A Compound Model for TCP Connection Arrivals for LAN and WAN Applications

Carl Nuzman Iraj Saniee Wim Sweldens Alan Weiss
Bell Laboratories, Lucent Technologies
600 Mountain Ave., Murray Hill, NJ 07974 USA

Abstract

We propose a two level model for TCP connection arrivals in local area networks (LAN). The £rst level are user sessions whose arrival is time-varying Poisson. The second level are connections within a user session. Their number and mean interarrival are independent and biPareto across user session. The interarrivals within a user session are Weibull, and across all users are correlated Weibull. Our model has a small number of parameters which are inferred from real traf£c collected at a £rewall. We show that traf£c synthesized with our model closely matches the original data. We extend this approach to a general model involving shot noise and show it is asymptotically consistent with more common fractal models used in data networks. Finally, we show that this model extends to the wide area network applications without alteration and it predicts smoothing of wide area network traf£c pro£le due to spatial aggregation, which we observe experimentally by synthetically creating large aggregate TCP load.

1 Introduction

The development of data networking in the 1960s, £rst as an academic exercise, and subsequently as a means of providing a host of new services, presents new challenges and opportunities to the well-established £eld of teletraf£c theory. Apart from a few pioneering investigations, such as Mandelbrot [20], some of the features that distinguish data traf£c from voice telephony were noticed as early as the late 1980s (Fowler and Leland [13], Meier-Hellstern et al. [21]). In their

pioneering study of LAN traf£c, Willinger et al. [17] presented data and argued for use of alternative models. By showing that suf£ciently aggregated data traf£c exhibited self-similarity over a wide range of time scales, the authors argued for the use of fractal models and, more explicitly, use of statistical processes exhibiting long-range dependence (LRD). In subsequent reports [26], the authors contended that the underlying cause of self-similarity was effectively unrelated to the mechanisms of data transmission, and were exclusively due to the nature of aggregated load. Unlike voice, individual streams of load in data networks followed distributions with heavy tails, the aggregation of which, it was argued, gave rise to its observed self-similarity. Cox [5, 4] had given a related theoretical model for long-range dependence, namely an $M/G/\infty$ queue with heavy tailed service times.

Many of the succeeding investigations (for Web data notably Crovella [6] and Paxson [22]) have followed the same general methodology. This consists of counting bytes and packets carried over identical non-overlapping time intervals and studying the general behavior of this process as the aggregation interval is increased both spatially and temporally. Some studies of the impact of self-similarity on the performance of switching systems have argued in favor of the self-similar approach [9], while others have argued against it [7, 25]. The latter studies argued that the short-term correlation in data traf£c could capture much, if not all, of the performance impacts attributed to self-similarity. Followup work in this direction has concentrated on demonstration and causes of multi-fractality in low aggregate Web traf£c [11] and potential performance impacts of multi-fractality [8].

TCP connection arrival processes from various data sets were classi£ed according to source application in [22]. Here it was observed that protocols such as TELNET that tend to generate a single connection per "user session" are well modeled by Poisson processes, whereas applications such as HTTP and X11 which may generate multiple connections per "user session" are not. The authors conjectured that if the underlying user sessions could be observed, such sessions would follow a Poisson arrival process. The concept of user sessions ("user equivelents") is also central to the SURGE model used to simulate the workload of HTTP servers [1].

In this report, we use high-resolution packet data, collected at one Bell Labs facility, to provide evidence for the correctness of the user session concept. We also show that this concept is useful for analyzing the traf£c and that it can be used to develop simple, accurate, and ef£cient models for TCP connection arrivals. In this analysis phase, we group TCP connections into user sessions, and then measure

distributions of measurements such as the number of connections per user session. We found that the user sessions arrivals followed a Poisson process, while per session measurements tend to be heavy tailed. We are then able to generate synthetic traf£c traces from a simple, two layer model, using a small number of parameters determined by the measurements in the analysis phase. After empirically verifying the accuracy of the model, we examine some of its theoretical properties. We show that it generalizes the $M/G/\infty$ queue in the shot noise framework, and demonstrate the way in which long-range dependence can arise.

In contrast to the count processes used extensively in the study of self-similar and multi-fractal phenomena, we look at the *arrival processes* at the user sesion and connection layers. The arrival processes can easily be converted to count processes as needed. The data we studied and the resulting model agree with the observation made in [10], that the interarrival times for TCP connections are subexponential and well-modeled by Weibull distributions (when signi£cant HTTP traf£c is present). In addition, we observe that the sequence of interarrival times is long-range dependent, in the data and in our model. As reported extensively in [3], the relative amplitude of this long-range dependence diminishes as traf£c is aggregated spatially; this is also true in the model we present here.

Our contributions consist of

- Inferring user session groupings from TCP packet headers, and verifying the appropriateness of a Poisson model for these sessions (Section 2).
- Providing a model of TCP connection arrivals within a user session, using a newly de£ned biPareto distribution (Section 2).
- Verifying that this compound model agrees with measurement (Section 3).
- Showing the model extends to characteristics of large-scale aggregation observed in traf£c in the network core (Section 3).
- Placing the model in a general shot noise framework (Section 4).

