Contradictory Relationship between Hurst Parameter and Queueing Performance

Ronn Ritke, Xiaoyan Hong, Mario Gerla
UCLA – Computer Science Department,
405 Hilgard Ave., Los Angeles, CA 90024
ritke@cs.ucla.edu, hxy@cs.ucla.edu, gerla@cs.ucla.edu

Key words: Performance analysis, Queuing, Time Series analysis, System Engineering.

ABSTRACT

Long Range Dependent (LRD) network traffic does not behave like the traffic generated by the Poisson model or other Markovian models. The main difference is that LRD traffic increases queueing delays due to the burstiness of the traffic over many time scales. LRD traffic has been measured in different types and sizes of networks, for different applications (eg. WWW) and different traffic aggregations. Since LRD behaviour is not rare nor isolated, accurate characterization of LRD traffic is very important in order to predict performance and to allocate network resources. The Hurst parameter is used to describe the degree of LRD and the burstiness of the traffic. In this paper we analyze UCLA Computer Science Department network traffic traces and compute their Hurst parameters. Queueing simulation is used to study the impact of LRD and to determine if the Hurst parameter accurately describes such LRD. Our results show that the Hurst parameter is **not** by itself an accurate predictor of the queueing performance for a given LRD traffic trace.

1 INTRODUCTION

Accurate characterization of Internet traffic is very important for precise modeling and network design decisions. Modeling of Internet traffic is based on the traffic characteristics and the resulting models often serve as input for simulations. The results of simulations are used for a number of network design decisions. For many years the Poisson model was widely used to model Internet traffic, but in the last few years new characteristics have been discovered in Internet traffic. Long Range Dependence (LRD) has been discovered in LANs (Willinger et al. 1996; Leland et al. 1997), WANs (Paxson and Floyd 1995) and MANs (Willinger et al. 1997). It has also been discovered in different services and applications: Aggregate traffic (Willinger et al. 1996; Leland et al. 1997), World Wide Web (Crovella and Bestavros 1995, Crovella and Bestavros 1996), Variable-Bit-Rate (VBR) video traffic (Beran et al. 1995) and different types of computer networks: Ethernet (Willinger et al. 1996; Leland et al. 1997), ATM (Willinger et al. 1997) and CCSN/SS7 (Duffy et al. 1994). The traffic with the LRD property is more bursty than traffic generated with the Poisson model. The Poisson model is Short Range Dependent and does not accurately model LRD traffic (Paxson and Floyd 1995). In comparison to LRD traffic, the use of the Poisson model (or other Markovian models) results in overly optimistic queueing performance. The queue length distribution decays much more slowly for LRD traffic. The queueing delay rises dramatically with increasing LRD (Erramilli et al. 1996) and the Hurst parameter quantifies this long range dependence.

(Leland et al. 1997) compared a current model (a compound Poisson process) with actual network traffic. They found that after aggregation over the seconds time scale you will see a smoothed traffic with a compound Poisson process. But in contrast real traffic did not smooth out and is bursty over many time scales (self-similar). They argue that the Hurst parameter quantifies the degree of self-similarity and can be used as a measure of the burstiness of the traffic (the higher the Hurst value the burstier the aggregate traffic).

Other researchers including (Crovella and Bestavros 1995; Crovella and Bestavros 1996) have also found that the H value declines somewhat when they use traffic from less busy hours as compared to busy hours, which is consistent with results found by (Leland et al. 1997).

Based on the study of the fractional Brownian motion model, Neidhardt and Wang discovered a further complexity - that burstiness depends on time scales (Neidhardt and Wang 1998). When comparing a high Hurst value process and a low Hurst value process (assuming both processes are exactly second-order self similar) the variance of the two processes match at a unique time scale t_m. There are time scales t_{qi} that are most relevant for queueing the arrivals of process i. If for both processes the queueing scales t_{qi} are greater than variance matching scale t_m, then the higher Hurst value process queue will result in worse queueing performance; if they are both smaller than t_m then the lower Hurst value process queueing performance is worse.

