Loss Correlation for Queues with Bursty Input Streams

Henning Schulzrinne
Dept. of ECE
James F. Kurose and Donald F. Towsley
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
hgschulz,kurose,towsley@cs.umass.edu

Abstract

The loss probability of a queueing system provides, in many cases, insufficient information for performance evaluation, for example, of data link layer protocols and applications with forward error correction.

This paper evaluates and characterizes the correlation between packet losses for two queueing systems in discrete time that are motivated by BISDN applications. The first, a two-class discrete-time queueing system, approximates the output queue of an ATM switch. The queue serves periodic foreground traffic and random background traffic. The background traffic is modeled as i.i.d. batches of arbitrary distribution. It is shown that the conditional loss probability (CLP) is independent of the buffer size if the buffer size is at least as large as the period of the foreground traffic. Example calculations indicate that losses occur essentially randomly as long as the foreground traffic uses less than 10% of the channel capacity.

The second analysis derives the CLP seen by a selected stream in a slotted finite-buffer system with a superposition of interrupted Poisson sources. Here, the total number of arrivals may be correlated from slot to slot. Traffic correlation is seen to have a strong influence on loss correlation, while buffer size is seen to have virtually none.

1 Introduction

With the advent of high-bandwidth wide area networks and their high ratio of propagation to transmission and queueing delay, loss probability rather than delay or throughput has become the dominant performance metric. However, since losses in queueing systems can be strongly correlated, the loss probability (LP) of a queueing system often provides insufficient information for performance evaluation, for example, of data link layer protocols and forward error correction. Bursty traffic, as generated by multimedia applications, has made concerns about loss correlation more urgent.

The investigation of loss correlation has a long history in the context of bit errors in data communication [1], but has only recently attracted stronger interest in the area of packet networks. We briefly summarize related work, extending the review in [2]. By rate conservation methodology, Ferrandiz and Lazar developed a theoretical framework for studying loss correlation in queueing systems [3–5]. Our focus is more on deriving computable results, relying on well-known Markov chain methods. Motivated by forward error correction methods, [6] computes the number of packets lost in a fixed-size packet group for both continuous-time and discrete-time systems.

A number of authors have quantified the influence of loss correlation for network performance. For example, it has been shown [7–10] how the throughput of go-back-N ARQ increases with positive loss correlation. Similar results for the N-packet transfer time and go-back-∞ and blast protocols are presented in [11]. All papers[1] assume a two-state Markovian error model with geometrically distributed run lengths for losses and successes, which will be seen to be appropriate for some, but not all queueing systems studied below. Thus, our results can be used in conjunction with the work cited here to directly predict ARQ performance.

In some simulation studies, cell loss correlation has been investigated [12,13] and methods for the compensation of correlated cell loss have been proposed [14,15].

This paper continues the work begun in [2], extending it to derive the loss correlation seen by packets of a selected stream in queues with *several* arrival streams. The goal of this work is to provide insights as to the structure and magnitude of loss correlation, so that the protocol and system architecture designer can judge when loss correlation will significantly effect system performance. The designer can judge whether either ameliorating measures are called for (in the case of real-time sequence reconstruction) or the benefits of loss correlation (in the case of retransmission algorithms) can be reaped.

Throughout the paper, our interest in broadband ISDN and ATM-like systems motivates a discrete-time perspective, with packet losses caused by buffer overflow[2]. Loss correlation is measured here by the conditional loss probability (CLP), the conditional probability that a packet from a stream is lost given that its immediate predecessor was also lost.

Due to space constraints, we will focus on the influence of traffic correlation (burstiness) on the loss correlation. In section 2, we briefly review results for a traffic model consisting of a periodic or general i.i.d. foreground source competing with a batch arrival process for queue resources. Details on these and related models and questions regarding service and buffering policies are addressed in a companion technical report [16]. In the main part of the paper, section 3, input traffic is modeled as a superposition of identical discrete-time interrupted Poisson (IPP) sources. IPP sources are commonly used to model voice and video traffic. We find that buffer size plays almost no role in the value of the loss correlation, while the burstiness affects both LP and CLP in a similar way.

2 Superposition of Periodic and Random Traffic

As a simple model of actual network traffic, we consider a discrete-time system similar to the one studied in [2]: single, first-come, first-serve queue with batch arrivals at every time slot and unit service time. The stream of interest is modeled as periodic (period τ), superimposed on background traffic with general, but i.i.d., batch arrivals. The background traffic is envisioned as the superposition of a large number of sources. For

This work is supported in part by the Office of Naval Research under contract N00014-90-J-1293, the Defense Advanced Research Projects Agency under contract NAG2-578 and a National Science Foundation equipment grant, CERDCR 8500332.

