On the Nonstationarity of Internet Traffic

Jin Cao, William S. Cleveland, Dong Lin, and Don X. Sun
{Statistics, Statistics, Networked Computing, Statistics} Research
Bell Labs, Murray, Hill, NJ
{cao, wsc, dong, dxsun}@bell-labs.com

ABSTRACT

Traffic variables on an uncongested Internet wire exhibit a pervasive nonstationarity. As the rate of new TCP connections increases, arrival processes (packet and connection) tend locally toward Poisson, and time series variables (packet sizes, transferred file sizes, and connection round-trip times) tend locally toward independent. The cause of the nonstationarity is superposition: the intermingling of sequences of connections between different source-destination pairs, and the intermingling of sequences of packets from different connections. We show this empirically by extensive study of packet traces for nine links coming from four packet header databases. We show it theoretically by invoking the mathematical theory of point processes and time series. If the connection rate on a link gets sufficiently high, the variables can be quite close to Poisson and independent; if major congestion occurs on the wire before the rate gets sufficiently high, then the progression toward Poisson and independent can be arrested for some variables.

1. INTRODUCTION

Long-range dependence and heavy-tailed marginal distributions have been established as important, fundamental characteristics of Internet traffic ([18; 23]). Here, we show that nonstationarity needs to be added to this list of fundamentals because nonstationarity is pervasive and affects many traffic variables that are measured through time on an Internet link. The form of the nonstationarity is important for Internet engineering. We show that as the rate, ρ, of new TCP connections on a link increases, traffic tends *locally* toward Poisson and independent.

1.1 Traffic Variables

We study the traffic variables empirically and theoretically. The empirical study is based on analysis of packet headers from four databases: measurements of the variables resulting from capturing all packet headers on links and adding a timestamp to each packet. We study all traffic variables by application protocol because their behavior differs from one application to the next. Here we present results in detail just for HTTP, and we briefly describe results for SMTP. We first established the traffic characteristics by study of one of the databases, from a link at Bell Labs; we present results for this database in detail. We then studied other databases for confirmation; the confirmation results are briefly described.

Two types of traffic variables are studied, packet and connection. The packet variables are the following: (1) size; (2) inter-arrival time. Connection variables describe the characteristics of application requests to TCP for transfers, as well as information about the Internet environment at the time of the request. The time of a TCP connection is taken to be the timestamp of the initial SYN packet. Each measurement of a connection variable is associated with a specific connection and each connection has measurements of a collection of connection variables all of whose timestamps are the timestamp of the connection. The connection variables are the following: (1) server-side round-trip time measured by the time between the client SYN and server SYN/ACK; (2) client-side round-trip time measured by the time between the server SYN/ACK and client ACK; (3) client file size measured from sequence numbers of the connection; (4) server file size measured from sequence numbers of the connection; (5) connection inter-arrival time, that is, the time until the next connection.

1.2 Individual Measurements vs. Counts

Here, we study individual measurements. In the past, much of the understanding of Internet traffic has arisen from studying counts in fixed-length intervals. Such counts, an aggregation of individual measurements, do not reveal our findings.

Consider packet arrivals as an example. We study the individual arrivals by studying inter-arrival times, t_j, as a time sequence. A standard practice has been to study packet counts, an aggregation of the individual arrivals, in fixed-length time intervals such as 10 ms. We study t_j because devices see packets, locally, arriving one by one, not all at once as an aggregate over 10 ms. We study t_j because they are the individual measurements and characterize the arrival process, but packet counts are summaries that do not retain the information in the arrival process. Consider what happens when the rate ρ increases. If we study the characteristics of 999 consecutive t_j then we study the characteristics of 1000 arrivals no matter what the value of ρ. But for byte counts in 10 ms intervals, as ρ increases, each interval aggregates more and more arrivals, masking the individual behavior more and more. Here, we show that the t_j tend to independent and exponentially distributed, the properties of a Poisson process. But packet counts become normally distributed and retain the same correlation structure. Thus the counts do not reveal the local Poisson limit.

The Poisson limit appears to provide the relevant framework for queuing behavior. Here, we investigate queuing by an open-loop simulation in which we feed measured packet arrivals and sizes from an Internet wire into an infinite-buffer queue. We find that as ρ increases, the queuing tends toward that of Poisson arrivals

with constant service times. Elsewhere, we carry out a queuing analysis with a closed-loop simulation that takes TCP feedback into account, and reach the same conclusion [4].

1.3 Contents of the Paper

Section 2 of this paper presents previous results. Section 3 describes the four databases used in our analysis. Section 4 discusses our tools for statistical modeling. In Section 5 we study the properties of the HTTP packet variables by a statistical modeling that summaries statistical behavior. In Section 6 we study the properties of the HTTP packet variables by an open-loop queueing analysis. Section 7 analyzes the HTTP connection variables; for each variable, a model is built that is comprehensive in the sense that it can generate synthetic values that mimic the statistical characteristics of the measurements. Section 8 discusses confirmation of our results. Section 9 summarizes and discusses our results.

2. PREVIOUS RESULTS

The traffic variables characterized here have been studied widely in the literature, and long-range dependence and heavy-tailed marginal distributions have been established as important characteristics [18; 27; 28]. These two characteristics are in fact closely related: the heavy-tailed marginal distribution of file sizes [22; 1; 23; 9] are a major factor in long-range dependence [9; 21]. Connection and packet inter-arrival times have been found to be long-range dependent and have a distribution that is either exponential or has longer tails than the exponential [22; 23; 11; 14]. Packet sizes have been shown to be multi-modal, primarily due to the mix of control packets such as ACKs, and data packets with differing maximum segment sizes [26]. Multifractral wavelet models have been developed to generate traffic with possibly heavy-tailed marginals and long-range dependence [25; 24].

Some authors have observed changes in the statistical properties of traffic. It is well known that traffic rates exhibit weekly and daily patterns because network usage has such patterns [23; 15; 26]. These changes, however, are a straightforward, simple instance of nonstationarity. The nonstationarity we address here is far more profound and involves the form of the long-range dependence and marginal distributions. Several authors have made reference to changes in these, but without connecting the changes to a changing rate ρ. Feldman [15] used the Weibull distribution to model HTTP connection inter-arrival times over one-hour time blocks and reported a substantial change of shape parameters from one time period to the next. Crovella and Bestavros [9] noticed a change in the Hurst parameter as the network utilization changes.

There have been some theoretical treatments of the effect of a changing amount of superposition on queuing. Queue length changes with an increasing number of superpositions while holding the utilization fixed [5; 2; 13; 12]. However, in these treatments, the probability of exceeding a level that grows with k, the number of superpositions, is studied for a fixed interval of time; but it seems to us more appropriate to let the time interval get smaller with the rate since we would not expect buffers to get arbitrarily large with an increase in traffic load but rather would expect devices to get faster.

Also, changes in the means and variances as a function of traffic load are studied by Morris and Lin [19]. They show that the relationship between changes in mean bandwidth and changes in variance for aggregated Web traffic are the same as those for Poisson traffic.

The first extensive empirical and theoretical modeling of nonstationarity was described in [7] for one HTTP connection variable: inter-arrival time. The overall results cited here about HTTP inter-arrivals are from that source, and we confirm the results with new data. The inter-arrival statistical model presented there is altered here to include a transformation that greatly simplifies the structure. The purpose of the work here is to study a large number of traffic variables, build models for them, and demonstrate the ubiquity of such nonstationarity in Internet traffic.

3. DATABASES

The primary database on which we draw for our empirical study in this paper arises from packet header collection on a 100mb/s Ethernet link at Bell Labs that connects a network of 3000 hosts to the rest of the Internet. For HTTP, all clients are on the inside of the network and all servers are on the outside, so incoming packets are from servers and outgoing are from clients. Data collection began on 11/18/1998 and continues through the present. Altogether there are 20 billion packets that have been organized in the database into 1 billion TCP connection flows. The Bell Labs database was used to establish the results presented here. We used three other databases to confirm the results; the confirmation is discussed in Section 8.

