A MODEL FOR SELF-SIMILAR ETHERNET LAN TRAFFIC: DESIGN, IMPLEMENTATION, AND PERFORMANCE IMPLICATIONS

Ying Chen Zhiang Deng Carey L. Williamson
Department of Computer Science
University of Saskatchewan
Saskatoon, SK, Canada S7N 5A9
e-mail: carey@cs.usask.ca

Keywords: Network, Model Design, Validation, Performance Analysis, Telecommunications

ABSTRACT

This paper presents a synthetic workload model that is capable of modeling self-similar Ethernet traffic. The design of the model is presented, along with a brief description of its implementation in SimKit, a C^{++}-based simulation interface to support efficient parallel simulation on Shared Memory TimeWarp (SMTW).

Preliminary simulation results are presented comparing the cell loss and cell delay performance in a simulated ATM network, both with and without the presence of self-similar traffic. Results suggest that simple Poisson traffic models, as opposed to self-similar traffic models, may underestimate ATM cell loss by several orders of magnitude, depending on the network utilization. Self-similar traffic models thus provide a valuable tool for "worst-case" stress testing of ATM networks.

1 INTRODUCTION

Recent research in network traffic measurement has established the "self-similar" nature of aggregate Ethernet local area network (LAN) traffic (Leland *et al.* 1994), as well as the presence of self-similarity in Internet traffic (Paxson and Floyd 1994) and video traffic (Garrett and Willinger 1994) as well. This "non-Poisson" behaviour has serious implications on the performance of networks and systems that have been designed on the assumption of uncorrelated Poisson traffic.

Understanding the performance implications of self-similar traffic is thus important to network designers and planners, and is especially relevant in the context of ATM networks (Frost and Melamed 1994). ATM networks are based on a statistical multiplexing access method that assumes that the law of large numbers holds: given a large number of independent traffic flows, the aggregate bandwidth requirements for all the traffic flows is stable and predictable. Self-similarity suggests that the law of large numbers does not hold: aggregating many bursty sources together still results in a bursty aggregate stream.

The ATM-TN TeleSim project has thus incorporated a self-similar Ethernet traffic model into its set of synthetic workload models for the ATM-TN network simulator. The Ethernet model is capable of generating self-similar traffic with various degrees of self-similarity, and various network utilizations. This model is useful as a background traffic generator for network backbone traffic, as well as a "worst-case" stress test for ATM switches and ATM networks.

The remainder of this paper is organized as follows. Section 2 presents some mathematical background on self-similarity, and describes the design and implementation of our synthetic Ethernet traffic model. Section 3 describes model validation, and Section 4 presents a set of simple simulation experiments designed to compare self-similar and Poisson traffic models in the context of ATM networks. Finally, Section 5 concludes the paper.

2 ETHERNET TRAFFIC MODEL

2.1 Purpose

The Ethernet LAN traffic model in the ATM-TN TeleSim project is designed to represent the workload characteristics of aggregate data packet traffic on existing local area networks, such as campus LANs. Since one of the anticipated uses of ATM is as a backbone interconnect between existing LANs, LAN-LAN traffic may form a significant fraction of the background load on an ATM network.

The purpose of this traffic model is to model the aggregate LAN data traffic workload as accurately as possible. The model produces a unidirectional flow of ATM cells, with statistical characteristics matching those of aggregated traffic.

2.2 Background

Modeling studies of Ethernet performance often make many simplifying assumptions, such as assuming that the packet arrival process is Poisson. The results of some recent measurement studies, however, have shown that traffic is very much *not* Poisson (Leland *et al.* 1994; Paxson and Floyd 1994). In particular, Ethernet LAN traffic has been found to have a "fractal" self-similarity behaviour. Such behaviour has potentially serious implications on the design, control, modeling and analysis of high speed cell-based networks.