[1] except for [9], which allows a more general Markovian error model

[2] It appears to be generally accepted that buffer overflow will be the dominant loss mode in wide-area high speed networks.

brevity, the two streams are abbreviated as "foreground traffic" (FT) and "background traffic" (BT). If the foreground traffic is periodic, it also referred to as constant bit-rate (CBR) or continuous bit-stream oriented (CBO) traffic. This traffic model is motivated by voice and video applications, where the source of interests emits packets periodically during a talk spurt, a video scan line or even a whole image. To gain insight into the loss behavior, we require transient rather than steady-state results.

In this paper, we summarize some results for this system; details may be found in [16]. First, we can show:

Theorem 1 *The CLP for FT packets is independent of K for values of $K \geq \tau$, where K is the system size.*

This behavior is shown in Fig. 2, where LP and CLP are plotted as a function of K. Even below the threshold of τ, the CLP shows very little dependence for values of $K \geq \tau/2$. (This behavior was observed across a range of K and τ.) For interesting loss probabilities, the CLP is several orders of magnitude larger than LP, but the average number of consecutively lost packets is still very close to one (1.049 for geometric arrivals in this example).

Fig. 3 shows that CLP approaches LP for values of τ on the order of ten. Thus, as long as the periodic source does not command more than about ten percent of the total channel bandwidth, losses can be treated as approximately random.

Almost all queueing systems show a roughly exponential rise in loss probability as the system load approaches one. This behavior is also found in Fig. 1 for the loss probabilities of both traffic types as the background traffic is varied from $\lambda = 0.2$ to 0.9. The buffer size and FT period are held constant at $K = 10$ and $\tau = 10$, respectively.[3] The flattening of the loss curve seen here is also common to all finite-buffer systems. The more interesting observation in our context, however, is that the expected loss run length varies relatively little, staying loss to the value of one associated with random loss. Note the vastly different scales of the left and right ordinate in Fig. 1. While the loss probability ranges over eleven orders of magnitude, the mean loss run length changes by less than ten percent.

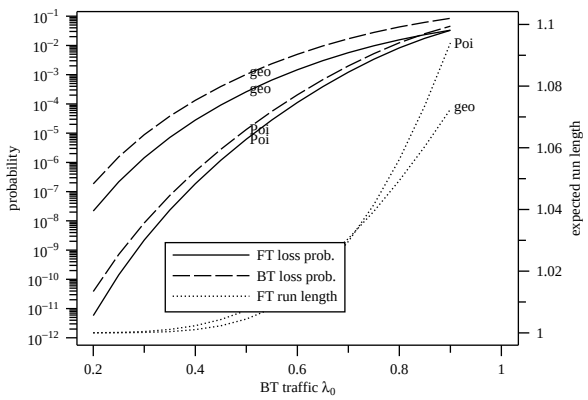

Figure 1: Loss probabilities and expected run length for Poisson and geometrically distributed background traffic BT and periodic traffic FT (random arrival), as a function of BT intensity λ_0; system size $K = 10$, period $\tau = 10$

3 Correlation in the $N \cdot \mathbf{IPP}/D/c/K$ Queue

There appears to be general agreement that Poissonian or i.i.d. batch arrival are not well suited for many non-data traffic sources, such as packet audio and video. Since the exact behavior

[3] Thus, the expected run length would remain the same for all values of K greater or equal to 10.

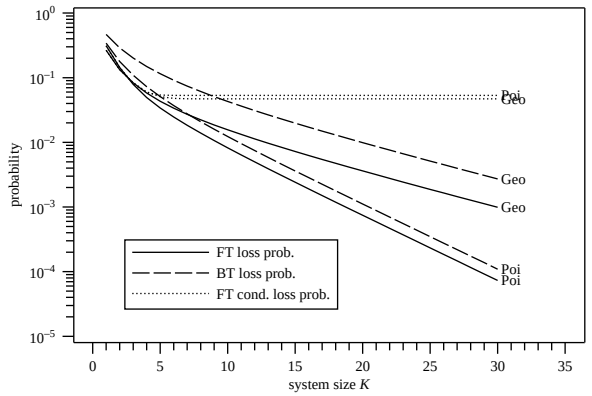

Figure 2: LP and CLP for Poisson or geometrically distributed BT with $\lambda = 0.8$ and periodic traffic FT with period $\tau = 10$, as a function of system size K