We used real traffic traces as input to trace-driven queueing simulations to examine the relationship between the Hurst parameter and queueing performance. We found some results similar to (Neidhardt and Wang 1998) in addition to some different results.

When we compared a lower H parameter traffic with a higher H parameter traffic we found that the lower H parameter traffic resulted in worse queueing performance for all the different time length traffic traces we used. We also showed that the Hurst parameter can differ for a given traffic trace over different time lengths. This is significant because it reduces the importance of the Hurst parameter and opens future research into new parameters and models for LRD traffic.

The rest of the paper is organized as follows: Information about the computer network traffic traces used in this paper is presented in Section 2. The definition for Long Range Dependence is given in Section 3. Section 4 contains the definitions for and generation of the time series used in this paper. The computation of the Hurst parameter and a discussion on the variance-time plot and examples of its use in estimating the Hurst parameter with UCLA network traffic is covered in Section 5. Section 6 describes the queue simulator. The performance of LRD traffic in a queueing simulation in order to determine the effectiveness of the Hurst parameter in predicting the resulting queueing performance is covered in Section 7. Lastly, in Section 8 we summarize our findings.

2 TRAFFIC TRACES

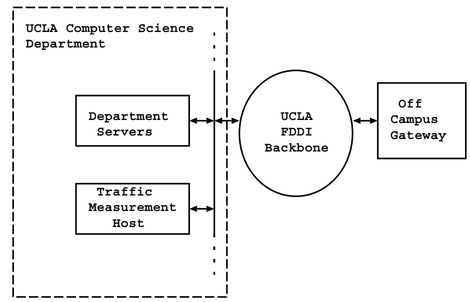

Figure 1: UCLA CSD Measurement Connection.

Network traffic traces were taken at UCLA Computer Science Department (CSD) over a 5 week period (Feb - March 1998). Network traffic information was collected at a host running Tcpdump (Jacobson et al. 1989). This host was connected (via a special link) to department servers and to the router that connects the CSD to the FDDI backbone (see Figure 1). The traces represent the network traffic in the Computer Science Department. The resulting output was processed to obtain the format needed to test for LRD (arrival time and packet length for each packet). Information for each traffic trace obtained is summarized in Tables 1 and 2. Note that the 'All' in the trace name signifies that it is an aggregate traffic trace.

Trace	Date	Start	Duration(s)	Arrivals
All1	3/6/98	1pm	955.698	500000
All4	2/26/98	10am	2017.56	1200000
All7	2/23/98	9am	4637.41	2000000
All8	3/23/98	12pm	995.695	680000

Table 1: UCLA CSD Traffic Trace Information.

Trace	Total Bytes	Bytes/Arriv	Arriv/Sec
All1	2.36126E+08	472.252	523.178
All4	8.40773E+08	700.644	594.779
All7	1.32400E+09	662.002	431.275
All8	2.02779E+08	298.204	682.940

Table 2: UCLA CSD Traffic Trace Information.

3 SELF-SIMILAR AND LRD PROCESSES

Our approach is to define LRD following the definitions given in (Willinger et al. 1995; Leland et al. 1997).

Let $X = (X_t : t = 0, 1, 2, \dots)$ be a covariance stationary stochastic process with mean μ, variance σ^2 and autocorrelation function $r(k), k \geq 0$. Assume $r(k)$ is of the form

$$r(k) \sim k^{-\beta}, \quad as \quad k \to \infty \tag{1}$$

where $0 < \beta < 1$.

