TECHNICAL RESEARCH REPORT

Cell Loss Probabilities in Input Queueing Crossbar Switches Via Light Traffic

by Y.B. Kim, A.M. Makowski

CSHCN T.R. 94-5
(ISR T.R. 94-72)

CENTER FOR SATELLITE AND HYBRID COMMUNICATION NETWORKS

Cell Loss Probabilities in Input Queueing Crossbar Switches Via Light Traffic Derivatives

Young B. Kim and Armand M. Makowski
Institute for Systems Research, Center for Satellite and Hybrid Communication Networks,
and Department of Electrical Engineering
University of Maryland
College Park, MD 20742
(301) 405-6554

Abstract

Under most system assumptions, closed form solutions of performance measures for input queueing crossbar switches are not available. In this paper, we present expressions and bounds for the derivatives of cell loss probabilities with respect to the arrival rate evaluated at a zero arrival rate. These bounds are used to give an approximation by Taylor expansion, thereby providing an economical way to estimate cell loss probabilities in light traffic.

INTRODUCTION

The recent deployment by NASA of its advanced communication technology satellite (ACTS) represents a key step in demonstrating the feasibility of packet-switched communications satellite with on-board processing and spot-beam operation. In these advanced satellite systems, a reliable on-board fast packet switch is essential for ensuring that packets are routed from different uplink beams to different downlink beams. Although the on-board and terrestrial (ATM) switches share many similar features and capabilities, the design of an on-board fast-packet switch needs to incorporate additional factors due to the unique satellite communication environment; these factors include size, power, reliability, fault-tolerance, multicasting and congestion control, to name a few.

Several candidate architectures could possibly support on-board fast-packet switching. Noteworthy among them are the various space-division switching fabrics developed in the past decade for terrestrial ATM networks; in particular, we focus here on non-blocking crossbar switches with input buffering. Having in mind satellite applications, we assume that the input buffers are finite, and typically small. In that context, the key performance measure we wish to evaluate is the cell loss probability (CLP).

To carry out this evaluation, we consider a simple discrete-time model for a synchronous $K \times K$ non-blocking crossbar switch where the input queues are of finite size b. Cells arrive at each input port according to a Bernoulli process of rate λ; cells that find a full queue are rejected. Output contention manifests itself through head-of-the-line (HOL) blocking (Karol et al. 1987), and is resolved by a simple randomized arbitration mechanism. Despite their simplicity, these rules of operation produce a very complex queueing behavior as input queues become correlated over time. This explains why the performance analysis is possible only under special conditions like infinite switch size and saturation assumptions (Karol et al. 1987). Therefore, under most model assumptions, closed form solutions of performance measures of interest are not available, nor can they be expected. Worse perhaps, when evaluating CLP, Monte-Carlo simulation techniques turn out to be of limited use owing to their vast computational cost as the desired CLP in ATM networks, being usually in the range of 10^{-6} to 10^{-12}, corresponds to rare events.

In this paper we address the problem of evaluating the CLP in light traffic, where light traffic refers to the the system being lightly loaded, or equivalently to the situation $\lambda \simeq 0$. If $P_b(\lambda)$ denotes the CLP when the arrival at each port is λ, we show how to evaluate, at least in principle, the light traffic derivatives $\frac{d^k}{d\lambda^k} P_b(0+) \equiv \lim_{\lambda \to 0+} \frac{d^k}{d\lambda^k} P_b(\lambda)$ for various values of k. In particular, we show that $\frac{d^k}{d\lambda^k} P_b(0+) = 0$ for $k = 0, 1 \ldots, 2b-1$, and spend most of our efforts on evaluating the first two non-zero derivatives. These formulae exploit the regenerative structure of a Markov chain associated with the queueing model. For lack of space, we only state the results, and refer the reader to Kim 1995 and Kim and Makowski 1994a for a complete discussion.

