

On Long-Range Dependence in NSFNET Traffic*

Steven M. Klivansky
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332
smk@cc.gatech.edu

Amarnath Mukherjee
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332
amarnath@cc.gatech.edu

Cheng Song
Advantis, Inc.
1311 Mamaroneck Ave.
White Plains, NY 10605
csong@vnet.ibm.com

GIT-CC-94-61

Last Revised: December 19, 1994

Abstract

Several recent studies indicate that (i) dependent packet arrivals can have a significant impact on buffer-overflow probabilities and delay characteristics at a network multiplexor, with long-range dependence showing a larger impact than short-range dependence, and (ii) that traffic in several different campus-level network environments appears to be long-range dependent. These two related issues are beginning to make a significant impact on traffic theory and network resource management.

The objective of this paper is to show that long-range dependence extends to wide-area networks as well, and to identify the factors contributing to it. In particular, we find the presence of long-range dependence in packet level traffic over the NSFNET wide-area network. The analysis is based on raw traffic data collected from multiple NSFNET core switches distributed across the US.

Investigation of the underlying processes for a subset of traffic created by TCP indicates that the (heavy-tailed) joint distribution of TCP conversation sizes and conversation transmission rates are the key factors contributing to long-range dependence. Since these distributions are primarily application-level characteristics and to a certain extent network independent, long-range dependence will be present in highly-multiplexed data networks that are running current TCP applications.

*The work of Steven M. Klivansky and Amarnath Mukherjee was supported in part by the National Science Foundation under grant NCR-9396299.

Contents

1	Introduction	3
2	Trace Gathering and Data Pre-Processing	5
2.1	Trace Gathering	5
2.2	Processing and Interpretation	7
2.3	Processing Mode C Traces	7
3	Background on Long-Range Dependence	7
3.1	Long-range and short-range dependence	7
3.2	Self-similarity	9
3.3	Implications of Long-Range Dependence	9
3.4	Estimation of H	10
4	Summary of Observations	10
4.1	Packet Level Statistics	10
5	Factors Contributing to Long Range Dependence in NSFNET TCP Traffic	12
5.1	Objective	12
5.2	Designed Experiments	13
5.2.1	Introduction	13
5.2.2	Potential causes	13
5.2.3	Selection of factors	14
5.2.4	Selection of levels	15
5.2.5	Simulation model	16
5.2.6	Experimental results	17
5.2.7	Analysis of variance	18
6	Conclusion	19
	References	22

1 Introduction

A stochastic process $\{X_k\}$, ($k = 0, 1, 2 \dots$), is said to be long-range dependent if it is wide-sense stationary, and the sum of its auto-correlations diverge. An important property of long-range dependence is that the variance of $(X_1 + \dots + X_n)/n$ decays slower than $1/n$ even as $n \rightarrow \infty$, or equivalently, its power spectral density diverges as the frequency, w , goes to 0. (See Section 3.1.) In contrast, for an independent or short-range dependent process, the variance of $(X_1 + \dots + X_n)/n$ decays as $1/n$ as $n \rightarrow \infty$, and its power spectral density is finite as $w \rightarrow 0$. Several recent studies have shown that buffer-overflow probabilities and delay statistics for a single queue with long-range dependent arrivals are significantly larger than that with independent or short-range dependent arrivals [1, 14, 16].

A special case of long-range dependence where the auto-correlation structure is invariant at different time scales (“second-order self-similarity”) was observed recently for packet-level traffic over the Bellcore Ethernet local area network [23]. The objectives of this paper are:

- (i) To show that packet-level traffic exhibits long-range dependence even at large multiplexors such as NSFNET core switches,
- (ii) To observe that TCP traffic, in general, exhibits stronger long-range dependence than the overall traffic mix at these switches, and
- (iii) To identify through a set of statistically rigorous experiments the relative impact of application-level factors in TCP traffic on observed long-range dependence.

(i) and (ii) are based on traffic measurements over several geographically dispersed core switches on the NSFNET. (iii) is based on a systematic study of the effect of TCP application parameters (*factors*) on long-range dependence through a set of designed experiments [6, 22].

The results of this paper support and extend the traffic results of Leland et al [23] and Paxson and Floyd [30] to larger sub-nets. Table 1 summarizes the main results.

A key conclusion is that long-range dependence in TCP traffic is primarily caused by the joint distribution of two TCP conversation parameters: the one-way number of packets per conversation and the conversation duration. Specifically, it is found that:

- (i) The (marginal) distribution of conversation duration, Z , is heavy tailed with $P(Z = k) \propto k^{-\beta-2}$, $0 < \beta < 1$. A theorem due to Cox and Isham[9] states that this, in conjunction with (a) Poisson conversation arrivals and (b) unit transmission rate per active conversation, results in second-order self-similarity (special case of long-range dependence) in the packet arrival process.
- (ii) The (marginal) distribution of the number of packets per conversation is heavy tailed.
- (iii) The average transmission rate of a conversation, defined as the ratio of number of packets transmitted to transmission duration, increases with number of packets with high probability. A relationship of the form $\log(\text{Rate}) \sim \log(\#\text{Packets}) + \epsilon$, with ϵ a normally distributed residual, appears to hold approximately, and
- (iv) Even if conversation inter-arrival dependencies are ignored, these two factors together can create significant long-range dependence in aggregate TCP traffic.

The designed experiments to determine the factors contributing to long-range dependence were organized as follows. The factors chosen were *distributions* for conversation inter-arrivals, packet inter-arrivals within a conversation, conversation item size (number of packets), and a variable corresponding

<p>Packet level traffic at NSFNET core switches exhibit long-range dependence. For some datasets, and for a subset of the overall arrival process, packet level traffic is also asymptotically second-order self-similar. For the overall arrival process, estimates of the Hurst coefficient, \hat{H}, (a measure of long-range dependence) were in the range 0.69 to 0.84 with a median of 0.74. For traffic specific to TCP applications, the \hat{H} estimates were in the range 0.78 to 0.89 with most values above 0.8.</p>
<p>The joint distribution of observed number of packets per conversation and the observed transmission rate for each conversation is a primary contributor to long-range dependence in TCP traffic over the NSFNET. Together, they explain about 80% of the response variation in \hat{H} in a full factorial experimental design. In comparison, the distribution of packet interarrival time per conversation explains about 15% and the distribution of the conversation arrival process explains about 0.1% of the variation.</p>
<p>Since the primary factors contributing to long-range dependence in TCP traffic are application-level characteristics (e.g., file-size distributions), long-range dependence may be expected to be present in most highly-multiplexed data networks.</p>
<p>Long-range dependent arrivals at a single backbone switch may be generated using a simple simulation containing accurate empirical distributions for the mean conversation transmission rate, number of packets in a conversation, packet interarrival time, and conversation interarrival time.</p>

Table 1: Key Findings.

to a conversation’s duration. (Through this last factor, we also attempt to capture the dependency between number of packets per conversation and conversation duration.) Alternative distributions for a factor were assigned to different levels (settings at which a factor is evaluated in a designed experiment), and simulation experiments were run for all possible combinations of levels and factors. The response variables of the experiments were chosen to be the Hurst coefficient (a measure of long-range dependence, defined in Section 3). The effect of each factor on each response variable was then quantified using Analysis of Variance with nested factors [6]. This led to the identification of the factors and their levels that lead to high \hat{H} estimates for TCP traffic.

The rest of the paper is organized as follows. Section 2 describes the network environment, the traffic measurement procedure, and post-processing of the data. Section 3 gives a brief background on long-range and short-range dependence. Section 4 presents statistical analysis of the trace data. Section 5 addresses factors contributing to long-range dependence in TCP traffic. Section 6 presents the conclusions of this study and outlines future work.

2 Trace Gathering and Data Pre-Processing

2.1 Trace Gathering

Each NSFNET Core Nodal Switching System (CNSS) location consists of several IBM RS6000 based routers. Each CNSS location is an MCI Point-of-Presence (POP) in the network. These routers can be linked via a variety of network media including T1, T3, Ethernet and FDDI. The NSFNET CNSS architecture underwent a series of changes during its lifetime. During the phase of this research when the traces were gathered, the routers within a backbone site were connected initially via T3 links and later via an FDDI ring.

Figure 1 gives a high level view of the router's architecture. IP packet routing and forwarding functions are distributed over a configurable number of adapter cards plugged into the IBM RS6000 MicroChannel I/O bus. Each card supports receive and transmit functions on at least one type of network medium. The system runs the routing daemon, and downloads its route table updates to all the cards. Each card has an Intel 80960 RISC processor and performs IP routing using the downloaded route table. Normally, IP packets are routed and forwarded from card to card without passing through the system.

The packet tracing point on the adapter cards and the packet sampling point in the system are shown in Figure 1. Tracing was done at the input port of the adapter card's network interface. The tracing code is part of the receive-path packet-processing module of the on-card microcode. It can be toggled by commands from the system, with parameters specifying the type of packets to trace by protocol family, size of each packet header to obtain, etc. In addition to the header, for each packet traced, we also collected a packet arrival time stamp, the total frame size including the link level header size, and a sequence number by which we can detect losses.

To ensure minimal interference with the real-time performance of packet routing and forwarding, the tracing code is not separately scheduled by packet arrival interrupt. Instead, it is run whenever the microcode kernel schedules the receive path processing module to process packets that have been received into packet buffers. The clock that the time stamp uses has a 14.4 MHz frequency, but the timestamp accuracy is estimated to be within a millisecond due to waiting for the scheduled processing time. We consider such accuracy to be sufficient for the purposes of our analysis since the smallest time scales we are concerned with are of the order of 10 milliseconds (as we shall see in Sections 4 and 5).

Before IP routing and card-to-card transfer of a received packet, the tracing code assembles a record in a buffer with the packet header, its time of origination, and size information. When the buffer is filled with a configurable number of such records, it is scheduled by the microcode kernel to be DMA transferred to the system memory. On the system, a sampling program sends tracing commands down to the interface card and reads the packet records DMA'ed back from the trace-activated card. The program runs as a user process and is given a low enough priority so as to not compete for CPU resources against other time-critical processes. The program selects specific packets to be included in the trace file by filtering based on various header fields.