For each m=1,2,3, ..., let $X^{(m)} = (X_t^{(m)} : t = 1, 2, 3, \dots)$ denote the new covariance stationary time series obtained by averaging the original series X over non-overlapping blocks of size m, i.e.,

$$X_t^{(m)} = (X_{tm-m+1} + \cdots + X_{tm})/m, \quad t \geq 1 \tag{2}$$

The process X is called (exactly) second-order self-similar if for all $m = 1, 2, 3, \dots,\ var(X^{(m)}) = \sigma^2 m^{-\beta}$ and

$$r^{(m)}(k) \sim r(k), \quad k \geq 0 \tag{3}$$

The process X is called (asymptotically) second-order self-similar if for all k large enough,

$$r^{(m)}(k) \to r(k), \quad as \quad m \to \infty \tag{4}$$

The key property of this class of self similar processes is the fact that the covariance does not change under block aggegation and time scale changes. The relationship between the Hurst parameter and β is $H = 1 - \beta/2$. Note that here $1/2 < H < 1$, since $0 < \beta < 1$. A self similar process with $1/2 < H < 1$ (i.e., $\beta < 1$) is long range dependent (LRD). Since $\beta < 1$ the function $\sum_k r(k) = \sum r^{-\beta} = \infty$. By contrast, a short-range dependent process (eg. Poisson Process) has fast decaying autocorrelation function (i.e., $\beta > 1$), hence, $\sum_k r(k) < \infty$. The Hurst parameter is thus a key indicator of LRD behavior. One immediate consequence of LRD behavior is that the traffic exhibits the same burstiness across many time scales see Figure 2. Having introduced the definitions for the Hurst parameter and LRD, in the following section we define some important time series and their relationship to the original traffic trace.

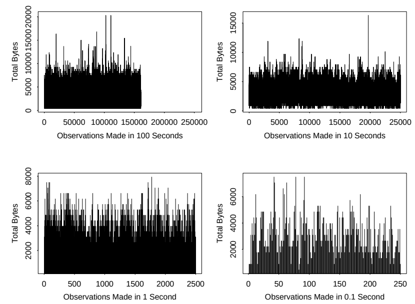

Figure 2: Bursty Traffic Over Many Time Scales.

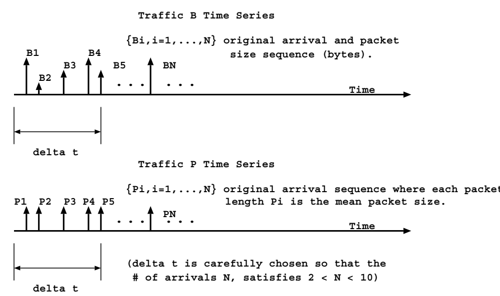

Figure 3: Time Series Diagrams.

4 TIME SERIES DEFINITION

This paper examines the correspondence between the queueing performance of self-similar traffic and the Hurst value. The original traffic trace is characterized by two variables: time (of arrival of a packet) and length (of the packet). From this trace, time series with only one variable must be generated in order to estimate the Hurst parameter for this variable. There are several methods for generating such single variable time series from data traces. Researchers from Bellcore (Willinger et al. 1996; Leland et al. 1997) have proposed the {Traffic B} and {Traffic P} time series. We derive the same two synthetic traffic traces ({Traffic P} and {Traffic B}) from the real trace to see the relationship between the Hurst value and queueing performance. The resulting two traces have comparable magnitude. Queueing experiments can be driven by them and produce comparable results. We will examine the degree of self-similarity of these two traces by calculating their Hurst values and observing the queueing behavior when both traces are fed into a FIFO queue.

The relationship of the two time series from the packet arrival times and packet lengths is displayed in Figure 3. Each of the two time series captures different aspects of the traffic trace. To compute {Traffic B} and {Traffic P}, we choose a time interval Δt which typically contains between 2 and 10 arrivals (see Figure 3). Within non-overlapping time intervals of size Δt we sum the number of bytes $\overline{B}_i$ arriving in each interval Δt_i and get the time series {Traffic B} $= \{\overline{B}_i, i = 1, 2, 3, \ldots\}$. Next, let $\overline{P}_{ave}$ be the mean packet size computed over the entire duration of the experiment. Consider the new traffic sequence where the actual packet size is replaced by the mean packet size $\overline{P}_{ave}$. Within non-overlapping time intervals of size Δt we sum the number of bytes $\overline{P}_i$ arriving in each interval Δt_i to get a time series {Traffic P} $= \{\overline{P}_i, i = 1, 2, 3, \ldots\}$.

