A Trace-Driven Simulation of an ATM Queueing System

Gilberto Mayor and John Silvester

CENG 96-17

Department of Electrical Engineering - Systems
University of Southern California
Los Angeles, California 90089-2562
(213) 740-4579

August 1996

A Trace-Driven Simulation of an ATM Queueing System

Gilberto Mayor and John Silvester
{gmayor, silveste}@usc.edu
Department of Electrical Engineering-Systems
University of Southern California
Los Angeles, California 90089-2562

Abstract

A series of studies have shown that local area network traffic can be modeled by a fractional Brownian noise (fBn) process. In this work, we compare the statistics of an ATM queueing systems driven by i) an fBn process and ii) an auto-regressive (AR) source, to a system driven by real network traffic. We show that the ATM queue system fed by realistic traffic suffers larger cell losses than the AR system's predictions. In this case, the fBn model is able to give more accurate performance results. Moreover, we verified that the unfinished work of the realistic queueing system exhibits long-range dependence (LRD) and the infinite variance syndrome (IFV). We also show that the busy period distribution has a heavy tail. We conclude that those phenomena are responsible for the large volume of cell losses experienced by this queueing system. Our results are obtained by a trace-driven simulation of an ATM queueing system.

1 Introduction

Bellcore researchers showed that local area network traffic can be better characterized by a self-similar process [1], *e.g.* a fractional Brownian noise (fBn), since they are able to reproduce the long-range dependence (LRD) exhibited by real network traffic. [1] In this work, we want to investigate if this better representation of the arrival traffic leads to more accurate queueing models. We compare the performance of an ATM queueing system driven by three different traffic sources: i) real network traffic, ii) a fBn process, and iii) an Auto-Regressive (AR) process. The real network traffic is a publicly available trace of actual Bellcore Ethernet traffic and has been shown to exhibit LRD. [2] The fBn sample was generated by an algorithm proposed by Chi [2, 3, 4]. The AR process is commonly used to model network traffic sources and is only able to model short-range dependence (SRD). All three arrival processes have the same mean and variance, the most distinguished feature among them is the strength of their long-range correlation structure.

Moreover, it has been shown that LRD can cause large cell losses in ATM systems [5, 4]. In fact, Duffield [6, 7] showed that the tail distribution of an ATM queue driven by a fBn process follows a Weibull distribution, *i.e.* it differs from standard Markov results. Here, we investigate the dynamics of this ATM system in order to explain these *unexpected* cell losses.

[1] The real valued-process $X(t)$ is *self-similar* with index $H > 0$ if for all $a > 0$, $X(at) \doteq a^H X(t)$.

[2] The samples are available with anonymous FTP from ftp.bellcore.com.

2 A Trace-Driven Simulation

We simulate a discrete-time, deterministic service, queueing (ATM) system, driven by three different arrival samples containing 1,000,000 points each. We changed the departure rate in order to achieve different utilization levels. We computed the cell loss probability and the average cell delay in the system.

2.1 Cell Loss Probability

In the first experiment, the buffer size is equal to 100 cells. Figure 1 shows that at lower utilization values, the AR system achieves similar results to the real system. In this case, we conclude that the SRD is responsible for the cell losses, *i.e* Markov models can be used to compute queueing statistics. On the other hand, the fBn generator is able to accurately reproduce the level of cell losses at higher utilization values. In this scenario, the AR system clearly underestimates cell losses. Therefore, the LRD is responsible for the cell losses at higher utilization.

In the second experiment, we increased the buffer size to 1000 cells. Figure 2 shows that the increase of the buffer was able to significantly reduce cell losses in the AR system. The buffer accommodates the SRD component of the arrival processes, so that the LRD (low-frequency) component is responsible for cell losses [9],. Therefore, although the real system achieves significant cell losses, the AR system clearly underestimates it. In this case, the fBn system can accurately predict the cell loss probability.

2.2 Average Delay

In this section, we measured the average cell delay in the system. The AR process is able to reproduce accurate cell delay for utilization levels below 45%. For higher utilization values, the AR queueing system clearly underestimates the average cell delay, see figure 3. In this case, the fBn queueing system gives more accurate results. In figure 4 we increased the buffer size to 1000 cells. The AR system underestimates the delay increase experienced by the real traffic. Again, the fBn system results in better delay prediction.