A packet can be traced as long as it is received and routed. However, a packet loss may occur if the DMA transfer momentarily overruns the sampling program on the system. In order for the sampling code to detect any such loss, each such buffer is tagged with a sequence number. Discontinuities in buffer sequence numbers indicate losses. During a few of the tracing runs, we observed a loss of traced packets in the range of one to five hundred out of one million packets (i.e., less than 0.05%). This is clearly not desirable if these traces are to be used in trace-driven simulations of buffer management, scheduling or equivalent bandwidth studies. However, given their relative infrequency of occurrence, we believe that they should not significantly affect the long range auto-correlations in the data, which is

Filter Mode	Location	Date and Start Time	Duration (sec)	Pkts or Conv	Apparent H
ftp-control, nntp, and smtp (Mode A)	Greensboro	09/03/93 09:45	4120.450282	293736	0.784
	Denver	09/03/93 09:45	2077.906371	326388	0.823
	Hartford	09/01/93 12:05	3208.185325	242806	0.778
	Hartford	09/03/93 09:45	4103.843185	500003	0.852
	San Francisco	09/03/93 09:45	2470.963851	500003	0.848
	Seattle	09/03/93 09:45	2768.917640	500005	0.870
	St. Louis	09/03/93 09:45	2357.191570	499993	0.853
ftp-data, telnet, and rlogin (Mode B)	Greensboro	09/03/93 11:05	2113.847107	393626	0.836
	Denver	09/03/93 11:05	764.634235	361225	0.790
	Hartford	09/01/93 11:15	2315.540459	269414	0.824
	Hartford	09/01/93 13:20	2148.679826	243826	0.889
	Hartford	09/03/93 11:05	2123.836526	500015	0.882
	San Francisco	09/03/93 11:05	1192.609671	500002	0.829
	Seattle	09/03/93 11:05	1321.228742	500012	0.883
	St. Louis	09/03/93 11:05	785.107739	500023	0.843
SYN or FIN flag set (Mode C)	Greensboro	09/02/93 11:05	18034.251051	115931	
	Greensboro	09/03/93 11:50	11352.876624	79887	
	Denver	09/02/93 11:05	15387.292089	75507	
	Denver	09/03/93 11:50	12120.272298	47298	
	Hartford	09/02/93 11:05	18027.765766	277057	
	Hartford	09/03/93 11:50	11588.213249	221835	
	San Francisco	09/02/93 11:05	16064.823477	262855	
	San Francisco	09/03/93 11:50	12180.219205	179392	
	Seattle	09/02/93 11:05	16092.873998	254269	
	Seattle	09/03/93 11:50	12234.860278	183643	
	St. Louis	09/02/93 11:05	13881.642478	250636	
	St. Louis	09/03/93 11:50	12072.583719	191590	
All packets (Mode D)	Hartford	05/17/94 14:20	518.557384	1000000	0.811
	Hartford	05/17/94 14:40	542.335767	1000020	0.762
	Cleveland	05/17/94 14:40	235.998821	1000020	0.841
	Washington	12/10/93 11:20	1978.396636	1000000	0.721
	Washington	12/10/93 14:36	1847.151878	1000000	0.722
	Chicago	12/10/93 14:55	1748.708039	647400	0.694

Table 2: Summary of Traces

the subject of this study.

As CNSS routers, the RS6000 machines were not configured with adequate disk space for prolonged tracing. In the initial phase of our measurements, the total number of sampled packets for each tracing session was limited to no more than 500,000 so that the data could fit on the disk. The generated data file had to be transferred to a repository machine before a new tracing session could begin since sending sampled data to a remote machine with tracing in progress would bias the traffic measurement results. Later enhancements allowed for 1 million packets to be captured.

The traces on which this paper is based were performed to study a series of network issues. Given the limited disk space available, we used filtering, especially at the initial stages when disk space was more constrained, to store only those packet headers that belonged to a particular set of applications. Table 2 summarizes information about the traces collected. The *Apparent H* column refers to the estimated Hurst parameter observed for the dataset. The *Filter Mode* column refers to the types of packets which were included in the trace file. To simplify future references, we will refer to the traces containing ftp-control, nntp, and smtp traffic as “Mode A,” the ftp-data, telnet, and rlogin traces as “Mode B,” the SYN or FIN traces that mark the beginning or end of a TCP conversation as “Mode C,” and the traces with all packets included as “Mode D.” The Mode D traces were collected after an upgrade to the trace system and consisted of approximately one million IP packets each, without any filtering at the application level.

Due to the time and resources involved in performing tracing at all CNSS sites, only the subset of NSFNET nodes shown in Table 2 was used.

2.2 Processing and Interpretation

The packet records in the traces we collected were binary records of header fields. To simplify subsequent processing, we translated the records into a format which `tcpdump` [34] could understand. We then employed `tcpdump`'s powerful command syntax to select only packets that fit certain criteria. While Modes A, B, and D traces were ready for analysis, Mode C traces required further filtration and validation tests.

2.3 Processing Mode C Traces

The primary objective of Mode C traces was to obtain TCP conversation-level statistics (e.g., conversation arrival times, their durations, number of bytes transferred, and TCP application type). We were concerned with the bidirectional nature of TCP conversations (for a related project), and thus tagged each packet with a conversation identifier based on the full TCP/IP address of both conversation endpoints. This meant assigning a unique key to a bidirectional TCP connection based on its source and destination IP addresses and transport-level port numbers, irrespective of whether the IP address and the transport port number appeared as the source or the destination. The effect of this scheme was to assign different connections between the same pair of hosts to different conversations. For example, consecutive `ftp-data` connections initiated via an `mget` command would appear as separate conversations since a different transport port is used on the client side for each item transferred. This tagging scheme enabled us to examine network traffic as an aggregation of individual conversations.

Several validation filters were run on the data. First, we needed to ensure that SYN-FIN pairs belonged to the same side of a conversation so that the difference in sequence number could be used to calculate conversation duration and the simplex number of bytes transferred examined in Section 5. In addition, only those conversations for which both a SYN and a FIN packet can be found in the trace are used, as it is impossible to determine conversation duration and size without both.

Another important concern for Mode C traces was to ensure that the conversation tagging scheme did not accidentally group separate conversations with a single tag. This could happen since higher layer entities could create the same transport level address over a trace duration. We expected such occurrences would be rare and indeed found that it took place in 74 out of 38465 conversations (0.19%) in all mode C traces combined. Among these 74 conversations, all but two were repeat attempts to establish an ftp conversation a few seconds apart from each other. These conversations were excluded from the analysis to simplify data processing.

3 Background on Long-Range Dependence

3.1 Long-range and short-range dependence

We begin with a brief overview of long-range and short-range dependence in time series data. The interested reader is referred to Cox's excellent review of the subject [9]. Readers familiar with this topic may skip this section.

Let $\{X_t\}$ ($t = 0, 1, 2, \dots$) be a wide-sense stationary stochastic process, i.e., a process with a stationary mean $\mu = E[X_t]$, a stationary and finite variance $v = E[(X_t - \mu)^2]$, and a stationary auto-covariance function $\gamma_k = E[(X_t - \mu)(X_{t+k} - \mu)]$, ($k = 0, 1, 2, \dots$) that depends only on k , and not on t . Let the auto-correlation of $\{X_t\}$ at lag k be denoted as ρ_k , and let its spectral density function be denoted by

$g(w)$, where, for $-\pi < w < \pi$,

$$g(w) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \gamma_k e^{-ikw}. \quad (1)$$

Observe that $v = \gamma_0$ and $\rho_k = \gamma_k/\gamma_0$. Now, for each m , let $\{X_j^{(m)}\}$ denote the sample mean of $X_{jm-m+1} \cdots X_{jm}$, i.e.,

$$X_j^{(m)} = \frac{1}{m} (X_{jm-m+1} + \cdots + X_{jm}). \quad (2)$$

Let v_m , $\{\gamma_k^{(m)}\}$ and $\{\rho_k^{(m)}\}$ denote the variance, the auto-covariance and the auto-correlation functions of $\{X_j^{(m)}\}$. Then $v_1 = v$ and v_m is given by [9]

$$v_m = E \left[\frac{1}{m} (X_{jm-m+1} + \cdots + X_{jm}) \right]^2 - \left[E \frac{1}{m} (X_{jm-m+1} + \cdots + X_{jm}) \right]^2, \quad (3)$$

$$= \frac{v}{m} + \frac{2}{m^2} \sum_{k=1}^m (m-k)\gamma_k, \quad (4)$$

$$= \frac{v}{m} + \frac{2}{m^2} \sum_{s=1}^{m-1} \sum_{k=1}^s \gamma_k. \quad (5)$$

Equation (4) follows from (3) by expanding the first term, applying the definition of γ_k , and simplifying. It follows from (5) that

$$\gamma_k = \frac{1}{2} \nabla^2 (k^2 v_k), \quad (6)$$

where ∇^2 denotes the central second-difference operator defined as $\nabla^2(y) = y_{k+1} - 2y_k + y_{k-1}$ for some vector $\{y\}$. The reader may verify (6) by expanding its right hand side, substituting for v_k with the right hand side of (5), and simplifying.

The auto-covariance function for $\{X_j^{(m)}\}$ is given by

$$\gamma_k^{(m)} = \frac{1}{2} \nabla^2 (k^2 v_k^{(m)}), \quad (7)$$

$$= \frac{1}{2} \nabla^2 (k^2 v_{km}). \quad (8)$$

The process $\{X_t\}$ is said to have short-range dependence if $\sum_k \gamma_k < \infty$. Equivalently, from (5), v_m , for large m , is asymptotically of the form v'/m , with v' finite.

The process $\{X_t\}$ is said to have long-range dependence if $\sum_k \gamma_k \rightarrow \infty$. Equivalently, from (4), $mv_m \rightarrow \infty$ as $m \rightarrow \infty$.

An example of a correlation structure that leads to long-range dependence is when, for $0 < \beta < 1$, the following (equivalent) relations hold:

$$\gamma_k \sim \gamma' k^{-\beta}, \quad \text{for large } k, \quad (9)$$

$$v_m \sim v' m^{-\beta}, \quad \text{for large } m, \quad (10)$$

$$g(w) \sim g' w^{-(1-\beta)} \quad \text{as } w \rightarrow 0. \quad (11)$$

Simple relations based on (1) and (5) connect γ' , v' and g' . It follows from (8) and (10) that

$$\rho_k^{(m)} = \frac{\gamma_k^{(m)}}{v_{km}} \rightarrow \frac{1}{2} \nabla^2 (k^{2-\beta}), \quad (12)$$

$$\doteq \frac{1}{2} (2-\beta)(1-\beta) k^{-\beta}, \quad (13)$$

Equation (13) holds for large k from the asymptotic equivalence of differencing and differentiation.

A hyperbolically-decaying auto-covariance function shown in (9) is one example of a stationary long-range dependent process because $\sum_k \gamma_k \rightarrow \infty$. Other forms are possible as long as $\sum_k \gamma_k \rightarrow \infty$. A key property of long-range dependence is that the variance of $\{X_j^{(m)}\}$ decays slower than m^{-1} even for large m .

In comparison, a process with an exponentially-decaying auto-covariance function, $\gamma_k \sim \gamma' e^{-\beta k}$ for large k , leads to a short-range dependent process because $\sum_k \gamma_k$ converges. An example process that exhibits this property is the auto-regressive moving average (ARMA) process where γ_k is a sum of exponentials.

3.2 Self-similarity

The process $\{X_t\}$ is said to be exactly (second-order) self-similar if $\rho_k^{(m)} = \rho_k$ for all m and all k , i.e., the correlation structure is preserved across different time scales. Such a process is frequently characterized by the Hurst coefficient, H , where $H = 1 - \beta/2$. Notice that, for $0 < \beta < 1$, we have $1/2 < H < 1$.

It was observed by Hurst [20] and by Mandelbrot and Wallis [27, 28] that many naturally occurring time series¹ exhibit $H > 1/2$. Examples include annual water flows of the Nile and several other rivers, sun-spot numbers, tree-ring indices, etc. The Leland et. al. study [23] on traffic traces from the Bellcore Ethernet also shows second-order self-similarity with $H > 1/2$. In contrast, traditional Poisson and related Markovian models used in modeling data traffic, lead to $H = 1/2$, because the auto-correlations for these processes vanish at larger time scales.