5 HURST PARAMETER ESTIMATION

A. Estimation of the Hurst parameter

We present two different views of the Hurst parameter: one is a visual view (variance-time plot), the other is a computed view (see next subsection). To visually estimate the Hurst parameter, we plot $var(X^{(m)})$ as a function of m. The variance-time plot draws the variance vs. m in a log-log scale, which shows the slowly decaying variance of a self-similar series. If the input traffic has the LRD property, the curve should be linear (for large m) with slope larger than -1. The 'Reference' line on the variance-time plot (Figure 4) represents the slope of the line of $\beta = 1$, that is $var(X^{(m)}) = m^{-1}$, then H = 1/2. Any line with a slope less than 0 and greater than this reference line exhibits LRD and has an H parameter value $1/2 < H < 1$. The variance-time plots for {Traffic B} and {Traffic P} for all four CSD traces are shown in Figure 4. The captions represent the curves top down on the graph. By inspection of the variance-time plots, it is apparent that both curves on each plot have an H value greater than 1/2 and less than 1, demonstrating that all the curves show the property of long range dependence.

B. Computation of the Hurst parameter

The Hurst parameter was also computed using Least-Squares Curve Fitting (Trivedi 1982), leading to an analytic equation for the curve. The resulting equation is in the form of $y = -\beta x + b$ where $-\beta$ is the slope of the curve. The Hurst value is then computed using the relation $H = 1 - \beta/2$. Table 3 contains the computed Hurst values of {Traffic B} and {Traffic P} for each of the four traffic traces used in this paper. Notice that all four computed {Traffic P} Hurst values are greater than their respective computed {Traffic B} values. Simulation is used in the next section to show the relationship between the Hurst parameter and queueing performance.

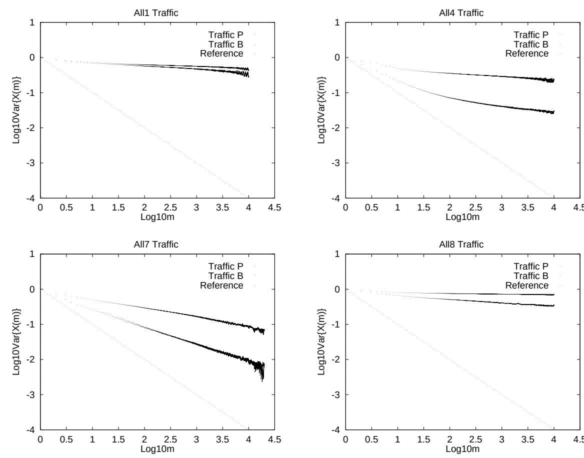

Figure 4: Variance-time plots for the 4 traces.

Traffic Trace	{Traffic B}	{Traffic P}
All1	0.9338	0.9652
All4	0.8991	0.9427
All7	0.7395	0.8475
All8	0.9542	0.9901

Table 3: Hurst values for the 4 traces.

6 THE QUEUE SIMULATOR

In this section, we introduce our queuing simulation which uses real traffic traces and then in the next section we discuss the simulation results. Previous studies on the queueing simulation are either driven by real traffic traces or by traffic models. As was done in many other modeling approaches, a single parameter (Hurst parameter) is used to describe the self-similar property of network traffic. Our paper differs from previous studies, however, in that we discuss the relationship between the Hurst parameter and queueing performance. We use UCLA traffic, which exhibits the long range dependency property, to drive a queueing simulation. It is well known that LRD traffic is burstier than the traditional Poisson model, and thus requires a much larger queue size. By observing the influence of such traffic on the queueing system, we will see that the Hurst parameter alone does not sufficiently quantify the LRD property, nor does it characterize the burstiness of real traffic.