2.3 Matching Delay and Cell Losses

The fBn queueing system gives optimistic results because the covariance structure of the fBn artificial sample decays faster than the real data's covariance structure, *i.e.* it has a weaker LRD component. In order to get more accurate results, we tried to get a better match of both covariance structures. We defined $Y_f(t) = kX_f(t)$ where $X_f(t)$ is the original fBn process and k is a constant, *i.e.* the covariance of $Y_f(t)$ is given by k^2 times $X_f(t)$'s covariance. We choose k empirically so that both covariance structures have similar values up to a certain lag s. Figure 5 shows that both queueing systems achieve similar cell loss probability. Therefore, the fBn model can accurately reproduce cell losses. Moreover, it is clear that the arrival model does not need to reproduce the entire tail of the real traffic's covariance structure in order to reproduce similar queueing behavior.

2.4 Explanation for the Cell Losses

Mandelbrot described a *infinite variance (IFV) syndrome* associated to a self-similar process [8]. Here, we show that if the arrival traffic *suffer* from the IFV syndrome, so does the queueing system. Let *Q(t)* be the number of cells in an infinite buffer system at time t, *i.e.* the unfinished work. Let $X(t)$ represent the aggregate arrival process at time t, *i.e.* the number of packets that joined the system until time t. Let $\lambda = E[X(t)]$ be the average arrival rate, c be the constant departure rate, and $\rho = \lambda/c < 1$ be the utilization. We can write [10]

$$Q(t) = X(t) - ct \tag{1}$$

Therefore $Var[Q(t)] = Var[X(t)] - Var[ct] = Var[X(t)]$, *i.e.* the unfinished work has similar variance to the arrival traffic. We measured the variance of *Q(t)* for a 1000 cell queueing system. Figure 6 shows that the unfinished work of the system fed by real traffic presents much higher variability than the AR system's unfinished work.

We also investigate if *Q(t)* exhibits LRD. We can compute the autocorrelation function by using equation [1]: $E[Q(t_1)Q(t_2)] = E[X(t_1)X(t_2)] - c\lambda(t_1+t_2) + c^2t_1t_2$. Therefore, whenever $c\lambda(t_1+t_2) < c^2t_1t_2$ the autocorrelation of the unfinished work is even stronger than the autocorrelation of the arrival process. We calculated the correlation coefficient of *Q(t)* for the real system and for the AR system at 60% utilization for a 1000 cell buffer. We can see in figure 7 that the real system has indeed LRD. In fact, its correlation structure decays even slower than the original arrival process's correlation. The explanation is that when the system is idle, the unfinished work does not change, generating this strong correlation structure. This phenomena has strong implication on queueing performance, *i.e.* we might experience very long *busy/idle* periods. Therefore, congestion might persist for very long periods of time. We ought to be able to re-route traffic from a *busy* link to another *available* link, on-demand. Another possible solution is to allocate bandwidth to Virtual Paths dynamically .

We also computed the unfinished work for a queueing system with different buffer sizes fed by the real traffic at 60% utilization, see figure 8. We verified that when massive cell losses occur the unfinished work presents very similar correlation coefficient to the arrival process' coefficient. In fact, for a 20 cell buffer, the correlation structure of the unfinished work is almost the same as the arrival process. On the other hand, when using larger buffers (even if losses occur) the correlation coefficient of the unfinished work is even higher than the arrival process' coefficient. For example, in figure 8, the loss probability for the 1000 cell buffer is $O(10^{-2})$ but it presents a very strong correlation structure.

2.5 The Busy Period

We investigated the busy period distribution for the real system. Since this distribution is related to the unfinished work distribution, we should expect that it also presents a large variance. In figures 9 and 10, we show the AR system's busy period distribution and the tail of the real traffic's busy period distribution respectively. By comparing both figures, we can see that the real traffic's busy period presents much larger variability than the AR's busy period. In fact, the real traffic's distribution has a very long tail, *i.e.* this system experiences very long busy periods that are responsible for most of the cell losses.

2.6 The Output Process

We checked if the output process of the queueing system presents LRD. Since a queueing system can be seen as a low-band pass filter, the low frequency should remain intact if the cell losses are negligible [11]. Nevertheless, we observed that the queueing system presents LRD, despite of the volume of cell losses, see figure 11. In fact for buffer sizes equal to 100 and 1000 cells, the output process's correlation coefficient structure is remarkably similar to the input process's structure, even though the smaller buffer achieves a $O(10^{-2})$ cell loss probability. On the other hand, a 20 cell buffer filters most of the high-frequency component of the input process, so that the resulting correlation structure has a weak SRD component. It also filters part of the low-frequency structure, but the remaining correlation structure still presents LRD. Therefore, a fBn process is also a good model for the output process of a queueing system fed by an input process with long-range dependence.