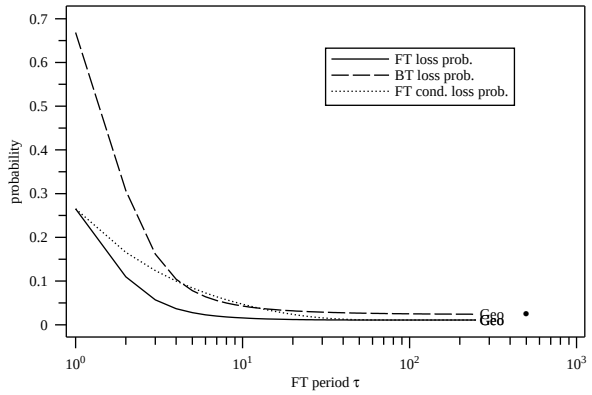

Figure 3: LP and CLP for geometrically distributed background traffic BT with $\lambda = 0.8$ and periodic traffic FT (random arrival), as a function of τ; system size $K = 10$

of these sources is ill understood and may depend strongly on the material being carried (voice vs. music; "talking heads" vs. entertainment video), several generic bursty-source models have been proposed and analyzed. These models differ not only in their generality, but also seem to be appropriate for different parts of the network. Note that altogether different models are to be used for aggregate traffic.

The first model, commonly referred to as an *on/off source*, features periodic arrivals during a burst and no arrivals during the silence periods separating bursts.

Multiplexors and other packet switch queues introduce delay jitter, thus, packet arrivals within a burst even of a on/off source will no longer be strictly periodic within the network. The interrupted Poisson process (IPP) source model attempts to capture this. In continuous time [17, 18], the on/off time periods are exponentially distributed and independent of each other. A fixed-rate Poisson source emits packets during the on periods. In discrete-time [19, 20], the on/off periods are taken to be geometrically distributed. A packet arrives at each slot during the on period with probability λ. (The term IPP is somewhat misleading in the discrete-time model, but maintained here for consistency.)

When the peak rate of sources or input channels exceeds the output rate, a batch arrival model is more appropriate. We can

generalize the IPP model by allowing batch arrivals at each time slot during the on period. No performance evaluations using this model are known to the authors.

We begin our analysis by formally defining the arrival process and the queueing system, proceed to find LP and finally, CLP. Numerical examples in section 3.4 illustrate the influence of buffer size and burstiness on loss correlation.

3.1 The Interrupted Poisson Process

An interrupted Poisson process (IPP) is defined as follows:

Definition 1 *A discrete-time interrupted Poisson process (IPP) alternates between active (on) and silence (off) periods. At each slot, the state moves from active to silent with probability* $1-\gamma$ *and from silence to active with probability* $1-\omega$*. Thus, silence and activity duration are geometrically distributed. In the active state, a packet is generated independently from slot to slot with probability* λ*.*

We assume that a state transition takes place just prior to the end of a time slot and that a packet is generated at the beginning of a slot with probability λ if the new state is the active state. [19] gives results for the single queue.

The average dwell times, T_i, for the off and on state, respectively, are given by $T_0 = 1/(1-\omega)$ and $T_1 = 1/(1-\gamma)$, or, conversely, $\omega = 1 - 1/T_0$ and $\gamma = 1 - 1/T_1$. We also define the activity factor $\alpha = \frac{T_1}{T_0+T_1} = (1-\omega)/(2-\omega-\gamma)$. The arrival rate is given by $\rho = \lambda\alpha$.

We analyze the $N \cdot \mathrm{IPP}/D/c/K$ queue operating in discrete time. N IPP sources, feed into the queue, which has K buffer positions, **not** including the customer being served. There are c identical servers. Each customer requires a service time of one time slot. All customers arrive at the beginning of a slot, while the departure, if any, occurs at the end of the slot (so-called *early arrival system* [21]).

We work towards the goal of computing the CLP (i.e., the conditional probability that customer n is lost given that customer $n-1$ from the same source was also lost) in several steps. First, immediately below, we characterize the aggregate arrival process. Then, we write down the transition matrix, which also yields the (unconditional) loss probability. Through a number of intermediate quantities, we finally arrive at the desired CLP.

We first describe the behavior of the aggregate arrival process, reviewing well-known results. The transition probability that j sources are active in the current slot, given that i were active in the previous slot can be computed by [20, Eq. (3)]

$$S_{i,j} = \sum_{m=0}^{i} \underbrace{\binom{i}{m}\gamma^m(1-\gamma)^{i-m}}_{i-m \text{ go off}, m \text{ stay on}} \cdot \underbrace{\binom{N-i}{j-m}(1-\omega)^{j-m}\omega^{N-i-(j-m)}}_{j-m \text{ go on}, N-i-(j-m) \text{ stay off}}$$
$$\text{for } i,j \in [0,N], 0 \le j-m \le N-i$$

The steady-state distribution s of the number of active sources is, as usual, computed as $s = s\mathbf{S}$, but can be written down immediately:

$$s_i = \binom{N}{i}\alpha^i(1-\alpha)^{N-i}.$$

The mean and variance of the number of active sources are $N\alpha$ and $N\alpha(1-\alpha)$. Because of batch-effects, the activity distribution seen by a random arrival, $\mathbf{s}'$, is a scaled version of $\mathbf{s}$, i.e., $s'_i = (is_i)/E[i]$ (random-incidence problem).