The paper by Leland *et al.* (Leland *et al.* 1994) demonstrates that Ethernet LAN traffic is statistically self-similar in the sense that there is no natural length of a burst: at time scales ranging from milliseconds to seconds to minutes, similar-looking traffic bursts are evident. Mathematically speaking, a stochastic time series process $X = (X(t) : t = 0, 1, 2, \ldots)$ is said to be self-similar with self-similar Hurst parameter H if the process is covariance stationary (that is, the process has constant mean and finite constant variance) and the corresponding aggregated process has the same correlation structure as the original process (exactly self-similar) or agrees asymptotically with the correlation structure of the original process (asymptotically self-similar) over large intervals. The Hurst parameter expresses the behaviour of the rescaled adjusted range statistic (R/S statistic) for large sample sizes n (Leland *et al.* 1994). For many naturally occurring time series, including Ethernet traffic, $E[R(n)/S(n)]$ is proportional to n^H, with H "typically" about 0.7. For models with only short-range dependence, H is 0.5. This discrepancy is called the Hurst effect (Leland *et al.* 1994).

Self-similarity manifests itself in a number of ways. One manifestation is a slowly decaying variance. That is, if you take the original time series X and compute a new smoothed time series by averaging every non-overlapping group of m observations to get a new time series $X^{(m)}$, then the variance of the sample mean for the new time series $X^{(m)}$ decreases more slowly than the sample size m (i.e., $Var(X^{(m)}) \approx am^{-\beta}$ as $m \to \infty$, where a is a constant, and $0 < \beta < 1$). On a log-log plot of the variance of $X^{(m)}$ versus m, the result is a straight line with (negative) slope β. The Hurst parameter H can be estimated using $H = 1 - \frac{\beta}{2}$. A second manifestation is a hyperbolically decaying autocorrelation function of the form $AF(k) \approx ck^{-\beta}$, for some constant c. If X is self-similar, then the autocorrelation function for $X^{(m)}$ is indistinguishable from the autocorrelation function for X, for any value of m. A third manifestation is a straight line of slope β on a log-log plot of the R/S statistic.

Ethernet traffic has been found to have this self-similarity property, and the degree of self-similarity tends to increase with the load level on the Ethernet (Leland *et al.* 1994). Furthermore, the self-similarity property has been found to hold both in the traffic within a LAN, as well as in the traffic leaving a LAN destined for the Internet. In short, the self-similarity property is pervasive enough to merit its inclusion in the traffic model for aggregate LAN packet traffic.

2.3 Modeling Methodology

A synthetic workload model for self-similar Ethernet traffic was constructed using a two step process.

The first step of the modeling process used an S (Venables 1990) program to generate a pseudo-self-similar time series[1] based on Fast Fourier Transform inversion. The S program produces as output a zero-mean time series with self-similarity controlled by an input parameter H (the Hurst parameter). Time series were generated for five different values of H (namely, $H = 0.5$, 0.6, 0.7, 0.8, and 0.9). The output time series from the S program were then translated and scaled (within S) to produce non-negative Ethernet packet count time series with a specified mean network utilization U. These synthetically generated time series were then used as input to the second step of the modeling process.

The second step of the modeling process used the TES (Melamed 1992) modeling methodology. TES provides a very general approach to modeling first-order (i.e., frequency histogram) and second-order (i.e., autocorrelation function) statistics of time series. In this case, the time series of interest is the "packet count per interval" time series from the first step above. The `testool` (Geist and Melamed 1992) program was used to construct an abstract TES model that had the same frequency histogram and autocorrelation function as the (synthetic) Ethernet time series. One TES model was constructed for each value of H considered. The resulting TES models were converted (automatically) to C^{++}, and incorporated into the ATM-TN SimKit model.

2.4 Parameters

The Ethernet traffic model in the ATM-TN TeleSim project has two parameters:

[1] A definition of pseudo-self-similarity is given in Appendix F of a paper by Paxson and Floyd (Paxson and Floyd 1994) (extended version, available by ftp). The source code for the S program (`ss.fourier`) is given in the same appendix.

- **H**: The Hurst parameter $0.5 \leq H < 1$ expresses the degree of self-similarity in the traffic. Data sets with higher self-similarity parameter H have higher "burstiness", in the intuitive sense.

  Empirical measurements of Ethernet traffic (both in (Leland *et al.* 1994) and at the University of Saskatchewan) suggest that the value of H for aggregated LAN traffic is $0.7 \leq H \leq 0.9$.

- **U**: The network utilization $0 \leq U \leq 1$ determines the average load level offered by the Ethernet LAN, as a fraction of its total 10 Mbps bandwidth. Most Ethernets operate at low to moderate levels of load (e.g., 1% to 30% utilization).