3 Conclusion

Our work has confirmed recent studies in showing that traditional Markovian queueing systems can underestimate the average cell delay and cell loss probability experienced by real ATM networks. We also showed that whenever the high-frequency component of the arrival process is buffered, the LRD dominates the bandwidth requirements. In this scenario, a traffic model needs to *emulate* part of the tail of the covariance structure exhibit by real network traffic in order to achieve accurate queueing statistics.

Moreover, we explained the large volume of cell losses experienced by real ATM queueing systems by showing that the unfinished work of those queueing systems present *long-range* dependence and the *infinite-variance* syndrome associated with the self-similar (fBn) arrival process. Moreover, we showed that the output process of a queueing system fed by a long-range dependence input process has also long-range dependence even in the presence of massive cell losses. Therefore, fBn processes are also good candidates for modeling the output process of an ATM queueing system.

We also concluded that in order to archive higher degrees of link utilization, ATM management protocols need to be able to either re-route traffic on demand whenever it sees a *busy* link, or to allocate bandwidth dynamically.

Acknowledgments

This work was supported in part by a grant from Intel Corporation. Gilberto Mayor is supported by a grant from CAPES/Brazil.

References

[1] W. Leland, M. Taqqu, W. Willinger and D. Wilson,"On the Self-Similar Nature of Ethernet Traffic (Extended Version)", IEEE/ACM Transactions on Networking, February 1994.

[2] M. Chi, E. Neal and G. Young, "Practical Application of Fractional Brownian Motion and Noise to Synthetic Hydrology", Water Resources Research, Vol. 9, 1523-1533, December 1973.

[3] B. Mandelbrot,"A Fast Fractional Gaussian Noise generator", Water Resources Research, Vol. 7, 543-553, 1971.

[4] G. Mayor and J. Silvester,"An ATM Queueing System with a Fractional Brownian Noise Arrival Process", IEEE ICC'96.

[5] A. Erramilli, O. Narayan and W. Willinger, "Experimental Queueing Analysis with Long-Range dependence Packet Traffic", submitted to IEEE/ACM Transactions on Networking.

[6] N. Duffield, J. Lewis and N. O'Connel, "Predicting Quality of Service for Traffic with Long-Range Fluctuations", IEEE ICC'95.

[7] I. Norros, "A Storage Model with Self-Similar Input", Queueing Systems 16, 1994.

[8] B. Mandelbrot, "Long-run linearity, locally Gaussian Processes, H-spectra and infinite variances. International Economic Review, Vol. 10, 1969.

[9] H.D. Sheng and S.Q.Li,"Spectral Analysis of Packet Loss Rate at a Statistical Multiplexer for Multimedia Services", IEEE/ACM Transactions on Networking, January 1994.

[10] V. Benes,"General Stochastic Processes in the Theory of Queues", Addison-Wesley, 1963.

[11] S. Q. Li et al., "Link Capacity Allocation and Network Control by Filtered Input Rate in High-Speed Networks", IEEE/ACM Transactions on Networking, Vol. 3, February 1995.

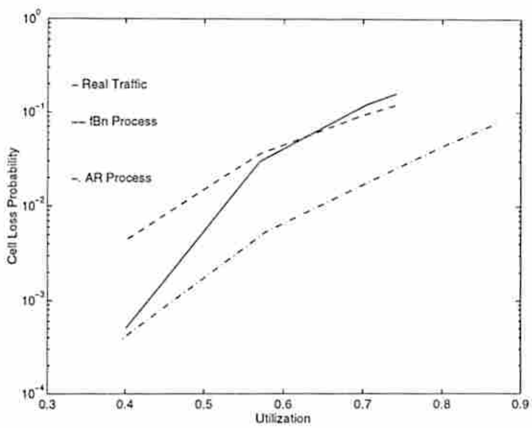

Figure 1: Loss probability for a 100 cell buffer.