Given that i sources are active, the number of packets generated during a slot is binomially distributed with probability mass function

$$a_{i,k} = \binom{i}{k}\lambda^k(1-\lambda)^{i-k} \text{ for } k = 0,1,\ldots,i$$

and zero otherwise. The distribution of batch sizes seen by a random observer is given by $a_k = \sum_{i=1}^{N} a_{k,i}s_i$, while the distribution of batch sizes seen by a random *arrival* is given by $a'_k = ka_k/E[A]$, where $E[A]$ is the expected batch size as seen by a random observer. Naturally, it equals ρ.

3.2 The Loss Probability (LP)

As a first performance metric, we derive the customer loss probability of this queueing system through the state transition matrix. The state of the system is described by the tuple (k,i), i.e., k cells in the system (waiting and in service), with i sources active and $N-i$ silent. For brevity, k and i will be referred to as the occupancy and activity, respectively. The state is tracked on slot boundaries, with state transitions occurring on the boundary and batch arrivals immediately afterwards. As in all finite-capacity, early-arrival systems with deterministic service time and batch arrivals, the first cell in a batch always enters the system. At most a single cell departs immediately prior to the slot boundary.

The $(N,K)\times(N,K)$ transition matrix is given by

$$\begin{array}{lcll} P_{(i,q);(j,r)} & = & S_{i,j}\,a_{i,r-q+1}, & q \in [0,K], r \in (0,K) \\ P_{(i,q);(j,K)} & = & S_{i,j}\sum_{m=K-q+c}^{\infty} a_{i,m}, & q \in (0,K] \\ P_{(i,q);(j,0)} & = & S_{i,j}\sum_{m=0}^{c-q} a_{i,m} & \\ P_{(i,q);(j,r)} & = & 0 & \text{otherwise} \end{array}$$

where we define $a_{i,k} = 0$ for $k < 0$ and $k > i$. We denote the corresponding steady-state distribution as $\mathbf{p}$ with elements $p_{i,k}$.

The probability that a customer is lost, $P[L]$, is most readily computed by conservation-of-flow arguments as the ratio of channel occupancy, called ν here, to arrival rate:

$$1 - P[L] = \frac{\nu}{\rho} = \frac{1-\sum_{i=0}^{N} p_{i,0}a_{i,0}}{\rho} \text{ for } c = 1.$$

For general values of c, the mean channel occupancy is given by

$$\nu = \sum_{i=1}^{N}\left[\sum_{k=1}^{c-1}\sum_{j=0}^{i} p_{i,k-j}a_{i,j} + c\sum_{j=0}^{i} p_{i,k-j}\sum_{m=j}^{i} a_{i,m}\right]$$

3.3 The Conditional Loss Probability (CLP)

The CLP probability is computed by tracing the state of the system from the time some random customer, called the test customer, is lost to the time when the next customer from that same source, the so-called test source, arrives. The computation is only possible because arrivals from a *single* source are i.i.d., even though the aggregate arrival process is correlated.

The computation proceeds in four major steps: (1) compute the activity seen by the test customer; (2) compute the activity without the test source at the next slot boundary immediately following the loss. (At both instants, the occupancy is known to be $K+1$ or K, respectively, since at least one loss occurred.) (3) Condition on the interarrival time A of the test source. For each value of the interarrival time, we compute the state, i.e., occupancy and activity, at the slot beginning of the slot containing the next arrival from the test source. (4) We arrive at the loss probability of the next arrival from the test source by computing the number of customers from other sources that arrive in that slot.

For the first step, we need the probability $P[i|L]$ that a lost customer sees i active sources (including itself), which is most readily derived by Bayes' rule,

$$P[i-1|L] = \frac{s_i' P[L|i]}{\sum_{i=1}^{N} P[L|i] s_i'},$$

where

$$P[L|i] = \sum_{k=0}^{K} \frac{p_{i,k}}{s_i} \sum_{j=K-k+1}^{i-1} b_{i,j},$$

and s_i' is the activity seen by a random arrival, with $s_i' = is_i/E[s]$.