  The utilization U has been empirically observed to have an effect on the degree of self-similarity observed in the traffic (Leland *et al.* 1994). The higher the network load, the higher the degree of self-similarity usually observed. Despite this observation in the literature, the Ethernet model in the TeleSim project treats H and U as two independent and orthogonal parameters.

2.5 Model Implementation

The model has been implemented in SimKit, a simulation language that provides an interface to Shared Memory TimeWarp (SMTW), an environment for efficient parallel simulation. At this point in time, only the sequential implementation of the model is complete. Parallel implementation of the model is still in progress.

The model produces a unidirectional flow of ATM cells between components (e.g., switches, ports, links) in the overall ATM-TN model. The model uses two logical processes (LPs): one for the Ethernet source, and one for the Ethernet sink. The Ethernet model contains 1905 lines of SimKit code, of which approximately 1400 lines are for the TES models. A standalone version of the model (in C^{++}, not SimKit) is also available.

3 MODEL VALIDATION

Throughout the modeling process, measurements from Ethernet LANs at the University of Saskatchewan have been used in the construction and parameterization of the model.

Measurements of Ethernet traffic at the University of Saskatchewan have verified the presence of the self-similarity property, and estimates of the Hurst parameter have been made (e.g., $H \approx 0.73$). The main methods used in the analysis are variance-time plots, and time-domain analysis based on the R/S-statistic. The latter method subdivides a given sample of N observations into K non-overlapping blocks and computes the rescaled adjusted range (R/S statistic) defined in (Leland *et al.* 1994). Graphical R/S analysis is used to estimate the Hurst parameter H from the graph's asymptotic slope. More detail about this method can be found in (Leland *et al.* 1994). The same analysis methods are used to assess the self-similarity of the synthetically generated Ethernet traces.

Figures 1, 2, 3, and 4 illustrate several characteristics of the synthetically generated Ethernet traces. Figure 1 shows a time series plot of the Ethernet trace for $H = 0.5$, along with its autocorrelation function. This trace, as expected, has no long range dependence, and no significant degree of self-similarity. These observations are illustrated by the variance-time plot and the R/S plot as well. That is, the variance-time plot has slope -1, and the R/S plot has slope 0.5 (the same as the dashed line in the figure).

Figure 2 shows a time series plot for the $H = 0.9$ Ethernet time series. This trace has a significant degree of self-similarity, as reflected by the slowly decaying autocorrelation function, the flatter variance-time plot, and the steeper R/S pox diagram. For example, the average slope of the R/S plot is approximately 0.9, slightly less than the slope of 1.0 shown by the solid line in Figure 2.

Finally, Figures 3 and 4 show the effects of aggregating (in the time domain) the Ethernet traffic time series, for different values of the Hurst parameter H, and different levels of aggregation M. These plots are shown in a format similar to Figure 4 in the paper by Leland *et al.* (Leland *et al.* 1994). When $H = 0.5$ in Figure 3 (i.e., the lefthand column), the time series aggregates to smooth traffic quite quickly as M increases from 1 to 1000. This behaviour is consistent with the law of large numbers for independent traffic arrivals (i.e., Poisson traffic). When $H = 0.9$ in Figure 4 (i.e., the right-hand column), the traffic remains bursty across three or four orders of magnitude[2] in the time domain, similar to (Leland *et al.* 1994). We thus believe that our synthetic workload model is a good generator of self-similar traffic.

4 PERFORMANCE IMPLICATIONS

What are the performance implications of self-similar traffic? As mentioned earlier, self-similar traffic implies that the law of large numbers does not hold, and thus

[2] One possible explanation for why our traffic model becomes smooth sooner than the traffic shown in (Leland *et al.* 1994) is that the traffic in (Leland *et al.* 1994) undoubtedly contains time-of-day changes in the network utilization, whereas our model uses a fixed value of network utilization (U) for the entire duration of the trace. Another possible explanation is that our model is based on pseudo-self-similarity (Paxson and Floyd 1994), not true self-similarity.