For a given set of observations $\{X_k, k = 1, 2, \dots, n\}$ with sample mean $\bar{X}(n)$ and sample variance $S^2(n)$, the re-scaled adjusted range or the R/S statistic is defined as

$$R(n)/S(n) = 1/S(n)[\max(0, W_1, \dots, W_n) - \min(0, W_1, \dots, W_n)],$$

with

$$W_k = (X_1 + \dots + X_k) - k\bar{X}(n), k = 1, 2, \dots, n.$$

If $\{X_k\}$ is a white noise process, or a process with a short memory correlation structure, it is known that $E[R(n)/S(n)]$ is asymptotically proportional to $n^{1/2}$. Hence a plot of $\log[R(n)/S(n)]$ versus $\log(n)$ should display an asymptotically linear trend with slope $1/2$ as $n \rightarrow \infty$. However, for long-range dependent processes, $E[R(n)/S(n)]$ is asymptotically proportional to n^H [9, 20]. Therefore, a plot of $\log[R(n)/S(n)]$ versus $\log(n)$ should display an asymptotically linear trend with slope H as $n \rightarrow \infty$.

3.3 Implications of Long-Range Dependence

In the context of network traffic, long-range dependence implies that, even if the number of packet arrivals were aggregated over larger and larger time-scales, its variance will not die out as one would expect with a finite-variance white-noise (i.e., independent) process such as Poisson traffic or a short-memory process such as Markov-Modulated Poisson Process [23]. Instead, the aggregate would show visually and statistically similar patterns of “bursty sub-periods and less bursty sub-periods” for small as well as large time-scales (see [23] for some instructive graphical illustrations).

Queuing analyses with long-range dependent input traffic is an open problem. Simulation results, however, implied a dramatic impact on queue statistics [1, 16, 18] if the process is long-range dependent.

¹The Central Limit Theorem for independent random variables with finite variance leads to $H = 1/2$. Mandelbrot and Van Ness [26] has shown that $H = 1/2$, even for short memory processes with finite variance.

It appears that two factors responsible for large queue statistics are the long-range dependence structure and the tail distribution of the arrival process.

3.4 Estimation of H

Several heuristic graphical techniques are available for “eye-balling” the presence of long-range dependence. These include:

- The *Variance-Time plot (VT-plot)* which validates Equation (10) by plotting $\log(v_m)$ versus $\log(m)$ for different aggregation levels, m . The slope of the linear least-squares fit line leads to an estimate of $-\hat{\beta}$. Since, by definition, $\hat{H} = 1 - \hat{\beta}/2$, and $0 < \beta < 1$ in (10), the test involves checking if the slope, $-\hat{\beta}$, satisfies the relationship $0 > -\hat{\beta} > -1$.
- The *Periodogram plot (PG-plot)* which validates Equation (11) by plotting $\log(g(w))$ versus w , and applying a linear least-squares fit to a section of the plot corresponding to small w . (A heuristic is to apply the test to the smallest 10% of the frequencies in the Periodogram.) The slope of this fit is set to $-(1 - \hat{\beta})$ (see (11)), from where \hat{H} is computed as $1 - \hat{\beta}/2$.
- The *R/S analysis* where $\{\log[R(n)/S(n)], \log(n)\}$ is computed for different values of n and then \hat{H} is estimated from the slope of the linear fit of $\log[R(n)/S(n)]$ versus $\log[n]$. Observe that for each value of n , there may be more than one set of $\{\log[R(n)/S(n)], \log[n]\}$ corresponding to non-overlapping intervals in the data.

The R/S statistic has good robustness properties, in particular, with respect to long-tailed distributions. However, its drawback is that it can lead to biased estimates [4]. The VT-plots and PG-plots are, similarly, useful starting points in data analysis. However, neither of these techniques are suitable for precise estimation [23].

More refined statistical analyses based on Maximum Likelihood Estimate (MLE) methods are used to estimate H when $\{X_t\}$ comes from a Gaussian distribution. These methods and their approximations are based on the spectral density of X_t and properties of the MLEs are discussed by a number of authors [4, 10, 37]. One specific method that has been used extensively is Whittle’s approximate MLE (e.g., see [23]). Most of the approximate MLEs are defined via quadratic forms which make them quite sensitive to deviations from normality. However, transformation of the data may sometimes alleviate these types of problems. Bayesian methods of estimating H are also proposed and discussed in [5].

4 Summary of Observations

4.1 Packet Level Statistics

Using the packet arrival timestamp in the NSFNET traces (Modes A, B, and D in Table 2) each trace was divided into 0.01 second slices and the number of packet arrivals in each slice was recorded. The slice value was chosen to provide details of network traffic at small time scales while retaining the ability to derive information for larger time scales via aggregation of neighboring slices. Figure 2 shows a section of traffic at time scales of 0.01, 0.1, 1.0 and 10.0 seconds. We observe that traffic remains bursty at all these time scales. However, they do not “look” self-similar.

The datasets were tested for long-range dependence and second-order self-similarity, first using the V-T, Periodogram, and R/S estimation methods discussed in Section 3.4, and then via Whittle’s approximate MLE. Figure 3 shows the results for the first Mode D Hartford trace. We observe that:

- The slope of the V-T plot is between -1 and 0 indicating that the variance v_m decays slower than $1/m$. For this dataset, the VT-plot's slope is estimated as $-\hat{\beta} = -0.471$. Therefore $\hat{H} = 1 - \hat{\beta}/2 = 0.765$.
- Several frequency components in the low-frequency domain in the Periodogram plot have high amplitude. Most of the high-frequency components, however, have relatively smaller amplitudes. Fitting a straight line on the lowest 10-percentile of frequencies yields $-(1 - \hat{\beta}) = -0.714$, giving $\hat{H} = 0.857$.
- From the R/S plots, the slope of the linear least squares straight line is 0.811 , which is the estimated value of \hat{H} from this figure. An interesting observation is that the points corresponding to fixed values on the x-axis shrink near $\log_{10}(n) = 2$, and then again around $\log_{10}(n) = 3$. \hat{H} -estimates based on Whittle's method also show a change in trend at these points.

Our experiments on all the packet-level traces in Table 2 indicate that \hat{H} is greater than $1/2$. The plot at the top of Figure 4 gives a graphical summary of \hat{H} values (estimated via R/S analysis) for the different datasets in the form of a box plot. Each box displays the three quartiles (the 25, 50 and 75 percentiles), the minimum, and the maximum. The lower and upper lines of the box are at the first and third quartiles respectively, and the whiskers extend to extreme values of the data for the box. The last column of Table 2 lists the apparent \hat{H} of each trace. The plot toward the bottom of Figure 4 summarizes \hat{H} estimates for the trace data according to the estimation method used. (The Whittle estimates are for $m = 20$.)

All traffic datasets show some level of long-range dependence. Mode A traces show \hat{H} in the range 0.78 to 0.87 with median 0.85 . Mode B traces show \hat{H} in the range 0.79 to 0.89 with a median of 0.84 . Mode D traces show \hat{H} in the range 0.69 to 0.84 with a median estimate of 0.74 . We observe that Mode D traces which consist of all packets show lower \hat{H} estimates than Modes A and B traces that are filtered by TCP protocol. A possible reason is that TCP applications have a larger variance in burst sizes than other application packets (primarily UDP and routing update packets). Observations in the data clearly show that the filtered Mode A and B traces are dominated by the more voluminous TCP protocols in those traces (such as `ftp-data` in Mode B or `smtp` in Mode A).

The above studies were performed using the heuristic tools of R/S analysis, V-T plots, and Periodogram plots. These were followed by rigorous estimation of H using Whittle's approximate MLE, and estimating σ^2_H through a corresponding Central Limit Theorem due to Dahlhaus (see [10, 23]). For each aggregation level m , the aggregate time-series $\{X^{(m)}\}$ yielded point estimates for $\hat{H}^{(m)}$, $\hat{\sigma}^2_{H^{(m)}}$ and the 95% confidence intervals $(\hat{H}^{(m)} \pm 1.96\hat{\sigma}_{H^{(m)}})$.

$\hat{H}^{(m)}$ and its 95%-confidence intervals were then plotted as a function of m for the packet level traffic (Modes A, B and D). The plots for four of the datasets are shown in Figure 5. We observe the following:

- For all datasets, $\hat{H}^{(m)} > 1/2$ for all m .
- For all datasets, $\hat{H}^{(m)}$ and its confidence intervals increase with m for small m (up to about $m = 10$, but the exact value depends on the dataset).
- For most datasets, $\hat{H}^{(m)}$ and its confidence intervals do not stabilize to a constant value; instead they show increasing as well as decreasing patterns, but remain significantly higher than 0.5 at the 95% confidence interval. This indicates that NSFNET packet level traffic is long-range dependent. However, their correlation structures are not necessarily self-similar at different time-scales.

- For several datasets (e.g., trace #3 of Modes A and B in Figure 5), $\hat{H}^{(m)}$ and its confidence intervals appear to stabilize around a fixed value for m larger than some m^* . In these figures, the different point estimates of $\hat{H}^{(m)}$ are within 95%-confidence intervals of each other, so they may be regarded as statistically indistinguishable, and hence asymptotically self-similar.
- For several datasets, $\hat{H}^{(m)}$ appear to approach 1 for large m . This implies a potential non-stationarity in the datasets when they are estimated for larger aggregation levels. For small m , it is possible to split a trace into multiple datasets and still have a large enough sample size per dataset. However, for large m , there is a trade-off between sample size and potential non-stationarity in the data.

The cause for the initial increase implies one or both of the following: (i) it is a reflection of the true nature of dependency in the dataset, and/or (ii) it is caused by the fact that Whittle’s method is more accurate for larger m . The conclusion from the Whittle plots is that most of the NSFNET traffic traces that we analyzed are long-range dependent, and a few pass the test for second-order self-similarity. Finally, for longer time periods, the data sets are most likely non-stationary.

5 Factors Contributing to Long Range Dependence in NSFNET TCP Traffic

5.1 Objective

Since long-range dependence appears to be present in several different networks, and has a serious impact on queueing performance [1, 14, 16], an important question is what factors are contributing to it. In what follows, we present preliminary results on this issue based on application level characteristics.

The arrival process at a large multiplexor such as an NSFNET core switch is a complex interaction of multiple factors. For TCP applications, this includes:

- the distribution and dependencies in the conversation arrival process,
- the joint distribution of the number of bytes (and packets) transmitted per conversation and the rate at which they are transmitted,
- the end-to-end flow control per conversation, and
- heterogeneous propagation delays and transmission speeds associated with different conversations.

The study in this section is based on statistical inference from a set of designed experiments. Due to tractability reasons, we study a reduced system where (a) only a subset of traffic belonging to TCP conversations are studied, and (b) conversation inter-arrival correlations are not considered. Extensions in both these directions are currently being investigated.

An important parallel result in [35, 36] shows that the distributions of ON and OFF periods of different hosts over the Bellcore Ethernet are heavy tailed and have large variances (the “infinite variance syndrome”). The study in [30] shows that the `tcplib` package [13] generates similar VT plots as measured traffic.

5.2 Designed Experiments

5.2.1 Introduction

In a designed experiment, the value of a set of **response variables** is studied as a function of the experimental parameters (**factors**), each set to some particular **level**. In a **full factorial design** all possible combinations of factor settings are investigated for each complete replication of the experiment. The variation in the response is then studied using a set of statistical analysis methods to understand how the response depends on the underlying phenomena that the factors represent (see [6, 19, 22]).

We developed a set of experiments to understand the underlying phenomena in TCP traffic that cause long-range dependence. The experiments were organized along a full factorial design using four factors with three levels per factor. Each experiment was repeated ten times using different random seeds for a total of $3^4 \cdot 10 = 810$ trials. The objective of this design was to determine the sensitivity of H on different application-level factors.