$P[L|i]$ depends in turn on the number of arrivals that precede a randomly chosen customer in the same slot, conditioned on the number of active sources i in a slot. Since a customer is equally likely to occupy every position within the batch, we have that $b_{i,k} = \sum_{j=k}^{i} \frac{a_{i-1,j}}{j+1}$ for $k < i$ and zero otherwise. For convenience, we define $b_{0,0} = b_{1,0} = 1$. Naturally, $P[L|i]$ is related to the loss probability:

$$P[L] = \sum_{i=1}^{N} P[L|i] s_i',$$

The initial state vector $\mathbf{q_0}$ seen by the test customer has elements

$$q_0(i,k) = \begin{cases} P[i-1|L], & \text{for } k = K \\ 0, & \text{otherwise.} \end{cases}$$

Note that here and in the following steps, we condition on the interarrival time so that the test source is known to be inactive up to the last slot of that interarrival time.

The transition matrix $\tilde{P}_0$ from the loss of the test customer to the next slot has to reflect the fact that the test source is inactive. Any arrivals that occur after the test customer are lost and do not affect the system state. Thus,

$$\tilde{P}_{0(i,k);(j,l)} = \begin{cases} \tilde{S}_{i,j}, & \text{for } k = l \\ 0, & \text{otherwise.} \end{cases}$$

Here, $\tilde{S}$ and $\tilde{P}$ are the source activity and queue state transition matrices for $N-1$ sources, respectively.

In the slot where the test source emits the next customer, the activity increases deterministically by one. Also, the arrivals before the customer from the test source increase the queue occupancy. Note that the occupancy reaches now from 0 to $K+1$. This $(N, K+1)\times(N, K+1)$ transition matrix is therefore written down as

$$\tilde{P}_\infty(i,k);(j,l) = \begin{cases} \sum_{j=l-k}^{N} b_{i+1,j} & \text{for } k \le l \le K+1, 0 \le i < N, j = i+1 \\ 1 & \text{for } 0 \le k \le K, i = j = N \\ 1 & \text{for } 0 \le i < N, j = i+1, k = l = K+1 \\ 0 & \text{otherwise.} \end{cases}$$

Given all these intermediate quantities, we then uncondition on the interarrival distribution $P[D = n]$ and obtain the state distribution vector seen by the next arrival from the test source as

$$\mathbf{q} = \sum_{n=1}^{\infty} \mathbf{q_0} \tilde{P}_0 \tilde{P}^{n-1} \tilde{P}_\infty P[D = n]. \quad (1)$$

The CLP is then $\sum_{i=0}^{N} q_{i,k}$. Note that the interarrival time distribution decreases monotonically, simplifying the truncation of the sum when numerically evaluating Eq. (1).

The computation cost is dominated by the matrix multiplication to compute $\tilde{P}^n$ and the matrix product. Since the computation time is proportional to $(N \cdot K)^3$, and, to make matters worse, larger N tend to increase the number of terms that need to be evaluated in Eq. (1), a direct computation is only feasible for relatively small systems. On the other hand, as pointed out in more detail below, the loss correlation depends very little on K, so that $K = 1$ or even $K = 0$ can be used. Also, non-random behavior is most pronounced when a small number of sources is superimposed. It remains to be determined whether bounding the CLP is possible, for example by assuming that the number of active sources remains constant in the interarrival interval. Choosing the distribution as seen by a lost customer should provide an upper bound, while choosing the distribution seen by an average arrival should bound the CLP from below. As a closer approximation, we could treat the number of active sources as a birth-death process, as is commonly done for voice multiplexors [22, 23], and then use numerical methods for sparse and block matrices. This approximation would probably improve in accuracy with larger N and longer state holding times.

While the discussion has focused on the $N \cdot \mathrm{IPP}/D/c/K$ system, a number of extensions are immediately apparent. We discuss them briefly in their rough order of complexity. A different ordering of sources in competing for buffer space (priority discarding) can be modeled by appropriately modifying $b_{i,j}$. Sources with more than two states are also readily handled at the expense of a larger state space. Note that as long as the loss probability for the individual source can be determined, only the aggregate batch arrival process of the remaining background sources enters into the calculation.

The case of an on/off source with *periodic* arrivals during the active periods can be approximately treated within this framework as well by making some simplifying assumptions. First, the interarrival time of the test source can be modeled as a constant since most of the consecutive losses will take place within the active period of the source. Then, the superposition of the remaining sources should have roughly the same impact on the test source as the superposition of the sources with Bernoulli arrivals, at least for large N. These assumptions remain to be validated by analysis. Also, the cases of several heterogeneous sources and of sources with two different Bernoulli arrival processes (MMBP) rather than on/off sources deserve closer scrutiny (see [24, 25] for a loss probability analysis).