We then propose to approximate the CLP $P_b(\lambda)$, at least for small values of the arrival rate λ, by a Taylor series expansion of $P_b(\lambda)$ near the origin, which here takes the form

$$P_b(\lambda) \simeq \frac{1}{(2b)!}\lambda^{2b}\frac{d^{2b}}{d\lambda^{2b}}P_b(0+) + \frac{1}{(2b+1)!}\lambda^{2b+1}\frac{d^{2b+1}}{d\lambda^{2b+1}}P_b(0+), \quad \lambda \simeq 0.$$

This approximation works well for small values of λ where the CLP is expected to be very small, a situation often handled by variance reduction techniques such as importance sampling (Kim and Makowski 1994b). The method proposed here thus provides a numerical alternative to these techniques, at least in light traffic.

In a next step, this light traffic information obtained here can be combined with medium to heavy traffic information to produce a "light traffic interpolation" as in Simon et al. 1988 and 1989, thereby providing an economical way to evaluate CLP in light traffic. This approach is taken in Kim 1995 and Kim and Makowski 1994a.

MODEL

All random variables (rvs) are assumed defined on a common probability triple $(\Omega, \mathcal{F}, \mathbf{P})$. Let K, the switch size, and b, the common buffer size, be positive integers held fixed throughout the discussion. We begin with $3K$ mutually independent collections of rvs $\{A^k_{t+1},\ t = 0,1,\ldots\}$, $\{\nu^k_{t+1}, t = 0,1,\ldots\}$, and $\{\varpi^k_{t+1}, t = 0,1,\ldots\}$, $k = 1,\ldots,K$, under the following assumptions:

1. For each $k = 1,\ldots,K$, the rvs $\{A^k_{t+1},\ t = 0,1,\ldots\}$ are i.i.d. rvs with $\mathbf{P}\left[A^k_{t+1} = 1\right] = 1 - \mathbf{P}\left[A^k_{t+1} = 0\right] = \lambda,\ t = 0,1,\ldots;$
2. For each $k = 1,\ldots,K$, the rvs $\{\nu^k_{t+1}, t = 0,1,\ldots\}$ are i.i.d. rvs with $\mathbf{P}\left[\nu^k_{t+1} = \ell\right] = \frac{1}{K},\ \ell = 1,\ldots,K;\ t = 0,1,\ldots;$
3. For each $k = 1,\ldots,K$, the rvs $\{\varpi^k_{t+1}, t = 0,1,\ldots\}$ are i.i.d. rvs with $\mathbf{P}\left[\varpi^k_{t+1} = \ell\right] = \frac{1}{K},\ \ell = 1,\ldots.$

As the switching fabric operates in synchronous mode, we divide time into contiguous slots of unit length; each such time slot is divided into two consecutive minislots. Loosely speaking, at each queue, new arrivals are completed by the end of the first minislot, at which time the address of the HOL cell is determined; contention management is carried out at the beginning of the second minislot, and this is followed with possible cell transmission across the switching fabric: Fix $i = 1,\ldots,K$ and $t = 0,1,\ldots$. Let Q^i_t and V^i_t respectively denote the number of cells present in the i^{th} input queue and the destination of the HOL cell in that queue at the beginning of time slot $[t, t+1)$; by convention $V^i_t = 0$ if $Q^i_t = 0$.

New cells which arrive into the system during a time slot are enqueued by the end of the first minislot, if buffer space is available. More precisely, A^i_{t+1} cell arrives at the i^{th} input port during time slot $[t, t+1)$. An arriving cell is accepted into the buffer and put at the end of the line only if it finds the i^{th} queue to be non–full, that is, if $Q^i_t < b$; otherwise the cell is blocked and rejected. Therefore, at the end of the first minislot in $[t, t+1)$, there are $Q^i_t + \mathbf{1}\left[Q^i_t < b\right] A^i_{t+1}$ cells in the i^{th} queue, and the HOL cell amongst them is eligible for transmission across the switch during the second minislot of the time slot $[t, t+1)$.

