

IMPROVED QUEUEING ANALYSIS OF SHARED BUFFER SWITCHING NETWORKS

Giuseppe Bianchi and Jonathan S. Turner
Washington University

Abstract

This paper describes several methods for analyzing the queueing behavior of switching networks with flow control and shared buffer switches. It compares the various methods on the basis of accuracy and computation speed, where the performance metric of most concern is the maximum throughput. The best of the methods accurately predicts throughput for multi-stage networks constructed from large switches (≥ 8 ports).

1 Introduction

In a widely cited paper [2], Jenq describes a method for analyzing the queueing behavior of binary banyan networks with a single buffer at each switch input. The method, while not yielding closed form solutions, does permit the efficient computation of the delay, throughput and packet loss performance of a network. Szymanski and Shaikh [5] extended Jenq's method to switching systems constructed from switches with an arbitrary number of inputs and an arbitrary number of buffer slots.

Turner [6] developed a similar method for switching networks with shared buffering. Recently, Pattavina and Monterosso [4] developed a new approach based on an exact model of a single shared-buffer switch element. This model, while highly accurate is computationally intractable for networks constructed from large switches. In this paper we describe a series of improvements to Turner's method which rival the accuracy of Pattavina and Monterosso's method while maintaining the computational effectiveness of Turner's method.

The analysis considered here is for delta networks (see [6]) constructed from switches with d input and

output ports. We use n to denote the number of network inputs and outputs and let $k = \log_d n$ denote the number of stages in the network. Each switch used to construct the network is assumed to have a single shared buffer with B buffer slots. Packets from any input can be placed in any available buffer slot and packets can proceed from any buffer slot to the desired output. Packets arriving at the inputs to the network are assumed to be assigned independent random output addresses.

Typically these systems are operated in a time-slotted fashion, with fixed length packets progressing synchronously from stage to stage. Low level flow control mechanisms regulate the flow of packets between stages to prevent packets from being lost due to buffer overflow. Consequently, we can think of the system as operating in two phases. In one phase, flow control information passes through the network from right to left. In the other, packets flow from left to right, in accordance with the flow control information. There are several types of flow control that can be used. We focus here on *local grant* flow control; other types are studied in [1]. Consider a shared buffer switch employing local/grant flow control and assume that the number of unoccupied buffer slots is x . The switch grants permission to send to $\min\{x, d\}$ of its upstream neighbors at the start of an operation cycle of the network. If $x < d$, we assume that x predecessors are chosen at random. Local/grant control is the easiest to implement and appears to have most practical interest while global/acknowledgement control provides the best performance.

Turner [6] models a network of shared buffer switches by modeling a shared buffer switch as a $B+1$ state Markov chain. We review this result briefly here. Let $\pi_i(s)$ be the steady state probability that a stage i switch contains exactly s packets and we let $\lambda(s_1, s_2)$ be the probability that a switch with s_1 packets during a given cycle contains s_2 packets in the subsequent

This work was supported by CEFRIEL, Ascom Timeplex, Bell Communications Research, Bell Northern Research, DEC, Goldstar Information & Communications, Italtel SIT, NEC, NTT and SynOptics.

cycle. In the steady state,

$$\pi_i(s_2) = \sum_{0 \leq s_1 \leq B} \pi_i(s_1) \lambda_i(s_1, s_2)$$

Let $p_i(j, s)$ be the probability that j packets enter a stage i switch that has s packets in its buffer and let $q_i(j, s)$ be the probability that j packets leave a stage i switch that has s packets in its buffer. Then

$$\lambda_i(s_1, s_2) = \sum_{0, s_2 - s_1 \leq h \leq d} p_i(h, s_1) q_i(h - (s_2 - s_1), s_1)$$

Let a_i be the probability that any given predecessor of a stage i switch has a packet for it. Then if we let $m = \min\{d, B - s\}$,

$$p_i(j, s) = \binom{m}{j} a_i^j (1 - a_i)^{m-j}$$

$$a_i = \sum_{0 \leq j \leq B} \pi_{i-1}(j) [1 - (1 - 1/d)^j]$$

We take a_1 to be the offered load. Note that the equation for a_i assumes that the destinations of the packets in an upstream switch are independent of one another.

