

A COMPUTATIONALLY EFFICIENT ALGORITHM FOR CHARACTERIZING THE SUPERPOSITION OF MULTIPLE HETEROGENEOUS INTERRUPTED BERNOULLI PROCESSES

Khaled M. Elsayed and Harry G. Perros

*Department of Computer Science, North Carolina State
University, Raleigh, NC 27695-8206, USA*

ABSTRACT

A computationally efficient algorithm for characterizing the superposition process of N heterogeneous and independent Interrupted Bernoulli Processes is introduced. The algorithm is then used to analyze a statistical multiplexer with finite buffer. Finally, numerical examples highlighting the algorithm accuracy are given.

1 INTRODUCTION

In an ATM environment many types of traffic, such as voice, data, and video, are to be efficiently transported by the same network. An ATM multiplexer receives cells (fixed size packets of 53 octets length) from a number of different incoming links and then transmits them out onto a single outgoing link. A finite buffer is provided in the multiplexer to accommodate the multiple arrivals of cells. Each arrival stream is modeled by a bursty and possibly a correlated process. The service time is deterministic and is equal to one slot of the outgoing link which is assumed to be long enough to transmit one cell.

The analysis of such a queueing system is quite complex due to the large number of arrival processes. A possible method for approximately analyzing the queue is to first characterize the superposition process of all arrival processes, and then analyze the queue with a single arrival process. In this paper, we con-

sider the case when the arrival processes are modeled as multiple independent heterogeneous Interrupted Bernoulli Processes (IBP). The IBP and its variants are popular models for bursty traffic sources in an ATM environment.

The problem of characterizing the superposition process of a set of arrival processes has been addressed extensively in the literature. One approach for obtaining the superposition process is to approximate it by a renewal process, see Albin [1], Whitt [16], Sriram and Whitt [12], and also Perros and Onvural [11]. An insightful discussion of the various time-scales affecting the accuracy of the approximate superposition of packet voice sources (modeled as a variant of IBP) was provided in [12]. Heffes and Lucantoni [6] considered the superposition process of packet voice sources modeled as an IBP where arrivals occur periodically. They approximate the superposition by a Markov Modulated Poisson Process (MMPP). The accuracy of the superposition is reasonable when the average delay in the multiplexer is the amount of paramount importance. However, the superposition does not provide a good estimate for the probability of loss. Several other authors (see [3, 10, 15]) considered the same problem and suggested alternative methods for characterizing the superposition process as an MMPP. The main objective of these papers was to improve the accuracy with regards to calculation of the cell loss probability. Heffes [5] obtained an MMPP approximation to the superposition of different MMPP arrival processes using a set of simple expressions.

An alternative method to model ATM multiplexers is the Uniform Arrival and Service (UAS) model (also known as fluid flow). In this case an on/off process produces a uniform flow of bits when in the on state. Cell departures are modeled as a uniform flow out of the queue. Anick, Mitra, and Sondhi [2] evaluated the system performance using elegant simple expressions for a multiplexer with infinite buffer space and homogeneous arrivals. Tucker [14] considered the finite buffer case. The methodology was generalized to the case of heterogeneous Markov Modulated Rate Processes in a series of papers [13, 4].

The work presented in this paper is related to the methodology presented Hong, Perros, and Yamashita [7] and Makhameh, McDonald and Georganas [9]. In [7], the authors used an aggregation method for approximating the superposition process. They first obtained the exact probability transition matrix of the Markov chain of the superposition process. This Markov chain has a dimension which is exponential in the number of input sources. Subsequently, the Markov chain was aggregated in order to obtain a superposition process with a small number of states (linear function of the number of sources). The method is limited to a small number of input sources. In [9], the authors used an ag-




Figure 1 The Interrupted Bernoulli Process.

gregation method to analyze an output-buffered ATM switch with correlated imbalanced traffic.

In this paper, we present an efficient algorithm for characterizing the superposition process of multiple heterogeneous IBP's. The algorithm is based on a step-wise aggregation scheme. It can be used to construct the superposition of a large number of IBP sources. The space complexity of the algorithm is $O(N)$, while the computational complexity is $O(N^4)$, where N is the number of superimposed IBP's. The superposition algorithm is then used to study the performance of a statistical multiplexer with finite buffer. Finally, we present a brief study of the range of parameters of the superimposed IBPs for which the accuracy of the algorithm is acceptable.