4. STATISTICAL MODELING

We build statistical models for the traffic variables in which unknown parameters of the models are estimated from the data. For each variable, the measurements are divided into a collection of time blocks; the connection rate, ρ, for the blocks ranges from smaller to larger. The model parameters are estimated for each block separately, and then the estimates are studied as a function of ρ. We found that simple statistical models do an excellent job of accounting for the statistical variability in the measurements. For each variable, the marginal distribution is modeled by fitting the empirical distribution. Then, taking the marginal distribution as fixed and known, a model is developed for the long-range dependence.

4.1 Marginal Distribution

The marginal distributions of the inter-arrival time variables are well approximated by Weibull distributions, where the parameters depend on the connection rate ρ. Let v_i be one of the inter-arrival variables, let $\lambda(\rho)$ be the shape parameter of the Weibull, and let $\alpha(\rho)$ be the scale parameter. Then

$$\left(\frac{v_i}{\alpha(\rho)}\right)^{\lambda(\rho)} = u_i$$

where u_i is a unit exponential. For packet sizes, we simply use the empirical distribution to model the marginal distribution because the distribution has atoms, or discrete points with many observations; approximately 2/3 of the values are 40 bytes or 1500 bytes. For round-trip times and file sizes we fit a semi-parametric model to a log transformation or power transformation of the variable. The transformed scale for a variable is broken up into intervals; the distribution has the same parametric form for all intervals but the parameters change. The parametric form is a Weibull for the round-trip times and a Pareto for the file sizes.

4.2 Long-Range Dependence

The long-range dependence of many of the traffic variables is described by a simple two-parameter model: a fractional sum-difference (FSD) model with transformation. Again, let v_i one of the variables. Let $F(v_i; \rho)$ be the cumulative marginal distribution function of v_i. (Note that in the notation we allow the possibility that

the marginal distribution depends on the rate ρ). Let $G(z)$ be the cumulative distribution function of a Gaussian, or normal, random variable with mean 0 and variance 1; that is, the distribution of the variable is $N(0,1)$. Let $z_i = H(v_i)$ be a transformation of v_i such that the marginal distribution of z_i is $N(0,1)$. Then $H(v_i) = G^{-1}(F(v_i;\rho))$. Let s_i be a $N(0,1)$ long-range dependent time series generated by a fractional ARIMA model:

$$(I-B)^d s_i = \epsilon_i + \epsilon_{i-1}$$

where ϵ_i is Gaussian white noise with mean zero and variance

$$\sigma_\epsilon^2(\rho) = \frac{(1-d)\Gamma^2(1-d)}{2\Gamma(1-2d)},$$

B is the backward shift operator

$$Bs_i = s_{i-1},$$

and

$$0 \le d < 0.5.$$

Let n_i be a $N(0,1)$ white noise time series independent of ϵ_i. Then the model is

$$z_i = \sqrt{1-\theta(\rho)}\; s_i + \sqrt{\theta(\rho)}\; n_i.$$

where

$$0 \le \theta(\rho) \le 1.$$

The term "sum-difference" is used because the fractionally differenced z_i is the sum of the output from applying a summation operator to a white noise series and the output of applying a difference operator to a white noise series independent of the first:

$$(I-B)^d z_i = \sqrt{1-\theta(\rho)}(I+B)\epsilon_i + \sqrt{\theta(\rho)}(I-B)^d n_i.$$

The autocorrelation function of z_i, from [17], is

$$a_z(k) = (1-\theta(\rho))\frac{(2k^2(1-d)-(1-d)^2)\Gamma(1-d)\Gamma(k+d)}{(k^2-(1-d)^2)\Gamma(d)\Gamma(k+1-d)},$$

for $k = 1, 2, \ldots$. The power spectrum is

$$p_z(f) = (1-\theta(\rho))\frac{(1-d)\Gamma^2(1-d)|1+e^{2\pi i f}|^2}{2\Gamma(1-2d)|1-e^{2\pi i f}|^{2d}} + \theta(\rho).$$

The variance of the component $\sqrt{1-\theta(\rho)}\; s_i$ is $1-\theta(\rho)$ and the variance of the component $\sqrt{\theta(\rho)}\; n_i$ is $\theta(\rho)$. We have allowed θ to depend on ρ but not d because, as we shall see, for the traffic variables studied here, θ changes with ρ but not d. As $\theta(\rho)$ tends to 1, the z_j tends to white noise.

To fit the FSD model with transformation to a set of measurements we take $F(v_i;\rho)$ to be the empirical marginal distribution function of the data. Then we fit the FSD models using ATS technology [8; 7]: the periodogram is averaged in blocks of size 5 and then the log of the FSD spectrum in the above equation is fitted to the log of the averaged periodogram using nonlinear least squares.

One of the connection variables, HTTP server file size, present opportunities to take aspects of the transfer into account to model the long-range dependence. The details of the tailored models will be given in Sections 7.

5. HTTP PACKET VARIABLES

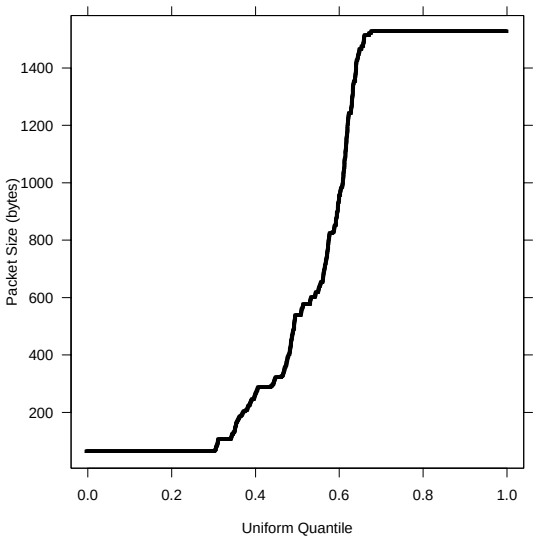

Figure 1: Empirical marginal distribution of 50000 packet sizes. The kth largest size is graphed against $(k-0.5)/50000$.

5.1 Data

We studied packet header data from the Bell Labs database to build statistical models for the HTTP packet variables and to carry out the queuing analysis in Section 6. The data cover the period 1/1/00 through 2/16/00.

The HTTP packets were taken to be those for which the server port is 80. We broke the time span of 46 days into 5 minute blocks; this block length is small enough to insure the requisite intra-block stationarity for most cases, but large enough to provide enough packets for fitting statistical models. For the analysis presented here, we selected a random sample of 500 blocks constrained so that the log rates were within a certain tolerance of being uniformly spaced from the minimum to the maximum log rate. The TCP connection rate $\hat{\rho}_b$ for each block b = 1 to 500 is the number of connections that started up in the block divided by the block length in seconds (300 sec). The $\hat{\rho}_b$ vary from 0.18 connections/sec (c/s) to 34 c/s. On the wire, packets travel in two directions: inbound (packets from the servers) and outbound (packets from the clients). We studied both, but present just inbound results here; outbound results are similar. The inbound packet rate for the 500 intervals varies from 1.7 packets/sec (p/s) to 452 p/s.

For each block we analyzed the packet size process and packet inter-arrivals. Each packet in a block is ordered from earliest to latest according to the timestamp. Suppose there are m_b packets in block b. The packet size variable for the block is the sequence of packet sizes, p_{bk} for k = 1 to m_b; p_{b1} is the size of the first arriving packet, p_{b2} is the size of the second arriving packet, and so forth. The inter-arrival variable is the time between successive packets, a_{bk} for k = 1 to $m_b - 1$; a_{b1} is the timestamp for the second packet minus the timestamp for first packet, a_{b2} is the timestamp for the third packet minus the timestamp for the second packet, and so forth.

