

Contradictory Relationship between Hurst Parameter and Queueing Performance (extended version)*

Ronn Ritke[†], Xiaoyan Hong and Mario Gerla
UCLA – Computer Science Department,
405 Hilgard Ave., Los Angeles, CA 90024
rritke@hotmail.com, hxy@cs.ucla.edu, gerla@cs.ucla.edu

Abstract

Long Range Dependent (LRD) network traffic does not behave like the traffic generated by the Poisson model or other Markovian models. From the network performance point of view, the main difference is that LRD traffic increases queueing delays due to its burstiness over many time scales. LRD behavior has been observed 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 commonly used to quantify the degree of LRD and the burstiness of the traffic. In this paper we investigate the validity and effectiveness of the Hurst parameter. To this end, we analyze the 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.

Key words: Hurst parameter, Long Range Dependence, queueing performance, trace-driven simulation.

*This research is supported by the National Science Foundation through grant NCR-9305376.

[†]Ronn Ritke is the corresponding author, Fax: 1 310 8257578

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 and network design. For many years the Poisson model was widely used to model Internet traffic, but in the last few years new characteristics have emerged. Long Range Dependence (LRD) in traffic arrival processes has been discovered in LANs [12] [19], WANs [14] and MANs [17]. This dependence spans several applications: World Wide Web [4] [5], Variable-Bit-Rate (VBR) video traffic [3] and also Aggregate traffic [12] [19]. It has been verified in measurements collected on different types of computer networks (see Figure 5) : Ethernet [12] [19], ISDN [8], ATM [17] and CCSN/SS7 [6]. LRD traffic is more bursty than traffic generated with the Poisson model. The Poisson model is Short Range Dependent [14] leading to overly optimistic queueing performance. In contrast, the queue length distribution decays much more slowly using the LRD traffic models. The queueing delay rises dramatically with increasing long range dependence of arrivals [7].

Leland, Taqqu, Willinger and Wilson [12] compared a traditional model (a compound Poisson process) with actual network traffic. They found that after aggregation of the Poisson arrivals over the seconds time scale the traffic is very smooth. In contrast, real traffic does 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 traffic burstiness (the higher the Hurst value the burstier the aggregate traffic). Other researchers including Crovella and Bestavros [4] [5] have also found that the H value declines somewhat in light load traffic conditions as compared to busy hours. This is consistent with results found by Leland [12].

Based on the study of the fractional Brownian motion model, Neidhardt and Wang gained further insight on the impact of the Hurst parameter on queueing performance [13]. When comparing a high Hurst value process and a low Hurst value process (assuming both processes

are exactly second-order self similar) the variances 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 the 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.

In our research, we used an experimental approach to examine the impact of the H parameter on queueing performance. Namely, we used real traffic traces as input to trace-driven queueing simulations and examined the relationship between the Hurst parameter and queueing performance. Some of our results confirm the findings in [13]. Others contradict them. As in [13], we found that the lower H parameter traffic can result in worse queueing performance. However, in contrast with [13], when the same comparison between low H and a high H parameter traffic was done over a broad range of time scales we found that the lower H parameter traffic resulted in worse queueing performance over all time scales. We also showed that the Hurst parameter can differ for a given traffic trace over different time lengths.

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 of 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

A. UCLA Computer Science Department Traffic Traces

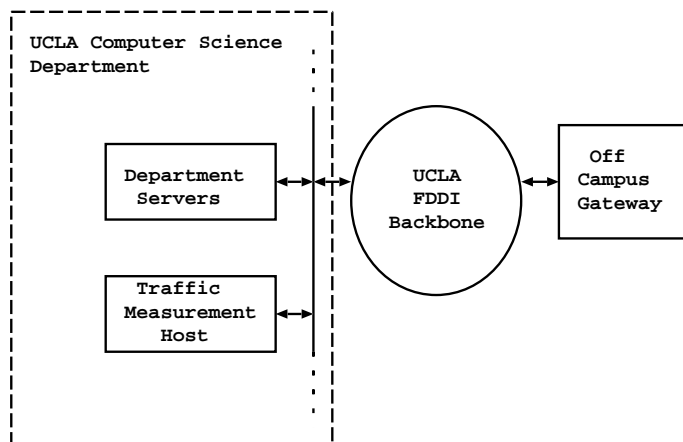


Figure 1: UCLA CSD Measurement Connection Diagram.

Network traffic traces were taken at UCLA CSD over a 5 week period (Feb - March 1998). Network traffic information was collected at a host running Tcpcdump [11] (see Figure 1). This host was connected (via a special link) to department servers and to the router that connects the CSD to the FDDI backbone. 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 is summarized in Tables 1 and 2. Note that the 'All' in the trace name signifies that it is an aggregate traffic trace (i.e., the trace aggregates several different applications).

Trace Name	Start Date	Start Time	Time Span (Sec)	Packets
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 Computer Science Dept Traffic Trace Information.