5.2.2 Potential causes

Intuition, in conjunction with the following result due to Cox and Isham (see [9]), guided the choice of factors and their levels.

Theorem 1 (*Cox and Isham’s Immigration/Death Model* [9]:) *Let $\{Y_t\}$ denote the arrival process of new immigrants into a population. Let the life-times of individual arrivals be denoted as $\{Z_n\}$, ($n = 0, 1, 2, \dots$), where Z_n is the life-time of the n th immigrant. Let $\{X_t\}$ denote the number of immigrants in the population at time t .*

If $\{Y_t\}$ is a stationary Poisson process, then for $\{X_t\}$ to be second-order self-similar (special case of long-range dependence),

$$\Pr\{Z_n = k\} \propto k^{-\beta-2}, \quad 0 < \beta < 1, \quad n = 0, 1, 2, \dots \quad (14)$$

This implies among other things that $\text{var}(Z_n) \rightarrow \infty$, (i.e., Z_n is a “heavy-tailed” distribution). $\{X_t\}$ is asymptotic second-order self-similar if (14) holds for large k .

Figure 6 illustrates two relevant features found in the trace data. The top graph shows a quantile-quantile (Q-Q) plot of conversation interarrival times against an exponential distribution. The bottom graph is a log-log plot of the conversation duration density function. Note the straight lines in both the Q-Q plot and the tail of the log-log density plot. The straight line in the Q-Q plot implies that the aggregate conversation interarrivals are close to exponential. The approximate straight line in the log-log density plot for large k (except at the higher tails where there are not enough sample points) implies that the conversation life-times are of the form in (14).

The assumptions in Theorem 1, of course, depart from data in several ways:

- Theorem 1 assumes that at time t , the contribution by each *live* immigrant to X_t is one. This does not apply to the packet arrival process because of widely varying transmission rates which lead to different contributions to X_t by different conversations;
- The conversation interarrivals, while exponential, are not completely uncorrelated.

The implications of variable transmission rates on H are explored in this study.

<i>Factor</i>	<i>Levels</i>	<i>Description</i>
Conv. Arrivals	Empirical	variates based on empirical distribution derived from traces
	Exponential	variates generated using exponential distribution
	Fixed	conversations arrive at regular fixed intervals
Packets	Empirical	variates based on empirical distribution derived from traces
	Pareto	variates generated using Pareto distribution
	LogNormal	variates generated using log-normal distribution
Rate	Empirical Resid.	variates based on empirical distribution derived from traces
	Normal Resid.	variates generated using normal distribution
	Fixed	all conversation have same rate regardless of packets
Packet Arrivals	Empirical	variates based on empirical distribution derived from traces
	Exponential	variates generated using exponential distribution
	Fixed	packets arrive at regular fixed intervals

Table 3: Factors and levels used in the experiments.

Figure 7 shows a 2-D histogram of the number of packets per TCP conversation and their durations. Observe that even after a log-transformation, the number of bytes distribution appears to be heavy-tailed. Experimental results that follow indicate that this joint distribution plays a vital role in long-range dependence in packet-level TCP traffic. For instance, when transmission rate and number-of-packets were independent variables, experiments yield a low value of H . On the other hand, if they are fully correlated (e.g., if they are related by a deterministic equation), experiments yield a higher value of H than what is observed in traces. The analysis of variance for the designed experiment leads to the conclusion that correctly modeling this relationship is vital to inducing realistic long-range dependence patterns in simulated TCP traffic.

5.2.3 Selection of factors

The factors chosen and the levels used per factor are summarized in Table 3. The technical rationale for these are as follows.

- (a) Conversations arrive from somewhere outside of the network (factor = **conv**). Each level of **conv** specifies a distribution for this factor, so that the relative impact of different distributions can be studied.
- (b) Each conversation transmits a certain number of packets (factor = **Pkts**). Different distributions are used as levels for this factor.
- (c) A conversation takes some X units of time to complete. Traffic data shows that X is correlated with number of packets. Over this duration, X , a conversation’s packet inter-arrivals are distributed according to end-to-end flow-control, intermediate queueing delays, user-interactions, etc. Since these are complex interactions to model, and our objective in this paper is primarily to determine its relative impact on H , a simple model that uses empirical and selected distributions is used as levels for the packet inter-arrival process. For technical convenience in generating simulation experiments corresponding to the factors chosen, the following method is adopted:
 - (i) each conversation duration is converted to a mean rate of transmission for this conversation, where mean rate is defined as number-of-packets/conversation-duration (factor = **Rate**), and

(ii) packets are spaced according to a packet-interarrival distribution with the computed rate as a parameter (factor = `pia`).

(d) When a conversation has transmitted all of its packets, it leaves the network.

Based on Theorem 1, the distribution of these factors appear to be important for presence (or absence) of long-range dependence in packet-level traffic. Therefore, the relative impact of different factors and different levels per factor are studied.

Traffic data showed that the correlation between $\log(\#Packets)$ and $\log(Rate)$, was 0.302. The impact of this dependence was tested in the design through a nested factors model (explained below) and the observation that $\log(Rate)$ was approximately a linear function of $\log(\#Packets)$ plus a Normally distributed residual:

$$Rate \sim a \cdot (\#Packets) + b + \epsilon, \quad (15)$$

where ϵ is a Normally distributed random variable with mean near 0.0 and variance 0.61. The Generalized Linear Model fit algorithm (`glm`) in the S+ package [6] gave $a = 0.39$, and $b = -0.38$. Figure 8 shows a sample of the transformed data along with the fit. The graph on the left of Figure 8 is the actual trace data; the graph on the right illustrates rates generated using Equation (15). While the fit and the scatter-plots look similar, the residuals were found to be somewhat correlated, so this model is only an approximation. We use this in this study because it allows for a separation of factors and facilitates experimenting with separate models for transmission-rate and number-of-packets. The alternative is to use a jointly-distributed random variate consisting of number-of-packets and transmission-rate as a single factor. This method will generate packets more accurately, and as such is our method of choice for our simulation model. The method adopted in this paper, however, is better suited for experimentation with and comparing impact of different factors on H .

5.2.4 Selection of levels

Any experimental design evolves with time as one obtains more experience with the sensitivities of the response variables. The design presented here is a compact version that highlights the sensitivity of different application factors on H . For each of the four factors, we selected three levels at which to evaluate the factor's effect on H . The levels were selected to include (i) an empirical description of the factor, (ii) a close analytical model based on off-line analysis of that factor in traffic data, and (iii) a (simple) level that helps determine if this factor makes any significant impact on H by providing contrast. The close analytical model as an alternative level selection was motivated by the need to build an accurate analytic workload (as part of a related project). The full set of experiments were replicated ten times to quantify the variation due to experimental errors. Table 4 gives a summary of level settings.

Observe that a given instance of rate is dependent on the number of packets per conversation, and the packet interarrival time in turn depends on the rate. While this is not acceptable in simple linear least-squares regression analysis, it is admissible in designed experiments using a nested factors model [6, Chapter 5]. When the levels of some dependent factor are only meaningful given the levels of some primary factor, the dependent factor is said to be *nested* within the primary factor. To analyze the variation in response due to nested factors, the primary factor is first fitted, and then the effects of the dependent factor are fitted in each level of the primary factor.

The main effects for the dependent factor are not meaningful in the case of nested factors. Thus, `rate` is not part of the factorial design formula, but `rate in pkt` is. Our analysis uses a nested factors model to account for the interrelation between the number of packets, transmission rate, and

<i>Factor</i>	<i>Parameter</i>	<i>Setting</i>
Conv. Arrivals	Fixed IA	0.0571 sec
	Exp. Mean IA	0.0831 sec
Packets	LogNorm Mean	0.4838
	LogNorm Shape	0.4004
	Pareto Shape	0.9020
	Pareto Scale	3.9425
Rate	Fixed Rate	0.7572 pkt/sec
	Fit Slope	0.3929
	Fit Intcp.	-0.3854
	Resid. Mean	$-6.613e^{-13}$
	Resid. Var	0.6129
Packet Arrivals	Fixed IA	1.6918

Table 4: Parameter setting used in the experiments

packet interarrival time. (For convenience in notation, `rate in pkt` is often referred to as `rate`, but its meaning is to be interpreted in the nesting sense.)

The formula we give to S+ [33] is specified as follows (the notation for nesting is changed slightly to make it more readable).

$$H \sim \text{conv} \star \text{pkt} \star (\text{rate in pkt}) \star (\text{piain}(\text{rate in pkt})). \quad (16)$$

In this notation, the \star is to be interpreted as follows: $A \star B$ is shorthand for a model with terms for the intercept, A , B and the interaction of A and B . The term `(rate in pkt)` says that `pkt` is the main effect and `rate` is nested within `pkt`. The main effects for the dependent factor are not meaningful in the case of nested factors. Thus, `rate` is not part of the factorial design formula, but `rate in pkt` is. Our analysis uses a nested factors model to account for the interrelation between the number of packets, transmission rate, and packet interarrival time in accordance with (16). (For convenience in notation, `rate in pkt` is often referred to as `rate`, but its meaning is to be interpreted in the nesting sense.)

Since several of the factors are nested, their interaction is not directly used as part of the model but is instead attributed to the parent factor. The first column in Tables 5 through 8 enumerate the predictor variables specified by the model.

5.2.5 Simulation model

We constructed a BONEs [3] simulation model to perform the experiments. It preserves the essential properties that we want to test: namely, each conversation has its own number of packets, a transmission rate determined by Equation (15) and packet interarrivals that are a function of the transmission rate and the level selected. (For example, when the level is deterministic, packet interarrivals are constant; when the level is exponential, it serves as the mean of the exponential, etc.) *The simulation is designed to maintain independent conversations that may overlap in time.*

The BONEs diagram of the simulation system is shown in Figure 9. We observe that this is similar to a flow chart with an added feature that BONEs implements it automatically and precisely according to the graphical specification shown. New conversations are created with the selected interarrival

distribution. Each new conversation is then assigned a random number of packets that it will transmit during its lifetime. A transmission rate is calculated via the model based on the `rate` level and the number of packets assigned (left side of Figure 9). The conversation then enters the central loop (right side of Figure 9) where it transmits a packet, decrements the number of packets left, and randomly picks the time until it transmits again (based on its rate and packet interarrival distribution). Simulation runs typically took about 5-10 minutes for the empirical level settings on a Sun SPARCstation 10/51 to generate a stream of approximately 1.2 million packets.

5.2.6 Experimental results

On the accuracy of simulation experiments

Figure 10 shows a section of \hat{H} estimates from (i) simulations (marked **Emp**, **LgN** and **Par**) and (ii) trace data filtered by TCP protocol (marked **Trc**). The first boxplot (labeled **Emp**) represents all \hat{H} estimates with number of packets set to the empirical distribution, and other factors taking all possible level-combinations. Similarly, the second and the third boxplots show \hat{H} estimates with number of packets set to Log-Normal and Pareto distribution, respectively, with other factors taking all possible level combinations. We observe that:

- The model with empirical distribution (**Emp**) for number of packets is fairly close to the TCP traces (**Trc**);
- Pareto is close to **Trc** but Log-Normal is a distant third in its ability to match real trace.