Let b_i be the probability that a successor of a stage i switch provides a grant and let $Y_d(r, s)$ be the probability that a switch that contains s packets, contains packets for exactly r distinct outputs. An output of a switch is called *active* if the switch contains some packet which is to be sent out using that output, so $Y_d(r, s)$ is the probability that a d port switch with s packets has r active outputs.

$$q_i(j, s) = \sum_{j \leq r \leq \min\{d, s\}} Y_d(r, s) \binom{r}{j} b_i^j (1 - b_i)^{r-j}$$

$$b_i = \sum_{0 \leq h \leq B-d} \pi_{i+1}(h) + \sum_{0 \leq h \leq d-1} \frac{h}{d} \pi_{i+1}(B - h)$$

We assume that network outputs can always accept packets and model this by letting $b_k = 1$.

Y is easily calculated, if we assume all distributions of s packets to the d outputs are equally likely. This is just a classical distribution problem. For the purposes of calculation, we can use the recurrence

$$Y_d(r, s) = \frac{r}{d} Y_d(r, s - 1) + \frac{d - (r - 1)}{d} Y_d(r - 1, s - 1)$$

when $0 < r \leq s$, $Y_d(r, s) = 1$ when $s = r = 0$ and $Y_d(r, s) = 0$ otherwise. Note that $Y_d(r, s)$ is independent of the stage of the switch in the network. For computational purposes, it is most convenient to

merely pre-compute a table with the values of Y required. We compute performance parameters by assuming a set of initial values for $\pi_i(j)$, then use these and the equations given above to compute $\lambda_i(s_1, s_2)$. These, together with the steady-state equations for the Markov chain are used to obtain new values of $\pi_i(j)$ and we iterate until we obtain convergence. We can now easily obtain the performance metrics of interest. The carried load is given by

$$\sum_{j=0}^d \sum_{s=0}^B j q_k(j, s) \pi_k(s)$$

The packet loss rate is the ratio of the offered load minus the carried load to the offered load. This type of analysis is referred to as the *scalar method*.

The scalar method is reasonably fast, even for switches with large values of B and d . The computation is dominated by the time required to compute the transition probabilities λ_i . If m iterations are required for convergence, this is proportional to $m(B/d)d^3 \log_d n$, while the time required for simulation is proportional to $m'n \log_d n$, where m' is the number of simulation steps required to get accurate results. Typically $m \ll m'$ and the elementary step in the analysis is faster to compute than the elementary step in the simulation (which typically involves making some routing decision on a packet data structure and moving it from one queue to another). Moreover, for networks with more than 3 stages $d^3 < n$, so the analysis' computational advantage grows with network size. While convergence is not guaranteed, our experience has shown convergence to be fairly rapid except when the offered load is approximately equal to the network's maximum throughput.

Notice that the calculation of $Y_d(r, s)$ given above, and the calculation of a_i rely on the assumption that the addresses of the packets stored within a switch's buffer are independent. This is not in fact the case. While it is true that the addresses of packets arriving at a switch are independent (given the input traffic assumptions), buffered packets are correlated as a result of having contended for outputs. The correlations are strongest when d is small and B large. The independence assumption causes the analysis to overestimate the maximum throughput. This is illustrated in Figure 1 which shows the maximum load that can be achieved by single stage networks for various values of d and B . Note that the throughput predicted by the scalar method can exceed that of the analysis by more than 10%. Figure 2 shows the packet loss experienced in several single stage networks as a function of offered load. Note that the scalar method grossly

