Non-Deterministic Periodic Packet Streams and Their Impact on a Finite-Capacity Multiplexer*

Randall Landry   Ioannis Stavrakakis
Department of Electrical Engineering and Computer Science†
University of Vermont
Burlington, VT 05405

Abstract

This paper considers the multiplexing of non-deterministic periodic traffic sources which deliver packets periodically according to an arbitrary i.i.d. batch arrival process. The proposed traffic sources are representative of the packet-stream periodicities commonly found in high-speed packet networks such as ATM. An exact analytical technique, which entails the solution of Markov chains whose dimensions are independent of the periods associated with the input sources, is presented for the study of a finite-capacity multiplexer fed by N of the proposed traffic models. Numerical results indicate that the source periodicities have a significant impact on the multiplexer's performance. In fact, for some regions of operating conditions, even a Poisson traffic assumption is shown to provide an extremely pessimistic estimate of the queue performance measured in terms of packet-loss probabilities.

1 Introduction

High-speed packet-switched networks based on technologies such as Asynchronous Transfer Mode (ATM) are widely believed to be the most promising candidates for providing services in the evolving Broadband-Integrated Services Digital Network (B-ISDN). In the ATM environment, sources generate information in the form of small fixed size packets which are multiplexed with the packet streams of other sources as they proceed from source to destination in a connection-oriented fashion. The main attraction of ATM is the potential statistical multiplexing gain which can be realized when integrating a variety of sources with very different statistical characteristics.

An effective manner in which to evaluate the merits of statistical multiplexing is to study the performance of multiplexing schemes under a variety of different, and hopefully realistic, packet arrival processes. Such studies can be found in [1]-[6] and the references therein, for input traffic models which are capable of capturing the burst-level correlations present in packetized voice. The multiplexing of variable bit rate (VBR) sources is considered in [7] and [8].

An important characteristic of high-speed networks, which is not explicitly considered in any of the previously mentioned studies, is the potentially large disparity in the relative speeds of the output link and the input source. This disparity, or *speed-up factor* – commonly appearing at the network edges where a number of sources are being multiplexed onto an output link (Fig. 1) – results in successive packets from an active source being spaced by the interval of time required to generate enough bits to form a packet. The speed-up factor $m^k = C/r^k$, where C is the link speed and r^k is the peak rate of source k (both expressed in units of bits/sec), is also relevant in inter-networking, where a number of slower packet networks of various speeds are connected to the B-ISDN through a common gateway. It has been shown in [9] that, under certain operating conditions, source periodicities at the network edge, like those mentioned above, give rise to packet streams within the network which also exhibit significant periodicities. This observation has been used in [10] to justify the adoption of periodic arrival streams for the simulation study of an ATM switch.

One way to address the speed-up issue, and the resulting packet stream periodicities, analytically is to consider the multiplexing of deterministic periodic sources, which have been studied in a number of papers such as [11]-[13]. However, a potentially more

*Research supported by the Advanced Research Project Agency (ARPA) under Grant F49620-93-1-0564 monitored by the Air Force Office of Scientific Research (AFOSR).

†email : rlandry@emba.uvm.edu / stavraka@emba.uvm.edu

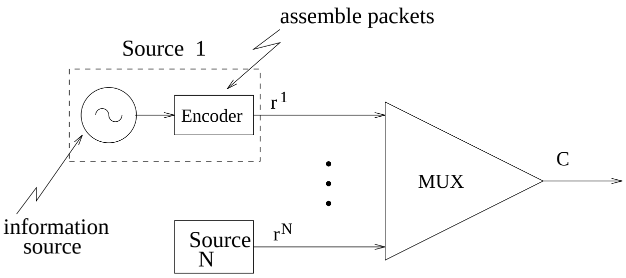

Figure 1: Multiplexing broadband packet sources at the network edge.

flexible traffic model for high-speed networks such as ATM may be a non-deterministic periodic (NDP) model in which packets are generated periodically according to a stochastic process, thereby allowing for a potential statistical multiplexing gain.

A relevant question that arises when considering the multiplexing of NDP sources is whether or not the aggregation of such sources can be approximated by any of the commonly used traffic models. For instance, it is widely accepted that the superposition of a large number of uncorrelated packet streams can be approximated by a Poisson process, which allows for a simple analytical solution to many queueing problems. However, this limiting assumption may not be valid when the arrival processes have a periodic structure. In fact, it will be shown that for the superposition of NDP sources, the Poisson assumption can be either pessimistic or optimistic, depending upon the operating conditions.