The addressing mechanism is random and uniform across input ports, and statistically independent of the generation of arrivals. Hence, there is no loss of generality in taking the viewpoint that each cell declares its destination address immediately upon reaching the head of line (at which time such a cell is called a *fresh* cell) and keeps its address until it begins transmission across the switching fabric (at the start of the second minislot of time slots). If $V^i_{(t+1)^-}$ denotes the address content of the HOL cell present in the i^{th} input queue at the start of the second minislot in $[t, t+1)$, then we have

$$V^i_{(t+1)^-} = A^i_{t+1}\mathbf{1}\left[Q^i_t = 0\right]\nu^i_{t+1} + (1 - A^i_{t+1}\mathbf{1}\left[Q^i_t = 0\right])V^i_t,\ i = 1,\ldots,K. \tag{1}$$

At the beginning of the second minislot in slot $[t, t+1)$, the switch controller mediates potential output contentions by randomly selecting one HOL cell amongst the HOL cells which have the same output address: Let $\mathcal{G}^\ell_{t+1}$ denote the set of input ports whose HOL cell has destination address ℓ at the beginning of the second minislot, so that

$$\mathcal{G}^\ell_{t+1} \equiv \{k \in \{1, \ldots, K\} : V^k_{(t+1)^-} = \ell,\ Q^k_t + A^k_{t+1} > 0\}, \quad \ell = 1, \ldots, K \tag{2}$$

and for convenience we set $\mathcal{G}^0_{t+1} \equiv \{k \in \{1, \ldots, K\} :\ Q^k_t + A^k_{t+1} = 0\}$. Whenever $|\mathcal{G}^\ell_{t+1}| > 0$, we define the $\mathcal{G}^\ell_{t+1}$-valued rv O^ℓ_{t+1} by

$$O^\ell_{t+1} = j \quad \text{w.p. } \frac{1}{|\mathcal{G}^\ell_{t+1}|}, \quad j \in \mathcal{G}^\ell_{t+1}, \quad \ell = 1, \ldots, K. \tag{3}$$

The rv O^ℓ_{t+1} selects an index in $\mathcal{G}^\ell_{t+1}$ at random, thereby indicating the input port with index in the set $\mathcal{G}^\ell_{t+1}$, whose HOL cell will be transmitted to output ℓ. Defining the binary rvs D^i_{t+1}, $i = 1, \ldots, K$, by

$$D^i_{t+1} = \sum_{\ell=1}^{K} \mathbf{1}\left[i \in \mathcal{G}^\ell_{t+1}\right] \mathbf{1}\left[O^\ell_{t+1} = i\right], \quad i = 1, 2, \ldots, K \tag{4}$$

we see that

$$Q^i_{t+1} = Q^i_t + A^i_{t+1}\mathbf{1}\left[Q^i_t < b\right] - D^i_{t+1} \quad \text{and} \quad V^i_{t+1} = (1 - D^i_{t+1})V^i_{(t+1)^-} + D^i_{t+1}\mathbf{1}\left[Q^i_{t+1} > 0\right]\varpi^i_{t+1}.$$

It is appropriate to view the pair $X_t \equiv (V_t, Q_t)$ as the state of the system at the beginning of time slot $[t, t+1)$, with $\{0, 1, \ldots, b\}^K$-valued rv $Q_t \equiv (Q^1_t, \ldots, Q^K_t)$ and $\{0, 1, \ldots, K\}^K$-valued rv $V_t \equiv (V^1_t, \ldots, V^K_t)$. Under the enforced assumptions, the rvs $\{X_t,\ t = 0, 1, \ldots\}$ form a Markov chain with finite state space $\mathcal{X} \equiv \{1, \ldots, K\}^K \times \{0, 1, \ldots, b\}^K$.