The rest of this paper is organized as follows. In section 2, we briefly describe the IBP model and discuss some of its statistical properties. In section 3, we discuss the step-wise aggregation algorithm for characterizing the superposition process. Section 4 presents the analysis of the statistical multiplexer and numerical results and a discussion of the algorithm accuracy are discussed in section 5. Section 6 concludes the paper.

2 THE IBP SOURCE MODEL

We consider a source that alternates between active and idle states according to a Markov chain (see figure 1). Arrivals occur in a Bernoulli fashion with probability γ only when the source is in the active state. No arrivals occur when the source is in the idle state. The transitions between active and idle states occur in a memoryless fashion. Let us assume that at the end of slot k the process is in the active (or idle) state. Then, in the next slot $k + 1$ it will remain in the active (or idle) state with probability α (or β), or it will change to idle (or active) state with probability $1 - \alpha$ (or $1 - \beta$).

The mean inter-arrival time \bar{t} and coefficient of variation CV^2 are given by:

$$\bar{t} = \frac{(1-\alpha) + (1-\beta)}{\gamma(1-\beta)}, \quad CV^2 = 1 + \gamma \left[\frac{(1-\alpha)(\alpha+\beta)}{[(1-\alpha) + (1-\beta)]^2} - 1 \right].$$

The triplet (α, β, γ) completely characterizes an IBP process.

3 THE APPROXIMATION ALGORITHM

We now present an efficient algorithm that can be used to characterize the superposition of $N \geq 2$ heterogeneous IBP sources. The sources are divided into G groups, where the sources within each group are identical. In practice, specially when N is large, this assumption is valid because traffic sources of a specific traffic type (e.g. data or voice source) would tend to have similar or even identical characteristics. The number of sources in group i is given by N_i , where $1 \leq N_i \leq N$ and $\sum N_i = N$.

I- The Superposition Process of Homogeneous IBP Sources: The first step in the approximation algorithm is to characterize the superposition of sources in group i , where $N_i > 1$, in terms of a stochastic process with $N_i + 1$ states. The superposition process is described in terms of the probability transition matrix between the $N + 1$ states, and the probability distribution of the number of arrivals per slot at each state.

The superposition process of N similar IBP's with descriptor (α, β, γ) is a Markov chain with state space $\chi = \{n, n = 0, 1, \dots, N\}$ which denotes the number of sources in the active state (or alternatively in the idle state). Let $A = [a(n_1, n_2)]$, $n_1, n_2 \in \chi$ be the probability transition matrix governing the transitions between the states of the superposition. Also, let $b(n, i)$ be the probability that i arrivals occur when the superposition is in state n , where $0 \leq i \leq n$. Consider two states n_1 and n_2 . Assume that in slot k the Markov chain is in state n_1 . The probability that in the next slot $k + 1$, the Markov chain would be in state n_2 is equal to the probability that l , $0 \leq l \leq n_1$, of the n_1 sources in the active state make a transition to idle state and that m , $0 \leq m \leq N - n_1$, of the sources in the idle state make a transition to active state, and that $n_1 - l + m = n_2$. The probability of the first event to occur is given by the binomial probability density function $\binom{n_1}{l} (1-\alpha)^l \alpha^{n_1-l}$. Similarly, the probability that the second event occurs is given by $\binom{N-n_1}{n_2-n_1+l} (1-\beta)^{n_2-n_1+l} \beta^{N-n_2-l}$ where $0 \leq n_2 - n_1 + l \leq N - n_1$ and $N - n_2 - l \geq 0$. We

then have:

$$a(n_1, n_2) = \sum_{l=0}^{n_1} \left[\binom{n_1}{l} (1-\alpha)^l \alpha^{n_1-l} \binom{N-n_1}{n_2-n_1+l} (1-\beta)^{n_2-n_1+l} \beta^{N-n_2-l} I(n_2 - n_1 + l \geq 0) I(N - n_2 - l \geq 0) \right]$$

where $I(x) = 1$ iff logical expression x is true and 0 otherwise.

When there are n sources each with probability γ of emitting a cell in a given slot, the the probability of having v arrivals is given by the binomial distribution $b(n, v) = \binom{n}{v} \gamma^v (1-\gamma)^{n-v}$, $0 \leq v \leq n$. For notational convenience, the quantities $b(n, v)$ are referred to by the lower triangular matrix $B = [b(n, v)]$.