Comparison of different levels for each factor

Figure 11 shows boxplots of \hat{H} estimates for different levels of each factor. We observe that:

- All levels of conversation inter-arrival distribution appear to yield the same distribution of \hat{H} estimates, implying that \hat{H} is not particularly sensitive to the choice of level for this factor (among the three we are experimenting with);
- A **Fixed Rate** model (representing homogeneous transmission speeds) yields a smaller \hat{H} than a model which accounts for heterogeneous rates through empirically distributed residuals or Normally distributed residuals.
- For number of packets, Pareto yields estimates close to Empirical, but remains smaller; Log-Normal yields \hat{H} estimates that are considerably smaller.
- Boxplots corresponding to \hat{H} estimates for different packet inter-arrival models show a different spread for each, but similar median values. This figure by itself does not lead to any specific conclusion, so further investigation is necessary to determine the role of this factor (see Section 5.2.7).

Relative impact of different factors on \hat{H}

Figure 12 compares the mean \hat{H} estimates for different factors and different levels per factor. The horizontal line represents the mean \hat{H} over all experiments.

For each factor, the mean \hat{H} estimates for its different level settings, are indicated on a vertical line. For each level of a given factor, this represents the mean over all experiments with this factor fixed, and other factors taking all possible level combinations. We observe the following:

- **Rate** has the largest spread, with Normal residuals slightly greater than Empirical residuals, and both substantially greater than Fixed, i.e., heterogeneity of speeds is a significant factor — and adds an extra dimension to \hat{H} over the Immigration/Death model in Theorem 1.
- **Pkt** has the second largest spread. Inside it, Empirical is larger than Pareto and both are considerably larger than Log-Normal.
- For packet inter-arrival distribution (**pia**), Fixed is larger than Exponential, and both are considerably larger than Empirical.
- For the conversation factor (**conv**), there is not much spread around the mean of all \hat{H} estimates (the horizontal line).

5.2.7 Analysis of variance

Tables 5 through 8 present the analysis of variance results. The results were obtained using the S+ [6, 33] Analysis of Variance algorithm (**aov**) to fit the model described by Equation 16. The column labels and their meanings are as follows:

- *Factor* : Represents either a single factor or the interaction of several factors as indicated by a “:” between the two factors;
- *Df* : The degrees of freedom for the factor;
- *Sum Sq*: The sum of squares term;
- *% Var*: The percent of total variation explained by this factor. This column gives a quantitative estimate of the relative sensitivities of different factors on the response variable.
- *Pr(F)*: The probability with which the mean-square regression (MSR) is less than mean square error (MSE) assuming that MSR/MSE has an F distribution (a consequence of Normality assumptions on error residuals). $1 - Pr(F)$ is the probability with which MSR is greater than MSE, and is the confidence level of the term corresponding to this factor.

Remarks:

- (a) The error residuals from the above nested factors model of the designed experiment yielded Normally distributed residuals with no apparent trend.
- (b) The response variable \hat{H} was estimated through R/S analysis, Whittle’s method, PG plot and VT plot. For all of these estimation methods, the variation in \hat{H} explained by the regression was more than 90%. The exact numbers are given in Tables 5-8.

The percent variation explained by each factor is illustrated graphically in Figure 13. In each plot, the bars correspond to the largest factors contributing to the variation in response are plotted from left to right. The last bar groups the remaining factors. The line over the bars indicates the cumulative amount of variation explained by the factors up to that point. We observe the following from these figures.

- If \hat{H} is computed using Whittle’s method or the R/S analysis, number-of-packets and rate (nested inside packets) explain most of the variation in \hat{H} . Together, they account for 77% of the variation using Whittle’s method and 82% percent of the variation using R/S analysis. (Because of nesting of factors, all interactions between the two are ascribed to the parent factor which in this case is number-of-packets.)
- Packet interarrivals are a distant third explaining approximately 15% percent of the variation.
- Conversation arrival distribution plays an insignificant role. (This does not, however, mean that conversation inter-arrival correlations do not play a role.)
- Interactions between factors (after the nesting) are negligible.
- The PG and VT estimators give different conclusions about which factors explain more of the variation in \hat{H} . More work will be needed to determine precisely why this is so. Given the robust statistical properties of the Whittle estimator [23], and the heuristic nature of the PG and VT estimators, the analysis of variance results computed for the response estimated by Whittle’s method are more believable.

6 Conclusion

An examination of packet traffic from geographically dispersed locations on the NSFNET T3 backbone indicates that:

- packet-level traffic over NSFNET core switches exhibits long-range dependence. A subset of these show asymptotic second-order self-similarity; and
- Hurst coefficient estimates for TCP-traffic, in general, appear to be larger than those for the overall aggregate traffic.

Further, a set of designed simulation experiments indicate that the primary reason for long-range dependence in TCP traffic is the joint distribution of the number of packets per conversation and conversation transmission rates. Since these are application-level characteristics which are largely network-independent, long-range dependence is expected to exist in most highly multiplexed data networks carrying similar TCP applications.

Intuitively, one might expect packet-level traffic to “smooth out” at larger time scales, especially when the number of conversations is large. However, empirical evidence suggests otherwise. The primary reason for this is the heavy-tailed nature of the afore-mentioned distributions and Theorem 1 [9]. Recall that the Central Limit Theorem holds only for i.i.d. variables with *finite* variance. Mandelbrot and Van Ness have shown that $H \rightarrow 1/2$ for short-range dependent processes as well. One ramification of these findings is that Markovian traffic models are inadequate for describing and analyzing highly-multiplexed data networks such as the NSFNET backbone because they inject exponentially-decaying auto-correlations that are characteristic of short-range dependent processes.

If the system under test needs to be studied for extended periods of time with some degree of accuracy, simulation models with empirically distributed parameters or trace-driven traffic models are more appropriate. For instance, long-range dependent traffic patterns for TCP traffic can be generated from within an event-driven simulation if conversation parameters (number of packets, transmission rate, etc.) are based on accurate empirical distributions. Currently, we are building a suite of simulation

<i>Factors for R/S</i>	<i>Df</i>	<i>Sum Sq</i>	<i>% Var</i>	<i>Pr(F)</i>
conv	2	0.008	0.068	0.00002
pkt	2	3.424	30.435	0.00000
rate in pkt	6	5.841	51.926	0.00000
conv:pkt	4	0.008	0.073	0.00011
pia in (rate in pkt)	18	1.698	15.097	0.00000
conv:(rate in pkt)	12	0.006	0.051	0.17400
conv:(pia in (rate in pkt))	36	0.012	0.104	0.57778
Residuals	729	0.253	2.246	NA

Table 5: Analysis of variance and related statistics for \hat{H} estimates from the $3^k r$ designed experiment. \hat{H} computed through R/S analysis.

<i>Factors for Whittle</i>	<i>Df</i>	<i>Sum Sq</i>	<i>% Var</i>	<i>Pr(F)</i>
conv	2	0.001	0.011	0.25344
pkt	2	4.356	35.564	0.00000
rate in pkt	6	5.832	47.611	0.00000
conv:pkt	4	0.024	0.195	0.00000
pia in (rate in pkt)	18	0.027	0.222	0.00000
conv:(rate in pkt)	12	1.621	13.238	0.00000
conv:(pia in (rate in pkt))	36	0.020	0.160	0.34293
Residuals	729	0.367	2.998	NA

Table 6: Analysis of variance and related statistics when \hat{H} is computed using Whittle's method.

<i>Factors for Periodogram</i>	<i>Df</i>	<i>Sum Sq</i>	<i>% Var</i>	<i>Pr(F)</i>
conv	2	0.029	0.135	0.00014
pkt	2	2.322	10.858	0.00000
rate in pkt	6	14.312	66.911	0.00000
conv:pkt	4	0.006	0.026	0.47530
pia in (rate in pkt)	18	3.505	16.388	0.00000
conv:(rate in pkt)	12	0.027	0.128	0.15210
conv:(pia in (rate in pkt))	36	0.016	0.077	0.99999
Residuals	729	1.171	5.476	NA

Table 7: Analysis of variance and related statistics when \hat{H} is computed using the Periodogram method.

<i>Factors for Variance-Time</i>	<i>Df</i>	<i>Sum Sq</i>	<i>% Var</i>	<i>Pr(F)</i>
conv	2	0.001	0.016	0.32404
pkt	2	6.297	76.287	0.00000
rate in pkt	6	0.644	7.807	0.00000
conv:pkt	4	0.011	0.135	0.00077
pia in (rate in pkt)	18	0.841	10.186	0.00000
conv:(rate in pkt)	12	0.028	0.344	0.00000
conv:(pia in (rate in pkt))	36	0.009	0.106	0.99901
Residuals	729	0.422	5.118	NA

Table 8: Analysis of variance and related statistics when \hat{H} is computed using Variance-Time method.

workloads tailored to specific performance studies (adaptive routing, flow control, switch scheduling, etc.)

With rapid development in high-speed wide-area communications and increasing concerns for quality of service guarantees, it is important to understand the statistical properties of traffic that these networks will carry. While higher speed networks are likely to usher in a new generation of applications, at least for the time being, current traffic models are the best source of information for network resource management. The results of this paper give a clearer picture of the statistical properties of today's medium-speed wide-area network traffic. A key conclusion from the analysis of variance study is that the long-range dependence characteristics are largely application dependent.

The principal findings are summarized in Table 1.

Acknowledgments

We wish to thank Sabyasachi Basu, Sugih Jamin, John Limb, Jean McManus, Biswanath Mukherjee and Walter Willinger for valuable feedback on earlier drafts of this paper. We would also like to thank Walter Willinger for providing us with his S functions for computing the Rescaled-Adjusted Range statistic and Jan Beran's S functions for estimating \hat{H} using Whittle's method.

References

- [1] A. Adas and A. Mukherjee, "On resource management and QoS guarantees for long-range dependent traffic," *Proc., IEEE Infocom '95*, Boston, 1995 (to appear).
- [2] D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic theory of a data-handling system with multiple sources," *Bell Sys. Tech. Journal*, pp. 1871-1894, Oct. 1982.
- [3] "The BONEs simulation package," *COMDISCO Systems Inc.*, 1992
- [4] J. Beran, "Statistical methods for data with long-range dependence," *Statistical Science*, vol. 7, pp. 404-427, 1992.
- [5] J. D. Carlin, P. Dempster, and A. B. Jones, "On methods and models for Bayesian time series analysis," *Journal of Econometrics*, vol. 30, pp. 67-90, 1985
- [6] J. M. Chambers and T. J. Hastie, *Statistical Models in S*, Wadsworth & Brooks/Cole, 1992
- [7] K. Claffy, H. W. Braun, and G.C. Polyzos, "Tracking long-term growth of the NSFNET backbone," *Proc., IEEE INFOCOM '93*, San Francisco, CA, August 1993.
- [8] K. Claffy, H. W. Braun, and G.C. Polyzos. "Application of sampling methodologies to network traffic characterization," *Proc., ACM SIGCOMM '93*, pp. 194-203, San Francisco, CA, September 1993.
- [9] D. R. Cox, "Long range dependence: a review," *Statistics: An Appraisal*, H. A. David and H. T. David (Eds.), The Iowa State University Press, Ames, Iowa, pp. 55-74, 1984
- [10] R. Dahlhaus, "Efficient parameter estimation for self-similar processes," *Annals of Statistics*, vol. 17, pp. 1749-1766, 1989.
- [11] P. B. Danzig, S. Jamin, R. Caceres, D. J. Mitzel, and D. Estrin, "An empirical workload model for driving wide-area TCP/IP network simulations," *Journal of Internetworking: Practice and Experience*, vol. 3, no. 1, March 1992.
- [12] P. B. Danzig, R. S. Hall, and M. F. Schwartz, "A case for caching file objects inside internetworks," *Proc., ACM SIGCOMM '93*, pp. 239-248, San Francisco, CA, September 1993.
- [13] P. B. Danzig and S. Jamin, "tcplib: A library of TCP/IP traffic characteristics," USC Networking and Distributed Systems Laboratory TR CS-SYS-91-01, October 1991. (Available as [jerico.usc.edu:/pub/jamin/tcplib/](http://jerico.usc.edu/pub/jamin/tcplib/)*)
- [14] A. Erramilli, O. Narayan and W. Willinger, "Experimental queuing analysis with long-range dependent traffic," submitted for publication, 1994.
- [15] A. Erramilli and L. J. Forys, "Oscillations and chaos in a flow model of a switching system," *IEEE Journal on Selected Areas in Communications*, pp. 171-178, February 1991.
- [16] H. J. Fowler and W. E. Leland, "Local area network traffic characteristics, with implications for broadband network congestion management," *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 1139-1149, 1991.