The finite state Markov chain $\{X_t,\ t = 0, 1, \ldots\}$ is irreducible and aperiodic, thus ergodic and there exists an $\{1, \ldots, K\}^K \times \{0, 1, \ldots, b\}^K$-valued rv X_∞ such that $X_t \Longrightarrow X_\infty$ (as $t \longrightarrow \infty$) with $\Longrightarrow$ denoting weak convergence. By invoking BASTA (Bernoulli arrivals see time average), we see that the cell loss probability $P_b(\lambda)$ is given by $P_b(\lambda) \equiv \mathbf{P}_\lambda[Q^1_\infty = b]$ (with $\mathbf{P}_\lambda$ standing for $\mathbf{P}$ if the arrival rate is λ). The Markov process $\{X_t,\ t = 0, 1, \ldots\}$ is a regenerative process which is positive recurrent, with the empty state $\mathbf{0}$ acting as regeneration state. The rv $\tau \equiv \inf\{t > 0 : Q_t = \mathbf{0}\}$ denotes the first return time to the empty state, and can be interpreted as the length of a regeneration cycle when $Q_0 = \mathbf{0}$. As well known, the steady–state measure $P_b(\lambda)$ can be expressed as the ratio $P_b(\lambda) = \Phi(\lambda)/\Psi(\lambda)$ with $\Phi(\lambda) \equiv \mathbf{E}_\lambda\left[\sum_{t=0}^{\tau-1} \mathbf{1}\left[Q^1_t = b\right]\right]$ and $\Psi(\lambda) \equiv \mathbf{E}_\lambda[\tau]$.

If the system is initially empty at time $t = 0$, then the process $\{X_t,\ t = 0, 1, \ldots\}$ is uniquely determined from the process $\{\xi_{t+1},\ t = 0, 1, \ldots\}$, where we have set

$$\xi_{t+1} \equiv (A_{t+1}, \nu_{t+1}, O_{t+1}, \varpi_{t+1}), \quad t = 0, 1, \ldots \tag{5}$$

With the processes $\{X_t, t = 0, 1, \ldots\}$ and $\{\xi_t, t = 1, 2, \ldots\}$, we associate the random elements $\mathbf{X} \equiv (X_0, \ldots, X_\tau)$ and $\Xi \equiv (\xi_1, \ldots, \xi_\tau)$. Because the process $\{X_t,\ t = 0, 1, \ldots\}$ is uniquely determined from the process $\{\xi_{t+1},\ t = 0, 1, \ldots\}$, there exist two functions $v_t(\cdot)$ and $q_t(\cdot)$ such that

$$V_t = v_t(\Xi) \quad \text{and} \quad Q_t = q_t(\Xi), \quad t = 1, \ldots, \tau. \tag{6}$$

If $\mathcal{S}_n$ denotes the set of all possible sample paths of Ξ which have a cycle of length n, i. e.,

$$\mathcal{S}_n \equiv \{(s_1, \ldots, s_n) : q_t(s_1, \ldots, s_n) \neq \mathbf{0},\ t = 1, \ldots, n-1, \text{ and } q_n(s_1, \ldots, s_n) = \mathbf{0}\}, \tag{7}$$

then $\mathcal{S} \equiv \bigcup_{n=1}^\infty \mathcal{S}_n$ represents the set of all possible sample paths. For every sample path $\mathbf{s}$ in $\mathcal{S}$, let $\ell(\mathbf{s})$ and $\#(\mathbf{s})$ denote the cycle length of $\mathbf{s}$, and the number of cells generated in $\mathbf{s}$, respectively.