B. Applications Present in UCLA CSD Traffic Traces

In order to determine which applications are present in the UCLA CSD network traffic we made use of the fact that (following RFC 1700 recommendations) a number of TCP and UDP

Trace Name	Total Bytes	Bytes/Sec	Bytes/Packet	Packets/Sec
All1	2.36126E+08	247072	472.252	523.178
All4	8.40773E+08	416728	700.644	594.779
All7	1.32400E+09	285505	662.002	431.275
All8	2.02779E+08	203655	298.204	682.940

Table 2: Detailed UCLA CSD Traffic Trace Information.

applications were assigned to well known port numbers. For example, HTTP network traffic was assigned to port 80. Each well known port has both a UDP and a TCP application assigned to it even if the application supports only the TCP implementation. Tcpcap allows for the collection of all the packet traffic (our All traces use this Tcpcap option) and it can be used to 'filter' the traffic so that only TCP or UDP packets are collected in the traffic trace. For a 7 day traffic trace taken earlier at UCLA CSD, Tcpcap was used to capture only TCP traffic. As a result, when a traffic analysis tool counted the number of packets or bytes sent to the well known ports there was no ambiguity whether it is TCP or UDP. Two different traffic summaries - packet and byte - are presented in Tables 3 and 4 for a 7 day trace. Both summaries are shown because some applications have a large number of packets but a small number of bytes per packet while other applications have a small number of packets but each packet has a large number of bytes. There are tens of thousands of different port numbers so to get an overview of the network traffic we look at the applications that make up the bulk of the measured traffic. Only the applications (with well known port numbers) that have over 0.9 percent of the total bytes or total packets are reported.

3 LRD processes Characterization

We follow the self similarity and LRD definitions given in [12] [18].

Let $X = (X_t : t = 0, 1, 2, \dots)$ be a covariance stationary stochastic process with mean μ ,

Port	Packets Percent	Application
6000	24.73	X-Windows
513	6.543	login
23	5.87	Telnet
80	3.849	http (World Wide Web)
119	3.42	nntp
514	1.855	Syslog
20	1.596	FTP-data
515	1.483	printer

Table 3: Summary of the Applications that have over 0.9% of the total packets.

Port	Bytes Percent	Application
6000	34.4	X-Windows
514	6.528	Syslog
119	2.739	nntp
515	2.62	printer
25	1.376	SMTP
20	1.253	FTP-data
513	0.998	login
23	0.952	Telnet
80	0.907	http (World Wide Web)

Table 4: Summary of the Applications that have over 0.9% of the total bytes.

variance σ^2 and autocorrelation function $r(k), k \geq 0$. Assume $r(k)$ is of the form

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

where $0 < \beta < 1$.

For each $m = 1, 2, 3, \dots$, 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.,

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

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

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

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

$$r^{(m)}(k) \rightarrow r(k), \quad \text{as } m \rightarrow \infty \quad (4)$$

The key property of this class of self similar processes is that the covariance does not change under block aggregation 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 $\sum_k r(k) = \sum k^{-\beta} = \infty$. 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. This property can be observed in the traces collected at UCLA and reported in 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.

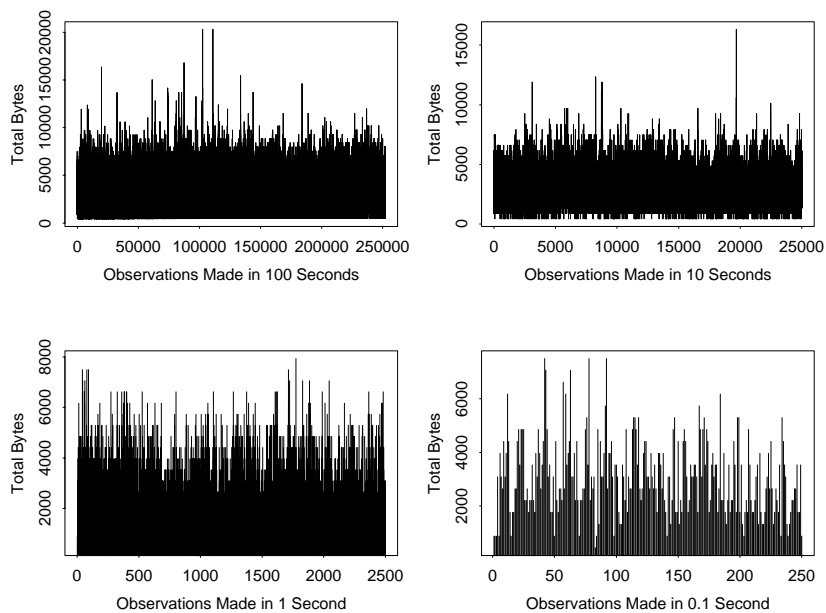


Figure 2: CSD Traffic Trace Over Many Time Scales.

4 Time Series Definition and Generation

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 Boston University [4] [5] have used the time series denoted as T_i and B_i , which are merely the arrival times and the packet size series. Researchers from Bellcore [12] [19] have proposed the {Traffic B} and {Traffic P} time series which combine the notion of arrival times and corresponding packet size.

The relationship of the four time series from the packet arrival times and packet lengths is displayed in Figure 3. Each of the four time series captures different aspects of the traffic trace. T_i is the inter-arrival time series. B_i represents the packet size sequence.

Both of the two timeseries on which the queueing simulations were based ({Traffic B} and

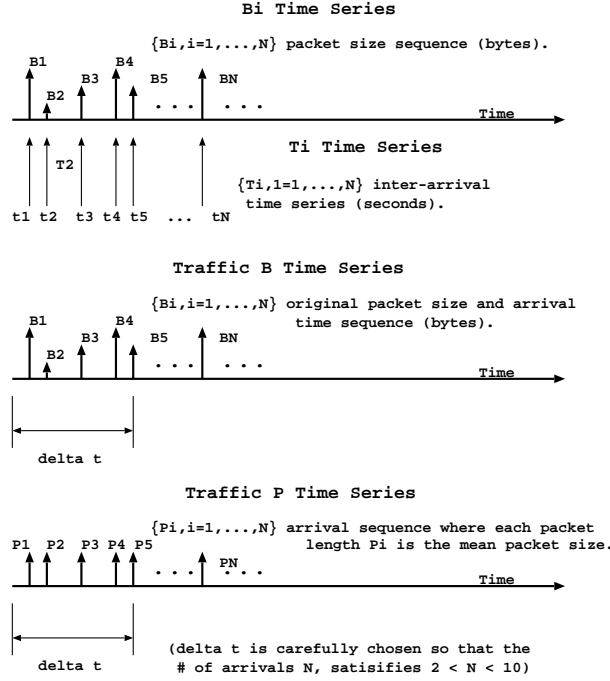


Figure 3: Time Series Diagrams.