In the Gilbert channel model [1], the system alternates between two states in a manner identical to the IPP source. In one state, every packet is lost, in the other, no losses occur. While the loss runs are indeed geometrically distributed, simulations indicate that the success runs (the sequences of customers from a single source without loss) are not. It remains to be investigated whether the model can be applied nevertheless. Naturally, the mean of the success runs can be computed from the loss probability and the loss run length (see [2]).

3.4 Numerical Examples

Compared to earlier systems, the large number of queueing system parameters (λ, γ, ω, N, K, c) covers a wide range of systems, but also makes it difficult to generalize from numerical examples.

The systems studied earlier had the property that the CLP depended only weakly, if at all, on the system size K. This property carries over into this queueing system, as indicated by Fig. 4, and other examples not shown here. The effect of K will be small as long as most of the probability mass of the interarrival time is concentrated in the first K slots [16].

The influence of the channel bandwidth is indicated by Fig. 5, where the ratio of the source on/off durations and the system load are kept constant, but the number of sources, N, is raised,

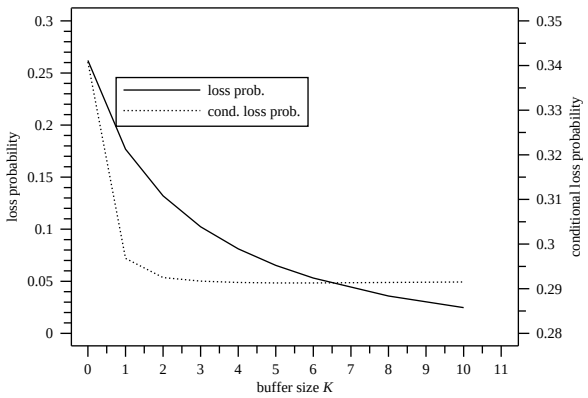

Figure 4: CLP and LP for $N \cdot \mathrm{IPP}/D/c/K$ system, as a function of the buffer size, K; $N = 4$, $\lambda = 0.8$, $\gamma = 0.7$, $\omega = 0.9$

with λ, γ and ω scaled correspondingly. Thus, the source correlation coefficient increases with N. This leads to a somewhat unexpected result: The CLP decreases with increasing N, indicating that the effect of burstiness is more than compensated for by the "thinning" during active periods.

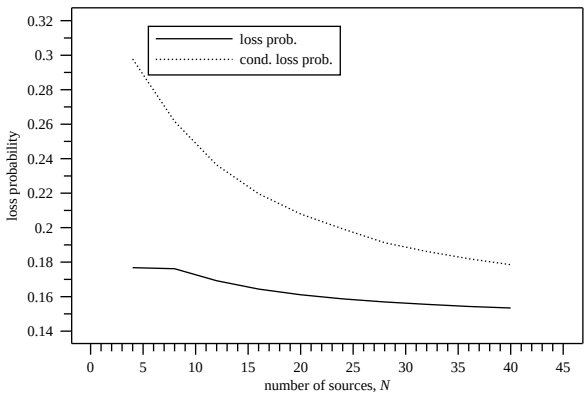

Figure 5: Conditional and unconditional loss probability for $N \cdot \mathrm{IPP}/D/c/K$ system, for constant activity, $\alpha = 0.25$, and load, $\rho = 0.8$, as a function of the number of sources, N; $K = 1$, $\lambda = 3.2/N$

As a final parameter study, we investigate the effect of different source behavior, again maintaining the same average load by adjusting λ and α. We also keep the sum of active and silent time constant. Thus, for high values of λ, a source emits a short burst of densely packed packets, followed by a long silence. A lower value of λ, conversely leads to a smoother traffic process. Fig. 6 shows that the loss probability is far more affected by the change in traffic burstiness than is the CLP.

To provide a real-life example, the loss correlation in a DS1 voice multiplexor was determined. The DS-1/T1 channel offers a user capacity of 1.536 Mb/s. Following Heffes and Lucantoni [26], each source encodes speech in 64 kb/s and generates a 128-byte packet every 16 ms. Talk spurts and silence periods are geometrically distributed with means 352 ms and 650 ms, respectively, for an activity factor of 35%. When active, this source produces a packet every 24 slots. The channel processes 1324.2 packets per second. These system characteristics translate, at a load of 0.8, to the model parameters $1 - \lambda = 0.04167$, $1 - \gamma = 2.145 \cdot 10^{-3}$, $\omega = 1.162 \cdot 10^{-3}$ and $N = 55$. At a buffer