II- The Combined Superposition Process of two Groups of Sources: Once the superposition of two or more separate groups of sources is constructed as described above, it is necessary to combine them in order to get the overall superposition process. Let us assume that the superposition process of two arbitrary groups, say groups 1 and 2, has been characterized in terms of matrices A_i and B_i , $i = 1, 2$, and that the number of states in process i is equal to $N_i + 1$. The superposition process is described by the state (n_1, n_2) , where n_i is the state of component process i . The number of states in the superposition process is equal to $(N_1 + 1)(N_2 + 1)$. The probability transition matrix of the superposition will be given by $A = A_1 \otimes A_2$ where \otimes is the Kronecker product operation of two matrices.

One of the useful properties of the Kronecker product operation is that $\bar{\pi}$, the invariant probability vector of A , is given by $\bar{\pi}_1 \otimes \bar{\pi}_2$ where $\bar{\pi}_i$ is invariant probability vector of A_i . This saves us from solving a system of linear equations for finding the invariant of P which is more time consuming than solving for $\bar{\pi}_1$ and $\bar{\pi}_2$ individually.

Note that the above discussion is also valid for more than two groups of sources. However, we focus here on the case of two groups since the approximation algorithm always superposes two processes at a time.

III- Aggregation of the Superposition Process: The above Markov chain which characterizes the superposition process of two groups of IBP sources is exact. However, the dimensionality of the resulting process may not be practical for studying the performance of a multiplexer with a large buffer. The dimensionality of the superposition process can be reduced by lumping all the states (n_1, n_2) of the superposition, where $n_1 + n_2 = n$, to a single state n . The Markov chain resulting from this aggregation represents what we call the *compact* superposition process.

Let the probability transition matrix of the compact process be $\tilde{A} = [\tilde{a}(n, m)]$. We then have

$$\tilde{a}(n', n'') = \frac{\left[\sum_{n'_1+n'_2=n'} \sum_{n''_1+n''_2=n''} \pi(n'_1, n'_2) a((n'_1, n'_2), (n''_1, n''_2)) \right]}{\left[\sum_{n'_1+n'_2=n'} \pi(n'_1, n'_2) \right]} \quad (25.1)$$

The remaining step is to find the probability distribution function of the number of arrivals in a particular state. This can be calculated from $B_i, i = 1, 2$ as follows:

$$b(n, v) = \frac{\left[\sum_{n_1+n_2=n} \sum_{m=0}^{\min(v, n_1)} \pi(n_1, n_2) \left[b_1(n_1, m) b_2(n_2, v-m) I(0 \leq v-m \leq n_2) \right] \right]}{\left[\sum_{n_1+n_2=n} \pi(n_1, n_2) \right]} \quad (25.2)$$

The resulting superposition process has a characterization identical to that of its component processes. This leads to the following simple iterative scheme.

IV- The Approximation Algorithm: Consider N heterogeneous IBP sources divided into G groups each having $N_i, i = 1, 2, \dots, G$, identical sources. Then, we have the algorithm outline as follows:

- ◊ Find $A_j, \tilde{\pi}_j$ and B_j for all groups $j, j = 1, \dots, G$
- ◊ Let $A_s = A_1$ and $B_s = B_1$
- ◊ For $j = 2$ to G do
 - ◊ $A' = A_s \otimes A_j$
 - ◊ $\tilde{\pi}' = \tilde{\pi}_s \otimes \tilde{\pi}_j$
 - ◊ Aggregate A' into A_s (use equation 25.1)
 - ◊ Calculate B_s (use equation 25.2)
 - ◊ Find π_s satisfying $\tilde{\pi}_s A_s = \tilde{\pi}_s, \sum_i \tilde{\pi}_s(i) = 1$

End

It can be shown that the computational complexity of the algorithm is $O(N^4)$ (The calculation of the Kronecker product of two matrices and the solution of a linear system of equations are the dominant factors.)

Hong, Perros, and Yamashita [7] tackled the same problem. However, their proposed algorithm constructs the complete probability transition matrix of the superposition process and then performs the aggregation. This causes the computational complexity and storage requirements of their method to be $O(2^N)$. This precludes the execution of the algorithm on a conventional computer. Even for a nominal value of N , say $N = 20$, it may not be possible to construct the superposition even using a supercomputer. We conducted a set of experiments where we used identical input mix and multiplexer parameters as the test cases reported in an earlier version of [7] and the achieved level of accuracy relative to simulation is the same. However, our method is considerably computationally more efficient.