{Traffic P}) use the original packet arrival times. But one timeseries {Traffic B} uses the original packet length and the other timeseries {Traffic P} uses the average packet length for the entire trace. Δt is used to aggregate the packet arrivals. 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 \bar{B}_i arriving in each interval Δt_i and obtain the time series {Traffic B} = $\{\bar{B}_i, i = 1, 2, 3, \dots\}$. Next, let \bar{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 \bar{P}_{ave} . Within non-overlapping time intervals of size Δt we sum the number of bytes \bar{P}_i arriving in each interval Δt_i to get a time series {Traffic P} = $\{\bar{P}_i, i = 1, 2, 3, \dots\}$.

The queueing simulation requires both the packet arrival times and the packet lengths. Although the variance-time plots for all four time series for CSD trace all7 are shown in the next section, we can not use the time series B_i (has only packet size information) and T_i

(has only the interarrival information) for the queueing simulation. As a result, we focus on {Traffic P} and {Traffic B} for the Hurst parameter estimation and for the queueing simulation.

We derive 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 the same total number of bytes. Queueing experiments can be driven by them and produce results which can be compared. 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.

5 Hurst Parameter Computation

A. Visual Representation 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 6) represents the slope of the line of $\beta = 1$, that is $var(X^{(m)}) = m^{-1}$. This corresponds to a process with independent interarrivals such as the Poisson process. For such a process, $H = 1/2$. Any process characterized by a slope less than 0 and greater than this reference line exhibits LRD and has an H parameter value $1/2 < H < 1$. Figure 4 shows the variance-time plots for all four time series derived from trace all7. Figure 5 shows the variance-time plots for traces from 3 different computer network types (FDDI, Ethernet and ATM). The ATM cell level traffic trace was taken on Friday March 19 1999 on the vBNS at Downer’s Grove (DNG) in Chicago. The trace represents the network traffic on the vBNS IP over ATM network (<http://www.vbns.net/>). The FDDI traces were taken at UCLA and represent the campus backbone traffic. The variance-time plots for {Traffic B} and {Traffic P} for all four CSD

traces are shown in Figure 6.

The captions represent the curves top down on the graph. By inspection of the variance-time plots, it is apparent that all the 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. Hurst Parameter Computation

The Hurst parameters were computed for the same traffic trace files that served as input for the variance-time plots. The Least-Squares Curve Fitting [15] was used to get an equation for the curves used for the variance-time plots. The resulting equation is in the Slope-Intercept form of the equation of a line $y = \rho(x) + b$ where m is the slope. From the definition in the previous section we know that $-\beta$ is the slope of the curve (so here $-\beta = \rho$). The Hurst value is then computed using the relation $H = 1 - \beta/2$. Table 5 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.

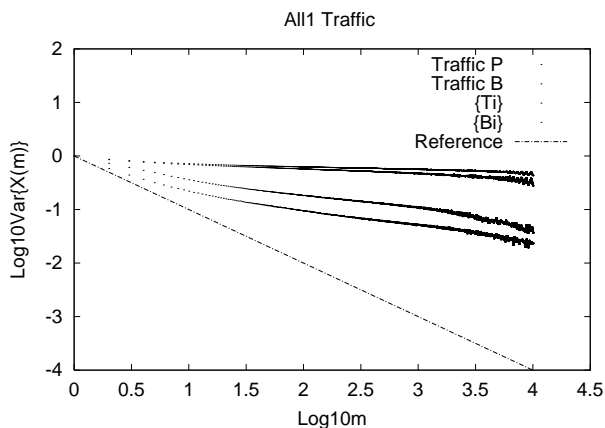


Figure 4: Variance-time plots for all 4 timeseries.

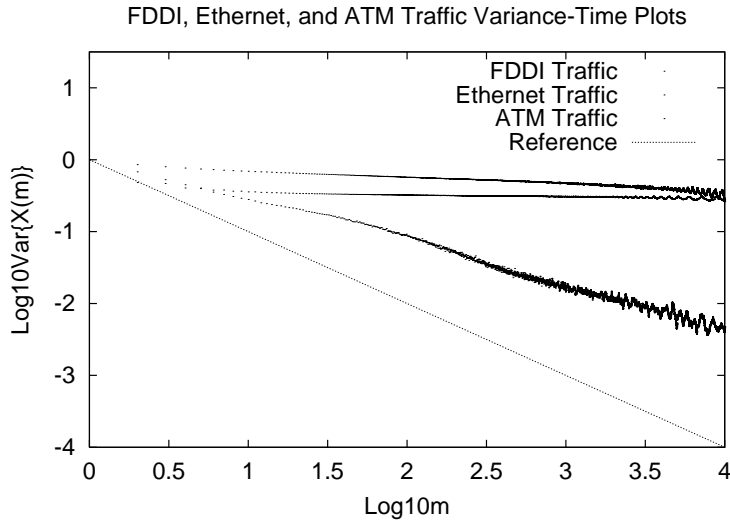


Figure 5: LRD Traffic in Three Different Network Types.

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 5: Hurst values for the 4 traces.