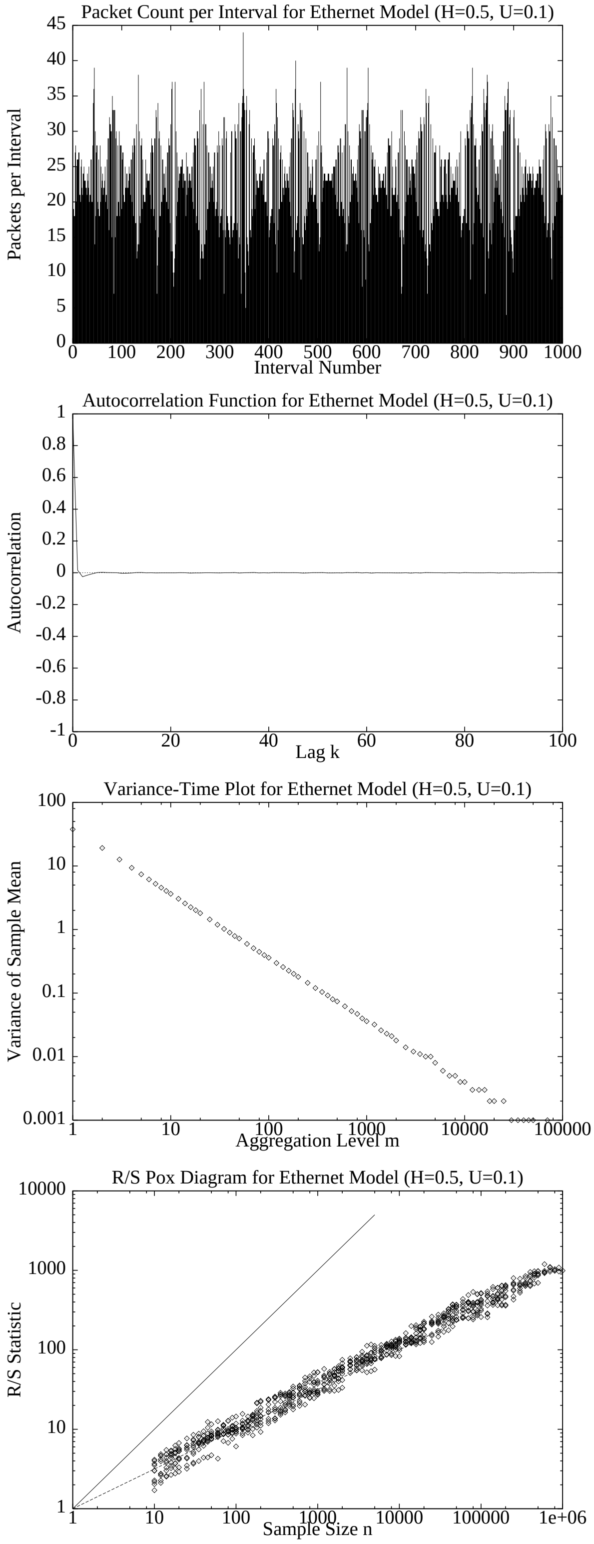

Figure 1: Ethernet traffic model for $H = 0.5$ ($U = 0.1$)

