Internet Traffic Tends To Poisson and Independent as the Load Increases

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

The burstiness of Internet traffic was established in pioneering work in the early 1990s, which demonstrated that packet arrival times are not Poisson, and packet and byte counts in fixed-length intervals are long-range dependent [17, 20]. Here we demonstrate that these results are one end of a continuum of traffic characteristics. At the other end are Poisson behavior and independence. Our study focuses on packets, what devices actually see; we study the statistical properties of packet inter-arrival times and packet sizes. As the traffic load increases — that is, as the number of simultaneous transport connections increases — arrivals tend to Poisson and sizes tend to independence. More specifically, long-range dependence of inter-arrivals and sizes decreases to independence, and the marginal distribution of inter-arrivals tends toward exponential; this happens (1) through time on a single link as the load increases due to daily variation, or (2) at a single point in time as the load increases going from lightly loaded links at the edges of the Internet to heavily loaded links at the core. Convergence is rapid; the packet traffic gets quite close to Poisson and independent loads far less than the maximum we observe.

1. FOCUS

This article focuses on packet traffic on Internet links. We study two traffic variables, packet size and packet inter-arrival time. Let q_j be measurements of the packet sizes on a link in a single direction where $j = 1$ corresponds to the first arriving packet, $j = 2$ to the second, and so forth. Let t_j be measurements of the inter-arrival times where $j = 1$ is for the time between packets 1 and 2, $j = 2$ is for the time between packets 2 and 3, and so forth. We will suppose that both variables are stationary time series in j over short blocks no larger than 5 minutes.

We focus on packet traffic variables because it is packets that network devices must send and receive, and the burstiness of traffic as seen by the devices is determined by the statistical characteristics of t_j and q_j. So our goal is to characterize as thoroughly as possible the statistical properties of these variables. It is for this reason that we do not study packet or byte counts in fixed-length time intervals. As we will demonstrate, it is not possible to readily determine the statistical properties of the packet variables from such counts.

We focus on empirical study to derive our results about changes in traffic with load, analyzing packet header traces for 6 links. We rely heavily on mathematical statistical models, but they are validated by the data. This empirical study is necessary to establish the results. Theory is important for understanding packet processes, and we will invoke theory, but it is not enough to determine the statistical characteristics of packet processes. Depending on the assumptions made about the source traffic, theoretical arguments lead either to burstiness or smoothness as the load increases; this was recognized in the very beginning in the initial work on long-range dependence of Internet traffic [17].

For each of the two packet variables t_j and q_j, we study the marginal distribution and the time dependence. We study marginal distributions using quantiles (percentiles). We study time dependence using the power spectrum, the Fourier transform of the autocovariance function; for the variable t_j, we first transform so the that the marginal distribution is $N(0, 1)$, that is, normal with mean 0 and variance 1, then we estimate the power spectrum using the transformed values $h_j = h(t_j)$. We study marginal distributions and dependence for many short blocks separately, and then determine how the statistical properties change with the traffic load. If arrivals are Poisson, then the t_j have an exponential marginal distribution, and the $h(t_j)$ are white noise, so the power spectrum is a constant. If the q_j are independent, then its power spectrum is constant as well. Closeness to Poisson arrivals and independent sizes will be of great importance in our study.

The driver of the changing statistical properties is superposition, or statistical multiplexing: an increased intermingling of packets from different transport connections as the number of simultaneous connections increases. For convenience, we measure load over a block of time, not directly by the average number of simultaneous connections, but rather by a surrogate measure, the packet rate over the block.

2. PREVIOUS RESULTS

Long-range dependence and non-Poisson arrivals have been established as important characteristics of Internet traffic under certain circumstances. Specifically, the q_j and t_j can be long-range dependent and the t_j can have a marginal distribution that has longer tails than the exponential. All of these properties, when they occur, contribute to an increased burstiness of packet traffic compared to what it would be if the arrivals were Poisson and the sizes independent. Realistic portrayal of traffic characteristics is important for designing Internet algorithms and protocols.

The discovery of long-range dependence as an element of Internet traffic began with two early articles [17, 20]. The first shows that packet counts and byte counts in time intervals of fixed length on an Ethernet LAN are long-range dependent. The second analyzes many traffic variables including packet inter-arrivals on two links connecting LANs to the rest of the Internet; it shows that the packet inter-arrival times have a marginal distribution whose upper tail is longer than the exponential, and confirms the long-range dependence of packet interval counts.

Since then, other writings such as [9, 12, 21] have corroborated and investigated the long-range dependence and non-exponential behavior with a variety of other Internet measurements. For example, [21] shows that inter-arrival times are long-range dependent and through wavelet modeling, investigates the properties of the dependence. Models of source traffic have been put forward to explain the traffic characteristics [9, 13, 18, 25]. The mathematics can get intricate but the intuition is straightforward. The sizes of transferred files utilizing a link vary immensely; to a good approximation, the upper tail of the file size distribution is Pareto with a shape parameter that is often between 1 and 2. This means there is a nonnegligible probability of the transfer of a file that is considerably larger than the sizes typically transferred. If this happens with high enough frequency for files large enough to have a substantial impact amidst the other aggregate link traffic, then the result is bursty activity. It has been argued that this burstiness increases the lengths of packet queues compared with Poisson arrivals and independent sizes, although there has been some debate about the magnitude and detailed causes of the effect [11, 14, 19, 22].

A critical issue is whether bursty behavior extends across today's entire Internet from lightly loaded LANs to heavily loaded core links. For a heavily loaded link, can the large transfers have sufficient impact to cause burstiness? So far, there has been insufficient empirical work to reach a conclusion. Specifically, there has been insufficient empirical study of the detailed statistical properties of packet inter-arrivals and packet sizes as a function of the load. There have been a number of theoretical discussions about possible implications of increases in load [1, 3, 10, 11]. The load depends on the amount of superposition (statistical multiplexing) of traffic sources, so one can study the implications of more and more superposition. The problem is that results depend heavily on the assumptions about the individual traffic sources being superposed, and theory can produce different burstiness results depending on the assumptions that are made. In fact, all of this was appreciated at the very onset of studies of long-range dependence; a discussion is given by [11], who conclude with the statement: "This is clearly an issue of practical importance, and there is considerable scope for further work."

Recently, the statistical properties of a number of traffic variables, measured on a single Internet link between a large local network and the rest of the Internet, have been related to load, establishing a pervasive nonstationarity in the variables in which fundamental statistical properties, such as shapes of marginal distributions and autocorrelations, change with the load [2, 7]. In this article we focus on packet variables and extend the nonstationary work by studying six Internet links with considerably higher maximum loads than in [2, 7].

3. DATA

We analyze measurements of packet sizes and inter-arrival times for six Internet links that come from three sources.

3.1 Bell Labs Database: One Link

One database on which we draw arises from packet header collection on a 100mb/s Ethernet link connecting a Bell Labs network of about 3000 hosts with the rest of the Internet. For HTTP, all clients are on the inside of the network and all servers are on the outside, so inbound packets are from servers and outbound 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.

For this database, we studied packets just for HTTP. 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. For the analysis presented here, we selected a random sample of 500 blocks constrained so that the log packet rates were within a certain tolerance of being uniformly spaced from the minimum to the maximum log rate. We studied both inbound and outbound packet processes separately, but present just inbound results here; outbound results are similar. The inbound load for the 500 blocks varies from 1.7 packets/sec (p/s) to 452 p/s.

3.2 Harvard Database: One Link

Our second data database is very similar to that arising from the Bell Labs link. The source in this case is a link connecting Harvard University to the rest of the Internet. The link speed is also 100 mb/s and we also study just incoming HTTP packets. The database consists of measurements for 100 blocks, each 60 sec long and collected from 12/21/2000 to 12/22/2000. The packet rate ranges from 101 p/s to 1224 p/s, so the maximum loads are higher than those available in the Bell Labs database.