6 The Queue Simulator

In this section, we introduce the queuing simulation environment. 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. We use UCLA traffic, which exhibits the long range dependency property, to drive a queueing simulation. As was done in previous modeling approaches, a single parameter (Hurst parameter) is used to describe the self-similar property of the UCLA traffic. Our paper differs from previous studies, however, in that it addresses the relationship between the Hurst parameter and queueing performance. It is well known that the LRD traffic burstiness, generally characterized by the Hurst parameter is higher than

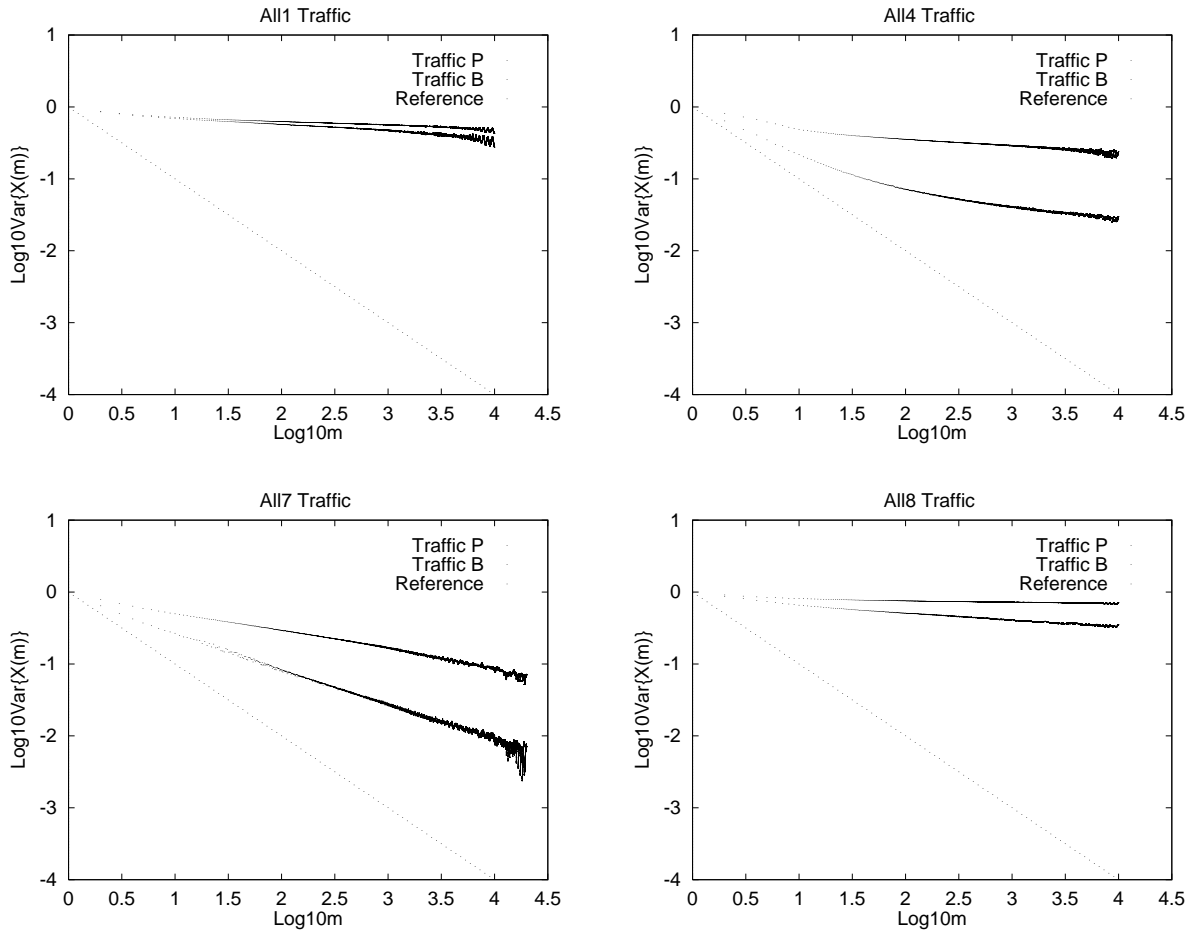


Figure 6: Variance-time plots for the 4 traces.

the traditional Poisson burstiness, and thus leads to much larger queue sizes. By observing the influence of long range dependence on the queueing system, we will see however that the Hurst parameter alone does not sufficiently quantify the LRD property, nor does it completely characterize the traffic burstiness.

The simulator used in our experiments was written in FORTRAN77. It implemented a single server model, infinite buffer size and FIFO discipline. Upon each packet arrival the queue length is calculated using a fluid model. That is, the queue length is counted in terms of waiting packets and also any in-service packet. The packet length of any in-service packet is only the un-serviced part. The queueing simulation is driven by {Traffic B} and {Traffic

P} sequences, which in turn were derived from real traces. Experiments were run with five different server loads (0.3, 0.5, 0.7, 0.9, 0.97). Except for trace All1 (all 5 utilizations are shown), only the experiments with utilizations (0.5, 0.7) are reported here. For each trace and each simulation utilization the Traffic P and B simulations were run for the same total duration, and reached steady state early in the run.

7 Simulation 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 [13] of a crossover point t_m where the variances of the two processes match is not critical in our application. Typically t_m is very small, while we are more interested in large time scale dependencies since these are the ones which determine the tail of the distribution.

In order to provide a reference, Figure 7 shows the queue length distributions for an M/M/1 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. Figure 8 compares the combined queue length distributions for the M/M/1 and LRD All1 trace. Figure 8 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.