2 Model-based Measurements

2.1 Data Set

Our data was collected by a system at Bell Laboratories in Murray Hill, New Jersey [3]. As pictured in Figure 1, the math and computer science research center

at Bell Labs is connected to the Internet through a £rewall. The £rewall connects to the center via a router; the data was collected from the link between this router and the £rewall. A packet sniffer recorded all the packet headers. Our data comes from a 27-hour period beginning the evening of November 17, 1998.

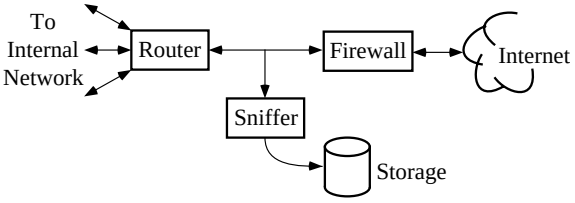

Figure 1: The data collection system

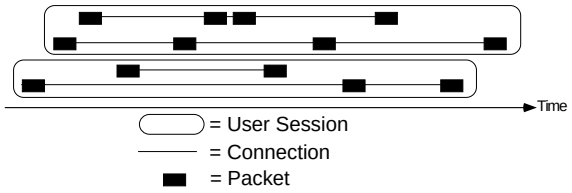

Figure 2: Sketch of user sessions, TCP connections, and packets.

We analyzed two subsets of the data:

1. All the TCP data. This amounted to 72% of the packets; almost all the rest were UDP.

2. HTTP data between 10 AM and 5 PM. This allows us to assume a stationary model and focus on a single type of traf£c. The majority of TCP packets in the trace are HTTP.

Note that the data represents carried traf£c (closed loop) while we model offered traf£c (open loop), ignoring the effects of vagaries in the Internet. Our justi£cation for this is that the £rewall, at least, was lightly loaded, so that any congestion was due to other links in the network. Also, it is very dif£cult to infer the effect of the TCP control algorithm on traf£c.

The model we develop in this paper has a Poisson process of user session arrivals at the top level. Based on this model, the important quantities that must be measured or estimated include the arrival rate λ of user sessions, the joint distribution of user session duration and rate, and the mechanism by which each user session generates connections. Each of these measurements are discussed in turn below.

2.2 De£nition and Characterization of User Sessions

The grouping of IP packets into TCP connections is determined by the TCP protocol. However, the way in which TCP connections should be combined into user

sessions is a matter of choice. Our de£nition of user session was based only on information contained in the TCP and IP packet headers. For each TCP connection, the host IP address initiating the connection was de£ned to be the user for that connection. The TCP connections corresponding to a given user were then grouped into user sessions as follows:

- A user session begins when a user who has been idle opens a TCP connection.

- A user session ends, and the next idle period begins, when the user has had no open connections for τ consecutive seconds.

This method, of course, has the problem that different users sharing the same machine cannot be distinguished, and this can lead to extremely long "user" sessions on shared machines.

We used $\tau = 100$ seconds, so that periods of HTTP activity separated by more than a couple of minutes were considered to be separate user sessions. We also tried, though don't report in detail, values of τ higher than 100 seconds, including 10 minutes and one hour. The qualitative results of the study do not depend on the particular choice of τ, although some quantitative factors such as the arrival rate of user sessions do vary with τ.

Using the above de£nition for user sessions, we determined that the user session arrival process is consistent with a Poisson process with a very slow (diurnal) variation in arrival rate. The arrivals came faster during the day than at night, but the rate was approximately constant over periods of a few hours. We found that the interarival times were exponentially distributed, and that various transformations of the interarrival process were uncorrelated.

2.3 Structure of Connection Arrivals

In the rest of work that will be reported here, we focused on approximately 4000 HTTP user sessions initiated between 9 AM and 5 PM. We split our model of connection arrival within a user session into two parts. This was done based on the data, not based on expectations we had of our model. In any user session we de£ne C as the number of TCP connections that are part of that user session. If $C > 1$ we de£ne two additional statistics. T is the total interarrival time; it is the time from the £rst connection arrival to the last connection arrival. The *average connection interarrival time* is de£ned to be $\mu = T/(C-1)$. Note that the timing

of connection arrivals is not addressed by C and μ; this will be addressed in Section 2.6.

2.4 Model Statistics for Connection Interarrivals

Figure 3 shows the tail distribution function of C, the number of connections generated by a user session. Figure 4 shows the tail distribution of μ. Note that on these log-log plots, a distribution of the form $1-F(x) = x^{-\alpha}$ would appear as a straight line with slope $-\alpha$. In Figure 4, there are two nearly linear regimes, with a smooth transition between them. In Figure 3, there appear to be two regimes, although the regime on the left is never linear.