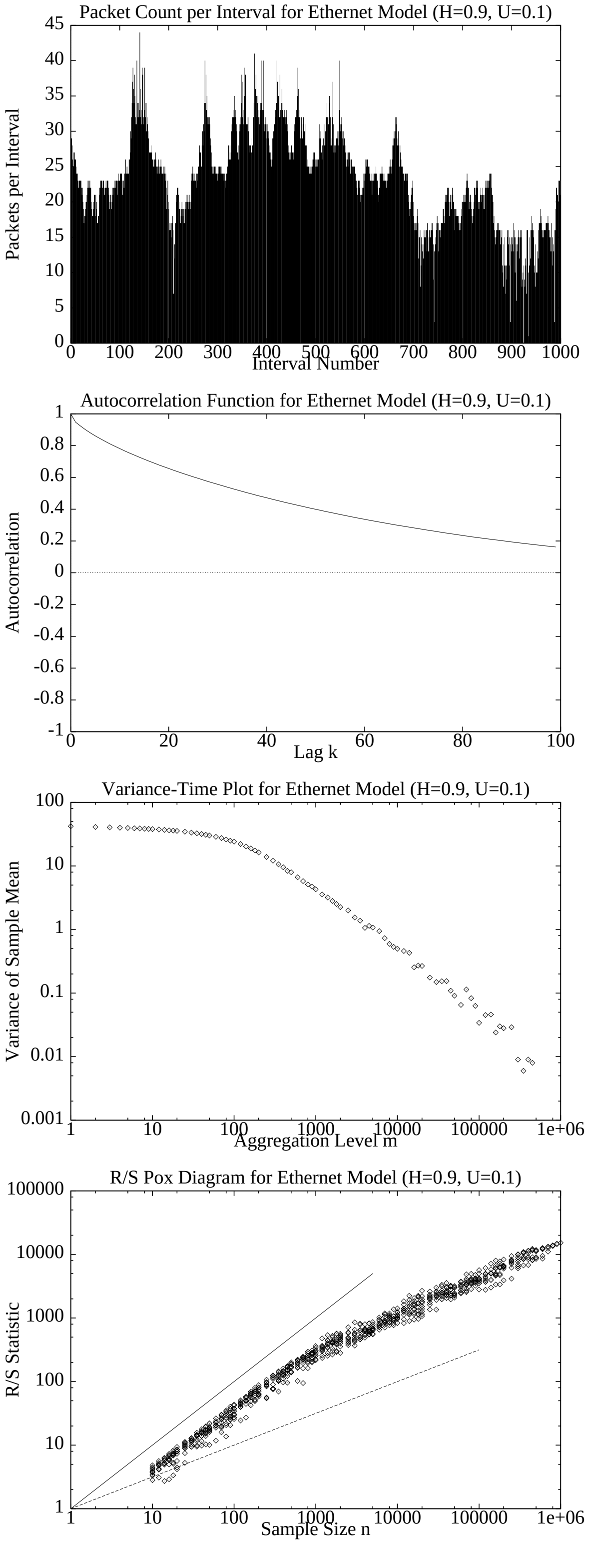

Figure 2: Ethernet traffic model for $H = 0.9$ ($U = 0.1$)