A. Queue Length Distributions for Whole Traces.

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 always lead to a larger queue. The following 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

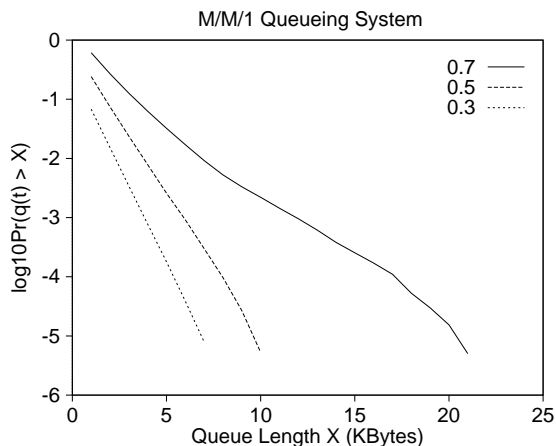


Figure 7: Queueing experiment for M/M/1 model using λ and μ from All1 trace.

were derived from the same traffic trace. Consider trace All1 (Figure 9 (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 5) are very close. But the tail of the queue for {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 9 (c) and Figure 11 (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 consistently 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.

Intuitively, we can easily explain why {Traffic P} leads to shorter queues. Recall that in defining {Traffic P} we have averaged our packet sizes over the entire trace. This averaging renders the offered load more regular than in the original trace and in {Traffic B}. Yet the Hurst parameter is high because of the high variance and interdependence of the interarrival times. This explanation does not dilute, however, our claim, that the Hurst parameter does not accurately predict queue behavior. In fact, this example shows that given a multivariable, complex, long range dependence experiment, it is possible to derive simplified time series which lead to contradictory Hurst parameter / queueing performance behavior.

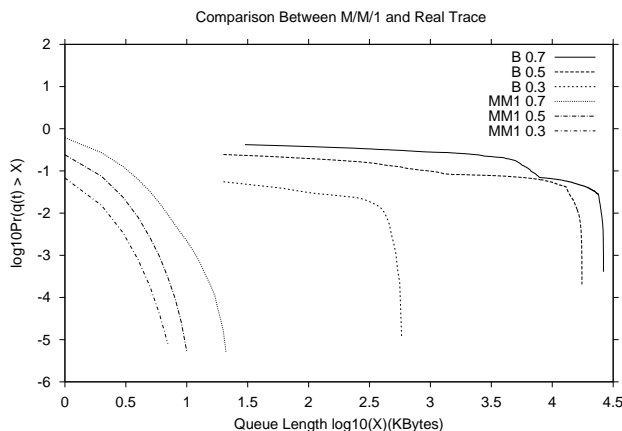


Figure 8: Queueing experiment for M/M/1 model and the LRD All1 trace.

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 6: Hurst values for the four segments of trace All7.

B. Queue Length Distributions for Trace Segments.

More support of our argument comes from the simulation experiments with segments from the same trace. Experiments on segments of the original trace were done earlier by Abry [1] to check if the Hurst parameter was constant across the segments in a test to see if the data was stationary. Our segmented experiments show that parts of the entire trace perform differently from the whole. 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 6 shows that the {Traffic P} and {Traffic B} segments have H values similar to each other. However, the queueing simulations (see Figure 10) show