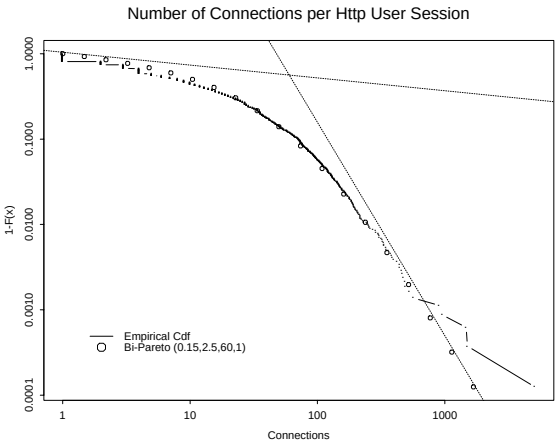

Figure 3: Empirical tail distribution and biPareto £t of C, the number of connections per user session.

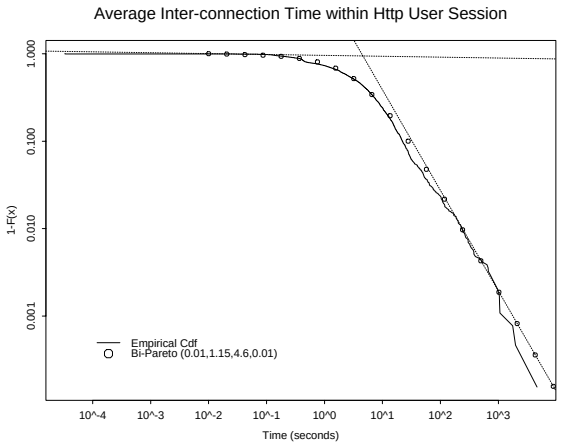

Figure 4: Empirical tail distribution and biPareto £t of μ, the mean connection interarrival time within a user session.

If the entire plot were almost linear, the distribution could be modeled by a Pareto distribution, for which the tail distribution function is given by

$$1-F(x) = \begin{cases} \left(\frac{x}{k}\right)^{-\alpha} & x > k \\ 1 & x \leq k \end{cases}$$

where the parameters $\alpha > 0$ and $k > 0$ are the decay exponent and scale parameter, respectively. The scale parameter is also the minimum possible value of the random variable.

To give the ¤exibility needed to model empirical distributions such as that in Figure 4 effectively, we de£ned the *biPareto* distribution, which has a tail distribution function given by

$$1-F(x)=\begin{cases}\left(\frac{x}{k}\right)^{-\alpha}\left(\frac{x+kb}{k+kb}\right)^{\alpha-\beta} & x>k\\ 1 & x\le k\end{cases}$$

The minimum possible value of such a random variable is the scale parameter $k>0$. From this point, the tail distribution initially decays as a power law with exponent $\alpha>0$. Then, in the vicinity of a breakpoint kb (with $b>0$), the decay exponent gradually changes to $\beta>0$. Some further properties of this distribution are given in Appendix B

The circles in Figures 3 and 4 show the tail distribution function of biPareto distributions matched to the empirical distributions. The parameters of the matching distribution are listed on the plots in the order (α, β,kb,k). The £t between the measured data and the biPareto distributions is excellent.

Some measured parameters of the traf£c trace seemed to be drawn from a *censored* biPareto distribution, namely the distribution that results from throwing out realizations of a biPareto distribution above a given threshold. An example is the duration T, in seconds, of user sessions. The lowest curve in Figure 5 is the empirical tail distribution function of user session duration. There are *three* linear regions in the log-log plot; the extreme tail appears to be asymptotic to a vertical line at $T\approx 600$ minutes, suggesting that the distribution does not allow observations above this level.

Suppose that a random variable X with tail distribution $\bar{F}(x)$ is censored by throwing out realizations greater than some threshold X_0. The resulting tail distribution is truncated to

$$\bar{F}^*(x)=\mathrm{P}(X>x\,|\,X<X_0)=\frac{\bar{F}(x)-\bar{F}(X_0)}{1-F(X_0)},\quad x<X_0.$$

Given the *censoring probability* $F(X_0)$, the original distribution $\bar{F}$ (on $x<X_0$) can be recovered from $\bar{F}^*$ as

$$\bar{F}(x)=\bar{F}(X_0)+(1-\bar{F}(X_0))\bar{F}^*(x).$$

Suppose that the observations recorded in the lower curve in Figure 5 are drawn from a censored distribution $\bar{F}^*$. The two upper curves in Figure 6 show two different estimates of $\bar{F}(x)$ based on the censoring probability estimate $\bar{F}(T_0)=$