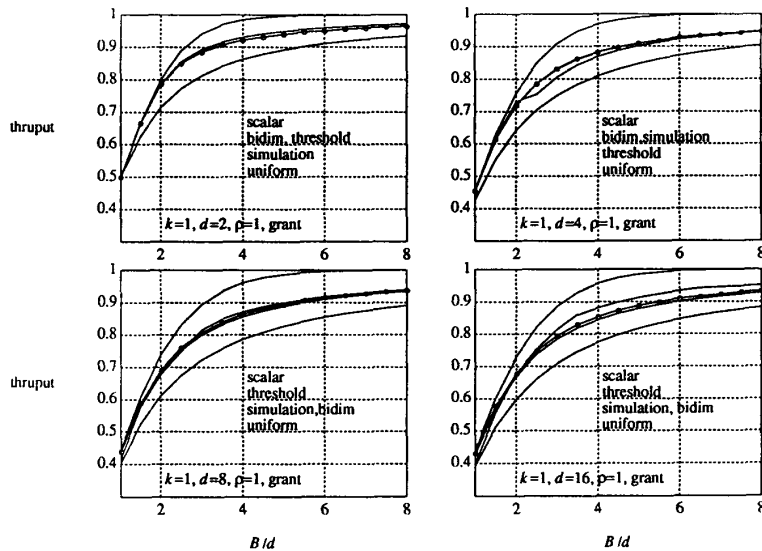


Figure 1: Throughput of Single Stage Networks

under-estimates the loss rate when compared to simulation.

Monterosso and Pattavina [4] have developed an exact model of a shared buffer switch which they call the *vector method*. In this approach, the state of a switch is represented by a vector $[s_1, \dots, s_d]$, where s_j denotes the number of stored packets for output j . The drawback of this method is of course that the number of states grows exponentially with d and the number of transition probabilities is roughly the square of the number of states. While Pattavina and Monterosso have taken care to avoid explicit representation of redundant states, a 16 port switch with 32 buffer slots still has about 30 thousand states and there are more than 10 million transition probabilities. In the scalar model, on the other hand, the same switch has 33 states and fewer than 1000 transition probabilities. These considerations limit the applicability of the vector method to networks constructed from small switches ($d \leq 4$). The objective of this paper is to develop an analytical method which rivals the accuracy of the vector method while retaining the computational effectiveness of the scalar method. We in fact describe three techniques. The first is a variant of the scalar method, which we call the *uniform scalar method*. The second method models the state of a switch using two variables and is called the *bidimensional method*. As we will see, the bidimensional method is the most accurate but also more computationally expensive. Consequently, we have developed

a third method called the *interval method* which is intermediate in both accuracy and speed.

2 The Uniform Scalar Method

The scalar method of queuing analysis computes the number of active outputs that a switch has using the assumption that if s packets are stored in a switch's buffer, that the outputs those packets are to take are independent of one another. To understand the implications of this assumption, it's helpful to compare the vector and scalar methods for binary switches ($d = 2$). Figure 3 shows the Markov chain corresponding to the vector model of a two port switch, with five buffer slots. In the illustration, a state (i, j) represents a switch in which i packets are destined for output 0 and j are destined for output 1. In the scalar method, a state corresponds to the sum $s = i + j$, so that the states of the scalar method correspond to sets of states in the vector method that lie along a common diagonal, as shown in the figure.

If a switch is in state (i, j) (vector method) where i or j but not both are equal to zero, then there is one active output. We refer to these as boundary states. The scalar method calculates the probability of these boundary states relative to the non-boundary states along a given diagonal by assuming that the packet destinations are independent. This can be interpreted as assigning probabilities to the states along a given




Figure 2: Packet Loss in Single Stage Networks

diagonal according to a binomial distribution. In particular, given that there are s packets in a switch, the probability assigned to state (i, j) is $\binom{s}{i}(1/2)^s$ (where $s = i + j$).

Comparison of the binomial probabilities with the frequency distributions observed during simulation runs (see [1]) reveals that the binomial assignment grossly under-estimates the probability of the boundary states ($i = 0$ and $i = 14$). It also makes clear that the relative frequencies of the states on the diagonal are well-approximated by a uniform distribution, that is equal values for all probabilities along a given diagonal. An intuitive justification for this can be found by considering the steady-state situation for a heavily loaded switch. In this case, one would expect that the state of a switch would move back and forth along a diagonal as packets come and go. While there would be some movement across diagonals as well, the predominant movement would be along a diagonal. Since there is equal probability of moving in either direction along a diagonal, one would expect the distribution along a diagonal to be approximately uniform. This reasoning leads to the *uniform scalar method* in which, when determining the number of active outputs that a switch has, we assume that the set of states of the underlying vector method corresponding to a given state of the scalar method are equally likely.