3.3 NLANR Database: Four Links

Our third database draws on a very large 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. These traces were selected for each site so that the log packet rates would have a reasonable spacing on the log scale.

Each trace is about 80 sec long, and contained packets from two links (or two directions), so altogether we have four links. Unlike the Bell Labs and Harvard databases, we study sizes and inter-arrivals of all packets for the NLANR database. We broke the data for each trace and each link into 5-second blocks. For Colorado, the result is 86 blocks for each link; loads for the first link range from 913 p/s to 13255 p/s, and for the second, from 1012 p/s to 12224 p/s. For Columbia, the result is 87 blocks for each link, with loads ranging from 452 p/s to 6858 p/s and from 759 p/s to 6620 p/s on the two links.

3.4 Database Summary

Altogether, we have 946 blocks of packet sizes and packet inter-arrival times for 6 links. The packet rates $\hat{\rho}_b$ for b = 1 to 946 range from 1.7 p/s to 13255 p/s, which covers a wide range of loads. We will study each of the links separately. But there are two distinct groups of very similar links. One group is Bell Labs and Harvard where the links are 100 mb/s Ethernet, and we study incoming HTTP packets. The second group is NLANR, where the links are OC3 ATM, and we study all packets. Because of the differences in protocols, and the different packet processes studied, we expect somewhat different results for the two groups.

4. RESULTS

For each of the 946 blocks of size and inter-arrival time measurements, we carry out analyses of the marginal distribution and time dependence of t_j and of q_j. For each of the 6 links separately, we then relate the block results to the block packet rates, or loads, $\widehat{\rho}_b$. But our methods of presenting the data also show how the block results depend on load overall for all links collectively.

For low rates, TCP dynamics can affect the statistical properties of inter-arrivals when loads are low. For the Bell Labs blocks, which include many blocks with quite low rates, TCP effects were evident in the estimates of spectra at the lower rates. Most effects were minor. However, 66 blocks revealed significant enough TCP effects to set them aside for their own analysis, although we used the same methods of analysis as for the “standard” blocks. We present results from their analysis separately in Section 4.7. The same overall result of the standard blocks, convergence to Poisson and independence, occurs for these 66 blocks as well, but the pattern of convergence is somewhat different, requiring the separate analysis. Significant TCP effects did not appear for blocks at other links, because the rates are too high for individual TCP connections to show a significant effect.

4.1 Marginal Distributions

We do not report on the marginal distribution of the q_j in detail here because it does not display marked changes with the rate; it does change with the link, however, reflecting the nature of the source traffic. For example, if a link has mostly web servers on the send side and web clients on the other side, then there will be a greater fraction of packets with 1460 bytes of data, and a smaller fraction of ACK packets with 0 bytes of data, than if sender and receiver are reversed.

We used quantile plots [4] to study the structure of the marginal distribution of the t_j, particularly Weibull quantile plots [7]. Let w be a random variable that has a Weibull distribution with shape parameter λ and scale parameter α. Then $(w/\alpha)^\lambda = u$ where u is a unit exponential. So if $\lambda = 1$, w has an exponential distribution. For $\lambda < 1$, the upper tail of w is longer than that of the exponential.

The Weibull quantile plots showed that the Weibull provides an excellent fit to the empirical distribution of the t_j. For the bottom 5% of the data, the values are higher than expected for a Weibull because the minimum inter-arrival time must be greater than the transit time of the packet with the minimum size. But the discrepancy is minor.

We estimated the shape parameter for each block using maximum likelihood [16], and found the estimates ranged across blocks from 0.30 to 1.16. Figure 1 graphs the shape estimates against the log of the block packet rates for each link.

The shape estimates in Figure 1 increase with the log packet rate for each link. For the Columbia and Colorado links the shape estimates are in the vicinity of 1 for the highest log rates; so at the highest rates overall, the packet inter-arrival distribution is well approximated by the exponential.

4.2 Dependence: Power Spectrum

We use the power spectrum to characterize the time dependence — in particular, the long-range dependence — of the sizes and inter-arrivals. Let $c(k)$ be the covariance function of either series at lag k. Then the log power spectrum, on the decibel scale, is

$$10 \log_{10} \left(\sum_{k=-\infty}^{\infty} c(k) \cos 2\pi k f \right),$$

where the frequency f ranges from 0 to 0.5. The units for f are cycles/inter-arrival for the inter-arrivals, and cycles/packet for the sizes.

We first transform t_j for each block to a time series, $h_j = h(t_j)$, whose empirical marginal distribution is $N(0,1)$, and then carry out the spectrum analysis using h_j. Suppose $F_t(u)$ is the sample distribution function of the t_j in a block. Let $G_n(u)$ be the quantile function of a $N(0,1)$ random variable, that is, u is the probability that the random variable is less than or equal to $G_n(u)$. Then the transformation is $h_j = G_n(F_t(t_j))$. The reason for the transformation is the following. The spectrum conveys the behavior of second moments; for a stationary normal time series, the dependence is characterized by second moments, so the power spectrum is a complete description of dependence. For example, if the autocorrelations are all 0, then the series is independent. The transformation is invertible so we can for any needed purpose transform back to the original scale.

For q_j, because the distribution has atoms (accumulations of many observations at specific values), we do not transform to normality; because the deviation from normality is discreteness rather than a highly skewed distribution such as that for q_j, second moments do an adequate job of characterizing dependence. However, just for convenience of interpretation and plotting results, we do adjust the sizes to have sample mean 0 and sample variance 1 by subtracting their mean and then dividing by their sample standard deviation. For q_j, 0 correlations do not strictly speaking imply statistical independence, but we use the term “independence” anyway since we expect near independence when no autocorrelation is present.

In the following we let $H(f)$ and $Q(f)$ denote the log power

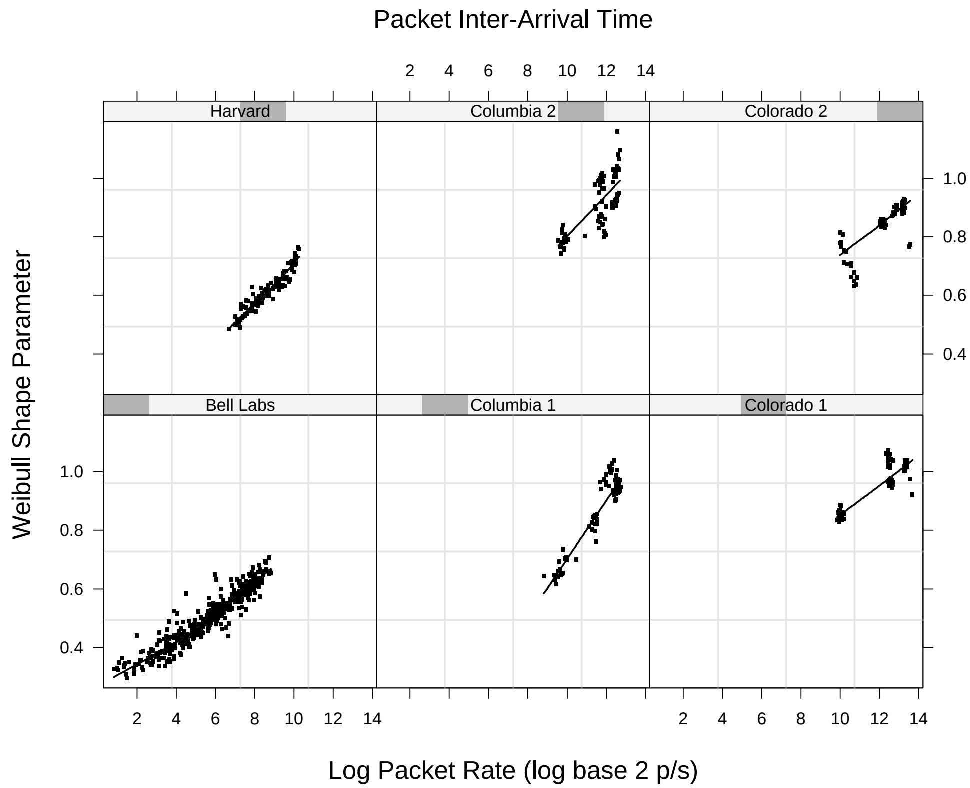

Figure 1: Inter-Arrival: Estimates of the Weibull shape parameter are plotted against the log packet rates for the standard blocks.

spectra of h_j and q_j, respectively.