A correlated traffic model in which packets are generated at a minimum spacing of s^k slots is considered in [14] and [15], where the packet arrival process is governed by an underlying 2-state Markov chain. The speed-up factor s^k, which is equal to the unit step transition time of the Markov chain for source k, is shown to have a significant impact on multiplexer performance. Specifically, as s^k increases, the inter-packet correlations present at the source become less dominant and the multiplexer's performance improves accordingly. These observations suggest that, for sufficiently large values of s^k, it may be possible to drop the correlated source model in favor of a simpler, uncorrelated model without significantly impacting on the model's accuracy.

In this paper, a non-deterministic periodic traffic model is considered in which packets from source k are generated independently every $m^k \geq 1$ time slots according to an arbitrary distribution. The model is motivated in part by the need to provide traffic models which account for the potentially large disparity between source and link speeds, as well as other periodicities that are commonly encountered in high-speed packet networks. An analytical technique is presented which allows for the exact solution of queue length distributions and packet-loss probabilities in a finite capacity multiplexer, where the input consists of a superposition of the proposed source models. The technique entails the decomposition of a potentially large (increasing with m^k) irreducible Markov chain into smaller irreducible Markov chains whose sizes are independent of the source periods m^k. The consideration of heterogeneous sources, which may have different periods and/or different packet generating distributions, introduces no additional numerical complexity into the problem.

In the next section, the source model, as well as the superposition of N such sources, is introduced and parametrically characterized. The analytical technique, which is shown to be numerically tractable under realistic operating conditions, is also described. Numerical results are presented in section 3, where the validity of the Poisson assumption mentioned above is discussed in detail, and concluding remarks can be found in section 4.

2 Traffic Characterization and Queueing Analysis

The system under consideration consists of a discrete-time, single-server multiplexer (Fig. 2) with a finite capacity of B packets, and an output link capacity of C bits/sec. The basic unit of time is taken to be the transmission time of a packet, or the duration of a time slot on the output link, which is equal to T seconds. All packets have a fixed size of $T \cdot C$ bits and are served in a First-Come First-Served (FCFS) manner.

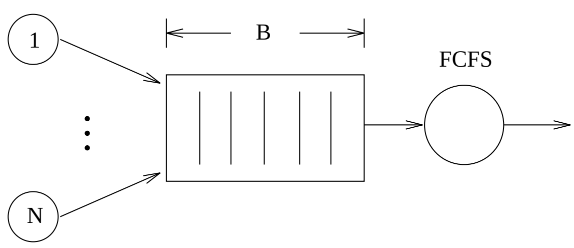

Figure 2: The finite-capacity queueing system.

2.1 The NDP Traffic Source Model

The input to the multiplexer described above consists of N NDP packet sources, each of which can be completely described in terms of the parameters $\{m^k, c^k, f^k(\cdot)\}$ defined as follows for $1 \leq k \leq N$.

- m^k: The speed-up factor associated with the k^{th} source, also referred to as the period of the source. It is assumed that $m^k \in \boldsymbol{Z}^+$, where $\boldsymbol{Z}^+$ denotes the positive integers.
- c^k: A constant which denotes the time instant (referred to here as the initialization time) of the first potential packet generation for source k, assuming that $1 \leq c^k \leq m^k$. Note that all N sources are synchronized when $c^1 = c^2 = \cdots = c^N$.
- $f^k(\cdot)$: The probability mass function (PMF) for the number of packets generated by the k^{th} source every $m^k \cdot T$ seconds (or m^k time slots) beginning at time $c^k \cdot T$; the maximum number of packet arrivals over a slot is $N^k < \infty$.