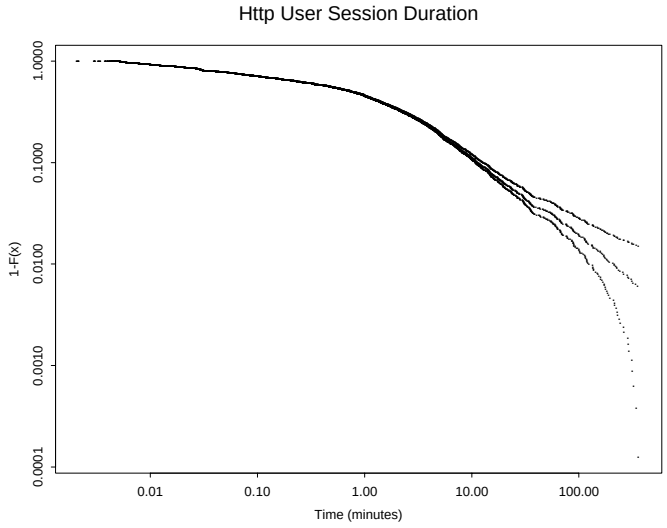

Figure 5: Empirical tail distribution of user session duration D and two hypothetical uncensored distributions. (Bottom curve: data, Middle curve: assuming $\bar{F}(T_0) = 0.006$, Top curve: assuming $\bar{F}(T_0) = 0.012$)

0.012 for the upper curve and $\bar{F}(T_0) = 0.006$ for the middle curve. The middle curve is asymptotically straight, indicating that a biPareto £t could be appropriate. The empirical $\bar{F}^*$ based on the estimate $\bar{F}(T_0) = 0.006$ is then plotted against a biPareto £t in Figure 6. In the upper curve of Figure 5, the tail is convex, indicating that 0.012 is too high an estimate of the censoring probability.

When working with heavy-tailed data, a censored model may sometimes be appropriate, even when no obvious truncation is apparent. Suppose for example that we observe N realizations of a random variable, and that the tail distribution shows a clear power-law decay, down to observable probabilities on the order of $1/N$. Based on the observations, we could form a model with an in£nite power law tail, or we could form a model censored to some value greater than the largest observed value. Suppose that we use the model to generate K synthetic realizations. If $K << N$, the two models would most likely generate similar values. However, if $K >> N$, the uncensored model will produce a number of realizations which are much larger than any of the observed values. In this case, the choice of model should depend on physical intuition; is it reasonable for the tail to extend in£nitely, or are there mechanisms in the system that would limit maximum values? Possible censoring mechanisms could include £nite data collection windows, hardware speed and memory limitations, and thresholds speci£ed by

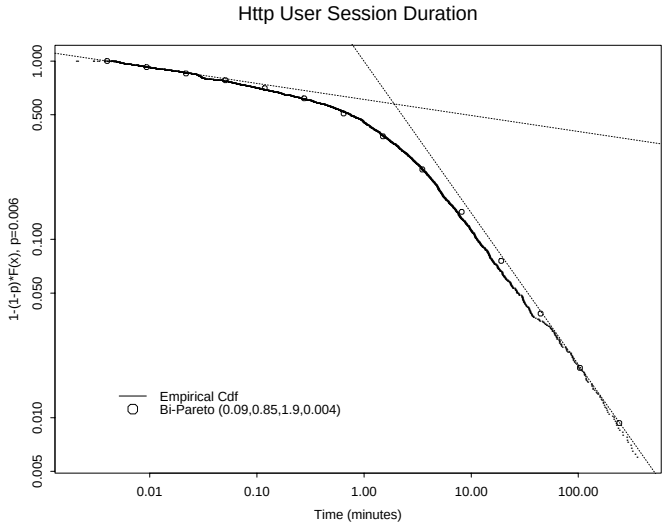

Figure 6: Empirical estimate of the uncensored tail distribution of user session duration D, and a biPareto £t, assuming a censoring probability of 0.006.

communication protocols.

2.5 Independence of number of connections (C) and mean interarrival time (μ)

To fully characterize the user session parameters such as C, μ, and T, we need to describe not only their marginal distributions, but also their joint distributions. Although T is strongly correlated with both μ and C, the relationship between μ and C is much weaker. The correlation coef£cient between μ and C was measured to be only -0.013 in our data set. We also examined the mean of μ, conditioned on various ranges of C. Here again we found a weak inverse relationship between μ and C.

In the mathematical model, we de£ne C and μ to be mutually independent random variables. This simpli£cation appears to have a relatively minor effect on the effectiveness of the model. If more careful modeling is required later, one approach would be to de£ne a few different condition distributions for μ conditioned on ranges of C.

2.6 Connection Arrivals within a User Session

In the model we have described so far, user sessions arrive according to a Poisson process, and the i-th user session is assigned a number of connections C_i and a total connection arrival time $T_i = \mu_i(C_i - 1)$, where C_i and μ_i are drawn independently from biPareto distributions. The last step is to model the way in which the connections are distributed within a user sessions.

Physical intuition and qualitative examination of the data suggest that the arrival of connections within a user session will be complex. For example, when opening a Web page, typical browsers will open as many as four simultaneous TCP connections for text and images, generating very closely spaced connection arrivals. Along with these machine-driven dynamics, there are also slower dynamics based on the behavior of the human user.