The equations for the uniform method are similar to those for the original scalar method. The only equa-

tions that must be changed are the ones for a_i (the probability that the predecessor of a stage i switch has a packet for it) and $q_i(j, s)$ (the probability that j packets leave a stage i switch when in state s).

$$a_i = \sum_{0 \leq s \leq B} \pi_{i-1}(s)(1 - \sigma(s, d-1)/\sigma(s, d))$$

$$q_i(j, s) = \sum_{j \leq r \leq d, s} \binom{d}{r} \frac{\sigma(s-r, r)}{\sigma(s, d)} \binom{r}{j} b_i^j (1 - b_i)^{r-j}$$

where $\sigma(s, d)$ is the number of ways d distinct non-negative integers can add to s .

$$\sigma(s, d) = \binom{s + d - 1}{d - 1}$$

The values of σ can be pre-computed and stored in a table, so the amount of computation required for the uniform method is essentially the same as for the original scalar method. Figure 1 shows the maximum throughput for single stage switches computed using the uniform method. The uniform method *under-estimates* the maximum throughput of these networks by a substantial margin and so is only marginally more useful than the original method. Nonetheless, as we shall see in the next section, it provides the basis for another method that is considerably better.




Figure 3: Comparison of Vector and Scalar Methods for Binary Switches

3 The Bidimensional Method

The two scalar methods attempt to use a single number to summarize the state of a switch. To obtain better accuracy, the bidimensional model keeps track of both the number of packets in a switch and the number of active outputs.

Let the state of a switch be a pair (s, c) where s is the number of packets and c is the number of active outputs. Then

$$\pi_i(s_2, c_2) = \sum_{(s_1, c_1)} \pi_i(s_1, c_1) \lambda_i((s_1, c_1), (s_2, c_2))$$

$$\lambda((s_1, c_1), (s_2, c_2)) = \sum_{0 \leq j \leq d} p_i(j, s_1) q_i(s_1 + j - s_2, c_1) R(s_1, c_1, j, s_1 + j - s_2, c_2)$$

where $R(s_1, c_1, j_1, j_2, c_2)$ is the probability that if a switch is initially in state (s_1, c_1) and then j_1 packets arrive and j_2 leave, that the new state is $(s_1 + j_1 - j_2, c_2)$. To compute R we must make some assumption about how the packets are distributed among the various outputs. Based on the reasoning discussed in the last section, we use a uniform assumption; that is, that all ways of distributing the s_1 packets among the c_1 active outputs are equally likely. We can compute R using a recurrence

$$R(s_1, c_1, j_1, j_2, c_2) = \sum_{0, c_1 - j_2 \leq h \leq c_1} R(s_1, c_1, 0, j_2, h) R(s_1 - j_2, h, j_1, 0, c_2)$$

Note that $R(s_1, c_1, 0, j_2, h)$ is independent of j_1 and that $R(s_1 - j_2, h, j_1, 0, c_2)$ is independent of both s_1

and j_2 . So if we define $S(s, c, j, c') = R(s, c, 0, j, c')$ and $T(c, j, c') = R(x, c, j, 0, c')$, the sum above is equal to

$$\sum_{0, c_1 - j_2 \leq h \leq c_1} S(s_1, c_1, j_2, h) T(h, j_1, c_2)$$

where

$$S(s, c, j, c') = \sum_{0, 2c - s \leq r \leq c} \binom{c}{r} \frac{\sigma(s - 2c + r, c - r)}{\sigma(s - c, c)} \frac{\binom{r}{c - c'} \binom{c - r}{j - c + c'}}{\binom{c}{j}}$$

$$T(c, j, c') = \frac{c'}{d} T(c, j - 1, c') + \left(1 - \frac{c' - 1}{d}\right) T(c, j - 1, c' - 1)$$

$$T(c, j, c) = (c/d)^j$$

It remains only to give the equations for p and q .