4 ANALYSIS OF THE MULTIPLEXER MODEL

Consider a FIFO finite buffer multiplexer serving $N \geq 2$ IBP sources. The multiplexer has $S \geq 1$ servers and can accommodate a total of $B \geq S$ cells at any time instant including those in service. The service time for all cells is constant and is equal to one time slot. The multiplexer can serve S cells every time slot. We assume that $N > S$, otherwise no queue will ever form in the multiplexer and the problem will be trivial to handle.

We seek the steady state probabilities, $\pi(n)$, $0 \leq n \leq B - S$, that there are n cells in the multiplexer's queue. From this, we can obtain other measures of interest such as the mean queue length, the probability of full buffer and the cell loss probability.

Let us first discuss the timing of events in our system. We follow an *early arrival* timing model as defined by Hunter [8]. That is, during an arbitrary time slot, the following sequence of events is possible: potential state transition in the superposition occurs, followed immediately by cell arrivals (if any), which is followed by service of waiting cells if there are any, and finally departure of cells that received service. If an arbitrary cell sees one or more empty servers upon its arrival, it is immediately admitted to one of the available servers

without waiting till the next slot. Hence, the multiplexer is effectively of the cut-through type.

The system state is described by the pair (n, h) where n is the number of customers in the system immediately after the end of a slot and h is the state of the superposition process at the current slot.

We first characterize the superposition process of the N IBP sources as described in the previous section. Then, the probability transition matrix of the Markov chain (n, h) is generated. Let $p[(n_1, h_1) \rightarrow (n_2, h_2)]$ be the probability that the Markov Chain (n, h) makes a transition from state (n_1, h_1) in slot k to state (n_2, h_2) in slot $k + 1$. Then

$$p[(n_1, h_1) \rightarrow (n_2, h_2)] = \sum_v b(h_1, v)a(h_1, h_2)$$

where $n_2 = (\min(n_1 + v, B) - S)^+$ and v , $0 \leq v \leq h_1$, is the random variable representing the number of arrivals when the superposition process is in state h_1 . Finally, we solve for the invariant probability vector of $P = [p[(n_1, h_1) \rightarrow (n_2, h_2)]]$ from which various performance metrics can be obtained. We used the Gauss-Seidel method when using a workstation, and the parallelized LAPACK routines for LU decomposition when using a KSR-1 parallel machine at the North Carolina Supercomputer Center.

5 VALIDATION AND NUMERICAL EXAMPLES

The approximation algorithm was validated by comparing it against simulation. The confidence intervals of our simulation were very narrow. Most of the figures are plotted without showing the confidence intervals for this reason. We only show the confidence intervals for the most stringent case of a multiplexer with an expected small cell loss probability. This will be addressed in example 6. The effect of the heterogeneity of the sources on the accuracy of the algorithm was investigated and metrics were defined to determine how heterogeneous a given set of sources is. These metrics can be used to determine in advance whether the superposition will be accurate or not.

In all the validation examples (except for example 6), we fix γ_j , the probability that source j emits a cell while in the active state, to be 1. This simplifies the

fitting procedure for characterizing an IBP from the first two moments of the inter-arrival time. The following additional notation is used.

- N_i = Number of sources at group i .
- ρ_i = Source throughput for a source in group i .
- R_ρ = $\max\{\rho_i\}/\min\{\rho_i\}$ is a measure of heterogeneity of the throughput of the input sources.
- CV_i^2 = Squared coefficient of variation of inter-arrival time for a source in group i .
- R_{cv^2} = $\max\{CV_i^2\}/\min\{CV_i^2\}$ is a measure of heterogeneity of the CV^2 of the input sources.
- λ = $\sum_{i=1}^G N_i \rho_i$ = Total throughput.
- B = Total buffer size.