Note that each source is non-deterministic and periodic in the sense that packets are generated periodically every m^k time slots according to the PMF $f^k(\cdot)$. The number of packets delivered by the k^{th} source during the j^{th} time slot is denoted by A_j^k, and described as follows.

$$A_j^k = \begin{cases} i \ \text{ w.p. } f^k(i) & \text{if } j \in \{c^k + nm^k\}_{n \geq 0} \\ 0 \ \text{ w.p. } 1 & \text{otherwise} \end{cases} \tag{1}$$

These arrivals are declared at the end of the slot over which they occur and the average number of packets per slot delivered by the k^{th} source (mean rate of source k) is given by

$$\lambda^k = \frac{1}{m^k} \sum_{i=0}^{N^k} i f^k(i) \quad . \tag{2}$$

The NDP traffic model considered above is potentially useful for describing a wide range of packet arrival processes in high-speed networks. For instance, most packet sources accessing the B-ISDN will have speed-up factors larger than 1 associated with them. This is the case whenever the peak rate r^k of source k is less than the link speed C. A Constant Bit Rate (CBR) source, which delivers packets at a constant rate (r^k) when active, can be described by the NDP model by letting $f^k(\cdot)$ be Bernoulli and setting m^k equal to C/r^k (assuming that this ratio is an integer). Note that r^k would be 64 kbps if source k were a voice source, which could result in very large values of m^k, depending upon the link speed C. It has been shown in [16] that even for a speed-up factor as small as 2, the queue length tail distribution in an infinite capacity multiplexer is significantly decreased over the equal source and link speed case ($m^k = 1$) when considering correlated on-off sources. In fact, as m^k increases, the queue performance improves further as *inter-packet* correlations at the source are no longer seen as *inter-slot* correlations at the server. For the infinite buffer case, the queueing behavior approaches that of the uncorrelated M/D/1 queue as m^k becomes large, which suggests that the presence of inter-packet correlations at a source does not necessarily require the inclusion of such correlations in the source traffic model. It is interesting to note that a further reduction in the effective (as seen by the server) source correlation is observed when considering a finite capacity multiplexer [14]. In fact, for some regions of buffer capacity, the impact of inter-packet correlations is shown to be negligible.

A NDP arrival process delivering at most one packet per period may also find applications in describing VBR video sources in which packets are prebuffered and equally spaced ("deterministic smoothing") over a frame (or some multiple of a frame). Such a scheme is considered in [17], where the Histogram Decomposition approximation originally presented in [18] is employed to study the multiplexing of deterministically smoothed VBR sources. This approach is motivated by the fact that due to both the relatively long duration of a video frame and the high positive correlations in the inter-frame packet arrival process, the periodic arrival state of each video source is sufficiently long to drive the multiplexer into steady-state. Consequently, the approximate steady-state queueing behavior is determined by performing a series of independent queueing analyses, each of which assume the superposition of periodic arrival processes as inputs; the period of each packet stream reflects the arrival state of the VBR source. Note that this approximation treats each VBR source as a "collection" of CBR sources (each of which can potentially be described by a NDP model) whose relative frequencies are determined from the arrival rate histogram.

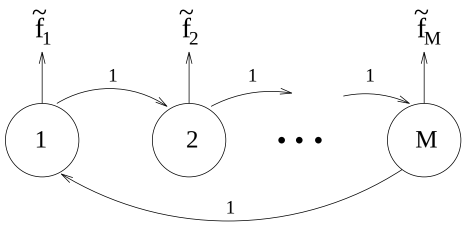

Figure 3: The aggregate source model.

2.2 The Superposition of NDP Sources

In order to study the queue response to N of the proposed traffic sources, their superposition is first considered. It turns out that the superposition of N NDP traffic sources results in another NDP source which can be determined by the parameters $\{m^k, c^k, f^k(\cdot) : 1 \le k \le N\}$. In fact, it is easily shown that the resulting source has period M given by,

$$M = LCM\{m^1, m^2, \ldots, m^N\} \quad , \tag{3}$$

where $LCM\{\cdot\}$ denotes the least common multiple of its argument.