4.3 Dependence: Spectrum Estimation

For each block b we form an estimate $\widehat{Q}_b(f)$ of $Q_b(f)$ and an estimate $\widehat{H}_b(f)$ of $H_b(f)$. The estimation is carried out by fitting local regression models [5, 6], but we alter the standard procedures to provide for long-range dependence, which results in a rapidly rising log spectrum as f tends to 0. The details are given in Section 4.8. Based on the estimates, we have observed three types of spectra for the 946 blocks — the standard spectrum, and two types of TCP-affected spectra. The TCP-affected spectra occur mostly for packet inter-arrivals at Bell Labs.

Figure 2 graphs $\widehat{H}_b(f)$ against f for 3 standard blocks, all for the Bell Labs link. Figure 3 graphs $\widehat{Q}_b(f)$ in the same way, again for 3 Bell Labs blocks. For both, the packet rates, shown at the top of each panel, increase from low to high in going from left to right. For all estimates, the low frequencies have the largest power and the spectrum decreases monotonically with f. This is a result of persistence in the inter-arrival and sizes — positive autocorrelations. The greater the amount of power at low frequencies, the further out the persistence extends; that is, as a function of lag, the autocorrelations decrease toward 0 more slowly. For both $\widehat{H}_b(f)$ and $\widehat{Q}_b(f)$, as $\widehat{\rho}_b$ increases, the relative amount of low-frequency power decreases; the estimates approach a flat log spectrum, the log spectrum of white noise. At high frequencies, the log spectrum increases toward 0 with $\widehat{\rho}_b$, and at low frequencies it decreases toward 0. This means the h_j and q_j are approaching independence. This description is, of course, qualitative, but shortly we will study quantitative measures of dependence.

For a single TCP connection, packet inter-arrivals can show very regular behavior due to TCP operation; the regular behavior results in peaks in the power spectrum of a single connection. If rates are low, the peaks can still come through in the spectrum estimates of the superposed traffic, but as the rates increase, the intermingling of packets of different connections eventually washes out the regular behavior. Two types of TCP behavior have been observed in the estimates $\widehat{H}_b(f)$. The first results in a minor departure from the behavior of the standard $\widehat{H}_b(f)$; the departure is a small peak at 0.2 cycles/inter-arrival, so we group the blocks with this TCP:0.2 behavior with the standard blocks. The second TCP type, however, a peak at 0.5 cycles/inter-arrival, can result in a major departure of $\widehat{H}_b(f)$ from the standard behavior; by major departure, we mean a peak with a significant amount of power. 66 blocks had such a significant peak at 0.5 cycles/inter-arrival, and we set them aside for their own analysis in Section 4.7.

Figure 4 shows $\widehat{H}_b(f)$ for a Bell Labs block with the TCP:0.2 effect. The estimates have the same overall decrease in power with frequency, but in addition, there is a peak at 0.2 cycles/inter-arrival. The cause is the HTTP 1.0 transfer of small files from one or more servers. The critical aspect of these small files is that they require only one packet to transmit because their size is less than the maximum segment size which is typically 1460 bytes. These small files are, for examples, “not-modified” or “not-found” messages resulting from cache validation requests, small GIF files, and some dynamically generated web pages. Also, the client request file is almost always contained within a single packet. Suppose we have a sequence of such transfers. In a typical TCP implementation, each tends to produce, from the server, 5 packets: (1) SYN/ACK in three-way hand shake; (2) ACK of client HTTP request packet; (3)

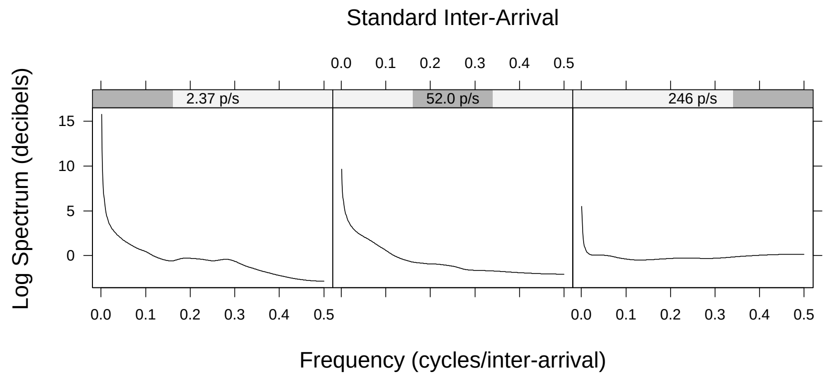

Figure 2: Inter-Arrival: Log power spectrum estimates for three standard blocks.

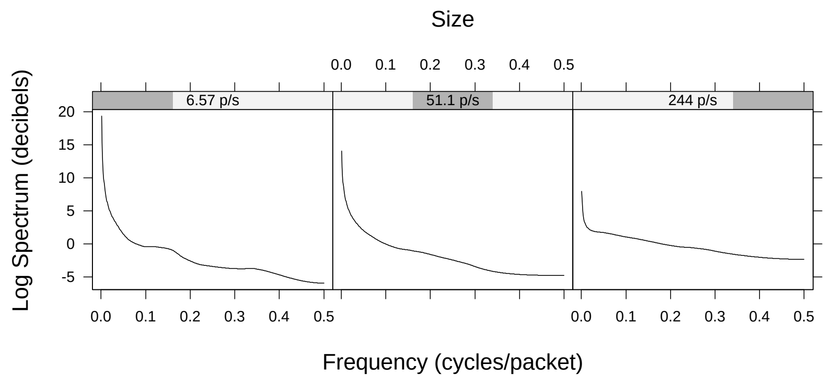

Figure 3: Size: Log power spectrum estimates for three standard blocks.

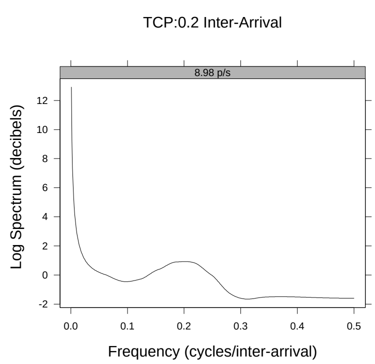

Figure 4: Inter-Arrival: Log power spectrum estimate for a TCP:0.2 block.

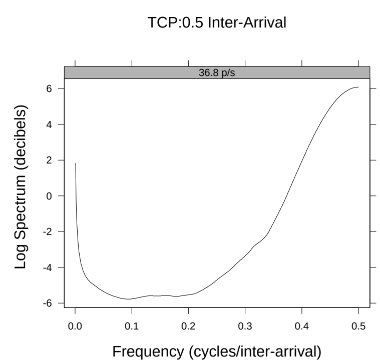

Figure 5: Inter-Arrival: Log power spectrum estimate for a TCP:0.5 block.

data packet for the small file (4) FIN; and (5) ACK of client FIN. The inter-arrival times between packets 1-2 and 4-5 are typically greater than or equal to the round-trip time, and between 2-3, and 3-4 are often close to 0, relative to the round-trip time. Also, the time between packet 5 of one connection and packet 1 of the next connection can be short or long. The exact nature of these packets and times may depend on the particular TCP implementation. But whatever they are, if the time patterns are consistent, then the result is a sequence of groups of inter-arrivals each with a similar pattern that will tend to induce a peak in the spectrum. For this Bell Labs block, the groups of 5 inter-arrivals from the 5-packet groups above give a peak at 1 cycle per 5 inter-arrivals, or 0.2 cycles/inter-arrival. There is ample opportunity for such regular sequences because on the Bell Labs network, about 60% of all HTTP connections are these small file transfers. However, while such peaks in the spectrum occur occasionally for low rates, the magnitude of the peaks is quite small, and the behavior of the spectrum is otherwise just like that of the standard behavior.