Rather than attempting to model these effects explicitly, we took a simple approach based on a renewal process. Given C and T, the $C-1$ interarrival times $\{X_j\}$ for a given user session were de£ned to be

$$X_j = T\frac{Z_j}{\sum_{k=1}^{C-1} Z_k}, \quad j \in \{1,\cdots,C-1\}$$

where the Z_k are positive, i.i.d. random variables. That is, C arrivals from a renewal process are scaled to span exactly T seconds. For the distribution function of the Z_k, we chose a Weibull distribution with parameter $c < 1$. The Weibull distribution was used because it is a commonly observed interarrival distribution [15, 10] and because its coef£cient of variation (ratio of the standard deviation to the mean) can be easily adjusted using the shape parameter, c. In the synthesis described below, $c = 0.48$ so that the coef£cient of variation matched the measured *average* coef£cient of variation for within-session connection interarrival times.

In our experience, the properties of the synthesized model are fairly robust to changes in the distribution chosen for the Z_k. For example, we found that increasing the coef£cient of variation of the Z_k tended to increase the coef£cient of variation of the overall interarrival times, but that large changes in Z_k were required to effect signi£cant changes overall.

3 Synthesized Connection Arrivals

To test the appropriateness of the connection arrival model, we used the model to generate synthetic connection arrivals. The rate of the user session arrival process,

the parameters of the biPareto distributions of C and μ, and the coef£cient of variation of the within-session interarrivals were all matched to measurements from the data as described in the previous section.

Here is the procedure we used to generate synthetic arrivals.

1. Generate a user session arrival at an independent exponentially distributed time from the previous user session arrival.

2. Generate C, the number of connections in this user session, as an independent biPareto random variable.

3. If $C > 1$, generate μ as an independent biPareto random variable.

4. If $C > 1$, generate $C-1$ independent Weibull random variables Z_j, and scale them to X_j as described in Section 2.6. Note that this uses the variable $T = (C-1)\mu$. The X_j represent the timings of connection arrivals.

We now show that our synthetic TCP connection arrivals are well-matched to the data both in terms of interarrival marginal and interarrival autocorrelation. The match of the two marginal distributions to each other is depicted in Figure 7. A straight line would indicate a perfect £t. In fact, the plot is curved "downward", indicating that the marginal of the simulated interarrivals is slightly more extreme than that of the data. In agreement with [10, 3], the distributions of the data and model are well-described by Weibull distributions.

The autocorrelation of the actual and simulated interarrivals are shown in Figure 8. Again the match is very good, though not exact. Because of the long-range dependence in the data, the empirical autocorrelation varies signi£cantly from realization to realization.

A measure of the accuracy of the model across different scales is depicted in Figure 9. In this £gure, we examine processes derived by counting the number of connections arriving in intervals of various durations. If the model matches the data well, then the marginal distribution of the count process derived from the model should match that of the data on all time scales. In this £gure, we have simply measured the square of the coef£cient of variation of the count processes as a function of scale. The line labelled "bi-level simulation" refers to the model. The line labeled "direct simulation" refers to a technique similar to that described in [3], which directly matches the marginal and autocorrelation of the observed interarrival sequence. In the line marked "i.i.d. simulation", the arrival process is a renewal process with interarrival distribution matched to the data. The bi-level model and direct simulation are roughly equivalent in their ability to match

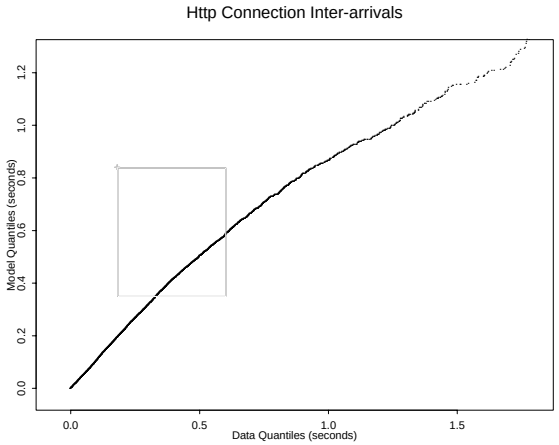

Figure 7: Fit of synthesized TCP connection interarrivals to data.

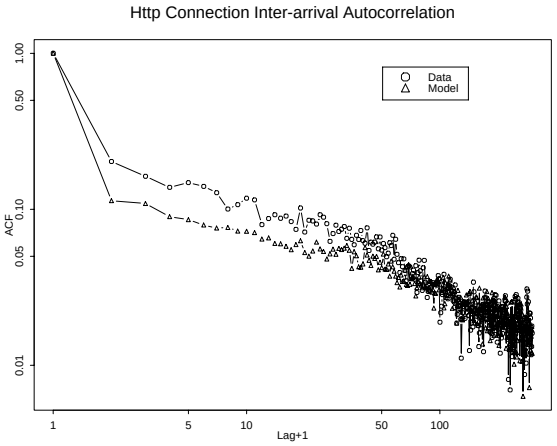

Figure 8: Empirical autocorrelation function of TCP connection interarrival sequence: from data and from model.

the data in terms of aggregated coef£cient of variation, while the i.i.d. simulation departs signi£cantly from the data as the time scale increases. The coef£cient of variation of the data has a dip at the 100ms scale that is not reproduced by any of the three models.