The queueing system utilized in the simulation has a single server, infinite buffer size and FIFO discipline. Experiments were run with five different server utilizations (0.3, 0.5, 0.7, 0.9, 0.97). Only the experiments with utilizations (0.5, 0.7) are reported here. The queueing simulation is driven by {Traffic B} and {Traffic P} sequences, which in turn were derived from real traces.

7 EXPERIMENTAL RESULTS

In the experiments, we measure the complementary distribution of the queue length. Let Q(t) be the number of bytes in the queue over time. In the plots we show P(Q(t) > x), the probability that the queue length is greater than x, in log scale. The longer the tail of the distribution, the burstier the traffic. The discovery by (Neidhardt and Wang 1998) of a crosssover point t_m where the variance of the two processes match in the very small time scales is not that important to researchers since they are more interested in larger time scales and the tail of the distribution.

In order to provide a reference, Figure 5 shows the combined queue length distributions for the M/M/1 and LRD All1 trace for queueing system with utilizations 0.3, 0.5 and 0.7. The M/M/1 queueing model uses the average interarrival time (0.0019 sec) and average packet size (472 bytes) extracted from the trace All1. In addition, Figure 5 shows the simulation results obtained by applying {Traffic B} derived from trace All1 to the FIFO queue (see also Figure 6 (a)). Figure 5 clearly shows that there is a large difference (3 orders of magnitude) in queue length distributions between the M/M/1 queue and the real traffic queue (corresponding to {Traffic B}) and as a result the Poisson process can not be used as a substitute for the real LRD traffic. The disparity between Poisson and Self Similar queues is revealed also by the analytical models. In particular, for the Poisson process (M/M/1 queue) the P(Q(t) > x) $\sim e^{-\eta x}$ (Beran et al. 1995) but for the self similar process the P(Q(t) > x) $\sim e^{-\alpha x^{2-2H}}$ (Beran et al. 1995, Erramilli et al. 1996).

A. Queue length distributions for the whole traces.

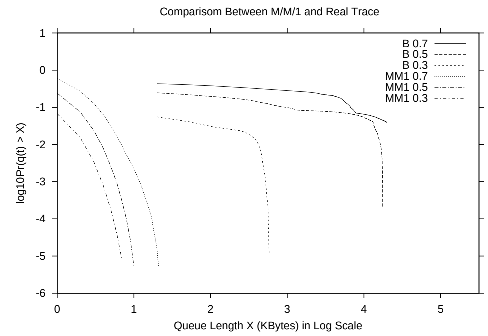

Figure 5: Queueing experiment for M/M/1 model and the LRD All1 trace.

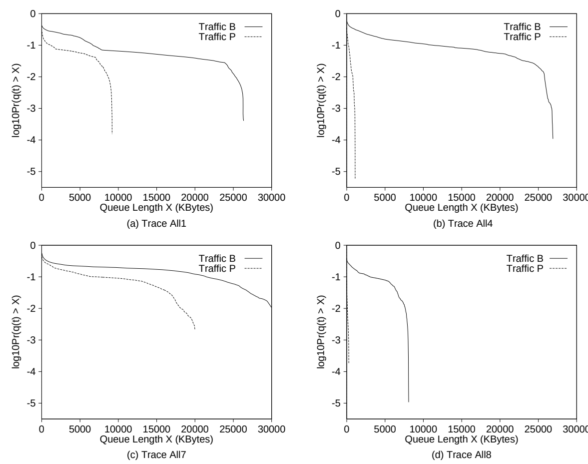

Figure 6: Queueing experiment of each trace with 0.7 utilization.

As confirmed by simulation, the burstier the traffic, the longer the tail of queue length distribution. However, our experiments show that a larger H value may not indicate a greater burstiness and a larger queue. Several queueing experiments and discussions of the queueing simulation results provide evidence from several different points of view to support such an argument.