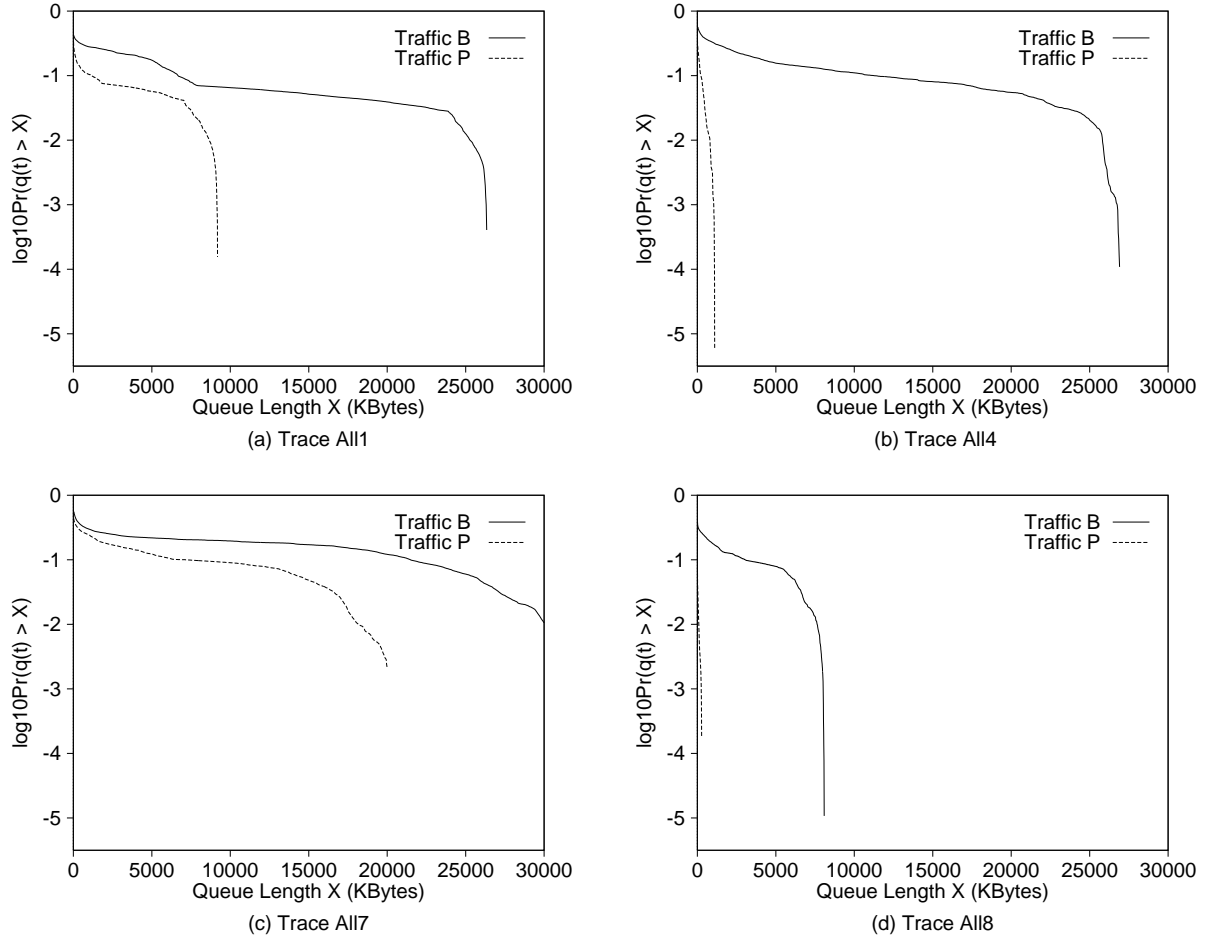


Figure 9: Queueing experiment of each trace with 0.7 utilization.

strong differences between queue length distributions among the four segments. For example, $\{\text{Traffic B}\}$ of Segment3 has the heaviest tail, yet 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 shorter than the tail for the entire trace All7 (see Figure 9 (c)). That is, the largest queue length of the four segments of $\{\text{Traffic B}\}$ is near 18MBytes (Figure 10(a)), but the probability of a queue length larger than 25MByte is still near 10 percent for $\{\text{Traffic B}\}$ of the entire trace (Figure 9 (c)). This is quite surprising, since the value of the Hurst parameter for All7 is smaller than that of the segments.

C. Queue Length Distributions for Different Utilizations.

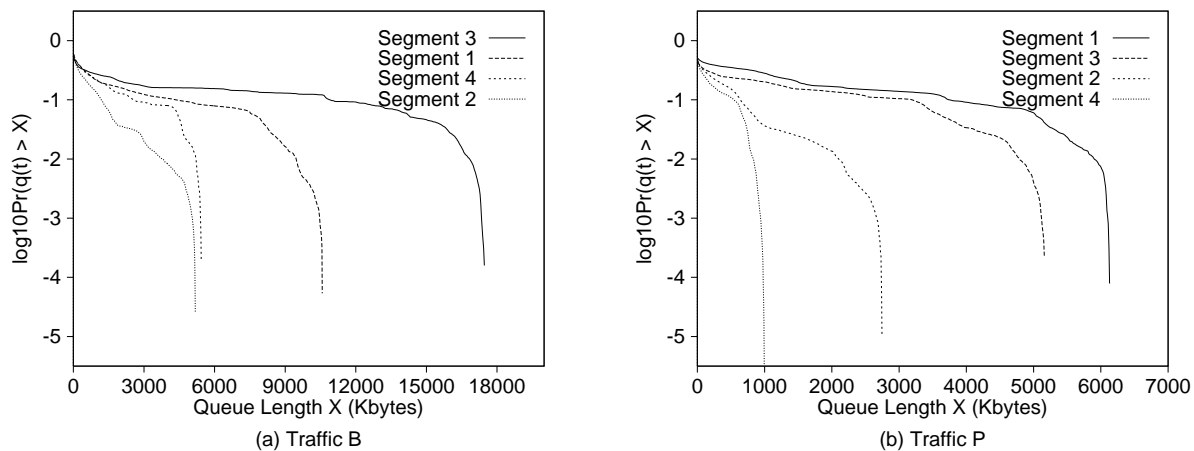


Figure 10: Queueing experiment of Segments from trace All7 with 0.7 utilization.

The queue length distributions of the traces when the system is at 50 percent load is shown in Figure 11. The distribution of queue buffer size decays faster than in 70 percent utilization. Just as Erramilli, et. al., showed in their paper [7], 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 9) is much longer than that of a light load (Figure 11). Therefore, our discussions have focused on the queueing 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 12). 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. Work is under way to develop more accurate predictors [9] [10].

