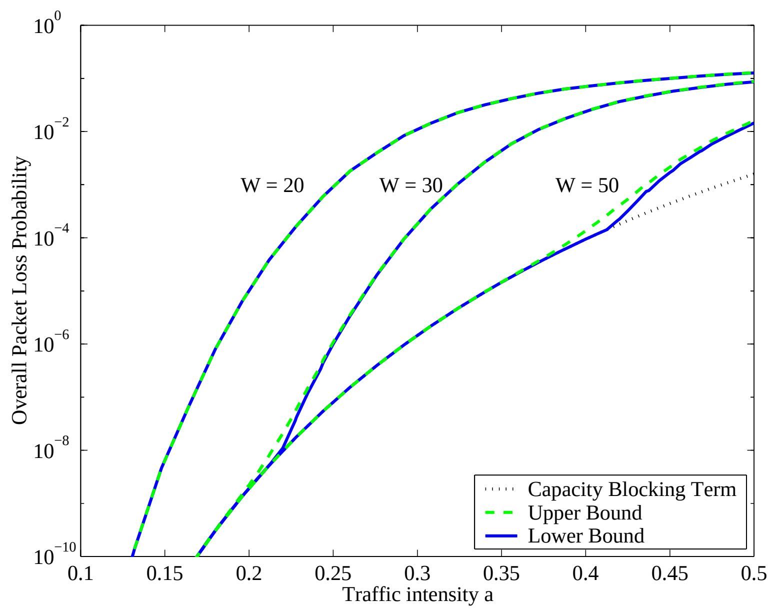

Figure 5: Bounds on overall packet loss probability as a function of traffic intensity a, for various converter pool sizes W. The bounds combine an analytic capacity blocking term (dotted line) with wavelength blocking curves computed by importance sampling. The switch has $n = 25$ fibres and $F = 20$ wavelengths.

the classical occupancy problem. The twisted measures for these problems are suggested by general sample-path large deviations analysis developed in [11]. Although the present importance sampling techniques have worked well, it would be worthwhile to analytically establish good properties such as asymptotic optimality. Extensions of the occupancy theory which would permit the direct estimation of $\mathbf{P_b}$ without resorting to the upper and lower bounds are also of interest.

References

[1] D. Bishop, C. Giles, S. Das, "The Rise of Optical Switching", *Scientific American*, **284** 1, pp. 88-94, Jan 2001.

[2] K.-C. Lee and V. Li, "A Wavelength-Convertible Optical Network", *J. of Lightwave Tech.*, **11** 5/6, pp. 962-970, May/June 1993.

[3] V. Eramo and M. Listanti, "Packet Loss in a Bufferless Optical WDM Switch Employing Shared Tuneable Wavelength Converters", *Lightwave Technology*, **18** 12, 2000.

[4] V. Eramo, M. Listanti, C. Nuzman, P. Whiting, "Analysis of an Optical Packet Switch with Shared Wavelength Conversion", Allerton Conf. on Communication, Control, and Computing, 2001.

[5] V. Eramo, M. Listanti, C. Nuzman, P. Whiting, "Optical Switch Dimensioning and the Classical Occupancy Problem", *International Journal of Communications Systems*, **15** 2/3, pp. 127-141, 2002.

[6] W. Feller, *An Introduction to Probability Theory and Its Applications*, Volume 1, John Wiley, New York, 1968.

[7] J. Bucklew, *Large Deviations Techniques in Decision, Simulation, and Estimation*, John Wiley, New York, 1990.

[8] P. W. Glynn and D. L. Inglehart, "Importance Sampling for Stochastic Simulations", *Management Science* **35**, pp. 1367-1392, 1989.

[9] Sandeep Junega, "Efficient Rare Event Simulation Using Importance Sampling: An Introduction", chapter in *Computational Mathematics, Modeling and Algorithms*, Narosa, New Delhi, pp. 357-396, 2003.

[10] S. Parekh and J. Walrand, "A Quick Simulation Method for Excessive Backlogs in Networks of Queues", *IEEE Transactions on Automatic Control*, **34**, 1, pp. 54-66, 1989.

[11] P. Dupuis, C. Nuzman, and P. Whiting, "Large Deviations for Occupancy Problems", manuscript under review, May 2001.