5.2 Empirical Study: Marginal Distributions

We analyzed the inbound packet size marginal distribution for each block by a quantile plot [6]. Let $p_{b(k)}$ for k = 1 to m_b be the sizes ordered from smallest to largest for block b. Then $p_{b(k)}$ is plotted against $(k-0.5)/m_b$. The study revealed a small change in the packet size marginal distribution; as $\hat{\rho}_b$ increases, the mean

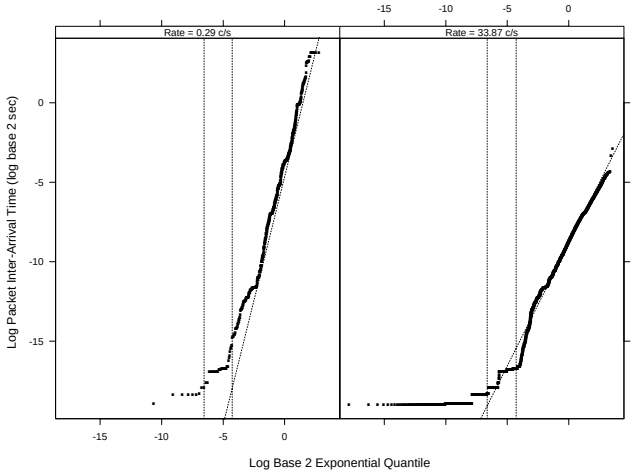

Figure 2: Quantiles of log inter-arrival times are graphed against log Weibull quantiles. The vertical lines show the 1% and 5% quantiles. The oblique line is drawn through the two quartile points.

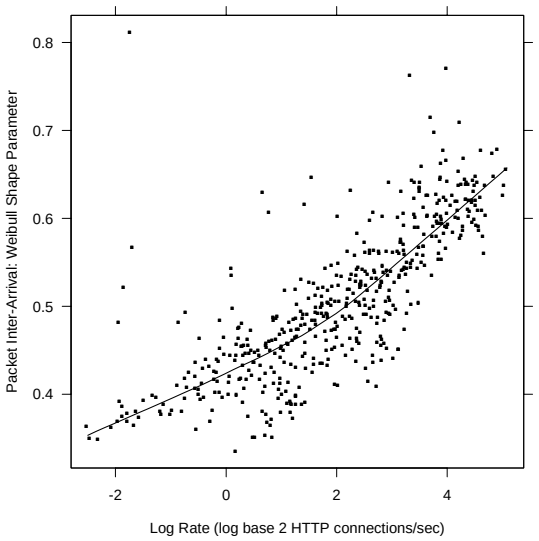

Figure 3: The maximum likelihood estimate of the Weibull shape parameter is graphed against the log connection rate for 500 blocks. The smooth curve is a loess fit with robust local linear fitting and a span of 0.75.

size increases but the increase is small compared with the range of the marginal distribution. Figure 1 shows the empirical marginal distribution using data from all of the blocks. In all of the 500 blocks there are 13.6 million packets which is far more than we need to characterize the marginal distribution so we randomly sampled 50000 sizes from each block to form the data used in the plot. 30% of the data have a size of 40 bytes, and 35% have a size of 1500 bytes. The remaining percentage is spread out between these two values although there are actually a few points of accumulation corresponding to infrequently used maximum segment sizes, but the usage is a very small percentage.

We study the inter-arrival time marginal distribution by a Weibull quantile plot. Let $a_{b(k)}$ for k = 1 to $m_b - 1$ be the inter-arrivals ordered from smallest to largest for block b. Let $\log_2$ denote log base 2, and let log denote the natural log. Then $\log_2(a_{b(k)})$ is plotted against the $\log_2$ of the quantile of probability $(k-0.5)/(m_b-1)$ of an exponential distribution, $\log_2(-\log(1-(k-0.5)/(m_b-1)))$. If the marginal distribution of the inter-arrival times is well approximated by a Weibull distribution with shape λ and scale α, then the pattern of points on the plot is approximately linear with slope λ^{-1} and intercept $\log_2(\alpha)$. For each block we estimate the shape and scale by maximum likelihood estimates, $\hat{\lambda}_b$ and $\hat{\alpha}_b$.

Figure 2 shows Weibull quantile plots for two blocks. In the left panel the connection rate is among the lowest, $\hat{\rho}_b = 0.29$ c/s, and in the right panel the connection rate is among the highest, $\hat{\rho}_b = 33.87$ c/s. On both plots the oblique line has slope $\hat{\lambda}_b^{-1}$ and intercept $\log_2(\hat{\alpha}_b)$. The leftmost vertical line is drawn at the 0.01 quantile of the $\log_2$ exponential distribution, which means that 1% of the points on the plot lie to the left of the line; the other vertical is drawn at the 0.05 quantile.

Figure 2 shows that the marginal distribution of the inter-arrivals for each block is well approximated by the Weibull. For the top 95% of the distribution the fit is excellent. For the bottom 5%, the data are larger than expected for the Weibull approximation. Part of the reason for the increase is the lower bound on the inter-arrival times of packets. The line speed on the Bell Labs wire is 100 mb/s. The inter-arrival time between packets k and $k-1$ in a block is no less than the transmission time of packet k which is $8p_{bk}10^{-8}$ seconds. For a 64 byte packet, the minimum size (40 bytes of TCP/IP header plus 14 bytes of Ethernet encapsulation), the transmission time is 5.12×10^{-6} sec, and for 1514 bytes it is 1.21×10^{-4} sec. But the departure from the Weibull involves only a small fraction of the data, and overall the fit is quite good.

The maximum likelihood estimate of the shape for the lower rate interval is $\hat{\lambda}_b = 0.378$, and for the higher is $\hat{\lambda}_b = 0.655$. We can see this in Figure 2 because the slope on the left is greater than on the right. This reflects the general trend. Figure 3 is a plot of $\hat{\lambda}_b$ against $\log_2(\hat{\rho}_b)$. The smooth curve is the fit from a nonparametric regression procedure, loess, with robust local linear fitting with a span of 0.75 [6]. The marginal distributions are tending toward the exponential, a Weibull with shape 1.

5.3 Empirical Study: Dependence

We fitted the FSD model with normality transformation (Section 4) to the measurements of the inbound inter-arrival variable of each block b, but we fitted the FSD model without transformation for the inbound packet size variable. Transformation makes sense for the continuous distribution of the inter-arrivals, but not for the sizes which have atoms in their distribution. We estimated the two FSD parameters, d and $\theta(\rho)$, for each block. For both the size and inter-arrival variables, the estimates of d remain nearly constant with the rate, and the median value for both is about 0.45. Thus we take the estimate for all blocks to be $\hat{d} = 0.45$ in both cases and re-estimate $\theta(\rho)$. Let $\hat{\theta}_b$ denote the new estimate (for either variable) for block b.

For the two blocks used in Section 5.2, the dots in Figure 4 show the $10\log_{10}$ of the averaged inter-arrival periodogram, and the curve shows the $10\log_{10}$ power spectrum of the fitted FSD model (with $d = 0.45$). The fits are quite good, so the FSD model appears to do a good job of fitting the second order dependence. This was the case for most of the 500 intervals. For the lower rate block ($\hat{\rho}_b = 0.29$ c/s), $\hat{\theta}_b$ = 0.58, and for the higher rate block ($\hat{\rho}_b$ = 33.87 c/s), $\hat{\theta}_b$ = 0.84, so the dependence decreases at the higher rate. The change is evident in the behavior of the power spectra. For the higher rate, the flat part of the spectrum at the higher frequencies extends far further down into the lower frequencies. This reflects the general trend. Figure 5 plots $\hat{\theta}_b$ against $\log_2(\hat{\rho}_b)$ for

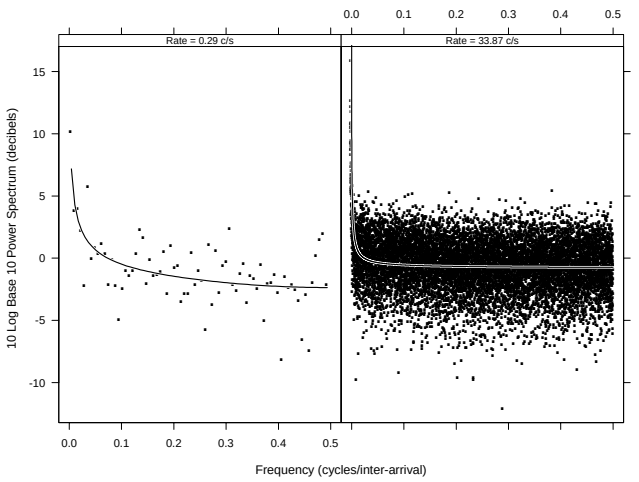

Figure 4: The log averaged periodogram is graphed against frequency for two blocks of packet sizes with low and high connection rates. The solid curve is the ATS estimate of the log power spectrum for the FSD model.