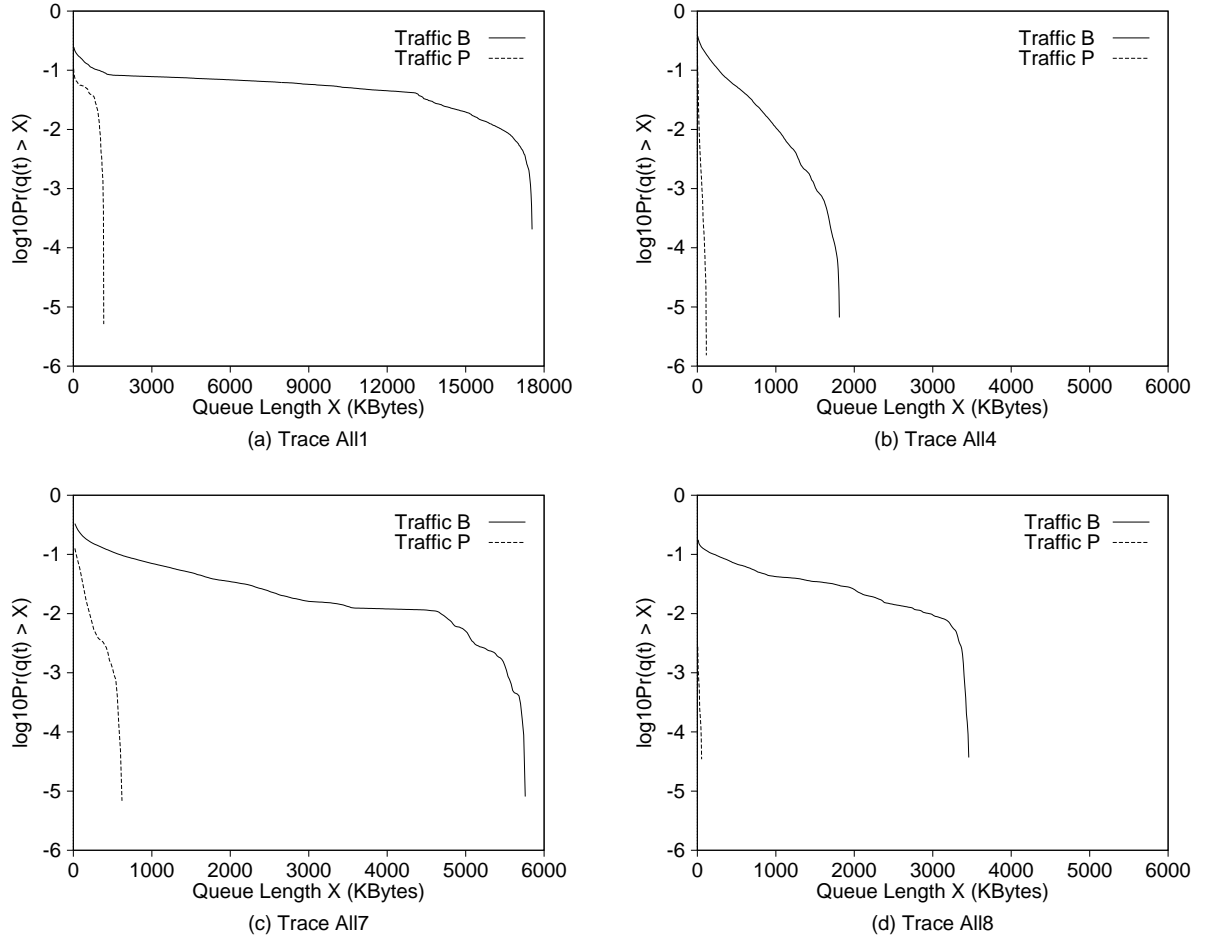


Figure 11: Queueing experiment of each trace with 0.5 utilization.

8 Conclusion

Using real traffic traces in our simulation we have shown that the H parameter is not a consistent, indicator of queueing performance. 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}. {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}. This contradictory relation also holds true for trace segments derived from the original traces. Namely, the value of the Hurst parameter does not fully

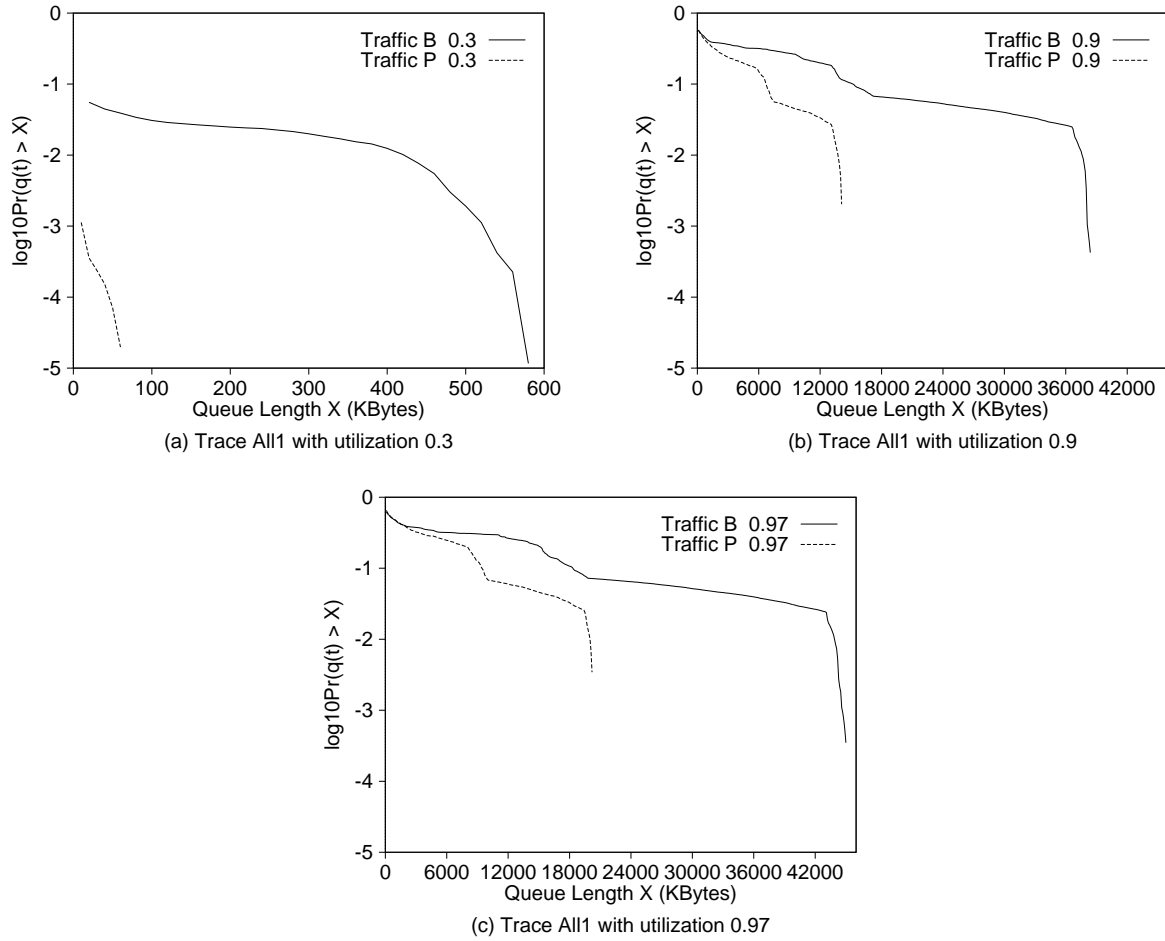


Figure 12: Queueing experiments for trace All1 with 0.3, 0.9 and 0.97 utilizations.