Under the independence assumptions enforced on the arrival processes $\{A_{t+1}^k,\ k=1,\ldots,K;\ t=0,1,\ldots\}$, the probability $\mathbf{P}_\lambda[\Xi = \mathbf{s}]$ that a cycle is realized along the sample path $\mathbf{s}$ in $\mathcal{S}$, can be expressed by

$$\mathbf{P}_\lambda[\Xi = \mathbf{s}] = c(\mathbf{s})\lambda^{\#(\mathbf{s})} \cdot (1-\lambda)^{(K\cdot\ell(\mathbf{s})-\#(\mathbf{s}))} \tag{8}$$

where the coefficient $c(\mathbf{s})$ is determined by the set of rvs $\{(\nu_t, O_t, \varpi_t), t=1,\ldots,\ell(\mathbf{s})\}$ pertaining to s. Hence, the quantity $\Phi(\lambda)$ can be computed by

$$\Phi(\lambda) = \sum_{\mathbf{s}\in\mathcal{S}} \mathbf{P}_\lambda[\Xi = \mathbf{s}]\,\phi(\mathbf{s}) = \sum_{\mathbf{s}\in\mathcal{S}} \phi(\mathbf{s})c(\mathbf{s})\lambda^{\#(\mathbf{s})} \cdot (1-\lambda)^{(K\cdot\ell(\mathbf{s})-\#(\mathbf{s}))}, \tag{9}$$

where we have set

$$\phi(\mathbf{s}) \equiv \sum_{t=0}^{\ell(\mathbf{s})-1} \mathbf{1}\left[q_t^1(\mathbf{s}) = b\right], \quad \mathbf{s}\in\mathcal{S}. \tag{10}$$

PRELIMINARIES

To proceed with the evaluation of the light traffic derivatives, we define several subsets of $\mathcal{S}$, namely

$$\mathcal{S}_\star \equiv \{\mathbf{s}\in\mathcal{S} : \phi(\mathbf{s}) > 0\} \quad \text{and} \quad \mathcal{T}_n \equiv \{\mathbf{s}\in\mathcal{S} : \#(\mathbf{s}) = n\}, \quad n = 0,1,\ldots.$$

Lemma 1

For any sequence s in $\mathcal{S}_\star$, we have $\ell(\mathbf{s}) \geq 2b$ and $\#(\mathbf{s}) \geq 2b$, while $\ell(\mathbf{s}) = 2b$ if s is an element of $\mathcal{S}_\star \bigcap \mathcal{T}_{2b}$.

Proof

Recall that the system is initially empty. During each time slot, at most one cell is fed into each input port, and at most one cell can be selected for transmission from each queue. Hence, a minimum of b time slots is required before the first queue gets full, and similarly once the first queue is full, at least b time slots must elapse for the system to empty. Therefore, we have $\ell(\mathbf{s}) \geq 2b$ for each sample path s in $\mathcal{S}_\star$. Moreover, $\phi(\mathbf{s}) > 0$ implies that at least b cells out of the $\#(\mathbf{s})$ cells generated in that cycle have been fed into the first input port since the buffer is of size b. In order to keep these b cells in the first queue, each time slot there must exist at least one cell residing in the other queues and these cells must always block the HOL cell in the first queue. Since at least one cell goes in each time slot, the minimum number of cells to be assigned to the other input ports is thus b. Therefore we have $\#(\mathbf{s}) \geq 2b$ for all s in $\mathcal{S}_\star$ and $\ell(\mathbf{s}) = 2b$ for all s in $\mathcal{S}_\star \bigcap \mathcal{T}_{2b}$.

Upon using Lemma 1 in conjunction with a well-known formula due to Leibniz for the higher-order derivatives of a product of two functions, we readily obtain the following results.