$$p_i(j, s) = \mathbf{B}(j, \min\{d, B - s\}, a_i)$$

$$a_i = \sum_{\substack{0 \leq s \leq B \\ 0 \leq c \leq d, s}} \frac{c}{d} \pi_{i-1}(s, c)$$

$$q_i(j, c) = \mathbf{B}(j, c, b_i)$$

$$b_i = \sum_{0 \leq c \leq d} \left[\sum_{0 \leq s \leq B - d} \pi_{i+1}(s, c) + \sum_{B - d < s \leq B} \pi_{i+1}(s, c) (B - s)/d \right]$$

where $\mathbf{B}(j, n, p) = \binom{n}{j} p^j (1 - p)^{n-j}$.

For computational purposes, S and T can be pre-computed and stored in a table. S is the larger of the

two and consumes 2 Mbytes for $d = 16$, $B/d = 4$. In fact, one can go further and pre-compute R . This will save time during the iterative computation of the steady-state probabilities but requires substantially more memory (32 Mbytes for $d = 16$, $B/d = 4$).

The maximum throughput predicted by the bidimensional method is shown in Figure 1. Notice that it closely tracks the simulation results, over-estimating slightly for small d and under-estimating slightly for large d . Figure 2 shows the loss predicted by the bidimensional method. Again note the close match with the simulation results.

4 Interval and Threshold Methods

While the bidimensional method offers a big improvement in accuracy over the scalar methods, it has a substantially higher computational cost. We can reduce the computational cost by keeping less detailed information about the number of active outputs. The *interval method* divides the range $[0, d]$ into two or more intervals and keeps track of the probability that the number of active outputs is in a particular interval.

Let (s, C) denote a switch state where $C = [\ell, h]$ is an interval. Then,

$$\begin{aligned} \pi_i(s_2, C_2) &= \sum_{(s_1, C_1)} \pi_i(s_1, C_1) \lambda_i((s_1, C_1), (s_2, C_2)) \\ \lambda_i((s_1, C_1), (s_2, C_2)) &= \\ &\sum_{c_1 \in C_1} Z(c_1, s_1, C_1) \sum_{0 \leq j \leq d} p_i(j, s_1) \\ &q_i(s_1 + j - s_2, c_1) R'(s_1, c_1, j, s_1 + j - s_2, C_2) \end{aligned}$$

where $Z(c, s, C)$ is the probability that there are exactly c active outputs, given that the state is (s, C) and where $R'(s_1, c_1, j_1, j_2, C_2)$ is the probability that if a switch initially has s_1 packets and c_1 active outputs and then j_1 packets arrive and j_2 leave, that the new state is $(s_1 + j_1 - j_2, C_2)$. We can compute R' and Z as follows.

$$\begin{aligned} R'(s_1, c_1, j_1, j_2, C) &= \sum_{c \in C} R(s_1, c_1, j_1, j_2, c) \\ Z(c, s, C) &= \binom{s}{c} \sigma(s - c, c) / \sum_{\ell \leq r \leq h, s} \binom{s}{r} \sigma(s - r, r) \end{aligned}$$

where R is defined as in the previous section. The equations for p and q are similar to the ones for the bidimensional method.

$$p_i(j, s) = \mathbf{B}(j, \min\{d, B - s\}, a_i)$$

$$\begin{aligned} a_i &= \sum_{\substack{0 \leq s \leq B \\ \forall C, \forall c \in C}} Z(c, s, C) \frac{c}{d} \pi_{i-1}(s, c) \\ q_i(j, c) &= \mathbf{B}(j, c, b_i) \\ b_i &= \sum_C \left[\sum_{0 \leq s \leq B-d} \pi_{i+1}(s, C) \right. \\ &\quad \left. + \sum_{B-d < s \leq B} \pi_{i+1}(s, C) (B - s) / d \right] \end{aligned}$$

If the number of intervals is not too large R' can be pre-computed and stored. We define the *threshold method* to be the particular case of the interval method in which there are just two intervals. In this case, the table for R' requires 4 Mbytes for $d = 16$, $B/d = 4$.