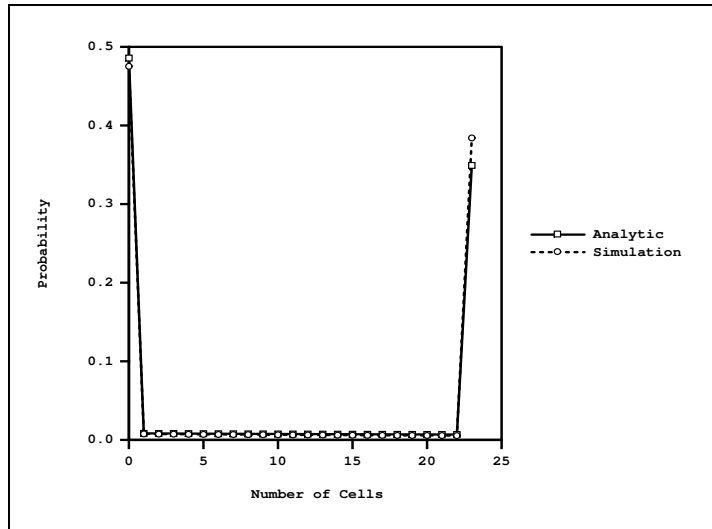
In all examples, the number of servers in the multiplexer was set to 1.

Example 1: In the first example, we have $G = 4$ and $B = 24$ and we use the input sources with parameters given in table 1.


i	N_i	ρ_i	CV_i^2
1	2	0.12	50.00
2	2	0.12	100.00
3	2	0.12	200.00
4	2	0.12	500.00

Table 1 Parameters for the first example.

In figure 2(a), we plot the distribution of the number of cells in the queue which is obtained from simulation and the approximation algorithm. The difference between the numerical and simulation values is shown in figure 2(b). The reason why we give the difference and not the relative error is that some of the calculated quantities are very small (in the order of 10^{-3}) which makes the relative error to be large, despite the fact that the overall accuracy is still acceptable. The probability distribution has two peaks, the first is for the probability that queue is full and the second peak for the probability that the queue is empty. This a well-known behavior in queues with bursty correlated input with a moderate to a high input traffic intensity. We note that the magnitude of error of the steady-state probability of a full buffer is high. This causes an underestimation of the cell loss probability. However, it should be stated that for this example the input sources are quite heterogeneous. The value of R_{cv^2} is 10, which is fairly large as it will be shown below.



(a) Probability distribution of the number of cells



(b) Error

Figure 2 Probability distribution of the number of cells and associated error, $B = 24$.

Example 2: In the second example, we use the input sources with parameters given in table 2, with $R_{cv2} = 4$. We first obtain the superposition process, and subsequently use it to generate the probability transition matrix for buffer sizes in the set $\{10, 20, 30, 40, 50, 60, 70\}$.

i	N_i	ρ_i	CV_i^2
1	2	0.1	10.00
2	2	0.1	20.00
3	2	0.1	30.00
4	2	0.1	40.00

Table 2 Parameters for the second example.


The cell loss probability and the mean queue length are shown in figures 3(a) and 3(b) respectively. It can be seen that the accuracy of the estimated mean queue length is very good while that of the cell loss probability is less accurate. This is due to the smoothing effect of the aggregation algorithm. It constructs a less bursty and correlated process than the actual superposition, which leads to underestimating the probability of observing a full buffer. However, we noted that the aggregation algorithm provides an upper-bound on the probability of having an empty buffer. Thus a balancing effect takes place and the overall mean queue length predicted by the simulation and numerical solution is close.

We also investigated the effect of increasing the number of sources in a group while keeping the overall group throughput fixed. The parameters of table 2 were modified as follows. N_i is increased to 4, ρ_i is halved to 0.05, for all $i = 1, \dots, 4$ and the values of CV_i^2 remained the same. For this new set of parameters, the probability of loss and mean queue length slightly increased over the values plotted in figure 3. The accuracy remained the same.


Example 3: The third example is in essence similar to the second one. We let $R_{cv2} = 1.6$ and the parameters of the input sources are given by table 3.

The probability of loss and the mean queue length are shown in figures 4(a) and 4(b) respectively. We note here that the accuracy of the approximation algorithm is very good. This is due to the fact that the amount of heterogeneity in the input sources is small as indicated by the relatively small value of R_{cv2} .

We know in advance that the more heterogeneous the sources are, the worse the accuracy of the iterative aggregation algorithm would become. However, we wanted to study the regions for the parameters of the sources in which the accuracy of the algorithm is acceptable. This is explored further in examples




(a) Probability of loss




(b) Mean number of cells

Figure 3 Probability of loss and mean number of cells.