Figure 5 shows $\widehat{H}_b(f)$ for a Bell Labs block with the TCP:0.5 effect. If one or more extremely large files are transferred, and the rate is otherwise quite low, then TCP will tend to send data packets in pairs, separated by a relatively longer time interval, so the result is an alternating sequence of short and longer inter-arrivals which creates a peak at the maximum frequency, 1 cycle per 2 inter-arrivals, or 0.5 cycles/inter-arrival. We screen for this behavior by selecting blocks for which the value of $\widehat{H}_b(0.5)$ minus the minimum of $\widehat{H}_b(f)$ is greater than 4 decibels. The result is 66 blocks.

4.4 Dependence: Entropy

A standard measure of the amount of randomness in a time series, Z_j, with power spectrum $p(f)$, is the entropy, the geometric mean of $p(f)$. The log entropy, on the decibel scale, is

$$2\int_0^{0.5} 10\log_{10}(p(f))\,df. \tag{1}$$

The entropy is the variance of the error of linear prediction of Z_j from the infinite past: $Z_{j-1}, Z_{j-2}, \ldots$. The entropy is less than or equal to the variance of Z_j. If it is equal to the variance, the series is uncorrelated (or independent if the series is normally distributed). If the entropy is 0, the series has so much dependence that it is perfectly predictable from the infinite past. Both h_j and q_j have variance 1, so their entropy is less than or equal to 1, and their log entropy is less than or equal to 0.

For each block b, we computed log entropy estimates for h_j and q_j from Equation 1 with $\widehat{H}_b(f)$ or $\widehat{Q}_b(f)$ in place of $10\log_{10}(p(f))$. Figure 6 graphs the entropy estimates against $\log_2(\widehat{\rho}_b)$ for the standard blocks and for h_j; Figure 7 makes the same plot for q_j. For both packet variables, the log entropy increases substantially with load, so in each case, dependence is substantially reduced. For the Columbia, Colorado, and Harvard links, the log entropy estimates for both series are in the vicinity of 0 for the highest loads, so h_j and q_j are close to independent, at least as measured by the entropy.

4.5 Dependence: Maxima and Minima

Figures 2 and 3, plots of $\widehat{H}_b(f)$ and $\widehat{Q}_b(f)$ against f, show that as the load increases, the power at low frequencies decreases toward 0 decibels, the power at high frequencies increases toward 0 decibels, and the estimates for each packet variable tend to the log power spectrum of white noise, which is 0 at all frequencies. To quantify this further, we study the maximum and minimum power for each estimate over a grid of frequencies.

We computed $\widehat{H}_b(f)$ and $\widehat{Q}_b(f)$ for each block at 500 equally spaced frequencies from 0.001 to 0.5, and then found the maximum and minimum of each estimate. The maxima of $\widehat{H}_b(f)$ are plotted against $\widehat{\rho}_b$ in Figure 8, and the maxima of $\widehat{Q}_b(f)$, in Figure 9. Similarly, the minima are plotted against $\widehat{\rho}_b$ in Figures 10 and 11. The maxima decrease and the minima increase substantially for some links, but less so for others where the spectrum estimates are already closer to flat at the lowest loads.

4.6 Dependence: An LRD Plus Noise Model

The behavior of the spectrum estimates, specifically, the increase in the minimum with load, suggests a component model for sizes and inter-arrivals that provides insight into the change in dependence with load.

Let Z_{bj} be either the h_j or q_j series for block b. Z_{bj} has mean 0 and variance 1. Let $p_b(f)$ be the power spectrum. Then the model is

$$Z_{bj} = \sqrt{1-\theta_b}S_{bj} + \sqrt{\theta_b}N_{bj},$$

where N_{bj} is a white noise series with mean 0 and variance 1, S_j is a long-range dependent series with mean 0 and variance 1, and θ_b is the minimum of $p_b(f)$, so

$$0 \le \theta_b \le 1.$$

The power spectrum of $\sqrt{\theta_b}N_{bj}$ is flat and equal to θ_b, and the power spectrum of $\sqrt{1-\theta_b}S_{bj}$ is equal to

$$p_b(f) - \theta_b.$$

An estimate, $\widehat{\theta}_{(h)b}$, of θ_b for h_j in block b is provided by the minimum of the spectrum estimate, $\widehat{H}_b(f)$; similarly, for q_j, an estimate, $\widehat{\theta}_{(q)b}$, is provided by the minimum of $\widehat{Q}_b(f)$. Thus Figures 10 and 11 graph $10\log_{10}(\widehat{\theta}_{(h)b})$ and $10\log_{10}(\widehat{\theta}_{(q)b})$ against load. As we saw, the estimates tend to 0 decibels.

An interesting question is whether we can, to a good approximation, take the statistical properties of S_{bj} to be the same for all b. The spectrum of S_{bj} is

$$\frac{p_b(f) - \theta_b}{1-\theta_b},$$

so the question is whether this spectrum is approximately the same for all b. To check this, we plotted estimates of this spectrum for h_j and q_j in all blocks against f; the estimates are

$$\frac{10^{0.1\widehat{H}_b(f)} - \widehat{\theta}_b}{1-\widehat{\theta}_b}$$

and

$$\frac{10^{0.1\widehat{Q}_b(f)} - \widehat{\theta}_b}{1-\widehat{\theta}_b}.$$

For each link separately we found the estimates for h_j to be similar to one another across b; the same was true for q_j, so for each link it appears reasonable to suppose that S_{bj} has the same statistics across b for h_j as well as for q_j.

The component model is a simple model that allows easy interpretation of the long-range dependence of h_j and of s_j, and how it

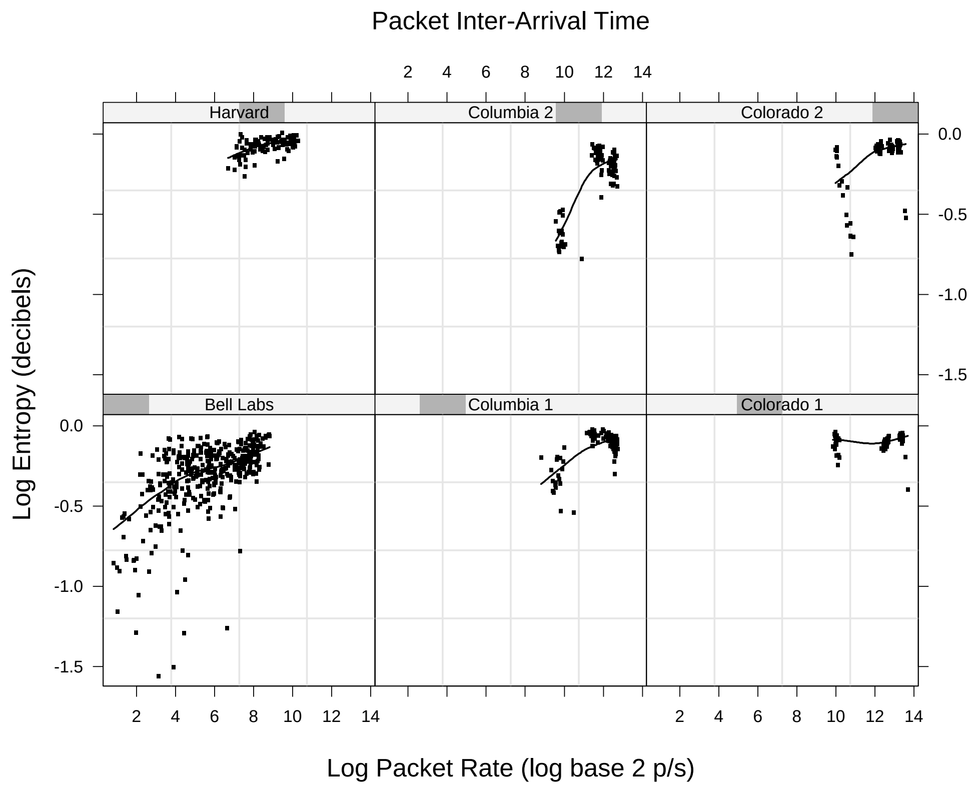

Figure 6: Inter-Arrival: Estimates of log entropy are plotted against the log packet rates for the standard blocks.