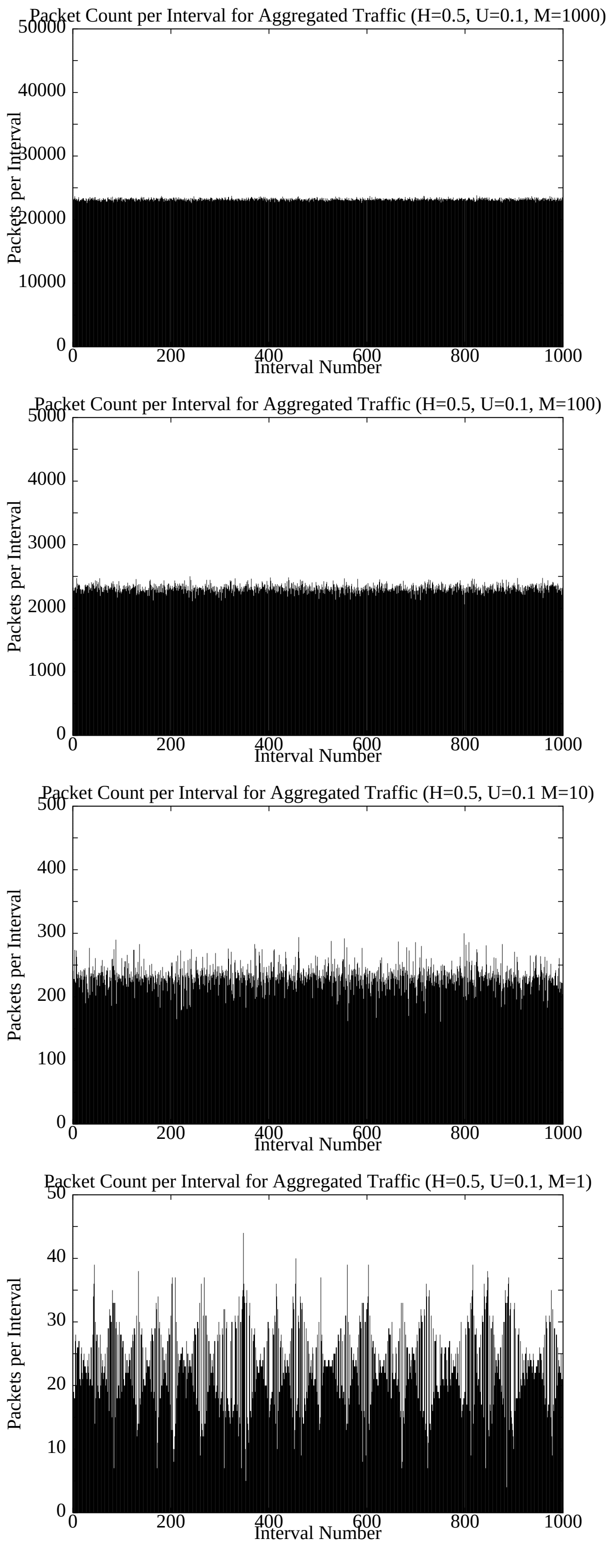

Figure 3: Effect of aggregation on traffic for $H = 0.5$ $(U = 0.1)$

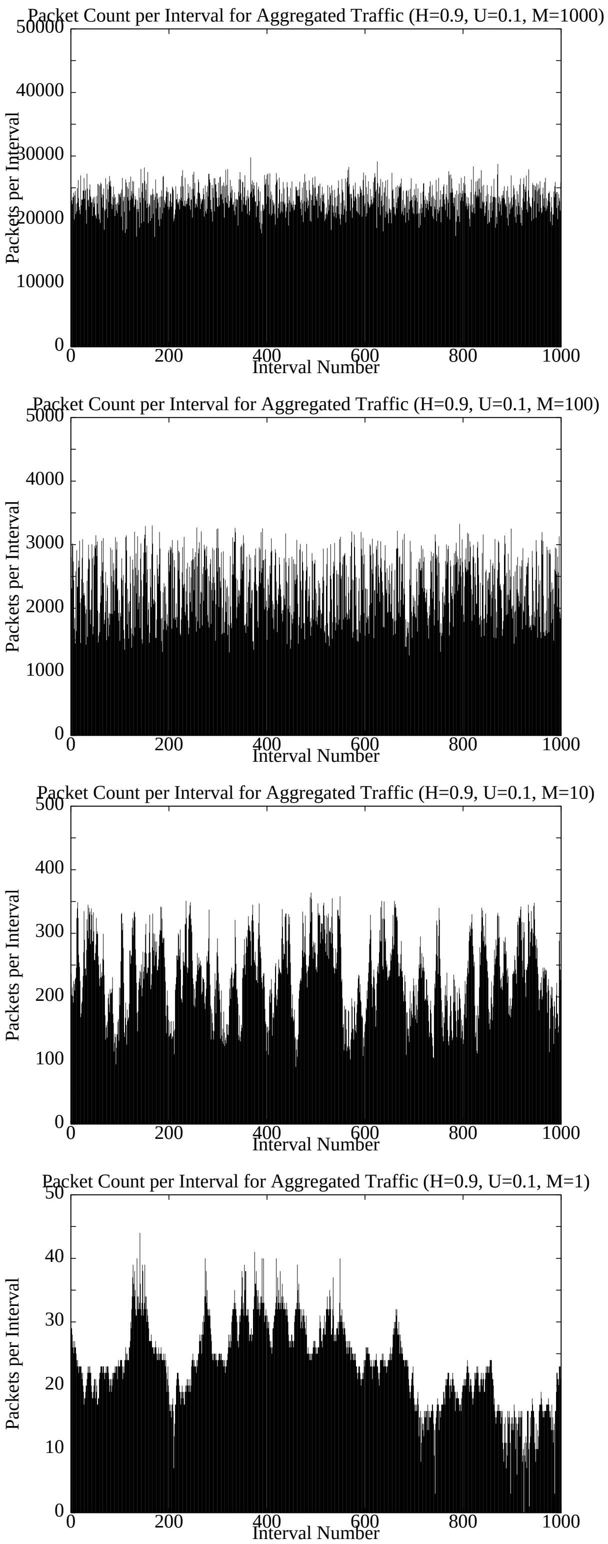

Figure 4: Effect of aggregation on traffic for $H = 0.9$ $(U = 0.1)$

aggregating bursty traffic streams together still results in bursty traffic. Thus cell loss in a network may be much higher than expected if network traffic is self-similar. This section describes a simple experiment that was conducted to assess the performance implications of self-similar traffic in terms of cell delay, cell loss, and buffer requirements, for a simulated ATM network.

The self-similar Ethernet traffic model was used to evaluate the cell delay and cell loss performance of a single simulated ATM switch. The simulated switch was a shared memory switch, with two input ports and two output ports. A single Ethernet traffic source was used to feed cells into the ATM switch, using a Permanent Virtual Channel (PVC). All Ethernet cells were destined for a low-bandwidth output link of the switch (1.5 Mbps). The switch has B buffers at the output port for this link, where cell queueing and cell loss may occur. No cells are lost within the switch fabric.