Let Φ_j denote the set of all sources for which state j, $1 \le j \le M$, is a potential packet-generating state, and let S_j denote the number of such sources; note that S_j represents the cardinality of Φ_j. Membership of the k^{th} source in Φ_j can be easily determined by $\{m^k, c^k\}$, $1 \le k \le N$. The PMF of the cumulative packet arrival process, given that the aggregate NDP input source is in state j, will be denoted by $\tilde{f}_j(i)$, $i \ge 0$, and is determined as follows:

$$\tilde{f}_j(i) = \begin{cases} 0 \;\text{ for } i > 0 & \text{if } \Phi_j = \emptyset \\ 1 \;\text{ for } i = 0 & \text{if } \Phi_j = \emptyset \\ S_j\text{-fold convolution of} & \\ f^k(\cdot) \text{ evaluated at } i, & \\ \;\forall k \in \Phi_j & \text{if } \Phi_j \neq \emptyset \;\; . \end{cases} \tag{4}$$

Given $\tilde{f}_j(\cdot)$, for $1 \le j \le M$, the new arrival process, which represents the superposition of N NDP sources, is completely specified by $\left\{M, \tilde{f}_j(\cdot) : 1 \le j \le M\right\}$ and can be represented as shown in Fig. 3. Note that while all M states may be packet-generating states, the superposition of N synchronized ($c^1 = \cdots = c^N$) sources with identical periods ($m^1 = \cdots = m^N$) results in a NDP source with only one packet-generating state. For instance, if $c^1 = \cdots = c^N = 1$ and $m^1 = \cdots = m^N = 10$, then the resulting source would have $\tilde{f}_1(0) < 1$ and $\tilde{f}_j(0) = 1$, for $2 \le j \le M = 10$. The consideration of un-synchronized sources with different periods and/or different PMF's $f^k(\cdot)$ (heterogeneous sources) simply results in more than one packet-generating state.

2.3 Performance Analysis of the Finite-Capacity Multiplexer

In this subsection a FCFS finite-capacity multiplexer fed by N NDP sources is analyzed in discrete time.

The system is observed at the end of each time slot and packet departures are assumed to occur before new arrivals are counted. The queueing system is described by the 2-dimensional state vector (q_k, x_k), where q_k represents the queue length at time k (defined at the end of time slot k), and x_k denotes the current state of the M-state input process. It is easy to establish that the stochastic process $\{q_k, x_k\}_{k\ge1}$ is a Markov chain with state space $\{(i,j) : 0 \le i \le B, 1 \le j \le M\}$, where i (j) denotes the current state of q_k (x_k). Let $p(i,j,i',j')$ denote the probability that the Markov chain $\{q_k, x_k\}_{k\ge1}$ moves from state (i,j) at time k to state (i',j') at time $k+1$. Due to the simple periodic structure of the input process $\{x_k\}_{k\ge1}$, a fraction of $\frac{M-1}{M}$ of the total number of elements in the transition matrix for $\{q_k, x_k\}_{k\ge1}$ are zero. In fact, when the probabilities $p(i,j,i',j')$, for $0 \le i,i' \le B$ and $1 \le j,j' \le M$, are mapped into 2-dimensional space through the mapping,

$$\ell = (j-1)(B+1)+i \quad \text{and} \quad \ell' = (j'-1)(B+1)+i' \quad ,$$

the resulting transition matrix, $\mathbf{P}$, which is of order $M(B+1)$ with elements $p(\ell,\ell')$, takes the form,

$$\mathbf{P} = \begin{bmatrix} 0 & \mathbf{P}_1 & 0 & \cdots & 0 \\ 0 & 0 & \mathbf{P}_2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \mathbf{P}_{M-1} \\ \mathbf{P}_M & 0 & 0 & \cdots & 0 \end{bmatrix} \quad , \tag{5}$$

which happens to be the canonical form for an irreducible periodic Markov chain of period M [19]. The square matrices $\mathbf{P}_j$, $1 \le j \le M$, are of order $(B+1)$ and their elements $p_j(i,i')$, $0 \le i,i' \le B$, are given below.
For $0 \le i \le B$ and $max(0, i-1) \le i' \le B-1$,

$$p_j(i,i') = \tilde{f}_j(i'-i+1) + \tilde{f}_j(0) \cdot \mathbf{1}_{\{i+i'=0\}} \quad , \tag{6}$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function.

For $0 \le i \le B$ and $i' = B$,

$$p_j(i,i') = \sum_{l=B-i+1}^{N_j} \tilde{f}_j(l) \quad , \tag{7}$$

where N_j denotes the maximum number of arrivals, as determined by (4), when the NDP input source is in state j.