- [17] M. W. Garrett, *Contributions toward real-time services on packet switched networks*, Ph.D. Dissertation, Columbia University, 1993.
- [18] M. W. Garrett and W. Willinger, "Analysis, modeling and generation of self-similar VBR video traffic," *Proc. ACM SIGCOMM '94*, pp. 269-281, September 1994.
- [19] W. W. Hines and D. C. Montgomery, *Probability and Statistics in Engineering and Management Science*, 3rd ed., John Wiley & Sons, 1990.
- [20] H. E. Hurst, "Long-term storage capacity of reservoirs," *Transactions of the American Society of Civil Engineers*, no. 116, pp. 770-779, 1951.
- [21] Network Information Services. InterNIC, data available on `nis.nsf.net:/nsfnet/statistics`.
- [22] R. Jain, *The Art of Computer System Performance Analysis*, John Wiley & Sons, 1991.
- [23] W. E. Leland, W. Willinger, M. S. Taqqu, and D. V. Wilson, "On the self-similar nature of Ethernet traffic," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1-15, February 1994.
- [24] W. E. Leland and D. W. Wilson, "High time-resolution measurement and analysis of LAN traffic: implications for LAN interconnection," *Proc., IEEE INFOCOM '91*, Bal Harbour, FL, pp. 1360-1366, April 1991.
- [25] M. Livny, B. Melamed, A. Tsiolis, "Impact of autocorrelation on queuing systems," *Management Science*, vol. 39, no. 3, pp. 322-339, March 1993.
- [26] B. B. Mandelbrot and J. W. Van Ness, "Fractional brownian motions, fractional noises and applications," *SIAM Review*, vol. 10, pp. 422-437, 1968.
- [27] B. B. Mandelbrot and J. R. Wallis, "Computer experiments with fractional gaussian noises," *Water Resources Research*, vol. 5, pp. 228-267, 1969.
- [28] B. B. Mandelbrot and J. R. Wallis, "Some long-run properties of geophysical records," *Water Resources Research*, vol. 5, pp. 321-40, 1969.
- [29] V. Paxson, "Empirically-derived analytical models of wide-area TCP connections: extended report," *IEEE/ACM Transactions on Networking*, vol. 2, no. 4, pp. 316-336, August 1994.
- [30] V. Paxson and S. Floyd, "Wide-area traffic: the failure of poisson Modeling," *Proc., ACM SIGCOMM '94*, pp. 257-268, September 1994.
- [31] V. Paxson, "Growth trends in wide-area TCP connections," (available as `ftp.ee.lbl.gov:/WAN-TCP-growth-trends.ps.Z`)
- [32] A. Schmidt and R. Campbell, "Internet protocol traffic analysis with applications for ATM switch design," *Computer Communications Review*, vol. 23, no. 2, pp. 39-52, April 1993.
- [33] The S+ package, Version 3.0, Statistical Sciences, Inc., September 1991.
- [34] V. Jacobson, C. Leres, and S. McCanne, `Tcpdump`, June 1989 (available via anonymous FTP from `ftp.ee.lbl.gov`).

- [35] W. Willinger, “Traffic Modeling for High-Speed Networks: Theory versus Practice”, in: *Stochastic Networks*, F.P. Kelly and R.J. Williams (Eds.), IMA Volume in Mathematics and its Applications, Springer-Verlag, 1994 (to appear).
- [36] W. Willinger, D.V. Wilson, W.E. Leland and M.S. Taqqu, “On Traffic Measurements that Defy Traffic Models (and vice versa): Self-Similar Traffic Modeling for High-Speed Networks”, *ConneXions*, Vol. 8, No. 11, pp. 14-24, 1994.
- [37] Y. Yajima, “On estimation of long-memory time series models,” *Australian Journal of Statistics*, vol. 27, pp. 303-320, 1985.

List of Figures

1	CNSS router architecture and trace points.	26
2	A section of traffic showing number of packets per time unit at four different time scales.	27
3	Variance-Time, Periodogram and R/S plots for the first Mode D Hartford trace showing potential self-similarity.	28
4	Summary of H values.	29
5	Estimates of \hat{H} and its 95%-confidence intervals using Whittle's method. Also shown are the \hat{H} estimates from V-T, PG and R/S analyses.	30
6	Two features found in the traffic which are relevant to Theorem 1. The top graph shows a quantile-quantile plot of conversation interarrival times against an exponential distribution. The bottom graph is a log-log plot of the conversation duration density function. Note the straight lines in both the quantile-quantile plot and the tail of the log-log density plot.	31
7	Two-dimensional histogram of log(bytes) versus log(duration)	32
8	Data and fit for rate vs. packets	33
9	BONeS diagram of simulation model. An intuitive way to understand the diagram is to follow the information flow indicated by the lines and arrows. Conversations are initiated and parameters set toward the left side of the diagram. Conversations then circulate in the loop toward the right side of the diagram emitting one packet per pass until they run out of packets.	34
10	Boxplots of \hat{H} for levels of the <code>pkt</code> factor. <code>Trc</code> describes values of \hat{H} estimated for the trace data.	35
11	Boxplots by isolated factor for Whittle estimates of \hat{H} . Each box contains all values where the factor was set to the indicated level.	36
12	A graphical summary of the experimental results. The horizontal line is the mean of all \hat{H} estimates. For each level of each factor, the average response is as indicated on the vertical line corresponding to that factor.	37
13	Percent variation in response explained by different factors. The first four bars correspond to the top four factors contributing to variation in response. A fifth bar groups all the remaining factors. The line over the bars indicates the cumulative amount of variation explained by the factors up to that point.	38

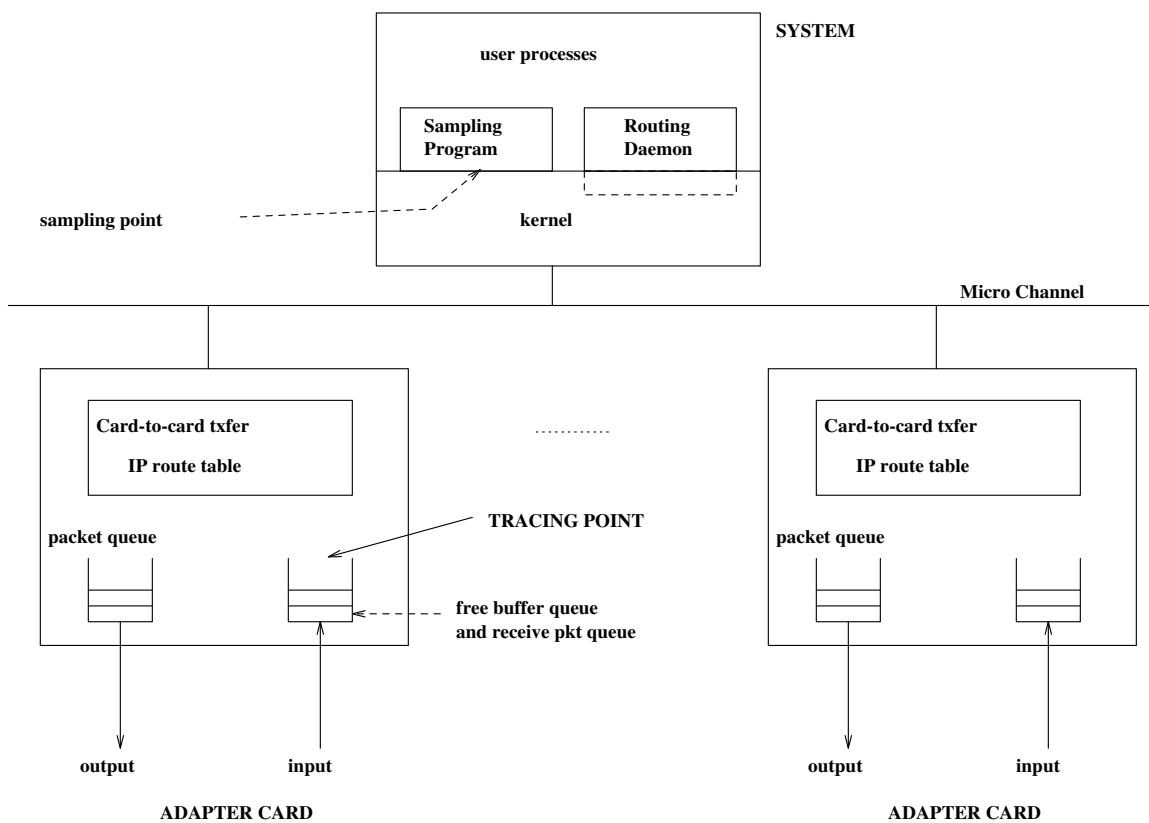


Figure 1: CNSS router architecture and trace points.