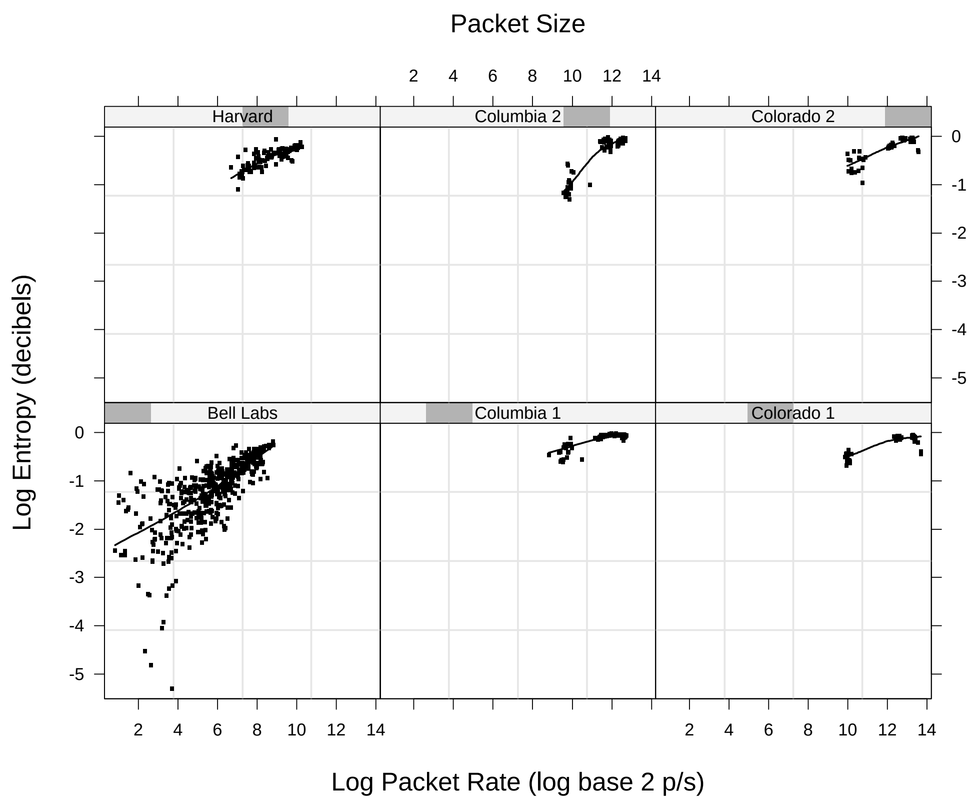

Figure 7: Size: Estimates of log entropy are plotted against the log packet rates for all blocks.

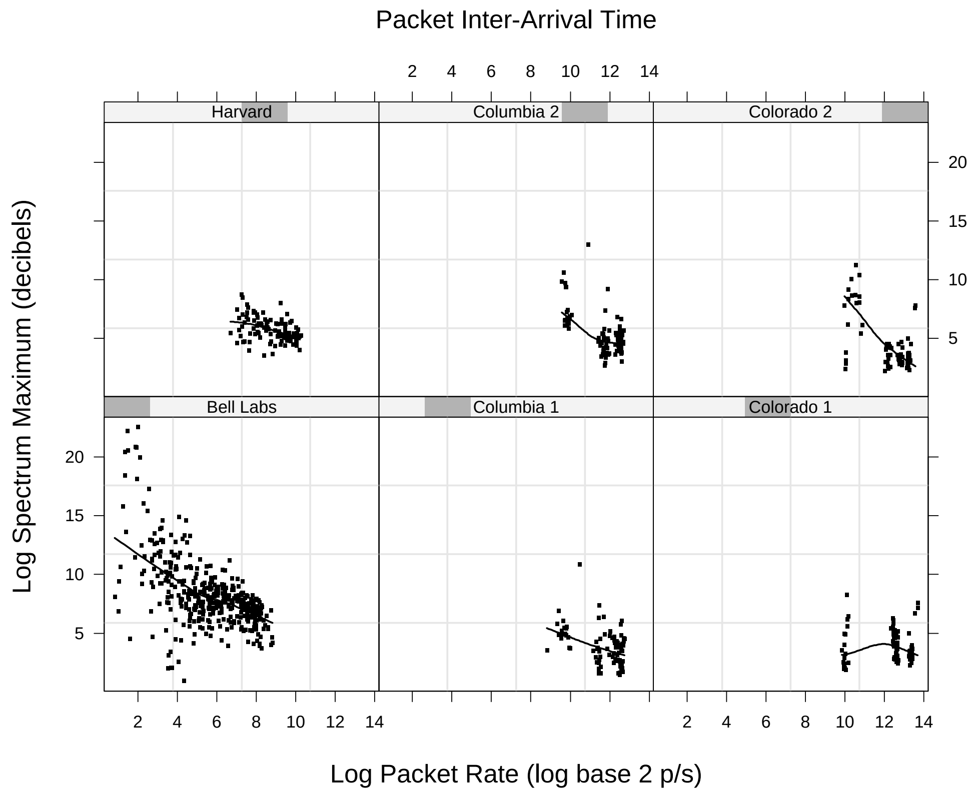

Figure 8: Inter-Arrival: The maxima of the log spectrum estimates are plotted against the log packet rates for the standard blocks.

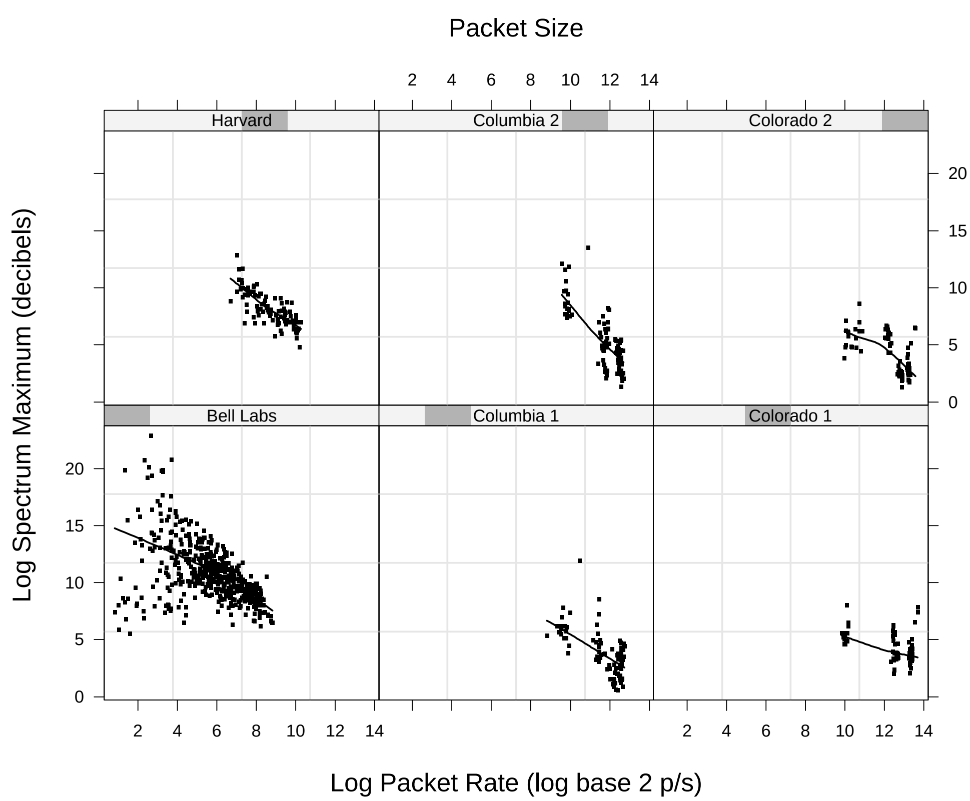

Figure 9: Size: The maxima of the log spectrum estimates are plotted against the log packet rates for all blocks.