We have, in Figures 10 and 11, a visual comparison of the number of connection arrivals observed in the data compared to a trace generated by the synthesized procedure. There are two time scales depicted: one is a count of arrivals per 100 ms; the other is a count per 100 seconds. Qualitatively, the left half of each plot (data) looks identical to the right half (synthetic). This is just a visual way to check that the synthetic data doesn't have any glaring anomalies.

Generating synthetic TCP connection using the model is computationally ef£cient; the slowest part of the algorithm is sorting together the connections from different user sessions, which may require something like $O(N \log N)$ operations, where N is the number of TCP connections.

3.1 The Effect of Spatial Aggregation

The proposed model captures the temporal aspects of TCP connections for data collected over a local area network. As shown in Figures 10 and 11, this corresponds to an average of (roughly) 10 TCP connections per second. We wish to

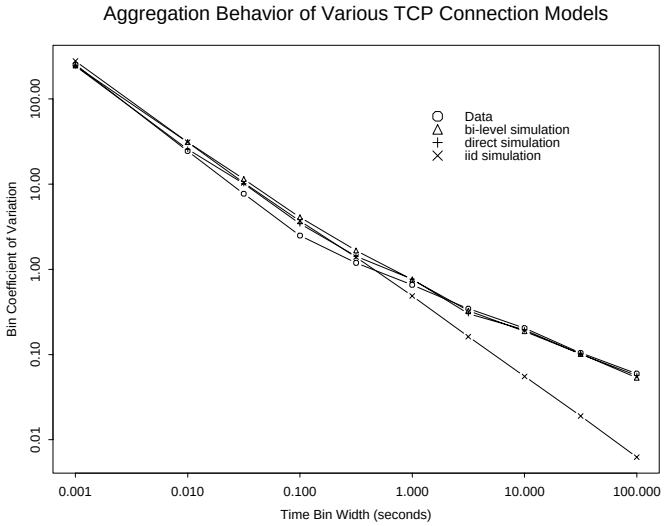

Figure 9: Square of the coef£cient of variation for a TCP connection count process as a function of bin length, from data and various models.

examine the validity of this model for wide area traf£c consisting of (say) 10-40gbps trunks carrying thousands of such connections per second. In the absence of explicit measurements from 10-40gbps trunks carrying primarily data traf£c, we propose to examine the effect of large-scale, or spatial, aggregation of TCP connections as follows. First, we synthesize spatial aggregates from the existing data. Next, we examine the temporal correlation structure of the resulting traces, as measured by the Hurst parameter, and compare it to that of the original trace. We then measure the asymptotic variability of the aggregate trace, as measured by its coef£cient of variation, and show how it relates to that of the local area network, the original trace. Finally, we show how the observed effects of large-scale aggregation are exhibited naturally in the proposed model.

Synthesis of Spatial Aggregation We use the technique of randomized shifts [9, 8] to create copies of the original trace which appear to be mutually independent, but whose correlation structures are virtually identically to the original trace. Each copy is generated by cutting the original trace at a random point and interchanging the beginning and ending segments. Multiple copies are then summed together to form the spatial aggregate. In order for this method to work, the observation interval should be much shorter than the total length of the trace. In our case, we took a trace consisting of about 100,000 counts, summed between 10 and 1000 shifted copies, and observed the results on a segment of 2000 counts, depicted in

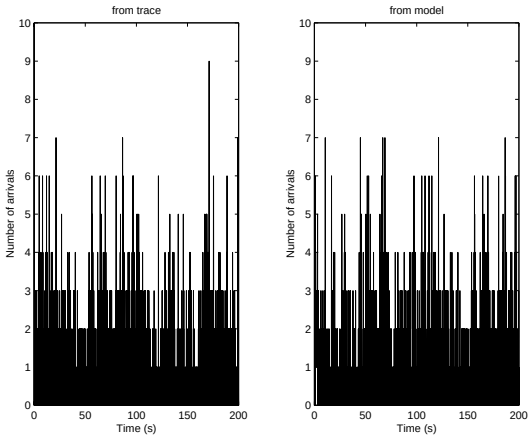

Figure 10: HTTP Connection arrivals in 100ms intervals

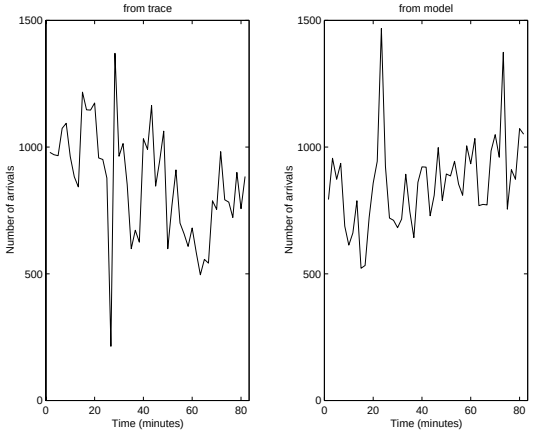

Figure 11: HTTP connection arrivals in 100 second intervals

the left side of Figure 12. This £gure shows estimates of the presumed counts over links carrying 10 to 1000 times more connections than the link we studied. On the right side of the £gure, we show corresponding counts generated from the compound model. The only difference between the model used in Figure 10 and the models used in Figure 12 is that the user session arrival rate ν is increased by a factors of 10, 100, and 1000 in Figures 12 (a), (b), and (c), respectively.