First, let us look at the behavior of the {Traffic B} and {Traffic P} synthetic traces which were derived from the same traffic trace. Consider trace All1 (Figure 6 (a)) for example. The computed Hurst parameters show that the H values of Traffic B (0.9338) and P (0.9652) of trace All1 (Table 3) are very close. But the tail of {Traffic B} is much longer than that of {Traffic P}. This suggests that {Traffic B} is much burstier than {Traffic P}. Moreover, trace All7 in Figure 6 (c) and Figure 8 (c) shows that while {Traffic P} has a greater H value than {Traffic B} (i.e. 0.8475 vs. 0.7395), the tail of the queue length distribution of {Traffic P} is shorter than that of {Traffic B}. These 'inversions' support our claim that the value of the Hurst parameter does not accurately reflect the relative burstiness between {Traffic B} and {Traffic P} from the same trace.

Traffic Trace	{Traffic B}	{Traffic P}
All7 Seg1	0.9171	0.9587
All7 Seg2	0.9387	0.9710
All7 Seg3	0.9462	0.9713
All7 Seg4	0.9574	0.9748
All7	0.7395	0.8475

Table 4: Hurst values for the four ALL7 segments.

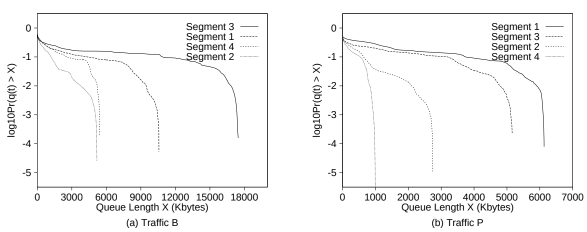

Figure 7: Queueing experiment of Segments from trace All7 with 0.7 utilization.

B. Queue length distributions for trace segments.

More support of our argument comes from the simulation experiments with segments from the same trace. Previous experiments on segments of the original trace were done by (Abry and Veitch 1998) to check if the Hurst parameter was constant across the segments in a test to see if the whole trace was stationary. Our segmented experiments show that parts of the entire trace perform differently from the original whole one. The original trace All7 is divided into four sections. Each section is treated as an individual trace. We derive {Traffic P} and {Traffic B} for each trace, estimate their H values, and feed them into the queueing system. The utilization is kept at 70 percent. The results given in Table 4 shows that the {Traffic P} and {Traffic B} segments have H values similar to each other. However, when the results of queueing experiments are considered (Figure 7), strong differences between queue length distributions among the four segments are observed. For example, {Traffic B} of Segment3 has a heavier tail than the other segments, but it doesn't have the largest H value. So the queueing performance is different for segments with similar H parameters. Furthermore, the tails of these distributions are much lighter than the tail of the queue length distribution of the entire trace All7 (see Figure 6 (c)). That is, the largest queue length of the four segments of {Traffic B} is near 18MBytes (Figure 7(a)), but the probability of a queue length larger than 25MByte is still near 10 percent for {Traffic B} of the entire trace (Figure 6 (c)). However, the value of the Hurst parameter for All7 is smaller than that of the segments. In other words the Hurst value can vary with different time scales.

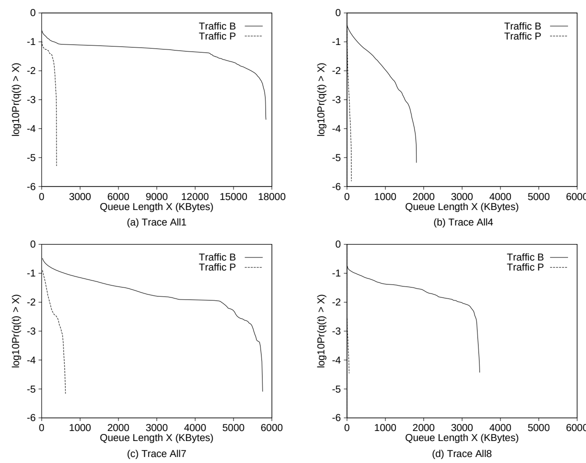

Figure 8: Queueing experiment of each trace with 0.5 utilization.