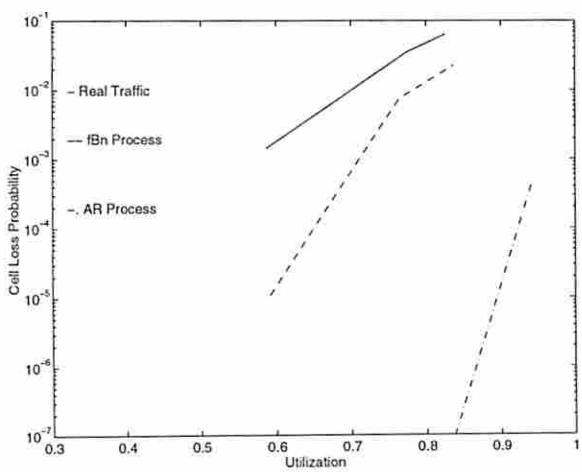

Figure 2: Loss probability for a 1000 cell.

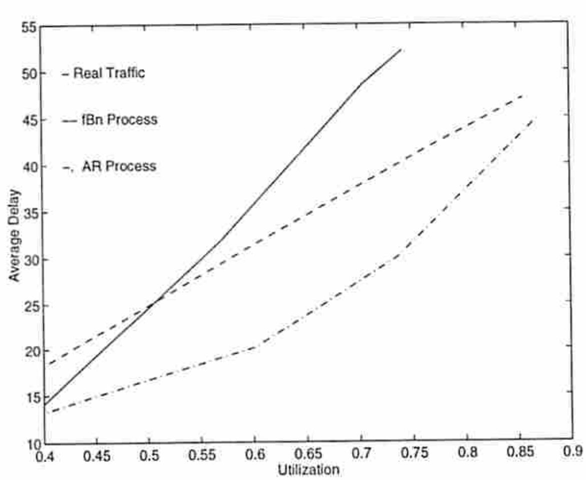

Figure 3: Average Delay for a 100 cell buffer.

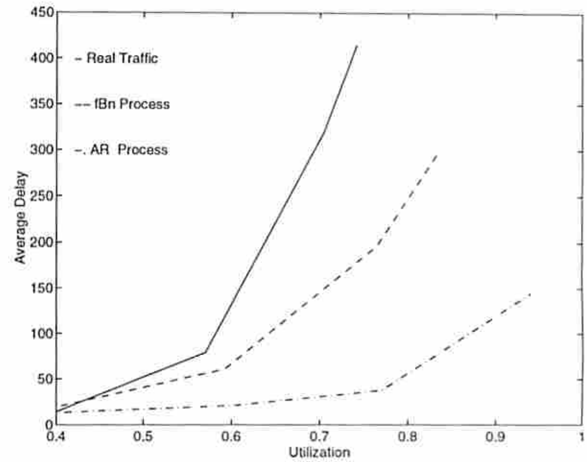

Figure 4: Average Delay for a 1000 cell buffer.

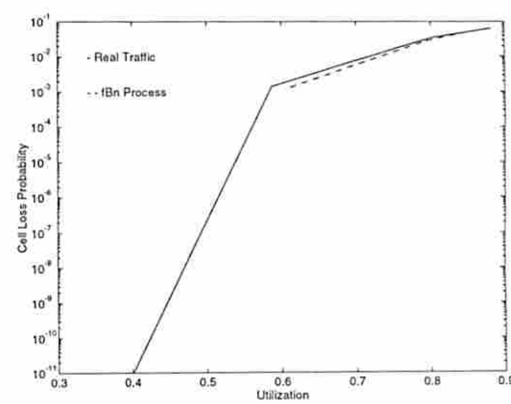

Figure 5: Cell Loss Probability for the modified fBn.

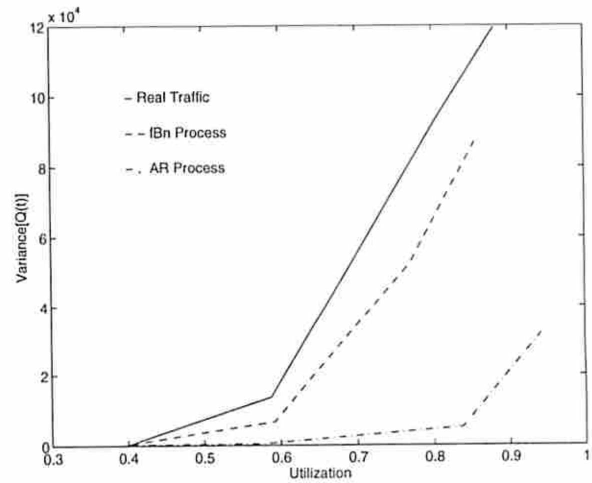

Figure 6: Variance of the unfinished work.