Figure 1 shows the maximum throughput predicted by the threshold method. To produce these results, a variety of thresholds were tried and the results for the thresholds that produced the closest match to the simulation results were plotted. It's surprising that a model with just twice the state complexity of the scalar models can produce results that so closely match simulation.

5 Closing Remarks

We have studied the running time required for the various analysis methods compared to the running time for simulation. For an n port network, let $k = \log_d n$ be the number of stages. Consideration of the equations for each of the methods studied shows that for large values of d and $B/d \geq 2$, the running time can be approximated by $Cmk \left(\frac{B}{d}\right) d^r$ where C is the time for a basic step in the computation, m is number of iterations required for convergence and $r = 3$ for the scalar methods, $r = 4$ for the threshold method and $r = 6$ for the bidimensional model, assuming that S and T are pre-computed. From experimental measurements, we've found that C is typically around 10 μ s on a SparcStation 2 and m is typically around 100 μ s.

We can do a similar analysis of the running time for a switch simulation. The simulation time takes the form $C'm'kd^k$ where C' is the time to simulate one time step at one switch element divided by d and m' is the number of iterations required for the simulation. For our simulation, we have found that C' is around 50 μ s and m' can be anywhere from 10^4 to 10^6 depending on the data of interest. (This of course, is for

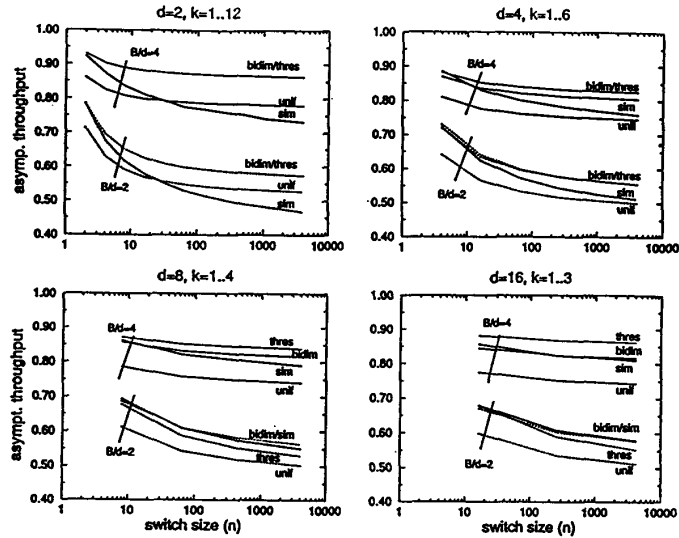


Figure 4: Throughput of Larger Networks

uniform random traffic with Bernoulli arrivals.) For throughput measurements, we have found $m = 10^4$ to be adequate for obtaining statistically reliable results. Using these equations to compare running times shows that for large networks, even the most computationally expensive of the analytical methods is substantially faster than simulation. See [1] for details.

Figure 4 shows the maximum throughput for multistage networks constructed from switches of various sizes. These curves show that as the number of stages in a network increases, all the analysis methods overestimate the throughput. This can be explained by the fact that none of the methods account for correlations between the states of switches in different stages. This is a common failing for this entire class of methods, including the vector method of Monterosso and Pattavina. These correlations are strongest for networks with lots of stages and small switches. For networks constructed from larger switches, the inter-stage correlations are relatively weak, allowing more reliable throughput predictions. In particular, we see that the bidimensional method gives acceptable results for switches with $d \geq 8$. For the threshold method, the thresholds used here were the ones found to give the best overall accuracy in the single stage case. The threshold method's accuracy is sensitive to the amount of buffering in a switch.

We have applied the bidimensional analysis method to determine packet loss in a network with mixed results. In realistic system configurations, a multistage

network is typically preceded by a collection of input buffers with flow control between the network and the input buffers. Packet loss can occur at these input buffers as the load offered to the network approaches its maximum throughput. The bidimensional method can accurately predict packet loss for single stage networks with no input queues (as shown in Figure 2) or very small input queues. However, when larger input queues are added (≥ 8 slots), the accuracy rapidly deteriorates. This is unfortunate since the ability to accurately predict packet loss with a fast analytical method would offer great advantages over simulation. We have considered two possible explanations for the inaccuracy of the bidimensional method in the presence of input queues. The first is that the extension of the bidimensional method to model the input queues, treats the different queues feeding a single switch as independent of one another, where in fact their states are strongly correlated. Unfortunately, an alternative model of the input queue which captures these correlations provides only a marginal improvement in the prediction of packet loss. The second (and more important) explanation was that the analysis ignores time correlations that become increasingly significant as the amount of input buffering is increased. Another model was developed to attempt to capture the effects of these time correlations and while the results were far more accurate, the complexity of the model makes it of little practical value.