The utilization of the traffic source was scaled relative to the output link capacity to achieve effective link utilizations ranging from moderate (60%) to high (90%) to very high (99%). Experiments were repeated for five different settings of the Hurst parameter for the traffic source, and for six different settings of the output port buffer size B, ranging from 10 to 1000.

Figure 5 shows the cell delay results for this simulation model. For this experiment, the number of buffers at the output port was fixed at $B = 100$. Figure 5 shows that the average queue size (in cells) at the switch output port increases with link utilization, as expected. Queue size results are also shown for five different values of H. The larger the value of H, the greater the degree of self-similarity in the traffic, and the larger the average queue size at the switch. Thus self-similar traffic implies higher average queueing delays within the network.

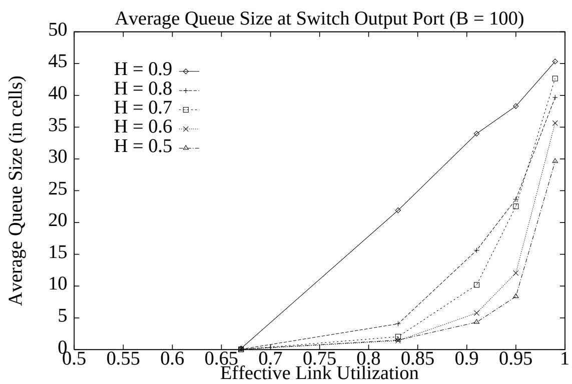

Figure 5: Average Queue Size at Switch Output Port for Ethernet Traffic

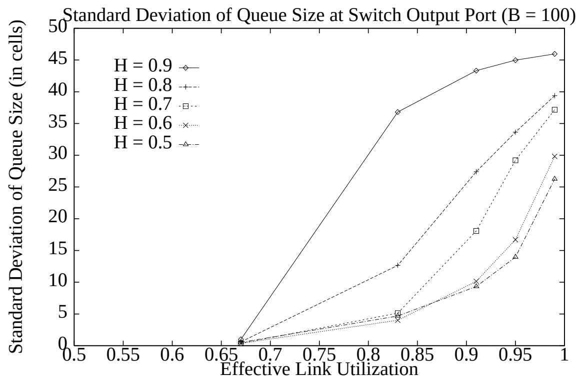

Figure 6: Standard Deviation of Queue Size at Switch Output Port for Ethernet Traffic

Self-similar traffic also affects cell delay variation (i.e., jitter). Figure 6 shows the standard deviation of the queue size at the switch output port for the same experiment as in Figure 5. The variation in the queue size increases with load, and increases with H. Thus the presence of self-similar traffic on a network can potentially affect the cell delay variation (jitter) of other traffic streams using the network.

Self-similar traffic has serious implications on cell loss performance in the network as well. Figure 7 depicts the average cell loss ratio experienced by the Ethernet traffic source, as a function of the number of buffers at the switch output port. These experiments are all conducted at a high link utilization (90%). Figure 7 shows that adding buffers at the output port always reduces the cell loss ratio, as expected. However, additional buffers are much less effective when the traffic source has a high degree of self-similarity. For example, for $H = 0.5$, increasing the number of buffers from $B = 10$ to $B = 100$ reduces the cell loss ratio from about 10^{-2} to below 10^{-4}. For $H = 0.9$, on the other hand, approximately 1000 buffers are needed to bring the cell loss ratio below 10^{-2}.

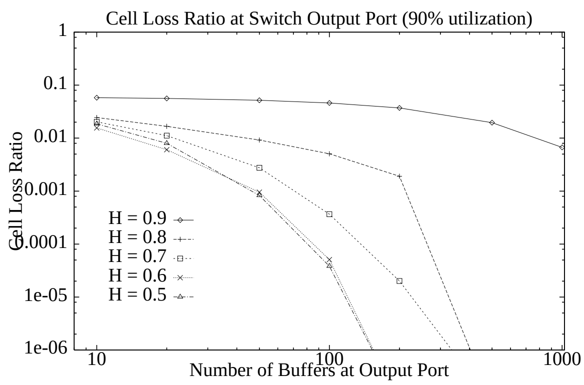

Figure 7: Cell Loss Ratio as a Function of Buffer Size

Table 1 summarizes the results from our simulations in tabular form. The table shows the mean queue size, the standard deviation of the queue size, and the cell loss ratio for different values of the Hurst parameter H. These results are for $B = 100$ and an effective link utilization of 90%. Note that simple Poisson traffic corresponds to $H = 0.5$. Cell loss and cell delay increase dramatically as the degree of self-similarity in the traffic increases.