Although the transition probabilities given above are easily determined, the order of the transition matrix $\mathbf{P}$, which affects significantly the complexity of the solution, increases linearly with respect to M. This suggests that the consideration of NDP input processes with large periods may be prohibitive. However, the following theorem provides for the exact steady-state solution of $\{q_k, x_k\}_{k\ge 1}$ by solving M Markov chains with transition matrices of order $(B+1)$. Its proof can be found in [20].

Theorem 1 *The finite irreducible periodic Markov chain, with transition matrix given by (5), has a unique stationary probability vector,* $\boldsymbol{\pi} = (\pi_{0,1} \cdots \pi_{B,M})'$*, given by*

$$\boldsymbol{\pi} = (1/M)(\boldsymbol{\pi}'_1 \boldsymbol{\pi}'_2 \cdots \boldsymbol{\pi}'_M)' \quad , \tag{8}$$

where $\boldsymbol{\pi}_i$ *is the unique stationary probability vector for* $\mathbf{A}_i = \mathbf{P}_i \cdots \mathbf{P}_M \cdot \mathbf{P}_1 \cdots \mathbf{P}_{i-1}$ $(i = 2, \cdots, M)$*, and* $\mathbf{A}_1 = \mathbf{P}_1 \cdots \mathbf{P}_M$*.* □

The above theorem significantly reduces the numerical complexity involved in calculating the stationary probability vector of the transition matrix in (5), since only matrices of order $(B+1)$ are operated upon.

Upon solving for $\boldsymbol{\pi} = (\pi_{0,1} \cdots \pi_{B,M})'$ using Theorem 1, the queue length distribution is easily evaluated as

$$p_q(k) = \sum_{j=1}^{M} \pi_{k,j} \qquad k = 0, \ldots, B \quad . \tag{9}$$

The average number of arrivals per slot for the cumulative NDP input process, or the total offered load, is given by

$$\lambda = \sum_{k=1}^{N} \lambda^k \quad , \tag{10}$$

where λ^k is the mean arrival rate of source k as defined earlier.

Under the departure-before-arrival order of service considered here, the server utilization, or system throughput ρ, can be calculated as

$$\rho = 1 - \sum_{j=1}^{M} \pi_{0,j} \cdot \tilde{f}_j(0) \quad . \tag{11}$$

Finally, using (10) and (11), the packet-loss probability L is given by,

$$L = 1 - \frac{\rho}{\lambda} \quad . \tag{12}$$

3 Numerical Results

In this section, some numerical results are presented for the analysis of the finite-capacity multiplexer fed by NDP input sources; all results are expressed in terms of packet-loss probabilities. Although the analysis described in the previous section allows for the superposition of input sources having different periods, the observations having the most direct relevance to high-speed networking are also facilitated by the consideration of sources with identical periods. Therefore, for added simplicity in the presentation and discussion of numerical results, all N input sources are assumed to have the same period denoted by $m = m^1 = \cdots = m^N$. Each source is also assumed to deliver packets according to a Bernoulli PMF with rate α, which results in a mean rate per source (see (2)) of $\lambda^1 = \cdots = \lambda^N = \alpha/m$; the total offered load is then given by (10) to be $\lambda = N\alpha/m$.

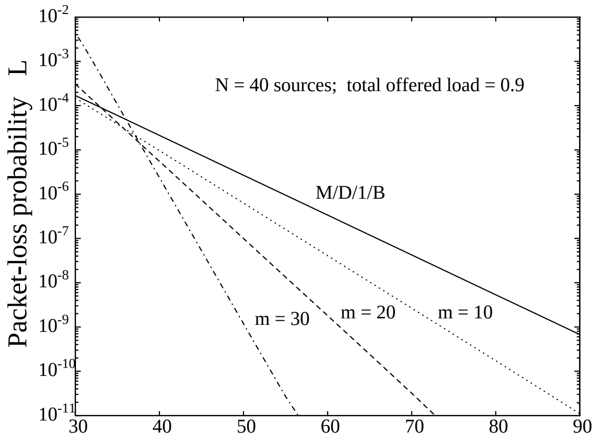

Figure 4: Loss probability vs. buffer capacity for various m.