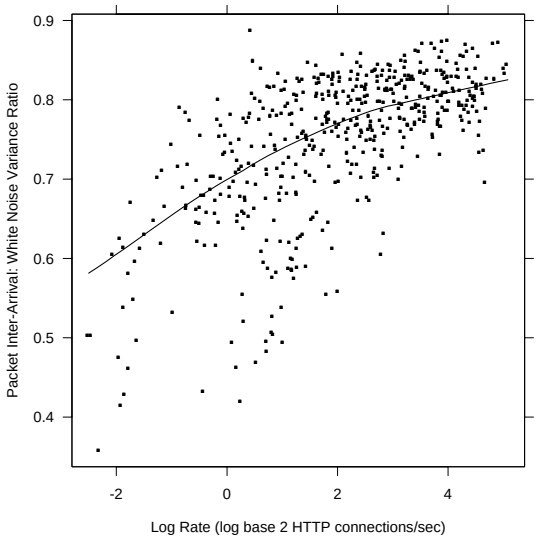

Figure 5: $\hat{\theta}_b$ is graphed against $\log_2(\hat{\rho}_b)$ for the packet inter-arrival variable. The smooth curve is a loess fit with robust local linear fitting and a span of 0.75.

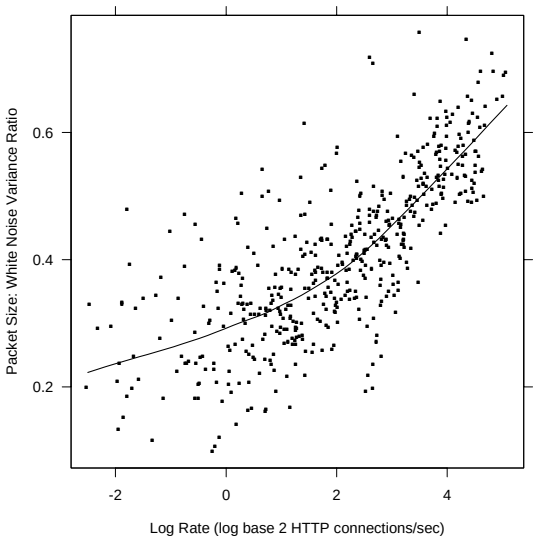

Figure 6: $\hat{\theta}_b$ is graphed against $\log_2(\hat{\rho}_b)$ for the packet size variable. The smooth curve is a loess fit with robust local linear fitting and a span of 0.75.

the inter-arrival variable. The smooth curve is a loess fit with local linear robust fitting with a span of 1 [6].

The results of fitting the FSD model to the measurements of the packet size variables are similar. Figure 6 shows the resulting plot of $\hat{\theta}_b$ against $\log_2(\hat{\rho}_b)$. Thus, for both packet size and packet inter-arrivals, the long-range dependence of the variables decreases with ρ and tends toward independence. There is more dependence and a greater change in dependence in the size variable.

5.4 Discussion of Results

Overall, the empirical results reveal a major change in the statistical properties of the inbound size and inter-arrival variables, a significant nonstationarity. The marginal distribution of size does not change much, but for the inter-arrivals, the Weibull distribution changes substantially, tending toward exponential. There is a major change in the long-range dependence of both variables. And as we will see in Section 6, these changes result in a major change in the queuing characteristics.

Increased superposition of HTTP flows provides a convincing theoretical explanation for the nonstationarity of both the size and the inter-arrival variables although the mathematical mechanism is somewhat different for the two.

Consider first the size variable. There are possibly two contributing factors here. First, within a connection, packet sizes are more similar than between connections. This is because two hosts establish a maximum packet size for each connection and different servers have different policies in terms of filling packets with data. Second, the TCP connection start up and tear down will create control packets in the two ends and data packets in the middle, and the sizes of control packets and data packets are markedly different. Because file sizes are heavy tailed, this will possibly generate long stretches of data packets of similar sizes within a connection. When $\hat{\rho}_b$ is low, sequences of packet arrivals tend to be from the same connection, so there tend to be bursts of packets of similar sizes. As $\hat{\rho}_b$ gets larger, packets from the individual connections tend to maintain similar inter-packet spacing, but for any particular connection, there are now many more packets to be interspersed among those of the connection. As the rate increases, the increased

intermingling of packets breaks up the bursts of similar packets, and the dependence decreases.

Now consider the inter-arrival variable. For low $\hat{\rho}_b$, the burstiness of HTTP connections ([14]; [7]), the heavy-tailed file sizes, the behavior of TCP, and the queuing of packets along the path to the link create bursts of packet arrivals. This results in a marginal distribution of the inter-arrivals that is skewed with respect to the exponential, and results in long-range dependence among the inter-arrivals. For a higher $\hat{\rho}_b$, the arrival process can be thought of as the superposition of arrival processes at a low $\hat{\rho}_b$; for example, if the high rate is 30 c/s, then the process can be seen as the superposition of 6 processes with rate 5 c/s. The theory of point processes shows that such superposition leads, locally, to a Poisson process. For example, Proposition 9.2.VI of [10] states that a process obtained by superposing k independent identically distributed stationary point processes and then dilating the time scale by a factor of k, will converge weakly to a Poisson process. This applies to a point process even if the inter-arrival times have a heavy-tailed distribution.

6. HTTP: QUEUING ANALYSIS

6.1 Experimental Design

We will explore further nonstationarity in the HTTP packet data described in Section 5 by carrying out 500 queuing simulation experiments, one experiment per block. Each experiment is a 2^3 factorial: three factors each at 2 levels, and 8 runs consisting of all possible combinations of the levels of the three factors. In Run 1, the packets are fed into the queue according to the observed inter-arrivals and sizes. In each of the other 7 runs, the sizes and inter-arrivals of the block are altered according to the factors.

The first factor is the form of the inter-arrival marginal distribution; we use either the original observed values or the values of an exponential distribution. The exponential case is as follows. As in Section 5, let $a_{b(k)}$ be the $m_b - 1$ packet inter-arrivals of block b, ordered from smallest to largest. Let $\bar{a}_b$ be the mean inter-arrival time. Let q_{bk} be the quantile of probability $(i - 0.5)/(m_b - 1)$ of an exponential distribution with mean $\bar{a}_b$. Then in the exponential case, the kth inter-arrival time is q_{bk}. Changing from the original inter-arrivals to exponential inter-arrivals decreases the two tails of the distribution, which as we saw in Section 5 were Weibull and thus more skewed or stretched out than the exponential. We expect this change to shift the queuing distribution to smaller values. The second factor is the order of the inter-arrival times in the block; we use either the original order or a randomized order. For the random order, the expected correlation is zero. Because the randomization removes the long-range dependence in the packet inter-arrival variable, we expect it to shift the queuing distribution to smaller values. The third factor is the packet size. We take the sizes to be either the original ones or to be constant and equal to the mean size of the packets in the block. Because moving to a constant removes the long-range dependence in the sizes we expect it to shift the queuing distribution to smaller values. Table 1 shows the eight combinations of the factors for the 8 runs.