Correlation Structure of Trace and Its Spatial Aggregate Figure 13 shows variance-scale log-log plots for the original trace and for spatial aggregates by the factors 10, 100, and 1000. The linearity of the plots indicates self-similar behavior, with the Hurst parameter H being proportional to the slope. The graph shows that the spatial aggregation procedure does not affect the value of the Hurst parameter, which is estimated to be 0.72+/-0.03 in all four cases.

Reduction in Coef£cient of Variation without Change in H An important characteristic of the spatial aggregation is increased smoothness visible in Figures 12(a)-(c). This is a simple manifestation of the law of large numbers. Let N denote the number of independent realizations of a count process that are summed to form an aggregate. The mean and variance of the marginal distribution of the aggregate process both scale linearly with N, so that the coef£cient of variation tends to zero as $N^{-1/2}$. This highlights the importance of considering the autocorrelation amplitude as well as the decay exponent in evaluating the "strength" of long-range dependence. Even though all of the aggregates have the same Hurst parameter, the long-range dependence is more important in the traces with high coef£cient of variation.

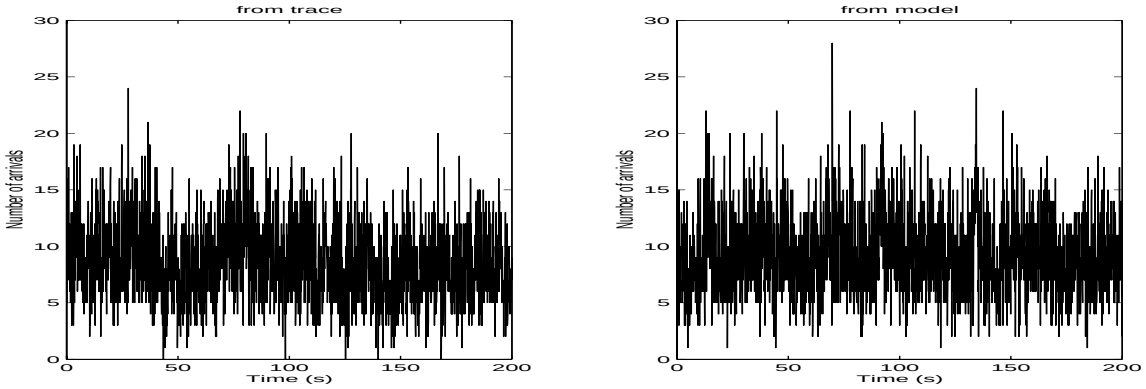

(a) trace and model spatially aggregated 10 times

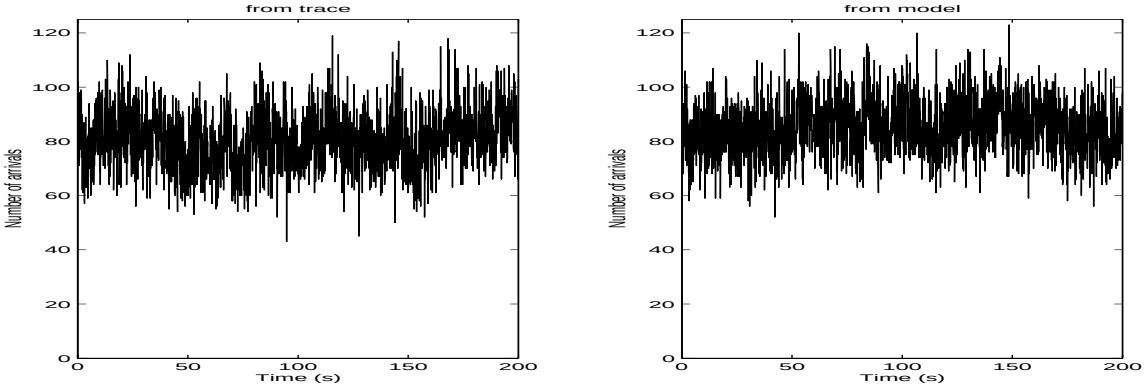

(b) trace and model spatially aggregated 100 times

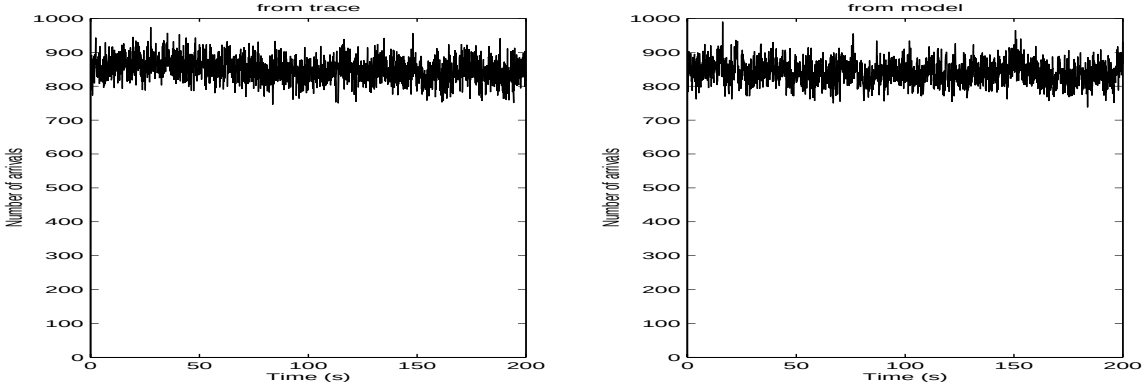

(c) trace and model spatially aggregated 1000 times

Figure 12: Incremental number of TCP connections per 100msec for the synthetically aggregated trace (left side) and corresponding model (right side).

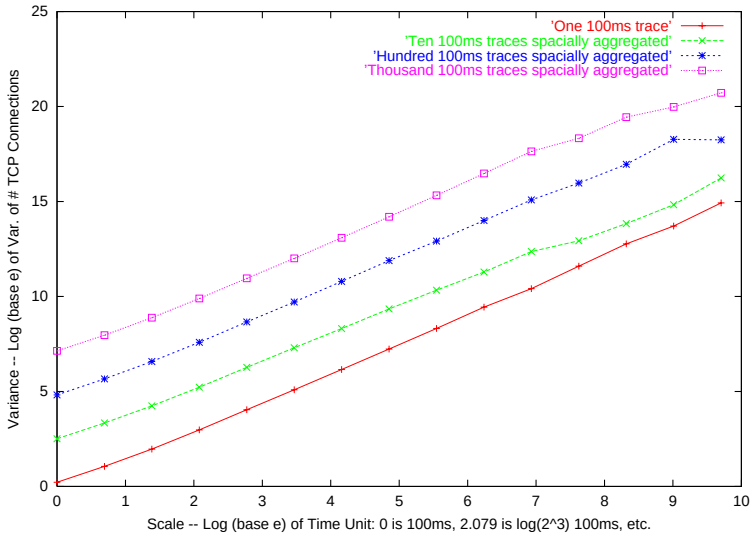

Figure 13: Variance-time log-log plots showing no change to H as a result of aggregation.

Comparing Figure 11 with Figure 12 (c) demonstrates the difference between spatial and temporal aggregation. The coef£cient of variation of a long-range dependent process falls more slowly under temporal aggregation than under spatial aggregation, because in the temporal case, the variables being aggregated are strongly correlated. For a short-term dependent process, the decay rate is the same in both cases.

4 General Framework

The model presented in the previous section is a special case of a more general approach that is presented below. This framework includes the approach of [4] and [19], and like these models provides a link to the more familiar fractal models used in data networking, for example [17] and [26].