Proposition 1

1. $\frac{d^k}{d\lambda^k} P_b(0+) = 0, \quad k = 0,1,\ldots,2b-1;$
2. $\frac{d^k}{d\lambda^k} P_b(0+) = \frac{d^k}{d\lambda^k}\Phi(0+), \quad k = 2b, 2b+1.$

Proposition 2

1. $\frac{d^{2b}}{d\lambda^{2b}}\Phi(0+) = (2b)! \sum_{\mathbf{s}\in\mathcal{S}_\star\cap\mathcal{T}_{2b}} c(\mathbf{s})\phi(\mathbf{s});$
2. $\frac{d^{2b+1}}{d\lambda^{2b+1}}\Phi(0+) = (2b+1)!\left[\sum_{\mathbf{s}\in\mathcal{S}_\star\cap\mathcal{T}_{2b+1}} c(\mathbf{s})\phi(\mathbf{s}) - 2b(K-1)\sum_{\mathbf{s}\in\mathcal{S}_\star\cap\mathcal{T}_{2b}} c(\mathbf{s})\phi(\mathbf{s})\right].$

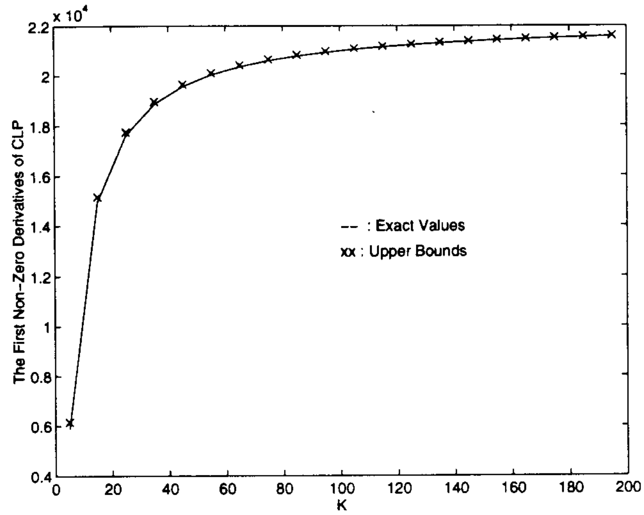

FIGURE 1. Comparison of 8th Derivatives When $b = 4$.

EVALUATION OF THE FIRST TWO NON–ZERO DERIVATIVES

For large values of b, it is not easy to evaluate the expressions $\sum_{\mathbf{s} \in \mathcal{S}_\star \cap \mathcal{T}_{2b}} c(\mathbf{s})\phi(\mathbf{s})$ and $\sum_{\mathbf{s} \in \mathcal{S}_\star \cap \mathcal{T}_{2b+1}} c(\mathbf{s})\phi(\mathbf{s})$ that appeared in Proposition 2. We cope with this difficulty by providing upper–bounds on these non–zero derivatives; these bounds turn out to be very tight when K is much bigger than b. Throughout we assume $K > b$.

To obtain the first bound, we consider a sample path $\mathbf{s}$ in $\mathcal{S}_\star \cap \mathcal{T}_{2b}$, so that $\ell(\mathbf{s}) = 2b$ and $\phi(\mathbf{s}) > 0$. Therefore, $2b$ cells are generated during the regeneration cycle to make the first input queue full (in other words, $Q_b^1 = b$), and exactly b cells should be assigned to the first input port and the remaining b cells to the other input ports such that every time there exists at least one cell (including the newly arriving cells) in the other queues, playing the role of blocking the HOL cell in the first queue. Using this remark, we show in Kim 1995 and Kim and Makowski 1994a that

$$\frac{d^{2b}}{d\lambda^{2b}} P_b(0+) \leq \frac{(2b)!}{K^b} \left\{ \sum_{\boldsymbol{x} \in \mathcal{A}} \prod_{t=1}^{b} \binom{K-1}{x_t} \frac{\sum_{k=1}^{t} x_k - t + 1}{\sum_{k=1}^{t} x_k - t + 2} \right\}, \tag{11}$$

with the notation $\mathcal{A} \equiv \{(x_1, \ldots, x_b) \in \{0, 1, \ldots, b\}^b : \sum_{t=1}^{b} x_t = b, \ \sum_{t=1}^{k} x_t - k \geq 0, \ k = 1, \ldots, b\}$. For $K \gg b$, the bound turns out to be tight enough as should be clear from Figure 1.