Let $\hat{r}_b$ be the throughput of the packets in block b, that is, the sum of the packet sizes (in megabits) divided by 300 sec. This, of course, is the throughput for Run 1 of the experiment for block b. The throughput is the same for all 8 runs of each block experiment. The sum of the packet sizes remains the same in the 8 runs, and all arrivals in each run are contained within the 300 sec period. Because the packets were collected on a 100mb/s link, $\hat{r}_b$ can never be more than 100mb/s, and is always much less. The values of $\hat{r}_b$ depends of course on the the connection rate of the block $\hat{\rho}_b$.

Run	Marginal	Order	Size	Combination
1	orig	orig	orig	orig.orig.orig
2	orig	rand	orig	orig.rand.orig
3	orig	orig	const	orig.orig.const
4	orig	rand	const	orig.rand.const
5	exp	orig	orig	exp.orig.orig
6	exp	rand	orig	exp.rand.orig
7	exp	orig	const	exp.orig.const
8	exp	rand	const	exp.rand.const

Table 1: Experimental Design

We suppose that when a packet arrives, space is allocated for it instantaneously, so the queue size is increased instantaneously at the time of arrival by the number of bytes in the packet. We also suppose that the queue buffer size is infinite, so packets are never dropped.

Let d_b be the output line speed (in mb/s), that is, the speed at which the queue drains. The utilization of the output link is r_b/d_b. For each run in each block we carry out a simulation for each of 4 utilizations: 0.1, 0.3, 0.5, and 0.7. That is, for each run we simulate four times, changing d_b to achieve each utilization. Altogether we run 16000 simulations: 500 blocks $\times$ 8 runs $\times$ 4 utilizations.

For each simulation we compute the queue length in bytes at the time of arrival of each packet, but do not include the arriving packet as part of the queue, so we measure what each packet sees in the queue when it arrives. To summarize the queue length marginal distribution of the simulation for a given block, run, and utilization, we compute the probability that the queue length exceeds each of 18 values: 0 bytes and 2^k bytes for k = 0 to 16.

Run 8, exp.rand.const, serves as a benchmark for all other runs. The inter-arrival times are independent and exponentially distributed, so they form a Poisson process. The packet sizes are constant. Thus we have a queue with Poisson arrivals and constant service time. For each utilization, the queuing distribution does not depend on $\hat{\rho}_b$; that is, the queuing distribution for Run 8 is the same across the 500 blocks for a fixed utilization. Furthermore, we expect the least amount of queuing for Run 8. So for each utilization we will compare the queue length distribution for Runs 1–7 with that of Run 8.

6.2 Previous Results and Goals

This queuing analysis is carried out in the same general mode as that of [13] — input observed packet traces into an infinite-buffer queue with fixed utilizations — but with an expansion of several aspects. Their experiment fed arrival times to the queue, but not packet sizes. Their experiment had one factor, inter-arrival order, and had one block, 30 minutes long. They use a single summary statistic of queue length, the mean.

The previous queuing analysis shows a substantial decrease in the mean queue length at higher utilizations for the randomized order. They then use this result as an argument for the importance of long-range dependence for network engineering. While we believe this type of queuing analysis might suggest characteristics of live queues on the Internet, it is open-loop and suffers from missing TCP feedback which has a large impact on queuing. Our purpose in carrying out the queuing analysis is not to make definitive conclusions about live queues. Rather, the primary purpose of the queuing analysis is to provide another method for studying the nonstationarity of the packet traffic variables. In the previous section we saw a substantial nonstationarity: parameters of statistical models fitted to the data change with increasing $\hat{\rho}_b$ in such a way that the packet

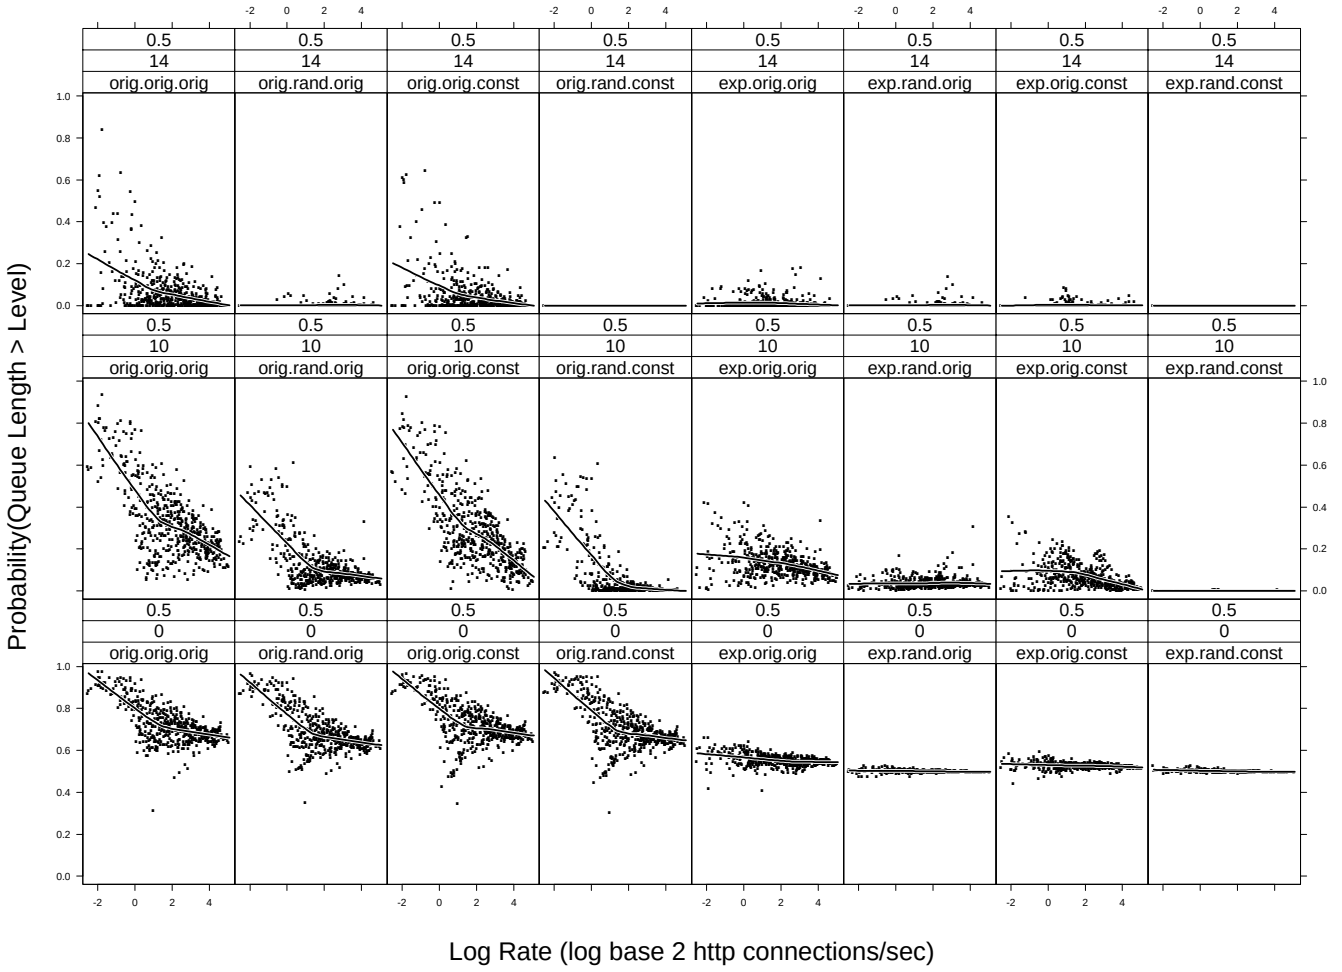

Figure 7: On each panel, the probability of exceeding a particular queue length for 500 blocks is graphed against the block HTTP connection rates. Each row of panels has the same queue length and each column has the same run type. The smooth curve on each panel is a loess fit with normal local linear fitting and a span of 0.75.

variables tend toward Poisson and independence. We seek to corroborate this result with another method of analysis. For example, we wish to see if the queue-length distributions of each of Runs 1–7 tends to that of Run 8 as $\hat{\rho}_b$ increases.