C. Queue length distributions for different utilizations.

The queue length distributions of the traces when the system is at 50 percent load is shown in Figure 8. The distribution of queue buffer size decays faster than in 70 percent utilization. Just as (Erramilli et al. 1996), showed in their paper, generally, a traffic load of 0.5 is near the "knee" of the delay-utilization curve. When the utilization is greater than 0.5, the queueing delay increases very fast. From the queue length distribution, this point can be clearly seen, where the tail of a heavy traffic load (Figure 6) is much longer than that of a light load (Figure

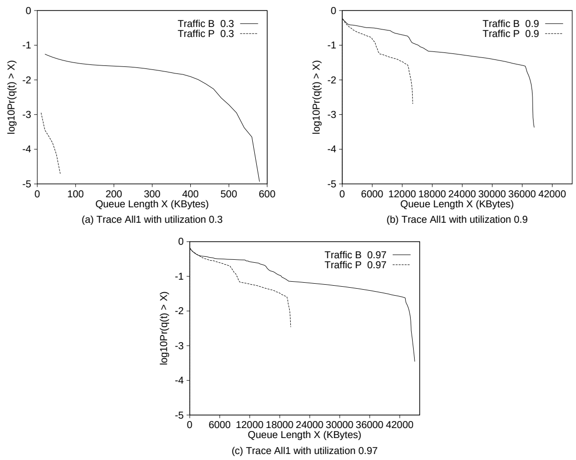

Figure 9: Queueing experiments for trace All1 with 0.3, 0.9 and 0.97 utilizations.

8). Therefore, our discussions have focused on the queuing performance under heavy load (0.7), which is of most interest. For the sake of completeness, however, we have compared the All1 simulation (both {Traffic P} and {Traffic B}) for loads 0.3, 0.9, 0.97 (see Figure 9). Queueing lengths for {Traffic B} are longer than for {Traffic P} for each traffic trace run at each utilization (0.3, 0.5, 0.7, 0.9, 0.97).

All these comparisons lead us to the same conclusion: that the Hurst parameter is not an accurate indicator of the traffic burstiness and queueing performance. Our queueing experiments clearly show that the Hurst parameter alone is not sufficient to predict the queueing performance.

8 CONCLUSION

Real traffic traces were used for the simulation and the results presented in this paper have shown that the H parameter is not a consistent, monotonic indicator of queueing perfomance. For example, our results show that the {Traffic B} and {Traffic P} time series derived from the same trace are such that {Traffic P} has a greater Hurst value than {Traffic B}, i.e. {Traffic P} should have a longer queue length distribution. Yet, the queue length distributions show an 'inversion' since {Traffic B} causes longer queues than {Traffic P}. These contradictory relations hold true for all four traffic traces. Namely, the value of the Hurst parameter does not fully reflect the relative queueing performance of {Traffic B} and {Traffic P}, even though they were from the same trace. Moreover, the differences in queueing performance among segments and between the segments and the entire trace confirm the inaccuracy of the H value in predicting queueing performance. The main conclusion of this paper is that the H parameter alone is not sufficient to fully describe the LRD property of a traffic source and to predict its queueing impact.

9 ACKNOWLEDGEMENTS

Authors would like to thank Jianbo Gao for some discussions and providing some software, Scott Seongwook Lee for his help in creating some software processing tools and many thanks to the UCLA CSD staff for their helpful cooperation.

10 BIBLIOGRAPHY

Journal:
Patrice Abry, Darryl Veitch, 1998: Wavelet Analysis of Long Range Dependent Traffic, *Trans. Info. Theory*, **44**, no. 1, pp. 2-15.

Journal:
Jan Beran, Robert Sherman, M. S. Taqqu and W. Willinger, 1995: Long-Range Dependence in Variable-Bit-Rate Video Traffic. *IEEE/ACM Trans. on Communications*, **2**, no. 4, pp. 1566-1579.