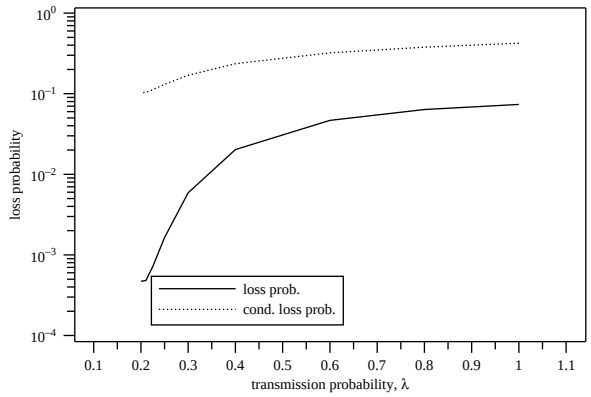

Figure 6: CLP and LP for $N \cdot \mathrm{IPP}/D/c/K$ system, for constant load ($\rho = 0.8$), cycle time ($T_0 + T_1 = 25$) and buffer size ($K = 1$), as a function of λ

size of $K = 3$, the CLP evaluates to 10.83%, or a mean loss run length of 1.12. This computation required about 20 hours of DECstation 5000 CPU time, evaluating 1483 terms of the total probability summation. A simulation confirmed these results (90% confidence on intervals 5.93 to 6.06% for loss and 10.78 to 10.99% for CLP) and showed that the numerical accuracy is satisfactory.

For a smaller multiplexing factor, consider video sources on an STS-3 channel, with a usable capacity of 150.336 Mb/s or 354566 ATM cells/s. Unlike voice sources, video sources have not been well characterized in terms of standard source models. We approximately model the CATV video coder of Verbiest and Pinnoo [27]. It has an average rate of 26.5 Mb/s and a standard deviation of 3.4 Mb/s. We assume that the coder transmits at channel capacity and that the active and silent period add up to the frame duration of 1/25 s. We choose model parameters $\lambda = 1$, $1 - \gamma = 3.623 \cdot 10^{-4}$, $1 - \omega = 8.755 \cdot 10^{-5}$, and, again for a load of approximately 0.8, $N = 5$. With these parameters and $K = 1$, the CLP and mean loss run evaluate to the significantly higher values (compared to the voice example) of 0.5658 and 2.303, respectively. Thus, appropriate packet interleaving or priority mechanisms must be used to reduce the incidence of noticeable display artifacts.

Li also investigated the loss behavior for a voice source model [28] related to the $N \cdot \mathrm{IPP}/D/c/K$ queue studied here. However, he considered the temporal properties of the *aggregate* loss rate for all sources, not the behavior of individual sources. However, he also arrives at the conclusion that the buffer size does not materially effect the loss behavior once an overload period sets in.

4 Summary and Conclusions

In this paper, we have attempted to provide both methodologies and traffic engineering insight. The tools provided should allow the computation of the loss correlation for a variety of sources of practical interest in high-speed networking.

Continuing the theme of [2], it was shown qualitatively that the buffer size has virtually no influence on the loss correlation. For periodic arrivals, we saw that as long as individual sources do not contribute more than, say, one tenth of the total bandwidth, the average loss run lengths are not significantly larger than one, even though the CLP can be orders of magnitude higher than the LP. As expected, the traffic burstiness and the number of multiplexed sources play significant roles in the loss correlation for the IPP system.

This investigation suggests a range of further explorations. Throughout this paper, we have used the CLP, or, equivalently, the average loss run length as an indicator of loss correlation. Another useful indicator could be the number of lost packets among a random sequence of packets with a given length. Also, the distribution of success runs deserves closer investigation. A discrete-time model where arrivals are periodic during the on-period is under investigation. Particularly for queues with correlated inputs, approximation methods need to be investigated that allow systems with large state spaces to be treated.

References

[1] E. N. Gilbert, "Capacity of a burst-noise channel," *Bell System Technical Journal*, vol. 39, pp. 1253–1265, Sept. 1960.

[2] H. Schulzrinne and J. F. Kurose, "Distribution of the loss period for some queues in continuous and discrete time," in *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, (Bal Harbour, FL), pp. 1446–1455 (12C.1), Apr. 1991.

[3] J. M. Ferrandiz and A. A. Lazar, "Consecutive packet loss in real-time packet traffic," in *Proceedings of the Fourth International Conference on Data Communication Systems*, (Barcelona), pp. 306–324, IFIP TC6, June 1990.

[4] J. M. Ferrandiz and A. A. Lazar, "A study of loss in $N/GI/1$ queueing systems." Department of Electrical Engineering and Center for Telecommunications Research, Columbia University, New York, NY, Jan. 1990.

[5] J. M. Ferrandiz and A. A. Lazar, "Rate conservation for stationary processes," *Journal of Applied Probability*, Mar. 1991.

[6] I. Cidon, A. Khamisy, and M. Sidi, "Analysis of packet loss processes in high-speed networks," EE Publication 786, Dept. of Electrical Engineering, Technion Israel Institute of Technology, Haifa, Israel, May 1991.