6.3 Results

Figure 7 shows the experimental results for the 500 simulations for which the utilization is equal to 0.5. There are 24 panels. Each panel conveys results for one queue length and one run type. For each row, the panels have the same queue length; from bottom to top the lengths are 1, 2^{10}, and 2^{14} bytes. For each column, the run type is the same; from left to right we go from Run 1 to Run 8. On each panel, the 500 block probabilities of exceeding the panel queue length for the panel run type are graphed against the 500 block connection rates.

The results in Figure 7 are representative of results at other utilizations. For each queue length, Run 8, the benchmark, has the lowest probabilities, and they are nearly constant as function of $\hat{\rho}_b$. For some other combinations of queue length and run, the probabilities are nearly constant and close to those for Run 8. In the remaining panels, the probabilities have an overall downward tread, tending toward the level for Run 8 as $\hat{\rho}_b$ increases. Of the three factors, the marginal-distribution factor has the largest effect, the dependence factor has the next largest, and the size factor has the least. However, for the longest queue length, the order factor has a slightly bigger effect than marginal-distribution factor.

However, the most important aspects of these results is that the probabilities for Run 1, those for the original data, tend toward the nearly constant probability for Run 8. In other words, as the rate increases, the queuing distribution of the live packet process tends toward that for Poisson arrivals and constant service time. This agrees with the result of Section 5 that the packet arrivals are tending toward Poisson and the packet sizes are tending toward independence.

7. HTTP CONNECTION VARIABLES

7.1 Data

For HTTP we studied the five connection variables — inter-arrival, client file, server file, client round-trip time, and server round-trip time — with data from the time period 1/1/00 through 2/16/00. We used HTTP connections for which the transferred numbers of data bytes from client to server and server to client are both greater than zero. We broke the data into 5-minute blocks and removed blocks for which there was gross nonstationarity. For the analysis presented here, we use the same 500 blocks as described in Section 5.

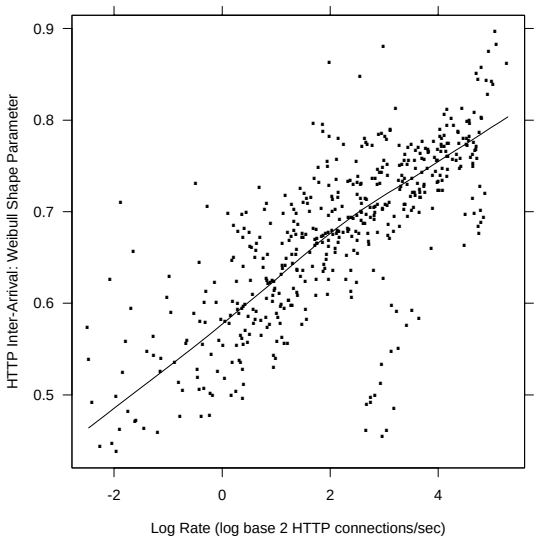

Figure 8: $\hat{\lambda}_b$ is graphed against $\log_2(\hat{\rho}_b)$ for HTTP inter-arrival time. The smooth curve is a loess fit with robust local linear fitting and a span of 0.75.

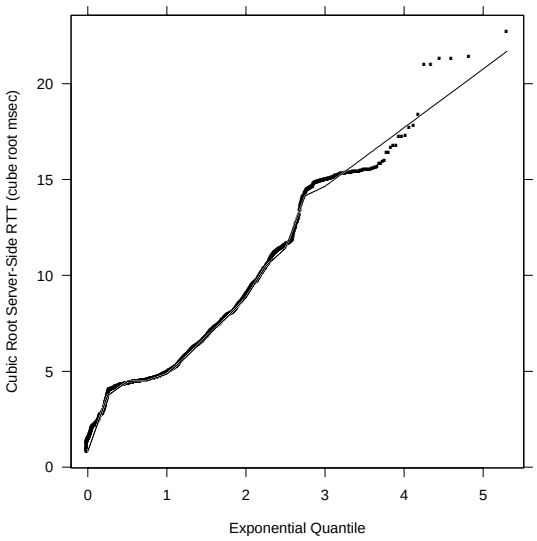

Figure 10: Quantiles of the cubic root of server side round-trip time are plotted against the quantiles of an exponential distribution. The curve is the semi-parametric fit to the marginal distribution.

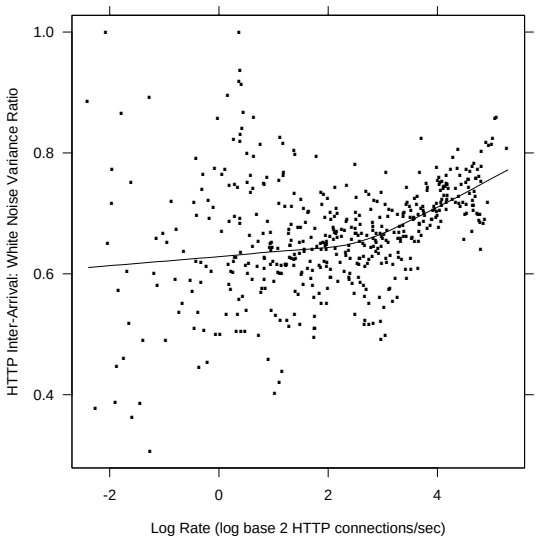

Figure 9: $\hat{\theta}_b$ is graphed against $\log_2(\hat{\rho}_b)$ for HTTP inter-arrival time. The smooth curve is a loess fit with robust local linear fitting and a span of 0.75.

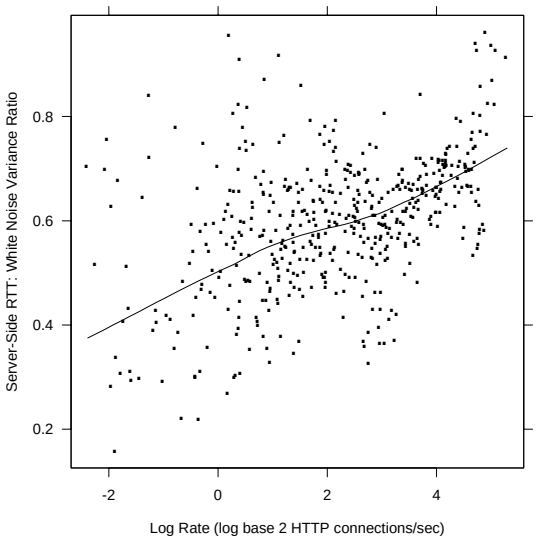

Figure 11: $\hat{\theta}_b$ is graphed against $\log_2(\hat{\rho}_b)$ for HTTP server round-trip time. The smooth curve is a loess fit with robust local linear fitting and a span of 0.75.

7.2 Inter-arrival Times

We fitted a Weibull marginal distribution to the inter-arrival times of each of the 500 blocks following the method set out in [7]. The results agree with this previous analysis, which used different data. Figure 8 graphs the estimates of the Weibull shape parameter against the log base 2 connection rate. As the rate increases, the shape changes from 0.46 to 0.80, an appreciable change. Since the exponential distribution is a Weibull with shape 1, this means the marginal distribution tends toward the exponential.

We fitted a FSD model to the inter-arrival times. In [7] an earlier version of the FSD model that had a much more complex transformation and back-transformation procedure was fitted; our new results are in general agreement with those of the previous ones. The estimated value of d in the new analysis did not change with the connection rate and was set equal to 0.4; $\theta_b(\rho)$ was estimated for this fixed value. Figure 9 plots $\hat{\theta}_b(\rho)$ against $\hat{\rho}_b$; the smooth curve was fitted by loess with local linear fitting and a smoothing parameter of 0.75. Overall, $\theta_b(\rho)$ changes substantially, from 0.61 to 0.77, so the long-range dependence is markedly reduced. [7] also provide an explanation of the dependence in the inter-arrival times and, and verify the results with several theoretical tools.