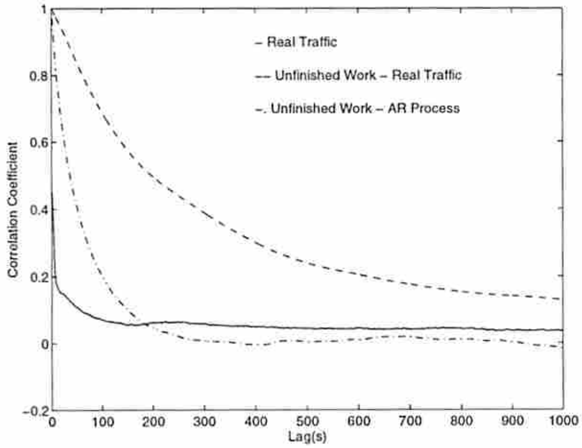

Figure 7: Correlation coefficient of the unfinished work for the real traffic and AR source.

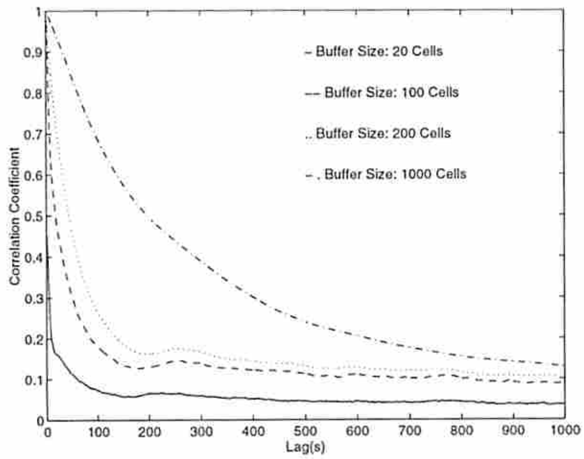

Figure 8: Correlation coefficient of the unfinished work.

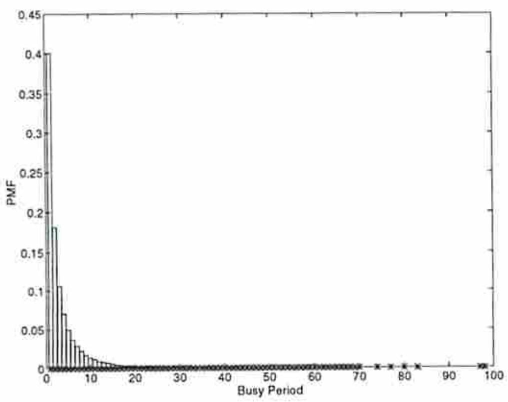

Figure 9: Busy Period for the AR source queueing system.

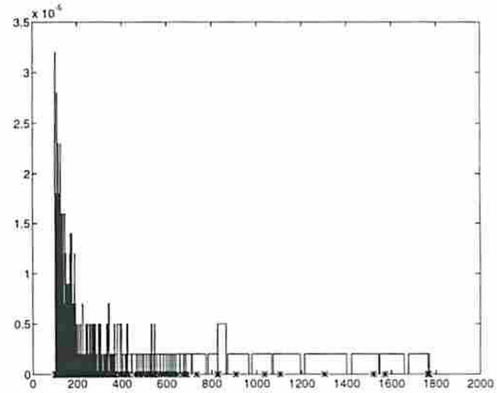

Figure 10: Busy Period for the Real Traffic queueing system.

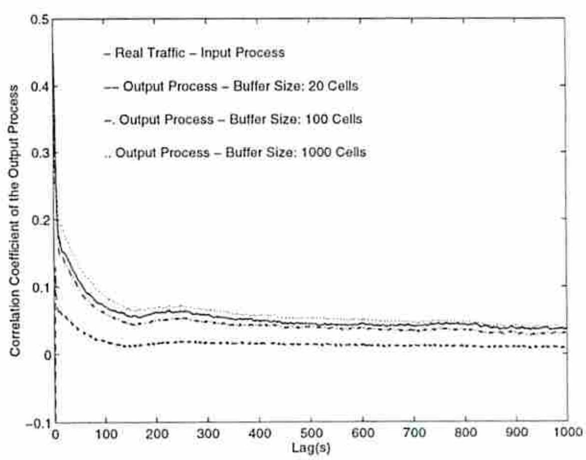

Figure 11: Correlation coefficient of the output process.