(a) Probability of loss



(b) Mean number of cells

Figure 4 Probability of loss and mean number of cells.

i	N_i	ρ_i	CV_i^2
1	2	0.1	10.00
2	2	0.1	12.00
3	2	0.1	14.00
4	2	0.1	16.00

Table 3 Parameters for the third example.

5 and 6. The conclusion of this study was that R_{cv2} is a major deciding factor in affecting the accuracy of the superposition process.

Example 4: In this example, we discuss the effect of the variation in the squared coefficient of variation of the inter-arrival time of the input sources on the accuracy. We constructed an input of four groups each having two sources, and fixed the source throughput to be 0.1. We vary the R_{cv2} to take values from the set $\{1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5\}$ respectively. We assume that $CV_1^2 = 10.0$, and for each value of R_{cv2} , we find the remaining CV_i^2 by solving the equations:

$$CV_i^2 / CV_{i-1}^2 = \theta, \quad i = 2, 3, 4$$

where θ is set equal to $\sqrt[3]{R_{cv2}}$.

We plotted the cell loss probability, obtained from the approximation algorithm and simulation, and the percentile error as a function of R_{cv2} in figures 5(a) and 5(b) respectively. We note that the relative error is an increasing pseudo-linear function of R_{cv2} . We observed that in general if $R_{cv2} > 3$, then the approximation algorithm does not perform satisfactorily. This seems to hold true independent of the maximum or minimum values of CV_i^2 . In some cases, the accuracy actually improves when the minimum and maximum value of the CV_i^2 are increased, while keeping the value of R_{cv2} fixed. This may be in contrast to what one would initially conjecture that the larger the absolute values of CV_i^2 the worse the algorithm would perform.

Example 5: In this example, we discuss the effect of variation in the throughput of the input sources. We consider four groups of sources each having a single source. The total throughput is fixed at 0.8 and each source has a squared coefficient of variation equal to 20. We let R_ρ take values from the set $\{1, 10, 20, 30, 40, 50\}$. The throughput of a single source $\rho_i, i = 1, \dots, 4$ is calculated by solving the equations:

$$\rho_i / \rho_{i-1} = \theta, \quad i = 2, 3, 4, \quad \sum_{i=1}^4 \rho_i = 0.8$$

where $\theta = \sqrt[3]{R_\rho}$.


We plotted the cell loss probability, obtained from the approximation algorithm and simulation, and the percentile error as a function of R_ρ in figures 6(a) and 6(b) respectively. We note here that the accuracy of the algorithm is less sensitive to variations in the throughput by many orders of magnitude than the coefficient of variation. Even for very large values of R_ρ , say 40, the relative error is about 10%. In addition, we have observed that the algorithm becomes less sensitive to R_ρ when there are more than one source in a given group. In other words, the case of one source per group is an extreme case in which the effect of heterogeneity in the throughput of the input sources is large.

Example 6: We now consider a realistic model of the multiplexer where the probability of loss is expected to be small. We have two groups of sources. The first group has $N_1 = 17$, and a source in the group is specified by a mean on period of 50 ($\alpha_1 = 0.98$), a mean off period of 200 ($\beta_1 = 0.995$), and the probability of an arrival during a slot in the on period, γ_1 , is equal to 0.1. The second group has $N_2 = 9$, and a source in the group is specified by $(\alpha_2, \beta_2, \gamma_2) = (0.96, 0.986666667, 0.2)$. a mean on period of 25, a mean off period of 75, and $\gamma_2 = 0.2$. The buffer size is increased from 20 to 200 in increments of 20.


We obtain the cell loss probability and the mean number of cells as obtained by the approximation algorithm and simulation. In this case, we show the confidence intervals for the simulation. We conducted a long simulation with 20 independent replications. A single run would be finished when all of the sources has undergone at least 100,000 state changes. The confidence intervals for the different variables are then evaluated using standard techniques. The results are plotted in figure 7. The results indicate the applicability of using the algorithm for operating regions with very small expected cell loss probability. In such regimes, simulation fails to provide reliable answers in a reasonable amount of time.

6 CONCLUSIONS

A computationally efficient algorithm for characterizing the superposition process of multiple possibly-heterogeneous groups of IBP sources was introduced. The accuracy of the algorithm depends to a large extent on the degree of heterogeneity of the sources. We found out that a good measure of the heterogeneity




(a) Probability of loss




(b) Percentile relative error

Figure 5 Probability of loss and relative error vs. R_{cv2} .