Fig. 4 illustrates the effect that the source period has on packet-loss probabilities as a function of the buffer capacity B when $N = 40$ sources are multiplexed at an offered load of $\rho = 0.9$ and the sources are assumed to be synchronized ($c^1 = \cdots = c^{40}$). Clearly, the source period m has a significant impact on queue performance when the total offered load is held constant. Note, however, that this impact is not necessarily negative as has been suggested in the past [10]. In fact, for "acceptable" loss probabilities on the order of 10^{-6} and smaller, the impact that increasing values of m have on the queue performance is positive. The curves for $m = 10, 20$, and 30 are also compared with the corresponding M/D/1/B approximation, and the results are quite enlightening. Fig. 4 reveals the interesting fact that the Poisson assumption can either overestimate or underestimate loss probabilities depending upon buffer capacity and the source period m. For a loss probability of 10^{-9}, the M/D/1/B approximation is extremely pessimistic, requiring a buffer capacity of $B = 90$, compared to a required buffer capacity of $B = 50$ when $m = 30$. In fact, the negative slope of the M/D/1/B curve is always smaller than the cases for $m > 1$, which implies that, beyond some value of B (which varies with m), the loss probabilities for the limiting Poisson traffic assumption upper bound the actual loss probabilities for the superposition of NDP sources.

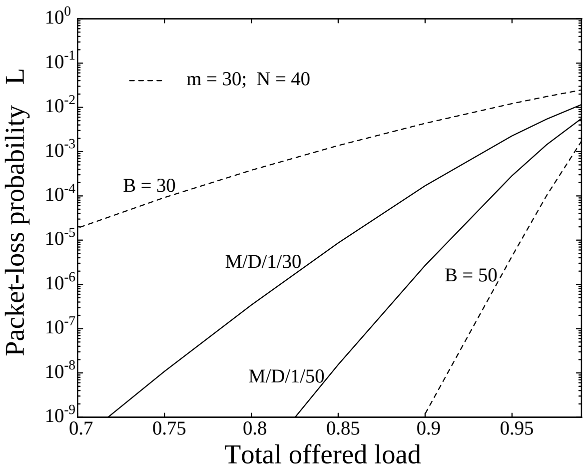

Figure 5: Loss probability vs. offered load for various B.

Fig. 5 demonstrates the relationship between packet-loss probabilities and offered load for a case in which the Poisson assumption is optimistic ($B = 30$) as well as a case in which the assumption is overly pessimistic ($B = 50$). The $N = 40$ NDP sources are assumed to have speed-up factors of $m = 30$ associated with them. When $B = 30$, the superposition of NDP input sources results in loss probabilities larger than 10^{-5} for the range of loads between 0.7 and 0.99. The corresponding M/D/1/30 queue, however, is overly optimistic and results in loss probabilities smaller than 10^{-9} for loads as high as 0.72. On the other hand, when $B = 50$, the M/D/1/50 approximation overestimates the loss probabilities for all loads. It is clear from Fig. 5 that the Poisson assumption performs poorly for low to moderate loads, and even for offered loads of 0.99 there is a substantial difference between the M/D/1/B queue performance and the corresponding case of NDP input sources.

Note that for the results presented in Figures 4 and 5, it is assumed that the NDP sources are synchronized ($c^1 = \cdots = c^N$), which represents a worst-case scenario in terms of queue performance. Unlike past analytical studies of deterministic periodic sources ([11]-[13]), the analytical technique employed here allows for the consideration of NDP sources with arbitrary initialization times. Although it is difficult to define a meaningful distribution of initialization times c^k ($1 \leq k \leq m$) for the NDP sources, it is reasonable to expect that uniformly distributing these times will improve the queueing performance. This is illustrated

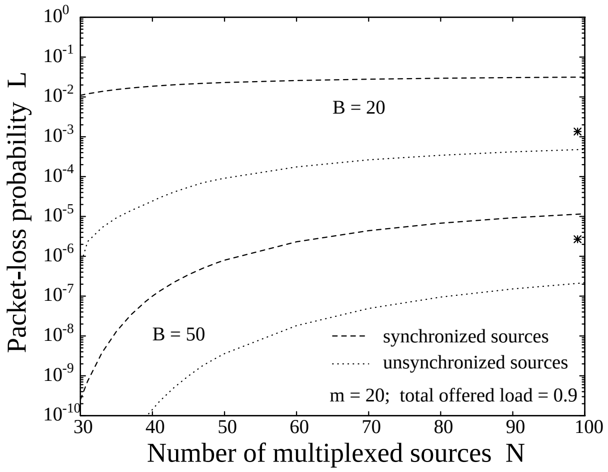

Figure 6: Loss probability vs. N for various B.