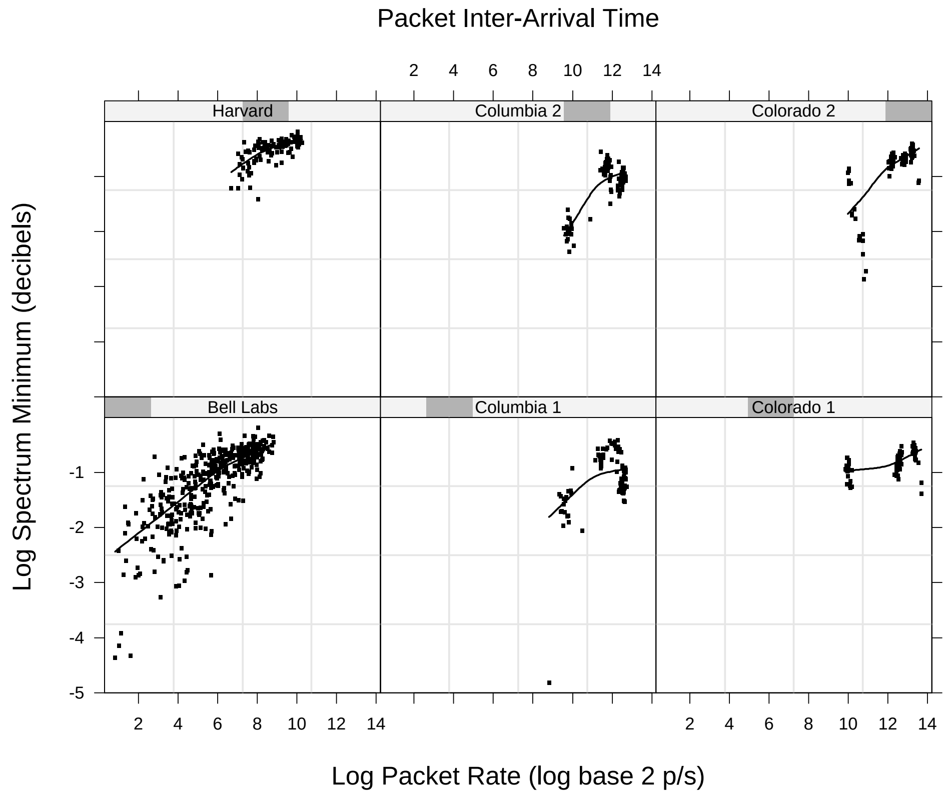

Figure 10: Inter-Arrival: The minima of the log spectrum estimates are plotted against the log packet rates for the standard blocks.

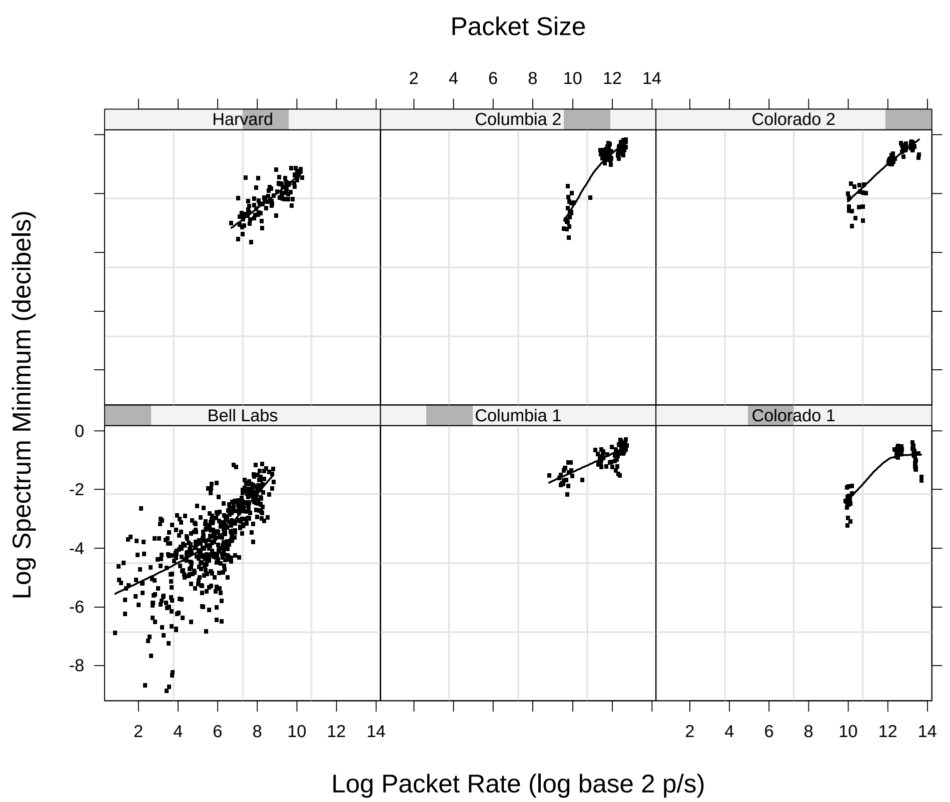

Figure 11: Size: The minima of the log spectrum estimates are plotted against the log packet rates for all blocks.

changes with load. For each link, each series is a mixture of a fixed long-range dependent series plus noise. As the load increases, the variance of the noise goes to 1, and the variance of the long-range dependent series goes to 0. One important implication is that the spectrum very close to the origin is always the same. But the behavior at the origin determines the Hurst parameter H [24]. This suggests that as the load increases, H remains constant, or at least close to it, even though the long-range dependence decreases. This says that H is not by itself an effective measure of long-range dependence, and that a measure in the spirit of θ is needed as well.

4.7 Dependence: TCP:0.5 Blocks

The results for the standard blocks just shown, also hold for the inter-arrivals of the 66 TCP:0.5 Bell Labs blocks. Figure 12 graphs estimates of the Weibull shape parameter against $\widehat{\rho}_b$; the shape increases toward 1 as for the standard blocks. Figure 13 graphs entropy on the decibel scale against $\widehat{\rho}_b$; the entropy tends to 0 as for the standard blocks. Figure 14 graphs the maximum spectrum estimate against $\widehat{\rho}_b$, and Figure 15 graphs the minimum against $\widehat{\rho}_b$; the spectrum tends to that of white noise as for the standard blocks.

4.8 Methods

A quantile of order u of a distribution, where $0 \leq u \leq 1$, is a number such that a fraction u of the distribution is less than or equal to (or approximately less than or equal to) the number. Suppose we have a set of n measurements x_i where x_1 is the smallest, x_2 is the next smallest and so forth. Then we take x_i to be the quantile of order $(i - 0.5)/n$. To determine whether the empirical distribution of the data is well approximated by a theoretical distribution such as the exponential, we plot x_i against the quantile of order $(i - 0.5)/n$ of the theoretical, and check to see if the pattern is linear. This quantile plot is an extremely powerful tool because it shows all of the information in the empirical distribution and allows us to effectively judge how well the theoretical distribution approximates the empirical, and shows where the problem lies when it does not. Quantile plots were used heavily to study the empirical distributions of t_j and q_j.