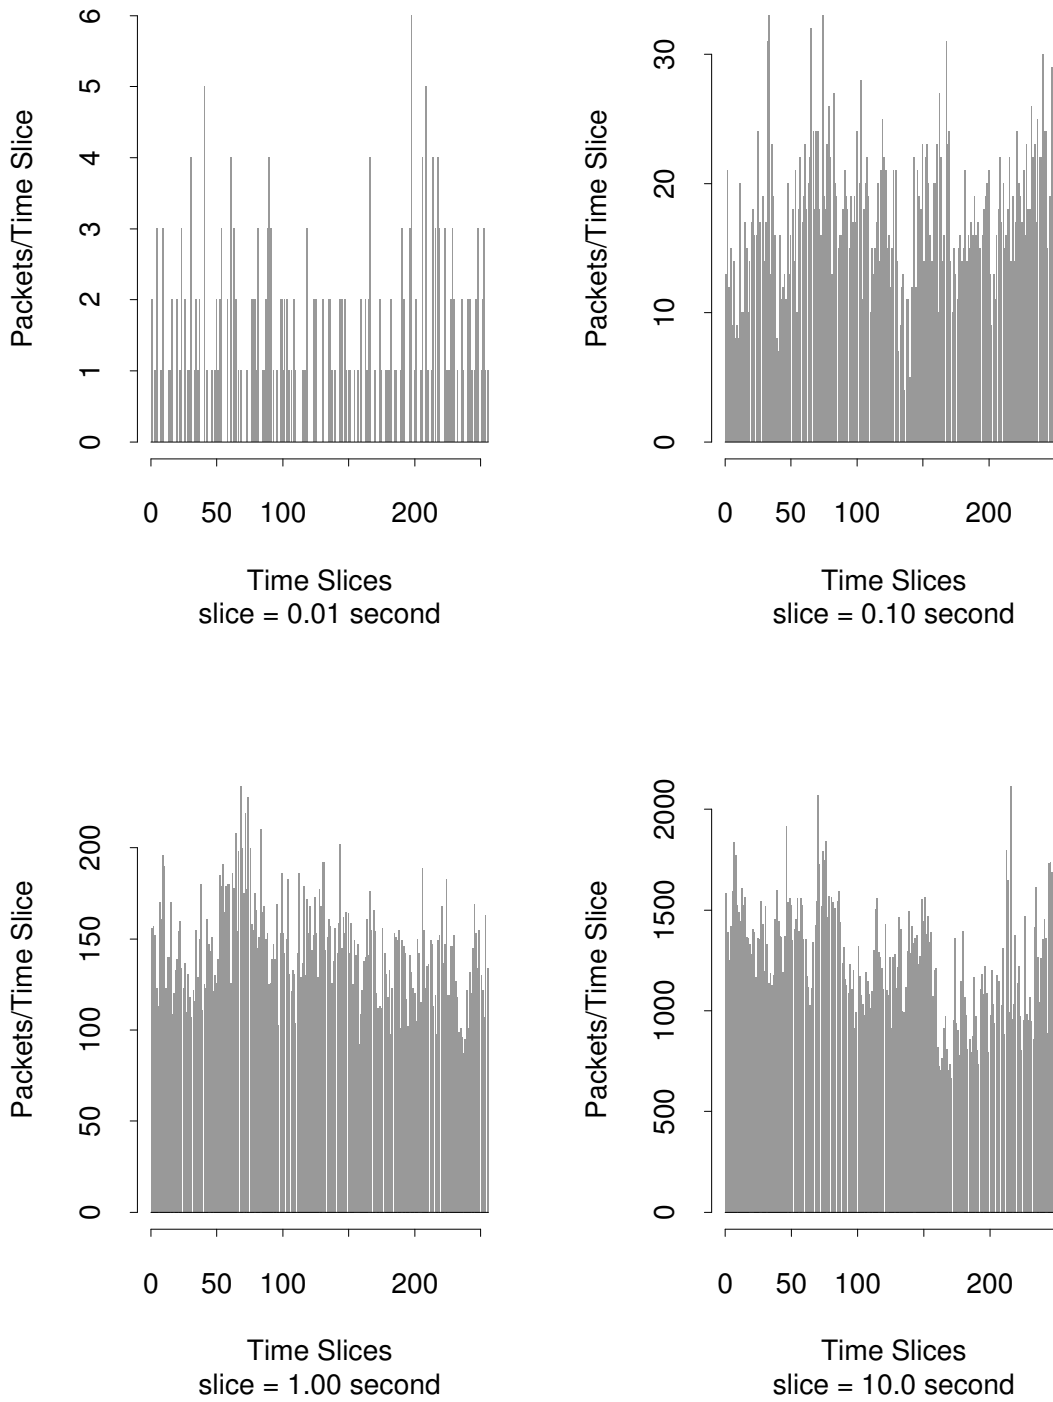


Figure 2: A section of traffic showing number of packets per time unit at four different time scales.

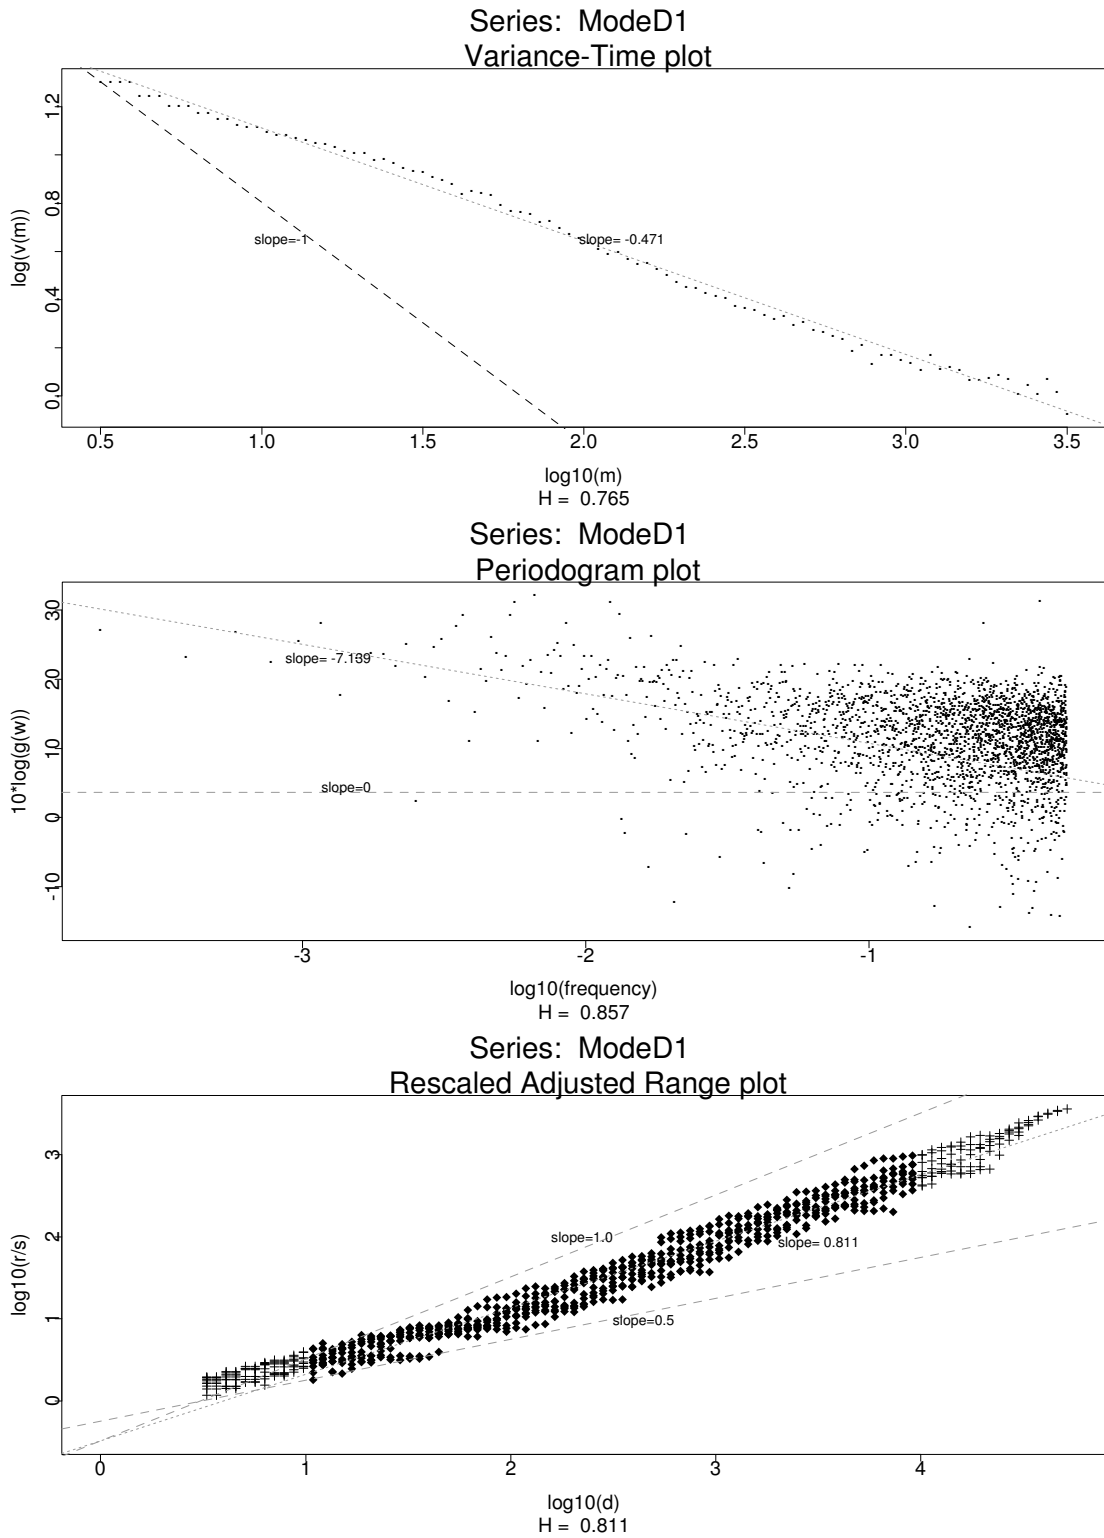



Figure 3: Variance-Time, Periodogram and R/S plots for the first Mode D Hartford trace showing potential self-similarity.

H Estimates by Trace Mode



H Estimates for Trace Data




Figure 4: Summary of H values.