Technical Report:
Mark E. Crovella and Azer Bestavros, 1995: Explaining World Wide Web Traffic Self-Similarity. *Technical Report TR-95-015, Boston University, August 1995.*

Conference Proceeding:
Mark E. Crovella and Azer Bestavros, 1996: Self-Similarity in World Wide Web Traffic Evidence and Possible Causes. *1996 ACM SIGMETRICS Int. Conf. Measurement, Modeling of Computer System*, pp. 160-169, May 1996.

Journal:
D. E. Duffy, A. A. McIntosh, M. Rosenstein, and Walter Willinger, 1994: Statistical analysis of CCSN/SS7 traffic data from working subnetworks. *IEEE J. Select. Areas Commun.*, **12**, no. 3, pp. 544-551.

Journal:
A. Erramilli, O. Narayan, and W. Willinger, 1996: Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Trans. on Networking*, **4**, no. 2, pp. 209-223, April 1996.

Misc:
Van Jacobson, Craig Leres and S. McCanne, 1989: The Tcpdump Manual Page. *Lawrence Berkeley Laboratory*, http://ee.lbl.gov/,

ftp://ftp.ee.lbl.gov/tcpdump.tar.Z.

Journal:
W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, 1994: On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Trans. on Networking*, **2**, no. 1, pp. 1-15, Feb. 1994.

Journal:
Arnold L. Neidhardt and Jonathan L. Wang, 1998: The Concept of Relevant Time Scales and Its Application to Queueing analysis of Self-Similar Traffic (or is Hurst Naughty or Nice?). *ACM Performance Evaluation Review*, **26**, no. 1, pp. 222-232, June 1998.

Journal:
Vern Paxson and Sally Floyd, 1995: Wide-Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Trans. on Networking*, **3**, no. 3, pp. 226-244.

Book:
Trishor S. Trivedi, 1982: Probability & Statistics with Reliability, Queueing, and Computer Science Applications. *Prentice-Hall Inc.*, Englewood Cliffs New Jersey, pp. 543-545.

Journal:
W. Willinger, M. S. Taqqu, et. al., 1995: Self-Similarity in High-Speed Packet Traffic: Analysis of Modeling of Ethernet Traffic Measurements. *Statistical Science*, **10**, no. 1, pp. 71-86, December 1995.

Journal:
W. Willinger, M. S. Taqqu, M. S. Sherman, and D. V. Wilson, 1996: Self-similarity through high-variability: Statistical analysis of ethernet LAN traffic at the source level (extended version). *IEEE/ACM Trans. on Networking*, **5**, pp. 71-86, December 1996.

Conference Proceedings:
Walter Willinger, Siddhartha Y. Devadhar, et. al., 1997: Measuring ATM Traffic Cell-by-Cell: Experiences and Preliminary Findings from BAGNet. *PM-CCN '97*, International Conference on the Performance and Management of Complex Communication, Tsukuba, Japan, November 17-21, 1997.

BIOGRAPHIES

Ronn Ritke received an M.S. from California State University Long Beach in 1994 in computer science, a second M.S. from and advanced to Ph.D. candidacy at the University of California at Los Angeles in 1996 and 1998, all in computer science. His research interests include network traffic measurement and analysis. He is currently working with San Diego Super Computer Center and NSF to help provide answers on vBNS performance for the PITAC Committee.

Xiaoyan Hong recieved a B.E. degree and an M.S. degree in computer science from the Zhejiang University, China during the 1980s. She has been a staff researcher at the Zhejiang University for eight years, and is currently a visiting researcher in the Computer Science Department at the University of California, Los Angeles. Her research interests include network traffic analysis and wireless Ad-hoc networks.

Dr. Mario Gerla was born in Milan, Italy. He received a graduate degree in engineering from the Politecnico di Milano, in 1966, and the M.S. and Ph.D. degrees in engineering from UCLA in 1970 and 1973, respectively. He joined the Faculty of the UCLA Computer Science Department in 1977. His research interests cover the performance evaluation, design and control of distributed computer communication systems and high speed computer networks. His current research projects are in the area of wireless ad-hoc networks (channel access, routing, multicast); internet (QoS routing, multicast, congestion control, IP over ATM), and; traffic modeling and measurements.