We developed a model to estimate the power spectra of the packet variables h_j and q_j in the presence of long-range dependence; for the size and inter-arrival processes, the long-range dependence causes rapid change in the spectrum near $f = 0$. The method is based on smoothing the periodogram. Suppose the number of observations of a packet variable is n. The following steps are carried out: (1) compute the periodogram $I(f)$ at the Fourier frequencies $f_k = k/n$, for $0 < f_k \leq 0.5$; (2) average the periodogram in non-overlapping blocks of size 5 (dropping the last block if there are fewer than 5 in the average), and average the frequencies in the same way, which results in averaged periodogram values $\bar{I}(\bar{f}_k)$ at m equally-spaced averaged frequencies $\bar{f}_k$, for k = 1 to m; (3) take log base 10 and multiply by 10 (so that units are in decibels), and then correct for expected bias by adding 0.448, yielding $L(\bar{f}_k)$; (4) smooth the $L(\bar{f}_k)$ using loess to form the log power spectrum estimate at each desired f. The loess estimate at f results from fitting the $L(\bar{f}_k)$ for f_k close to f by a parametric function, and evaluating the fit at f. It is important in using loess to choose a parametric family that approximates the true spectrum locally was well as possible. FARIMA models have an additive term in the log spectrum consisting of a constant times

$$g(u) = \log_{10}(4 \sin^2(\pi u)),$$

which introduces substantial curvature in the log spectrum near the origin, but is smooth elsewhere [15]. Since our packet spectra exhibit this behavior, the $L(\bar{f}_k)$ are smoothed by loess as a function of $g(\bar{f}_k)$, using a quadratic as the local regression function. The fit, though, is then plotted against $\bar{f}_k$. This results in a much better fit than smoothing directly as a quadratic function of $\bar{f}_k$. The loess smoothing parameter was 0.5, chosen by studying the residuals of the fitting; this smoothing parameter was the largest that did not introduce unacceptable distortion of the spectrum.

5. FIXED-LENGTH COUNTS

It has been common in the analysis of Internet traffic data to study packet counts and byte counts in fixed-length intervals [23]. Such counts would not readily reveal our results, which are based on study of the packet variables t_j and s_j. More generally, studying fixed-length counts is not sufficient to fully understand packet processes, which are defined by the packets sizes and inter-arrivals. We will address the general issue, and then discuss one specific technical issue regarding our results.

5.1 General

Network devices process packets as they arrive, and service times depend on packet size. The packet variables discussed here, t_j and s_j, are the defining variables of packet processes. They characterize packet processes in terms of what devices do with packets. Byte counts and packet counts in fixed-length intervals do not correspond to something typically done by network devices, which do not wait for a fixed length of time and then process what is there.

Byte counts and packet counts in a sequence of fixed-length intervals can be derived from t_j and s_j, but we cannot get t_j and s_j from a sequence of byte and packet counts except in the degenerate case where the intervals are infinitesimally small, and the counts are 0 or 1 in each interval, which would make the counts equivalent to t_j and s_j. This does not render byte and packet counts useless because it is possible to get some insight into the behavior of t_j and s_j from packet and byte counts. But if what we want is a fundamental modeling and description of packet processes, we will in the end need to study t_j and s_j for a full understanding.

5.2 Inter-Arrivals and Load

Counts would not readily reveal the marginal distribution of t_j and its change with load. It is simple and straightforward to study inter-arrival measurements directly, but surely not simple to infer their distribution through the filter of fixed-interval packet counts.

Counts would not reveal the change in long-range dependence of the inter-arrivals with load. We need some mild assumptions to argue this. Suppose the arrivals, $a_{(n)j}$, of packets on a wire are the n-fold superposition of n source arrival processes a_{kj} for k = 1 to n, where the source processes are stationary and generated independently of one another from the same population point process. Independence of the sources is quite a reasonable assumption if the link devices, or any other devices that carry a large fraction of the traffic on the link in question, do not experience more than minor congestion. We can think of the traffic sources as individual users or as groups of individual users.

Suppose we count arrivals in fixed-length time intervals. Let N_{ki} be the counts that arise from the k-th source process in interval i, and let $N_{(n)i}$ be the counts in interval i for the n superposed processes. Let μ be the mean of N_{ki}.

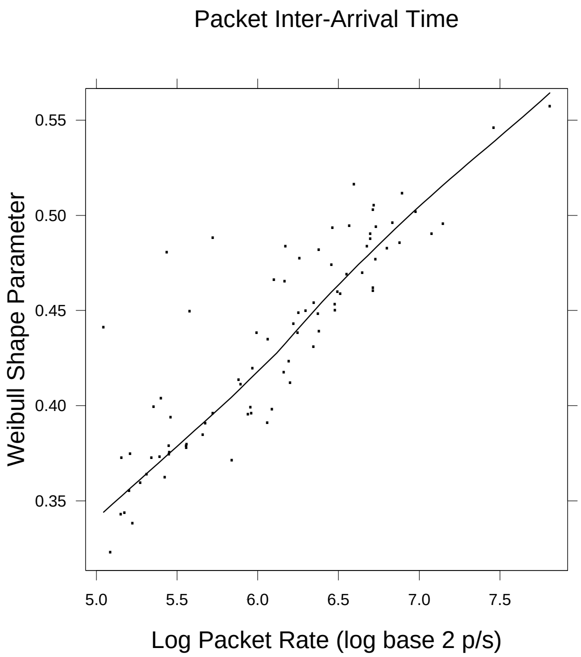

Figure 12: Inter-Arrival: Estimates of the Weibull shape parameter are plotted against the log packet rates for the TCP:0.5 blocks.

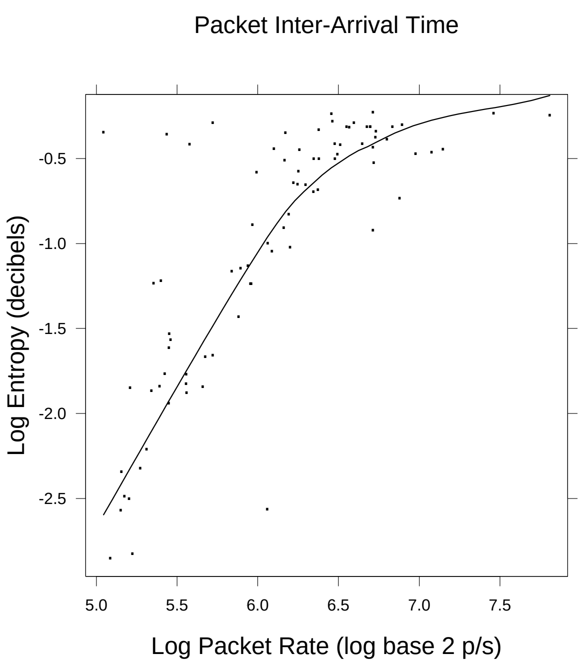

Figure 13: Inter-Arrival: Estimates of log entropy are plotted against the log packet rates for the TCP:0.5 blocks.

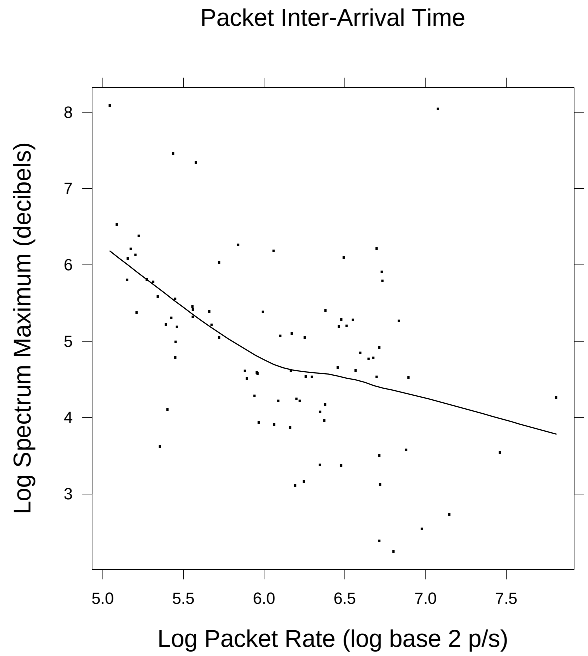

Figure 14: Inter-Arrival: The maxima of the log spectrum estimates are plotted against the log packet rates for the TCP:0.5 blocks.