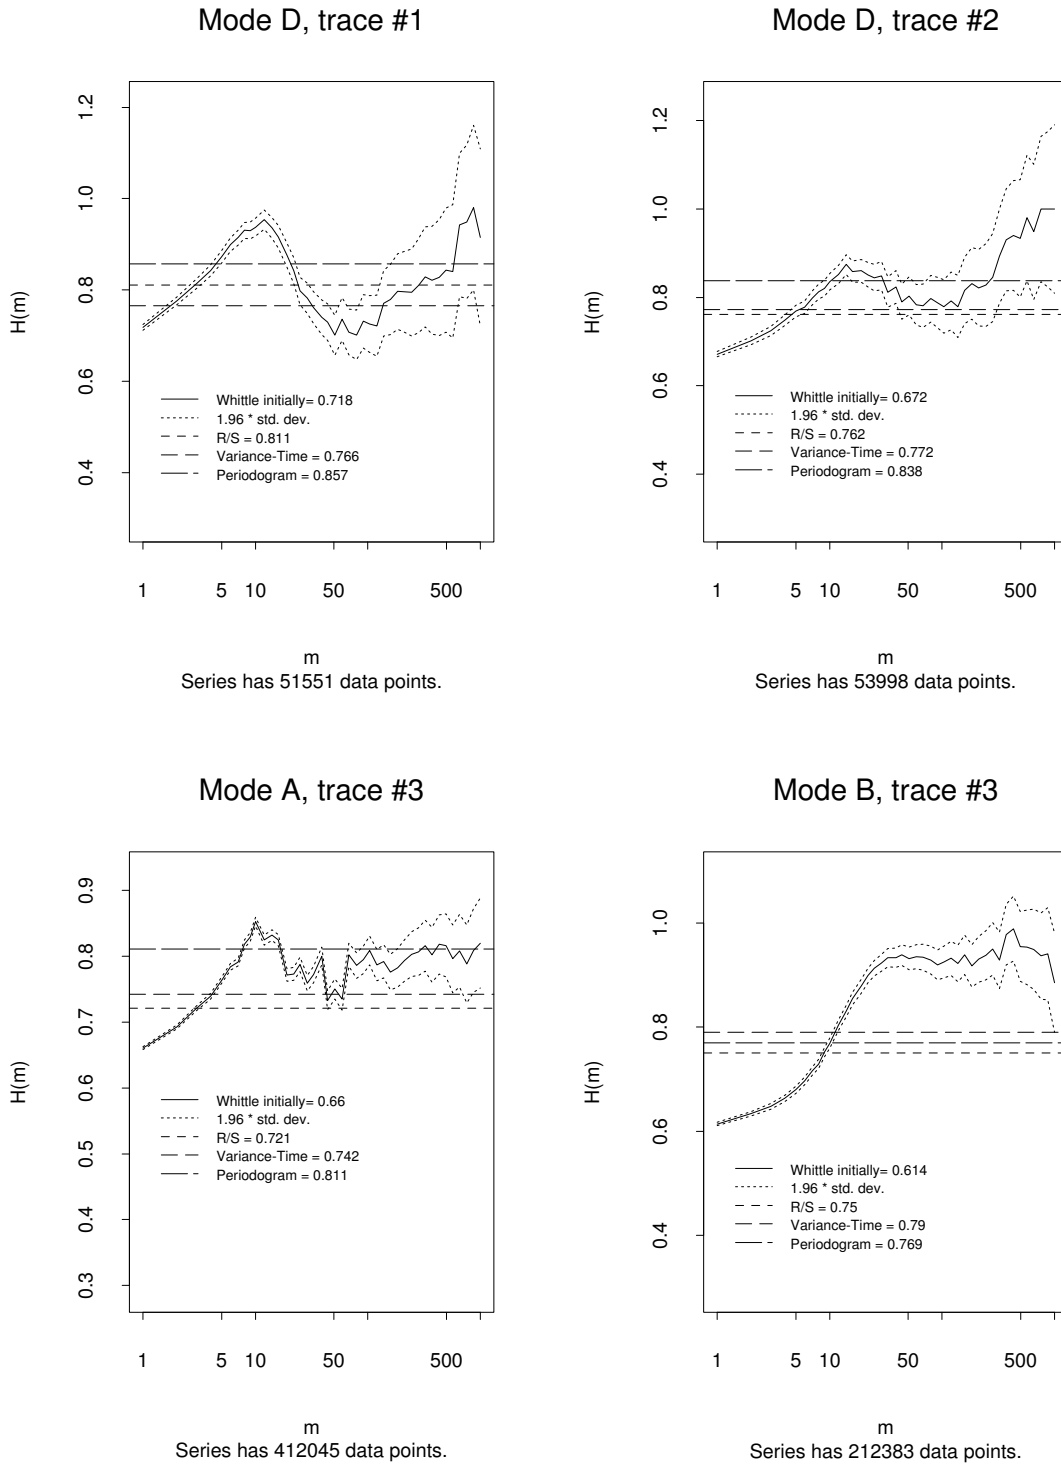



Figure 5: Estimates of \hat{H} and its 95%-confidence intervals using Whittle's method. Also shown are the \hat{H} estimates from V-T, PG and R/S analyses.

Conversation Interarrivals Q-Q Plot



exp (rate=11.41)

Log-Log Duration Distribution




Figure 6: Two features found in the traffic which are relevant to Theorem 1. The top graph shows a quantile-quantile plot of conversation interarrival times against an exponential distribution. The bottom graph is a log-log plot of the conversation duration density function. Note the straight lines in both the quantile-quantile plot and the tail of the log-log density plot.





Figure 7: Two-dimensional histogram of $\log(\text{bytes})$ versus $\log(\text{duration})$

Actual Rate vs. Packets with Fit



Modeled Rate vs. Packets with Fit




Figure 8: Data and fit for rate vs. packets

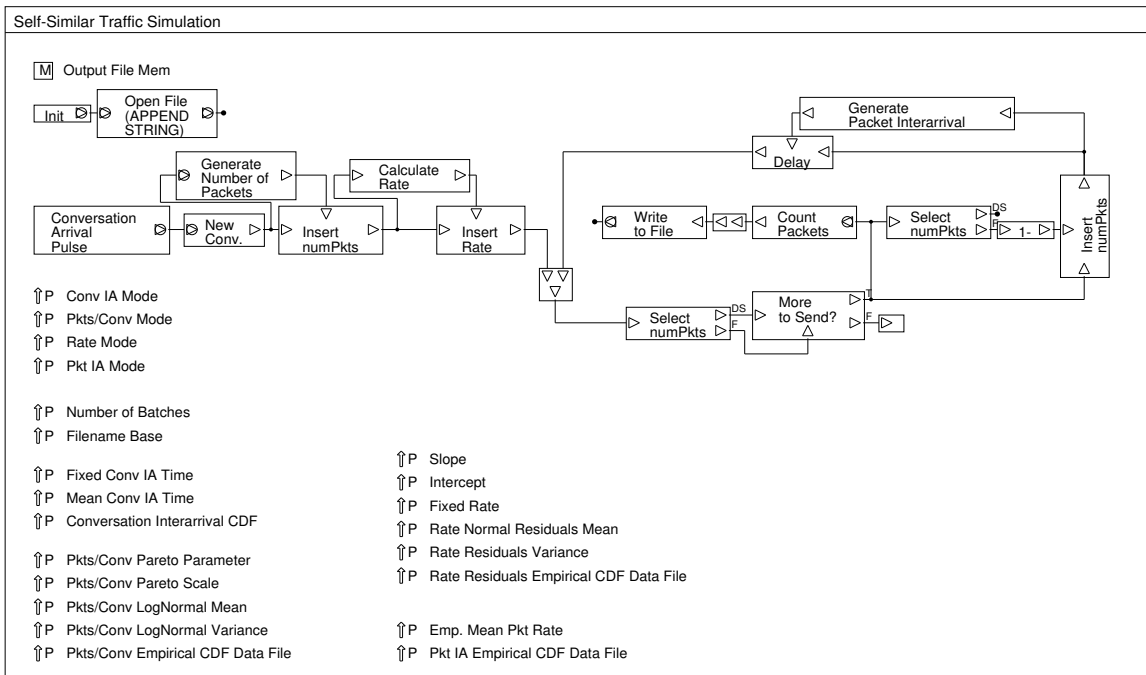


Figure 9: BONeS diagram of simulation model. An intuitive way to understand the diagram is to follow the information flow indicated by the lines and arrows. Conversations are initiated and parameters set toward the left side of the diagram. Conversations then circulate in the loop toward the right side of the diagram emitting one packet per pass until they run out of packets.




Figure 10: Boxplots of \hat{H} for levels of the `pkt` factor. `Trc` describes values of \hat{H} estimated for the trace data.

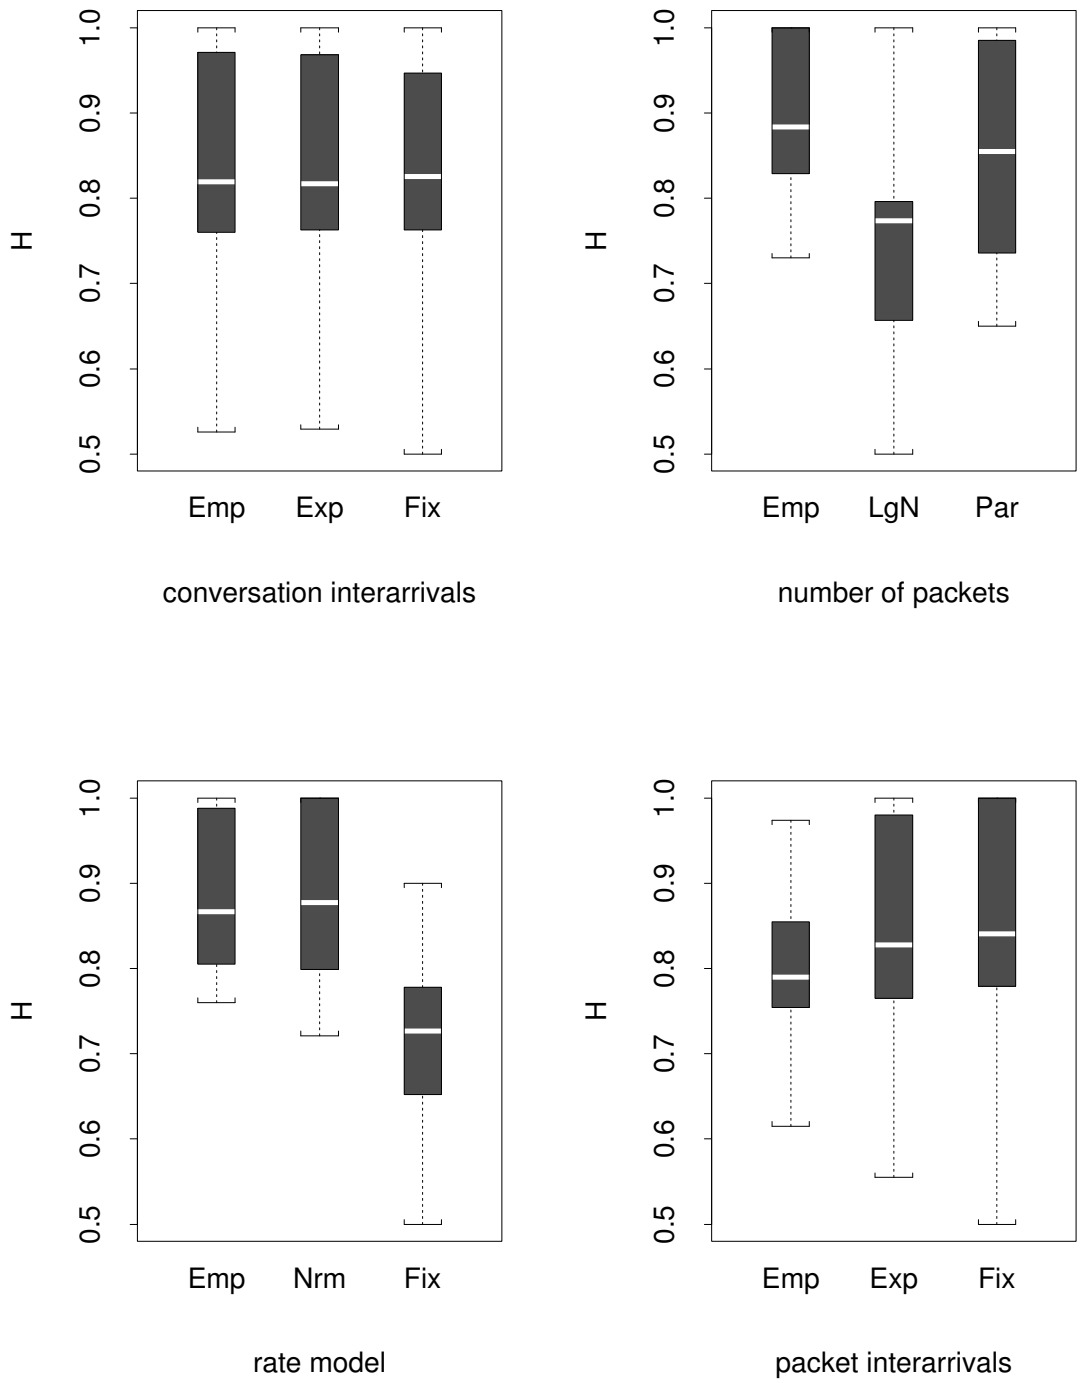


Figure 11: Boxplots by isolated factor for Whittle estimates of \hat{H} . Each box contains all values where the factor was set to the indicated level.




Figure 12: A graphical summary of the experimental results. The horizontal line is the mean of all \hat{H} estimates. For each level of each factor, the average response is as indicated on the vertical line corresponding to that factor.







Figure 13: Percent variation in response explained by different factors. The first four bars correspond to the top four factors contributing to variation in response. A fifth bar groups all the remaining factors. The line over the bars indicates the cumulative amount of variation explained by the factors up to that point.