7.3 Round-Trip Times

Both the client and server round-trip times have marginal distributions that do not change much through time. We built statistical models for these marginals that is nonparametric below a cut-off, and is Weibull above the cutoff. We again use a quantile plot. Let $r_{(k)}$ for k from 1 to m be samples of round-trip times from small to largest. We plot the power transform of $r_{(k)}$ of some α, $r^{\alpha}_{(k)}$, against the quantile of probability $(k-0.5)/m$ of an exponential distribution, $-\log_{10}(1-(k-0.5)/m)$. Figure 10 shows such an quantile plot of the cubic root of server round-trip times against quantiles of the exponential distribution. The power transform is used here to give a reasonable scale to model the distribution. Then we break up the exponential quantiles into n intervals using points $0 = q_1 < q_2 < \ldots < q_n < q_{n+1} = \infty$. In each interval $[q_i, q_{i+1})$, we fit the quantiles $r^{\alpha}_{(k)}$ of the same probability by a straight line, and require these lines to be continuous at points q_i. Since quantiles on last interval $[q_n, \infty)$ represent the tail of the distribution, this means we use the Weibull to approximate the tail of the round-trip time distribution. Such a piecewise linear fit for the server round-trip time is shown as the gray line in Figure 10, where we used $n = 13$ equally spaced points from 0 to 3. A similar semi-parametric fit is done for the client side round-trip time using a power transform of $\alpha = 1/5$.

Although the round-trip times do not show much change through time, the dependence is nonstationary and changes with the connection rate. We fitted FSD models to both variables for each of the 500 intervals. For both variables, the estimates of d do not depend on the connection rate, and are set to 0.4, which is close to the median estimate for each. The estimates of the variance ratios $\hat{\theta}_b$ for this fixed value are plotted against $\hat{\rho}_b$ in Figure 11 for server round-trip time. The ratios tend toward 1, which means the round-trip time variables tend toward independence. Client round-trip time shows similar results.

Why does dependence occur in these round-trip times, and why does it decrease with the connection rate? For HTTP 1.0, the dominant version of HTTP now in use, each click by a user on a link will likely to create a sequence of TCP connections for the same client and the same server, and continued clicks by the user for the same site creates a continuation of this sequence. Consider first the server round-trip time. The measurements for two separate servers can be quite different if the propagation delay is quite different, or if paths to them from the Bell Labs network experience very different amounts of congestion. This variation between the servers can be much larger than the variation of the measurements for either one of the servers separately. When the connection rate is low, the sequence of server round-trip times will tend to have runs from a single server, so we might get a run from one server of values that are very small and about the same, and then a run of values from another server that are vary large and about the same. The result is a large amount of time dependence. As the connection rate goes up, the run lengths of round-trip times from the same server are reduced because the connections of different servers intermingle more, and eventually, the server sequence tends toward random, so the server round-trip times tend toward random as well. The same consideration applies to the client round-trip times.

7.4 Server and Client File Sizes

The marginal distributions of the client and server file sizes do not change much with the rate and were fitted using the same semi-parametric method used for the round-trip times except that we used the log transformation instead of the power transformation. This means tails of the file size distributions are fitted by the Pareto distribution, which is consistent with the previous observations that file sizes are heavy tailed ([21]; [22]; [1]; [23]; [9]). We observe a time dependence in both the client and server file size time series that was fitted by FSD models. The estimates of the variance ratios $\hat{\theta}_b$ tend to 1 with $\hat{\rho}_b$, which means the size variables tend toward independence. The server size time series however has far less dependence than the client series.

We discovered that a specialized model fitted the server size variable even better than the FSD model. Servers respond to cache validation requests either by sending a file if there is a new version or sending an "unmodified" message. For the 500 blocks reported here, about 30% of the HTTP connections are in the latter category. The distribution of the file sizes, excluding the unmodified messages, do not differ markedly across servers. The unmodified messages are mostly less than 300 bytes, and, in fact, make up the majority of the messages less than 300 bytes. However, the format of the messages is the same for each server, and thus the file sizes are the same, but vary from server to server, which creates different server HTTP file distributions at the bottom end of the distribution. Because there are so many "unmodified" messages dependence is created.

Because of the special form of the dependence for the server file sizes, we developed a specialized model for this special case. First we generate a categorical time series each of whose values is either "modified" or "not-modified". We get the categorical series by generating alternating run lengths randomly generated from a distribution on the positive integers for "modified" and from a distribution on the positive integers for "not-modified". For example, the run series might be one "modified", four "not-modified", three "modified", one "not-modified", ... For a "modified" run we generate one value randomly from the portion of the file size distribution below 300 and then repeat it for the run. For a "not-modified" run we generate independent values, one for each element of the run, from the the portion of the file size distribution above 300.

The run-length distributions are from a family we call the discretized Weibull or dW distribution. A discretized Weibull with shape and scale parameter, λ, α, is the round up of a Weibull with the same shape and scale parameter. If the shape parameter is 1, then the run length distribution is geometric; when this occurs the the categorical series is a Bernoulli series: the values of the series are independent. This makes the file size variables independent. We fitted the parameters for the "not-modified" run lengths and

"modified" run lengths for the 500 blocks and graphs the shape parameters $\hat{\lambda}_b^*(\rho)$ against $\hat{\rho}_b$ for "not-modified" distribution. The graphs show that shape parameter for the "modified" run-length distribution is constant with a median value in the vicinity of 0.8, and the shape parameter for the "not-modified" run-length distribution increases with the rate and interestingly, converges on 0.8. The problem is a minor lack of fit of the discretized and not a convergence to a dependent file size time series. Actually, 0.8 produces results that are quite close to independent.

8. CONFIRMATION

We believe it is reasonable to hypothesize that our results will hold for other application protocols other than HTTP, for links other than the Bell Labs link, and for connection rates higher than those observed on the Bell Labs link. The reason is the cause, superposition, which is a universal Internet phenomenon. The characteristics of the nonstationarity are governed by the mathematics of superposition which is invariant across links and applications. However, the validity of the mathematics is dependent on our underlying assumptions. There are two critical ones: (1) the sources are homogeneous, and (2) the sources are independent. Assumption (1) is reasonable because different implementations and usages of individual Internet applications are not too disparate. Assumption (2) is reasonable for a link with connection rates below that where major congestion begins on the link.

8.1 SMTP

We carried out a study of SMTP for the same traffic variables as for HTTP using data from the Bell Labs database. We found very mild nonstationarity because even at very low rates, the SMTP variables are close to Poisson and independent.

8.2 Connection Variables

We carried out a confirmation study for HTTP connection variables using data from the Helios Next-Generation-Internet traffic analysis project [16]. Packet header collection is being carried out on the 1 gb/s Ethernet link connecting the Chapel Hill campus of the University of North Carolina to an OC48 fiber ring that carries UNC traffic to other local campuses and to the rest of the Internet. The ring is part of the NCNI gigapop [20]. The collection, management, and analysis of the 1gb/s Ethernet data is supported by DARPA under Federal Contract No. F30602-00-C-0034. Slightly less then 1 hour of data was collected to test the measurement operation. Since the test data appeared reliable we have used them to assess our results.

We carried out the same analysis of the Helios HTTP connection variables that was carried out in Section 7. For the Helios wire, there are HTTP clients on both sides, so there are two directions for each connection variable, inbound to UNC servers and outbound from UNC clients. We present here the results for the outbound traffic, since the monitored link is close to the the clients in this case, similar to the Bell Labs link. We divided the 54 minutes of data into 54 1-minute blocks; thus there are 54 fittings of our statistical models for each connection variable. The values of ρ range from 136.5 c/s to 172.7 c/s.

We found our extrapolation hypothesis to be true for the Helios data. In general, long-range dependence was reduced even further beyond that for our Bell Labs database, and inter-arrival distributions were closer to exponential.