In summary, we have developed several methods

of analyzing networks constructed from shared buffer switches. The bidimensional method yields excellent predictions of throughput for networks constructed from large switches. The interval method can approach the accuracy of the bidimensional method but is sensitive to the amount of buffering and the choice of intervals. We have concentrated on maximum throughput, as this is the performance metric of most importance in switching system design. In high speed switching applications, so long as a network is operated below its maximum throughput, the queueing delay will be satisfactory and packet loss can be reduced to acceptable levels by engineering the input buffers appropriately. Methods that can accurately predict packet loss for networks preceded by input buffers could be useful, but in practice, conservative engineering rules for these buffers can be applied with little impact on system cost.

These methods do not model inter-stage correlations which become significant in large networks constructed from small switches. This is one possible direction for future research. Another is in extending the methods for nonuniform traffic as has been done for other types of analysis in [3]. Improving the computational performance of these methods would be very useful. The bidimensional method, in particular, can be very time-consuming. One way to improve the computational performance is to modify the basic iterative algorithm. The basic algorithm computes the transition probabilities λ_i from the current state probabilities (π_i) and then computes new state probabilities from the previous values by a single application of the balance equations for π_i . By applying the balance equations multiple times each time we compute values of λ_i , we can reduce the number of iterations significantly. We have found that this variation of the basic method can reduce the overall computation time by an order of magnitude.

We have not fully explored the possibilities of the interval method. We suspect that by using three or four intervals one could match the accuracy of the bidimensional method with a lot less computational effort. There remains the problem of selecting the intervals. Our approach of using simulation to "tune" the method is not fully satisfactory. The interval method can also be extended so that the intervals are dependent on the number of packets in a switch. For example, when just two intervals are used (the threshold method), we would expect better results if the threshold increases as the number of stored packets increases. Ideally, the probability associated with the two intervals for any given value of s should be approximately

equal. This objective can be approximated by selecting the interval so as to divide the corresponding states of the underlying vector model equally. That is, for each value of s , we select a threshold t for which

$$\sum_{c < t} \sigma(s, c) \approx \sum_{c \geq t} \sigma(s, c)$$

Alternatively, the threshold for each value of s could be adjusted dynamically to make these probabilities as nearly equal as possible.

References

- [1] Bianchi, Giuseppe and Jonathan S. Turner. "Improved Queueing Analysis of Shared Buffer Switching Networks," Washington University Computer Science Department, wucs-92-19.
- [2] Jenq, Yih-Chyun. "Performance Analysis of a Packet Switch Based on a Single-Buffered Banyan Network," *IEEE Journal on Selected Areas in Communications*, vol. SAC-1, no. 6, 12/83, 1014-1021.
- [3] Kim, Hyong Sok and Alberto Leon-Garcia. "Performance of Buffered Banyan Networks under Nonuniform Traffic Patterns," *Proceedings of Infocom 88*, 4/88.
- [4] Monterosso, Alberto, and Achille Pattavina. "Performance Analysis of Multistage Interconnection Networks with Shared Buffered Switch Elements for ATM Switching," *Proceedings of Infocom 92*, 5/92.
- [5] Szymanski, Ted and Salman Shaikh. "Markov Chain Analysis of Packet-Switched Banyans with Arbitrary Switch Sizes, Queue Sizes, Link Multiplicities and Speedups," *Proceedings of Infocom 89*, 4/89.
- [6] Turner, Jonathan S. "Queueing Analysis of Buffered Switching Networks," *Proceedings of the International Teletraffic Congress*, June, 1991.