[7] A. K. Kaul, "Performance of high-level data link control in satellite communications," *COMSAT Technical Review*, vol. 8, pp. 41–88, Spring 1978.

[8] F. C. Fujiwara, M. Kasahara, K. Yamashita, and T. Namekawa, "Evaluations of error-control techniques in both independent and dependent-error channels," *IEEE Transactions on Communications*, vol. COM-26, pp. 785–793, June 1978.

[9] D. Towsley, "A statistical analysis of ARQ protocols operating in a nonindependent error environment," *IEEE Transactions on Communications*, vol. COM-29, pp. 971–981, July 1981.

[10] C. H. C. Leung, Y. Kikumoto, and S. A. Sorensen, "The throughput efficiency of the go-back-n ARQ scheme under Markov and related error structures," *IEEE Transactions on Communications*, vol. 36, pp. 231–234, Feb. 1988.

[11] G. R. Pieris and G. H. Sasaki, "The performance of simple error control protocols under correlated packet losses," in *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, (Bal Harbour, FL), pp. 764–772 (7C.1), IEEE, Apr. 1991.

[12] H. Ohta and T. Kitami, "Simulation study of the cell discard process and the effect of cell loss compensation in ATM networks," *The Transactions of the IEICE*, vol. E73, pp. 1704–1711, Oct. 1990.

[13] A. A. Lazar, G. Pacifici, and J. S. White, "Real-time traffic measurements on MAGNET II," *IEEE Journal on Selected Areas in Communications*, vol. 8, pp. 467–483, Apr. 1990.

[14] N. Shacham and P. McKenney, "Packet recovery in high-speed networks using coding and buffer management," in *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, (San Francisco, CA), pp. 124–131, IEEE, June 1990.

[15] T. Kitami and I. Tokizawa, "Cell loss compensation schemes in an asynchronous broadband ISDN," in *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, (San Francisco, CA), pp. 116–123, IEEE, June 1990.

[16] H. Schulzrinne, J. F. Kurose, and D. Towsley, "Distribution of the loss period for some queues in continuous and discrete time," Tech. Rep. TR92, Department of Computer and Information Science, University of Massachusetts, Amherst, MA, 1992.

[17] A. Kuczura, "The interrupted Poisson process as an overflow process," *Bell System Technical Journal*, vol. 52, pp. 437–448, Mar. 1973.

[18] J. H. Rath and D. Sheng, "Approximations for overflow from queues with a finite waiting room," *Operations Research*, vol. 27, pp. 1208–1216, November–December 1979.

[19] M. Murata, Y. Oie, T. Suda, and H. Miyahara, "Analysis of a discrete-time single-server queue with bursty inputs for traffic control in ATM networks," *IEEE Journal on Selected Areas in Communications*, vol. 8, pp. 447–458, Apr. 1990.

[20] Y. Ohba, M. Murata, and H. Miyahara, "Analysis of interdeparture processes for bursty traffic in ATM networks," *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 468–476, Apr. 1991.

[21] J. J. Hunter, *Mathematical Techniques of Applied Probability – Discrete Time Models: Techniques and Applications*, vol. 2. New York, NY: Academic Press, 1983.

[22] K. Sriram, P. K. Varshney, and J. G. Shanthikumar, "Discrete-time analysis of integrated voice/data multiplexers with and without speech activity detection," *IEEE Journal on Selected Areas in Communications*, vol. SAC-1, pp. 1124–1132, Dec. 1983.

[23] S. Li and J. W. Mark, "Performance of voice/data integration on a TDM system," *IEEE Transactions on Communications*, vol. COM-33, pp. 1265–1273, Dec. 1985.

[24] J. J. Bae, T. Suda, and R. Simha, "Analysis of a finite buffer queue with heterogeneous markov modulated arrival processes: A study of the effect of traffic burstiness on individual packet loss." submitted to *Infocom 1992*, May 1992.

[25] J. J. Bae, T. Suda, and R. Simha, "Heterogeneous arrival streams, burstiness and packet discarding: A study of individual packet loss," Technical Report 91-58, Dept. of Computer and Information Science at the University of California, Irvine, Irvine, CA, July 1991.

[26] 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*, vol. SAC-4, pp. 856–867, Sept. 1986.

[27] W. Verbiest and L. Pinnoo, "A variable bit rate codec for asynchronous transfer mode networks," *IEEE Journal on Selected Areas in Communications*, vol. 7, pp. 761–770, June 1989.

[28] S. Li, "Study of information loss in packet voice systems," *IEEE Transactions on Communications*, vol. 37, pp. 1192–1202, Nov. 1989.