in Fig. 6, where loss probabilities are plotted versus the number of multiplexed sources with $m = 20$ and $\rho = 0.9$ for synchronized and unsynchronized (initialization times uniformly distributed) sources. The uppermost curves represent results for a buffer capacity of $B = 20$, and the lower curves for $B = 50$. Similarly, the upper star on the right hand vertical axis denotes the M/D/1/20 result while the lower star denotes the M/D/1/50 result. For both cases ($B = 20$ and $B = 50$), the queue performance, in terms of packet-loss probabilities, approaches a limit for large values of N. However, even when the NDP sources are unsynchronized, the limiting results are significantly different than the M/D/1/B results.

The results presented thus far indicate that all of the parameters (B, m, N and ρ) affect significantly the queue performance, and consequently, the validity of the M/D/1/B approximation. For each set of numerical results there seem to be two distinct regions of operation; one in which the source periodicities have a negative influence on the queueing behavior, relative to a Poisson input source, and one in which the periodicities have a favorable impact. These operating regions are examined more closely in Fig. 7, where the ratio of the loss probability for the Poisson traffic source (L_P) to the loss probability for the NDP sources (L_{NDP}) is plotted versus buffer capacity B. When the ratio $\frac{L_P}{L_{NDP}}$ is less than (greater than) 1, the M/D/1/B approximation underestimates (overestimates) the actual loss probability. In Fig. 7, $N = 75$ NDP sources are multiplexed at a total offered load of 0.9 for source periods of 1, 10 and 50. When $m = 1$, the input to the multiplexer consists of $N = 75$ Bernoulli sources, and the ratio $\frac{L_P}{L_{NDP}}$ is close to 1 for all values of B. This result supports the assertion that the superposition of a large number of Bernoulli traffic streams can be approximated by a Poisson process (asymptotic behavior of the Binomial distribution), regardless of the buffer capacity. However, as the source periods increase, the two regions become more distinct. When $m = 10$ and $B = 10$, L_{NDP} is 5 times greater than L_P for the M/D/1/10 queue, and when $B = 100$, L_{NDP} is 5 times smaller than L_P. For this case, the M/D/1/B results are overly pessimistic for $B \geq 50$. When $m = 50$, the two regions are even more well-defined, and the Poisson assumption does not perform well in either region. It is interesting to note that, as was the case in Fig. 4, the region of operating conditions which produces loss probabilities considered acceptable in ATM (less than 10^{-6}) is the one in which the M/D/1/B queue provides overly pessimistic results ($L_P > L_{NDP}$).

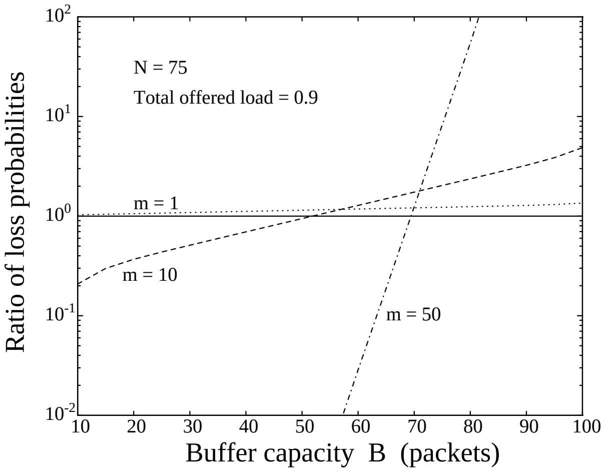

Figure 7: Ratio of loss probabilities for Poisson and NDP sources.

4 Conclusions

In this paper, an exact analysis was presented for a finite-capacity multiplexer subjected to an input of NDP packet streams.

In terms of packet-loss probabilities, the packet generation period, or speed-up factor, associated with each traffic source was shown through numerical results to have a significant impact on the multiplexer's performance. This implies that the speed-up factor could potentially play a prominent role in network resource allocation, thereby supporting its adoption as an important traffic descriptor.