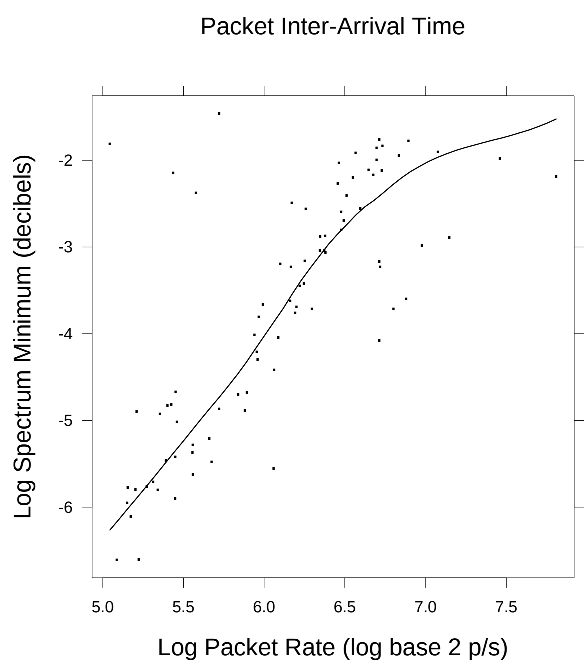

Figure 15: Inter-Arrival: The minima of the log spectrum estimates are plotted against the log packet rates for the TCP-0.5 blocks.

The autocovariance of N_{ki} at lag r is

$$c_r = \frac{\mathrm{E}(N_{k(i+r)} - \mu)(N_{ki} - \mu)}{\mathrm{E}(N_{ki} - \mu)^2}.$$

Now

$$N_{(n)i} = \sum_{k=1}^{n} N_{ki}$$

and the N_{ki} are independent across k, so the autocovariance of $N_{(n)i}$ at lag r is also c_r.

Let $t_{(n)j}$ be the inter-arrival times of the n-fold superposition process. The marginal distribution and dependence structure of $t_{(n)j}$ can change with n and does so with a certain regularity. In fact, suppose that the source processes were renewal processes, which would mean that the inter-arrival times for each would be independent. From standard results in the superposition of renewal processes $a_{(n)j}$ would tend to Poisson [8] with n. This means that the $t_{(n)j}$ would tend to exponential and independent as n gets large. Of course, source traffic packet inter-arrivals are not independent, because, as we have seen, inter-arrivals are dependent at low loads. But this standard result for renewal processes can be expected to hold even when the inter-arrivals are dependent. In fact, we have demonstrated this empirically in our paper. But, as we have just demonstrated, the packet count process in fixed-length intervals cannot even hint at the result, provided our mild assumptions are true.

This might seem like a contradiction because it suggests that for renewal processes we have two results, one that suggests a Poisson limit, and one that suggests a dependent point process from the counts. The resolution is that the counts in each interval grow without bound with n, but the standard renewal results are local and apply to a fixed number of arrivals. We would argue that for Internet devices, which process packets one by one, and not in blocks that grow without bound, the fixed number of arrivals are the right conceptual framework. This is borne out by queueing results which show that open-loop queueing of empirical packet traces tends to that of Poisson with independent service times [2].

6. SUMMARY

The statistical properties of 946 blocks of two packet variables — inter-arrival time and size — are studied for 6 links: two 100mb/s Ethernet links and four OC3 ATM links. The packet rates range from 1.7 packets/sec to 13255 packets/sec.

At low rates, the inter-arrival times have marginal distributions that are well-approximated by Weibull distributions with longer upper tails than the exponential, and the inter-arrivals and sizes are long-range dependent. For each link, as the packet rate increases, the Weibull marginal distribution of the inter-arrivals tends to an exponential, and the dependence of each of the two packet variables decreases. While dependence appears in the variables even at the higher rates of each link, the amount of dependence is very small. Thus at the higher rates, the packet inter-arrivals are nearly a Poisson process, and the packet sizes are nearly independent. Thus packet queueing at higher rates can be expected to behave like that of Poisson arrivals with independent service times.

The inter-arrival marginal distribution was studied, in detail, by quantile plots. Having established through the plots that the Weibull was a good approximation, the shape parameter was estimated for each block by maximum likelihood. Then the shapes were related to the packet rates to establish that the marginal distribution converged to the exponential as the rate increased.

The time dependence of the inter-arrivals and sizes was studied by estimating the power spectrum of each for each of the 946 blocks. A method of spectrum estimation based on a local fitting model was developed to ensure that the estimate did not distort the extreme curvature near the origin that occurs due to long-range dependence. Checks for distortion were carried out by graphing residuals against frequency. Dependence was characterized qualitatively by graphing log spectrum estimates against frequency. Dependence was characterized quantitatively by three measures based on the log spectrum estimates: entropy, maximum, and minimum. A simple LRD-plus-noise component model is developed that provides easy interpretation of the long-range dependence and how it changes with load.

7. REFERENCES

[1] 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.

[2] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun. On the Nonstationarity of Internet Traffic. Technical report, Bell Labs, Murray Hill, NJ, in preparation.

[3] G. L. Choudury. Squeezing the Most Out of ATM. *IEEE Transactions on Communications*, 44(2):203–217, 1996.

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

[5] W. S. Cleveland and S. J. Devlin. Locally-Weighted Regression: An Approach to Regression Analysis by Local Fitting. *Journal of the American Statistical Association*, 83: 596–610, 1988.

[6] W. S. Cleveland, E. Grosse, and M. J. Shyu. Local Regression Models. In J. M. Chambers and T. Hastie, editors, *Statistical Models in S*, pages 309–376. Chapman and Hall, New York, 1992.

[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 2000*, pages 166–177. ACM, 2000.

[8] D. R. Cox. *Renewal Theory*. Chapman and Hall, 1962.

[9] M. E. Crovella and A. Bestavros. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. In *Proc. ACM SIGMETRICS '96*, pages 160–169, 1996.

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

[11] 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.

[12] A. Feldman, A. A. Gilbert, and W. Willinger. Data Networks as Cascades: Explaining the Mulifractal Nature of Internet WAN Traffic. In *Proc. ACM SIGCOMM '98*, pages 42–55, 1998.

[13] S. Floyd and V. Paxson. Why We Don't Know How to Simulate the Internet. Technical report, LBL Network Research Group, 1999.

[14] D. P. Heyman and T. V. Lakshman. What Are the Implications of Long-Range Dependence for VBR-Video Traffic Engineering? *IEEE/ACM Transactions on Networking*, 4:301–317, 1996.

[15] J. R. M. Hosking. Fractional Differencing. *Biometrika*, 68:165–176, 1981.

[16] N. L. Johnson and S. Kotz. *Distributions in Statistics: Continuous Univariate Distributions*. Houghton Mifflin Company, Boston, 1970.

[17] 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.

[18] 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.

[19] K. Park, G. Kim, and M. Crovella. On the effect of traffic self-similarity on network performance. In *Proc. SPIE Intl. Conf. Perf. and Control of Network Systems*, 1997.

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

[21] 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.

[22] B. K. Ryu and A. Elwalid. The Importance of Long-Range Dependence of VBR Traffic in ATM Traffic Engineering: Myths and Realities. In *Proc. ACM SIGCOMM '96*, pages 3–14, 1996.

[23] W. Willinger, M. S. Taqqu, and A. Erramilli. A Bibliographic Guide to Self-Similar Traffic and Performance for Modern High-Speed Networks. In *Stochastic Networks: Theory and Applications*, pages 339–366, 1996.

[24] 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.

[25] 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.