8.3 Packet Variables

We carried out a confirmation study for the packet variables from two packet-header databases. The first results from measurements from a 100 mbps Ethernet link connecting Harvard University to the rest of the Internet. We studied just incoming HTTP packets. The second database draws on a database of traces available at the National Laboratory for Applied Network Research (www.nlanr.net) and developed under the auspices of the National Science Foundation NLANR/MOAT Cooperative Agreement (No. ANI-9807479). We obtained 5 traces measured on an OC3 ATM link at Colorado State University in Ft. Collins, Colorado, and 5 traces measured on an ATM link at Columbia University in New York City, from the period 09/01/2000 through 11/05/2000. The results strongly confirmed our results and are reported in [3].

9. RESULTS

We study the traffic variables described in Section 1 by analyzing measurements from many time blocks of packet traces from nine Internet links, and then relating the statistical properties to the new connection rates ρ of the blocks.

9.1 Conclusions

The empirical distribution of packet inter-arrivals are well approximated by a Weibull distribution with a shape parameter λ that is less than 1 for low values of ρ. This means that the inter-arrival distribution has a longer upper tail than the exponential. But as ρ increases, λ tends to 1, so the distribution tends to an exponential. The inter-arrivals, as a time sequence, are long-range dependent. The packet sizes have a marginal distribution that does not change with ρ. The size time sequence, like the inter-arrivals, is long-range dependent. As ρ increases, the dependence in both time sequences decreases, and the sequences tend toward independence. This means that the arrivals tend toward a Poisson process. A open-loop queuing study supports these results; the queuing distribution of the live packet traces tends toward that for Poisson arrivals and constant service times. The very simple statistical FSD model given in Section 4 does an excellent job of modeling the long-range dependence of the inter-arrivals and sizes.

Similar results hold for the five HTTP connection variables. HTTP inter-arrivals are long-range dependent and have a Weibull distribution but tend toward a Poisson process as ρ increases. The two file size variables and the two round-trip time variables are long-range dependent and tend toward independence as ρ increases, but their marginal distributions stay the same with ρ. Again, the FSD model does an excellent job of modeling the long-range dependence.

9.2 Discussion: The FSD Model

One important feature of the FSD model is that it is simple, with only two parameters d and θ. Another is that it shows that the use of the Hurst parameter $H = d + 0.5$ by itself to characterize long-range dependence is quite inadequate. For each variable we found d did not vary appreciably with ρ, but θ increased with ρ, sometimes dramatically. In other words, the long-range dependence decreased dramatically but d stayed constant. For our variables then we need both d and θ to characterize the long-range dependence.

9.3 Discussion: Counts vs. Inter-Arrivals

The packet arrivals on a link can be thought of as the superposition of independent point processes. As ρ increases, the amount of superposition increases. If we look at 1000 consecutive arrivals for any block of measurements no matter what the value of ρ, then we study the arrival process locally; as ρ increases, the 1000 arrivals tend to cover a smaller range of time. If we study counts of arrivals

in fixed-length intervals, the standard for studies of long-range dependence of Internet traffic, then the measure becomes less and less local because more and more arrivals occur. It is easy to see that the correlation structure of the counts of the superposed processes stays the same with the amount of superposition, so the fixed-interval counts retain their long-range dependence as ρ increases. But the theory of point processes shows that the fixed-number arrivals tend to Poisson. Thus studies of packet counts do not reveal the results shown here.

10. REFERENCES

[1] M. F. Arlitt and C. L. Williamson. Web Server Workload Characterization: The search for Invariants. In *Proc. ACM SIGMETRICS*, pages 126–137, 1996.

[2] D. D. Botvich and N. G. Duffield. Large Deviations, the Shape of the Loss Curve, and Economies of Scale in Larger Multiplexers. *Queueing Systems*, 20:293–320, 1995.

[3] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. Internet Traffic Tends to Poisson and Independence as the Load Increases. Technical report, Bell Labs, Murray Hill, NJ, 2001.

[4] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. PackMime: An Internet Traffic Generator. Technical report, Bell Labs, Murray Hill, NJ, 2001.

[5] G. L. Choudury, D. Lucantoni, and W. Whitt. Squeezing the Most Out of ATM. *IEEE Transactions on Communications*, 44(2):203–217, 1996.

[6] W. S. Cleveland. *Visualizing Data*. Hobart Press, Summit, New Jersey, U.S.A., 1993.

[7] W. S. Cleveland, D. Lin, and D. X. Sun. IP Packet Generation: Statistical Models for TCP Start Times Based on Connection-Rate Superposition. In *Proc. ACM SIGMETRICS '00*, pages 166–177, 2000.

[8] W. S. Cleveland, C. L. Mallows, and J. E. McRae. ATS Methods: Nonparametric Regression for Nongaussian Data. *Journal of the American Statistical Association*, 88:821–835, 1993.

[9] M. E. Crovella and A. Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proc. ACM SIGMETRICS*, pages 160–169, 1996.

[10] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes*. Springer-Verlag, New York, 1988.

[11] S. Deng. Empirical model of www document arrivals at access link. In *Proceedings of ICC/SUPERCOMM*, 1996.

[12] N. G. Duffield. Economies of Scale in Queues with Sources Having Power-Law Large Deviation Scaling. *Queueing Systems*, 33:840–857, 1996.

[13] A. Erramilli, O. Narayan, and W. Willinger. Experimental Queueing Analysis with Long-Range Dependent Packet Traffic. *IEEE/ACM Transactions on Networking*, 4:209–223, 1996.

[14] A. Feldmann. Characteristics of TCP Connection Arrivals. Technical report, AT&T Labs Research, 1998.

[15] A. Feldmann, A. C. Gilbert, P. Huang, and W. Willinger. Dynamics of IP Traffic: A Study of the Role of Variability and the Impact of Control. In *Proc. ACM SIGCOMM*, pages 301–313, 1999.

[16] Helios Next Generation Internet Project. www.anr.mcnc.org/projects/Helios/Helios.html. Technical report, North Carolina Network Initiative, Research Triangle Park, NC, 2000.

[17] J. R. M. Hosking. Modelling Persistence in Hydrological Time Series Using Fractional Differencing. *Water Resources Research*, 20(12):1898–1908, 1984.

[18] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the Self-Similar Nature of Ethernet Traffic. *IEEE/ACM Transactions on Networking*, 2:1–15, 1994.

[19] R. Morris and D. Lin. Variance of Aggregated Web Traffic. In *Proceedings of INFOCOM*, 2000.

[20] NCNI Network. www.ncni.org. Technical report, North Carolina Network Initiative, Research Triangle Park, NC, 2000.

[21] K. Park, G. Kim, and M. Crovella. On the Relationship Between File Sizes, Transport Protocols, and Self-Similar Network Traffic. In *Proceedings of the IEEE International Conference on Network Protocols*, 1996.

[22] V. Paxson. Empirically-Derived Analytic Models of Wide-Area TCP Connections. *IEEE/ACM Transactions on Networking*, 2:316–336, 1994.

[23] V. Paxson and S. Floyd. Wide-Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Transactions on Networking*, 3:226–244, 1995.

[24] V. J. Ribeiro, R. H. Riedi, M. S. Crouse, and R. G. Baraniuk. Simulation of NonGaussian Long-Range-Dependent Traffic Using Wavelets. In *Proc. ACM SIGMETRICS*, pages 1–12, 1999.

[25] R. H. Riedi, M. S. Crouse, V. J. Ribeiro, and R. G. Baraniuk. A Multifractal Wavelet Model with Application to Network Traffic. *IEEE Transactions on Information Theory*, 45(3):992–1019, 1999.

[26] K. Thompson, G. J. Miller, and R. Wilder. Wide-Area Internet Traffic Patterns and Characterics. *IEEE Network*, Nov. 1997.

[27] W. Willinger, M. S. Taqqu, W. E. Leland, and D. V. Wilson. Self-Similarity in High-Speed Packet Traffic: Analysis and Modeling of Ethernet Traffic Measurements. *Statistical Science*, 10:67–85, 1995.

[28] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level. *IEEE/ACM Transactions on Networking*, 5:71–86, 1997.