Depending upon the operating conditions at the finite-capacity multiplexer, the period of the NDP input streams was shown to be capable of having either a positive or negative impact on queueing performance, relative to that of a M/D/1/B queue. In fact, for each set of operating conditions, there were two distinct regions of performance; one in which the NDP sources impacted favorably on the finite queue and one in which they had a negative impact (again, relative to the M/D/1/B queue). This implies that operating conditions should be considered carefully before M/D/1/B results are used for bounding loss probabilities when studying the queue response to NDP input processes. Finally, the asymptotic (large N) behavior of NDP input sources was shown to be poorly approximated by the Poisson process when $m > 1$.

References

[1] H. Heffes, D. Lucantoni, "A Markov Modulated Characterization of Packetized Voice and Data Traffic and Related Statistical Multiplexer Analysis", IEEE J. Select. Areas Commun., Sept. 1986, pp. 856-868.

[2] H. Bruneel, "Queueing Behavior of Statistical Multiplexers with Correlated Inputs", IEEE Trans. Commun., Vol. 36, No. 12, Dec. 1988.

[3] K. Sriram, W. Whitt, "Characterizing Superposition Arrival Processes in Packet Multiplexers for Voice and Data", IEEE J. Select. Areas Commun., Sept. 1986, pp. 833-846.

[4] J. Daigle, Y. Lee, N. Magalhaes, "Discrete Time Queues with Phase Dependent Arrivals", IEEE INFOCOM'90, San Francisco, CA 1990.

[5] I. Stavrakakis, "Analysis of a Statistical Multiplexer under a General Input Traffic Model", IEEE INFOCOM'90, San Francisco, CA 1990.

[6] S. Q. Li, "A General Solution Technique for Discrete Queueing Analysis of Multimedia Traffic on ATM", IEEE Trans. Commun., Vol. 39, No. 7, July 1991.

[7] B. Maglaris, D. Anastassiou, S. Prodip, G. Karlsson, J. Robbins, "Performance Models of Statistical Multiplexing in Packet Video Communications", IEEE Trans. Commun., Vol. 36, July 1988.

[8] C. Blondia, O. Casals, "Performance Analysis of Statistical Multiplexing of VBR Sources", IEEE INFOCOM'92, Florence, Italy, 1992.

[9] V. Ramaswami, G. Latouche, "Modeling Packet Arrivals from Asynchronous Input Lines", Proceedings of the ITC 12, 1988.

[10] T. Eliazov, V. Ramaswami, W. Willinger, G. Latouche, "Performance of an ATM Switch: Simulation Study", IEEE INFOCOM'90, San Francisco, CA, 1990.

[11] A. Eckberg, "The Single Server Queue with Periodic Arrival Process and Deterministic Service Time", IEEE Trans. Commun., Vol. 27, March 1979.

[12] F. Hubner, "Analysis of a Finite-Capacity Asynchronous Multiplexer with Deterministic Traffic Sources", 7^{th} ITC Seminar, N.J., 1990.

[13] J. Roberts, J. Virtamo, "The Superposition of Periodic Cell Arrival Streams in an ATM Multiplexer", IEEE Trans. Commun., Vol. 39. No. 2, Feb. 1991.

[14] H. Kroner, "Statistical Multiplexing of Sporadic Sources - Exact and Approximate Performance Analysis", *Teletraffic and Datatraffic in a Period of Change*, ITC-13, Elsevier Science Publishers, 1991.

[15] I. Stavrakakis, "Statistical Multiplexing of Correlated Slow Traffic Sources", GLOBECOM'92, Orlando, FL, 1992.

[16] J. Daigle, *Queueing Theory for Telecommunications*, Addison-Welsey, 1992.

[17] N. Shroff, M. Schwartz, "Video Modeling within Networks using Deterministic Smoothing at the Source", IEEE INFOCOM'94, Toronto, Canada, 1994.

[18] P. Skelly, S. Dixit, M. Schwartz, "A Histogram-Based Model for Video Traffic Behavior in an ATM Network Node with an Application to Congestion Control", IEEE INFOCOM'92, Florence, Italy, 1992.

[19] J. Hunter, *Mathematical Techniques of Applied Probability*, vol. 1, *Discrete Time Models: Basic Theory*, Academic Press, 1983.

[20] J. Hunter, *Mathematical Techniques of Applied Probability*, vol. 2, *Discrete Time Models: Techniques and Applications*, Academic Press, 1983.