The evaluation of $\sum_{\mathbf{s} \in \mathcal{S}_\star \cap \mathcal{T}_{2b+1}} c(\mathbf{s})\phi(\mathbf{s})$ is more complicated but similar to that of $\sum_{\mathbf{s} \in \mathcal{S}_\star \cap \mathcal{T}_{2b}} c(\mathbf{s})\phi(\mathbf{s})$. The details are omitted for the sake of brevity and the reader is referred to Kim 1995 and Kim and Makowski 1994a.

AN EXAMPLE

We have applied our result to the estimation of CLP in crossbar switches of size $K = 20$ (See Figure 2). As we can observe in Figure 2, although the simulation results for extremely small input rates are very poor due to the limitation of Monte Carlo simulation techniques (or employing other variance reduction techniques such as importance sampling) in this light traffic regime, the two results fit well together as the input rate becomes larger.

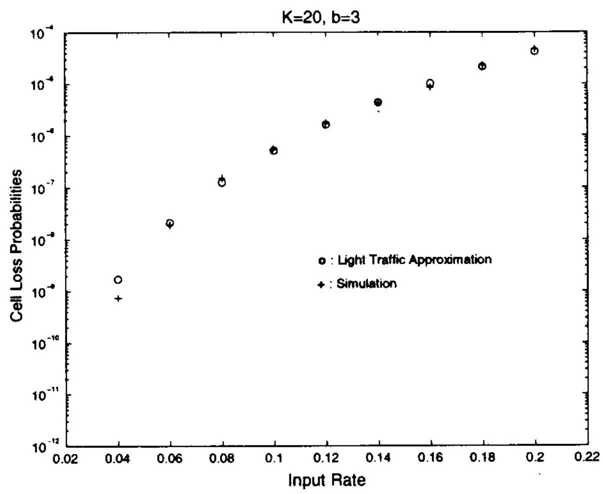

FIGURE 2. Comparisons between Light Traffic Approximation and Simulation Results.

CONCLUSIONS

In this paper, we have derived expressions for the non–zero derivatives of CLP with respect to the arrival rate evaluated in light traffic. These light traffic derivatives obtained are incorporated into an approximation of the CLP via a Taylor expansion, thereby providing an economical way to get a quick evaluation of CLP in light traffic. This result may be further extended by interpolating heavy and medium traffic values to yield a global configuration of input rate versus CLP; this is discussed in Kim 1995 and Kim and Makowski 1994a.

Acknowledgments

This work was supported partially through NSF Grant NSFD CDR-88-03012, through NASA Grant NAGW277S and through a grant from TRW, Inc., Redondo Beach, CA.

References

Karol, M. J., M. Hluchyj, and S. P. Morgan (1987) "Input Versus Output Queueing on a Space-Division Packet Switch," in *IEEE Transactions on Communications*, COM-35(12): 1347-1356.

Kim, Y.-B. (1995) *Performance Evaluation of Non–Blocking ATM Switches*, Ph.D. Thesis, Electrical Engineering Department, University of Maryland, College Park, MD.

Kim, Y.-B. and A. M. Makowski (1994a) "Cell Loss Probabilities in Input Queueing Crossbar Switches: Light Traffic Interpolations," Technical Report CCDS-TR-94-5, Center for Satellite and Hybrid Communication Networks, University of Maryland, College Park, MD.

Kim, Y.-B. and A. M. Makowski (1994b) "Cell Loss Probabilities in Input Queueing Crossbar Switches: Importance Sampling Techniques," Technical Report CCDS-TR-94-6, Center for Satellite and Hybrid Communication Networks, University of Maryland, College Park, MD.

Reiman, M. I. and B. Simon (1988) "An Interpolation Approximation for Queueing Systems with Poisson Input," in *Operations Research*, 36: 454–469.

Reiman, M. I. and B. Simon (1989) "Open Queueing Systems in Light Traffic," in *Operations Research*, 14: 26–59.