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. Based on the same idea, Dr. Rubin and Jianbo Gao at UCLA have developed a new model using two parameters [9] [10]. 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 *Acknowledgments*

The authors would like to thank Jianbo Gao for stimulating discussions and for providing some of the software tools used in our experiments. Scott Seongwook Lee for his help in developing some data reduction tools. Many thanks also to the UCLA CSD staff for their helpful cooperation.

References

- [1] Patrice Abry, Darryl Veitch, Wavelet Analysis of Long Range Dependent Traffic, *Trans. Info. Theory*, 44, no. 1, (1998) 2-15.
- [2] R. G. Addie, M. Zukerman, and T. Neame, Performance of a Single Server Queue with Self Similar Input. *IEEE INFOCOM'95 Seattle*, (1995).
- [3] Jan Beran, Robert Sherman, M. S. Taqqu and W. Willinger, Long-Range Dependence in Variable-Bit-Rate Video Traffic. *IEEE/ACM Trans. on Communications*, 2, no. 4, (1995) pp. 1566-1579.
- [4] Mark E. Crovella and Azer Bestavros, Explaining World Wide Web Traffic Self-Similarity. Technical Report TR-95-015, Boston University (1995).
- [5] Mark E. Crovella and Azer Bestavros, Self-Similarity in World Wide Web Traffic Evidence and Possible Causes. *1996 ACM SIGMETRICS Int. Conf. Measurement, Modeling of Computer System*, 1996, pp. 160-169.
- [6] D. E. Duffy, A. A. McIntosh, M. Rosenstein, and Walter Willinger, Statistical analysis of CCSN/SS7 traffic data from working subnetworks. *IEEE J. Select. Areas Commun.*, 12, no. 3, (1994) 544-551.
- [7] A. Erramilli, O. Narayan, and W. Willinger, Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Trans. on Networking*, 4, no. 2, (1996) 209-223.

- [8] H. J. Fowler and W. E. Leland, Local area network traffic characteristics, with implications for broadband network congestion management. *IEEE J. Select. Areas Commun.*, 9, (1991) 1139-1149.
- [9] Jianbo Gao and Izhak Rubin, Multifractal Modeling of Counting Processes of Long-Range Dependent Network Traffic. *ATS 99*, San Diego 1999.
- [10] Jianbo Gao and Izhak Rubin, Multifractal Analysis and Modeling of Long-Range Dependent Network Traffic. *IEEE ICC 99*, Proceedings of the ICC 99, Vancouver. British Columbia, Canada, (1999), 382-386.
- [11] Van Jacobson, Craig Leres and S. McCanne, The Tcpcdump Manual Page. Lawrence Berkeley Laboratory, (1989) <http://ee.lbl.gov/>, <ftp://ftp.ee.lbl.gov/tcpdump.tar.Z>.
- [12] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Trans. on Networking*, 2, no. 1, (1994) 1-15.
- [13] Arnold L. Neidhardt and Jonathan L. Wang, 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, (1998) 222-232.
- [14] Vern Paxson and Sally Floyd, Wide-Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Trans. on Networking*, 3, no. 3, (1995) 226-244.
- [15] Trishor S. Trivedi, Probability & Statistics with Reliability, Queueing, and Computer Science Applications. *Prentice-Hall Inc.*, Englewood Cliffs New Jersey, 1982, pp. 543-545.
- [16] Darryl Veitch and Patrice Abry, A Wavelet Based Joint Estimator of the Parameters of Long-Range Dependence, Presented at GRETI Grenoble, (1997), <http://universe.serc.rmit.edu.au/darryl/>.

- [17] Walter Willinger, Siddhartha Y. Devadhar, et. al., Measuring ATM Traffic Cell-by-Cell: Experiences and Preliminary Findings from BAGNet. PMCCN '97, International Conference on the Performance and Management of Complex Communication, Tsukuba, Japan, November 17-21, 1997.
- [18] W. Willinger, M. S. Taqqu, et. al., Self-Similarity in High-Speed Packet Traffic: Analysis of Modeling of Ethernet Traffic Measurements. *Statistical Science*, 10, no. 1, (1995) 71-86.
- [19] W. Willinger, M. S. Taqqu, M. S. Sherman, and D. V. Wilson, Self-similarity through high-variability: Statistical analysis of ethernet LAN traffic at the source level (extended version). *IEEE/ACM Trans. on Networking*, 5, (1996) 71-86.

BIOGRAPHIES

Ronn Ritke received an M.S. from California State University Long Beach in 1994 in computer science, a second M.S. and a Ph.D. from the University of California at Los Angeles in 1996 and 2000, all in computer science. His research interests include network traffic measurement and analysis. He is currently working with the San Diego Super Computer Center and NSF to help provide answers on vBNS performance for the President's Information Technology Advisory Committee (PITAC).

Xiaoyan Hong is a Ph.D. student in Computer Science Department at the University of California, Los Angeles; Recieved a B.E. degree and a 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 was a visiting researcher for two years 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.