(a) Probability of loss




(b) Percentile relative error

Figure 6 Probability of loss and relative error vs. $R\rho$.



(a) Probability of loss



(b) Mean number of cells

Figure 7 Probability of loss and mean number of cells.

of the sources is the ratio of the largest to the smallest coefficient of variation of the inter-arrival time. This ratio can be used to forecast the accuracy of the algorithm without the need to compare with simulation. A good rule of thumb is to use the algorithm for cases when this ratio is in the range $[0, 2.5]$. This, however, still needs further testing for large buffer sizes and under various ratios between the buffer size and the largest average on period of the sources (note that this ratio plays a major role in affecting system performance.)

An interesting study would be to compare our method with other approaches that characterize the superposition process as a 2-state MMPP, e.g. the work in [3, 10, 15]. We plan to use the results of this model in studying call admission control and comparing the different methodologies and bounds available in the literature.

REFERENCES

- [1] S. L. Albin. Approximating a point process by a renewal process, 2: Superposition of arrival processes to queues. *Operations Research*, 32:1133–1162, 1984.
- [2] D. Anick, D. Mitra, and M. M. Sondhi. Stochastic theory of a data-handling system with multiple sources. *Bell Sys. Tech. J.*, 61:1871–1894, 1982.
- [3] A. Baiocchi, N. B. Melazzi, M. Listani, A. Roveri, and R. Winkler. Loss Performance Analysis of an ATM Multiplexer Loaded with High-Speed On-Off Sources. *IEEE Journal on Selected Areas in Communications*, 9:388–393, 1991.
- [4] A. I. Elwalid, D. Mitra, and T. E. Stern. Statistical multiplexing of Markov modulated sources: Theory and computational algorithms. In *Proceedings of 13th International Teletraffic Congress*, pages 495–500, Copenhagen, June 1991.
- [5] H. Heffes. A class of data traffic process: covariance function characterization and related queueing results. *Bell Sys. Tech. J.*, 59:897–929, 1980.
- [6] H. Heffes and D. M. Lucantoni. A Markov Modulated Characterization of Packetized Voice and Data Traffic and Related Statistical Multiplexer Performance. *IEEE Journal on Selected Areas in Communications*, 4:856–867, 1986.

- [7] S. Hong, H. G. Perros, and H. Yamashita. A discrete-time queueing model of the shared buffer ATM switch with bursty arrivals. *Telecommunication Systems*, 2:1–20, 1993.
- [8] J. Hunter. *Mathematical Techniques of Applied Probability. Volume 2: Discrete Time Models: Techniques and Applications*, chapter 9. Academic Press, 1983.
- [9] I. I. Makhamreh, D. McDonald, and N. D. Georganas. Analysis of an Output-Buffered ATM Switch with Imbedded Bursty Traffic. Submitted for publication, 1994.
- [10] R. Nagarajan, J. F. Kurose, and D. Towsley. Approximation Techniques for Computing Packet Loss in Finite-Buffered Voice Multiplexers. *IEEE Journal on Selected Areas in Communications*, 9:368–377, 1991.
- [11] H. G. Perros and R. Onvural. On the superposition of arrival processes for voice and data. In *Fourth International Conference on Data Communication Systems and Their Performance*, pages 341–357, Barcelona, June 1990.
- [12] K. Sriram and W. Whitt. Characterizing Superposition Arrival Processes in Packet Multiplexers for Voice and Data. *IEEE Journal on Selected Areas in Communications*, 6:833–846, 1986.
- [13] T. E. Stern and A. I. Elwalid. Analysis of separable Markov-modulated rate models for information-handling systems. *Adv. Appl. Prob.*, 23:105–139, 1991.
- [14] R. C. F. Tucker. Accurate method for analysis of a packet-speech multiplexer with limited delay. *IEEE Transactions on Communications*, 36:479–483, 1988.
- [15] S. S. Wang and J. A. Silvester. A Fast Performance Model for Real-Time Multimedia Communication. In *Proceedings of the Fifth International Conference on Data Communication Systems and Their Performance*, Raleigh, NC, October 1993.
- [16] W. Whitt. Approximation of a point process by a renewal process, 1: Two basic methods. *Operations Research*, 30:125–147, 1982.