Table 1: Performance Results for Self-Similar Traffic ($B = 100$, 90% link utilization)

Hurst Parameter	Mean Queue Size	Std Dev Queue Size	Cell Loss Ratio
0.5	4.348	9.482	0.0000383
0.6	5.729	10.281	0.0000508
0.7	10.156	17.991	0.0003661
0.8	14.301	26.693	0.0050707
0.9	32.875	43.217	0.0460415

5 CONCLUSIONS

This paper has presented a synthetic workload model that is capable of modeling self-similar Ethernet traffic. The model captures the autocorrelation and long range dependence that has been observed in recent measurements of Ethernet LAN traffic (Leland *et al.* 1994).

The simulation results presented in this paper suggest that self-similar traffic has potentially serious implications on ATM networks. Cell delay, cell jitter, and cell loss may be significantly higher than predicted by simple traffic models. For example, simple Poisson traffic models, as opposed to self-similar traffic models, may underestimate ATM cell loss by two or three orders of magnitude, depending on the network utilization. Self-similar traffic models are thus a valuable tool in the arsenal of network simulation, providing for "worst-case" stress testing of ATM networks.

ACKNOWLEDGMENTS

Funding for this research was provided by the ATM-TN TeleSim project at the University of Calgary (via CANARIE and the industry sponsors of the TeleSim project) and by NSERC Research Grant OGP00121969. The authors are grateful to the other members of the TeleSim project team for their constructive comments and feedback on our traffic modeling work throughout the stages of the TeleSim project.

REFERENCES

Frost, V. and B. Melamed. 1994. "Traffic Modeling for Telecommunications Networks". *IEEE Communications*, Vol. 32, No. 3 (Mar.), pp. 70-81.

Garrett, M. and W. Willinger. 1994. "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic". In *Proceedings of the 1994 ACM SIGCOMM Conference* (London, UK, Aug. 31-Sept. 2), pp. 269-280.

Geist, D. and B. Melamed. 1992. "TEStool: An Environment for Visual Interactive Modeling of Autocorrelated Traffic". In *Proceedings of the 1992 International Conference on Communications* (Chicago, Illinois, June), pp. 1285-1289.

Leland, W.; M. Taqqu; W. Willinger; and D. Wilson. 1994. "On the Self-Similar Nature of Ethernet Traffic (Extended Version)". *IEEE/ACM Transactions on Networking*, Vol. 2, No. 1 (Feb.), pp. 1-15.

Melamed, B. 1992. "TES Modeling of Video Traffic". *IEICE Transactions on Communications*, Vol. E75-B (June), pp. 1292-1300.

Paxson, V. and S. Floyd. 1994. "Wide-Area Traffic: The Failure of Poisson Modeling". In *Proceedings of the 1994 ACM SIGCOMM Conference* (London, UK, Aug. 31-Sept. 2), pp. 257-268.

Venables, W. 1990. "S: A Programming Environment for Data Analysis and Graphics". Technical Report, Department of Statistics, University of Adelaide (Dec.).

Ying Chen is an M.Sc. student in the Department of Computer Science at the University of Saskatchewan. Her interests include Ethernet traffic characterization, workload modeling, and network traffic self-similarity.

Zhiang Deng is an M.Sc. student in the Department of Computer Science at the University of Saskatchewan. His interests include video traffic characterization, network traffic self-similarity, spectral analysis, and mathematical modeling of traffic flows in high speed networks.

Carey Williamson is an Associate Professor in the Department of Computer Science at the University of Saskatchewan. A Ph.D. graduate of Stanford University in 1992, his research focuses on network simulation and computer systems performance evaluation. Of particular interest are performance issues in high speed ATM networks, such as traffic modeling, admission control, congestion control, and quality of service guarantees. He is currently supervising six graduate students working in these areas.