We will construct a stochastic, instantaneous rate process for arrivals which will be used to modulate a point process, giving a doubly stochastic point process. This point process could be Poisson, as in [18], or a burstier renewal process, e.g. the process with Weibull interarrivals discussed in the previous section. Here we focus on properties of the modulating rate function.

4.1 Rate process construction

We construct the rate process using a shot noise construction for which the roots go back as least as far as [24] and [14]. In a general shot noise setting, events arrive in a poisson stream with rate ν. Each arrival A_i generates a pulse $h(t-A_i,X_i)$ parameterized by an i.i.d. sequence $\{X_i\}$ of random vectors which parameterize the pulses and which are independent of $\{A_i\}$.

The entire shot noise is the sum of pulses,

$$\lambda(t) = \sum_i h(t-A_i,X_i). \tag{1}$$

Some examples: Lowen and Teich [19] de£ne fractal shot noise as

$$h(t-A_i,X_i) = X_i f(t-A_i)$$

where

$$f(t) = \begin{cases} ct^{\alpha/2-1} & A < t < B \\ 0 & \text{otherwise} \end{cases}$$

Here X_i is a scalar quantity which may be random or deterministic, and B may be in£nite.

The power law decay in the impulse response of these shot noises leads to asymptotic long-range dependence and fractal behavior. These models are particularly appropriate for a variety of physical processes ranging from neuron £rings and Cherenkov radiation from charged particles [19],[18].

In the case of computer traf£c traces, long range dependence manifests itself in random distributions with power-law decay, rather than in slowly decaying impulse responses. If $X_i = T_i$ is a random time length parameter, and h is the function

$$h(t,T_i) = 1_{(0,T_i)}(t),$$

depicted on the left side of Figure 14, then we obtain the M/G/∞ "queue". That is, $\lambda(t)$ is the number of servers busy at time t when the arrival process is Poisson, T_i is the random service time, and there is no limit on the number of active servers. When the distribution of the user session service times is heavy-tailed, long-range dependence is induced in $\lambda(t)$.

In the traf£c analysis reported in previous sections, we noted that different users generated traf£c at widely different rates. We also observed that to £rst order the average interarrival time for connections within a user session was independent of the total amount of traf£c it generated. We de£ne $X_i = (C_i, \mu_i)$, where C_i is

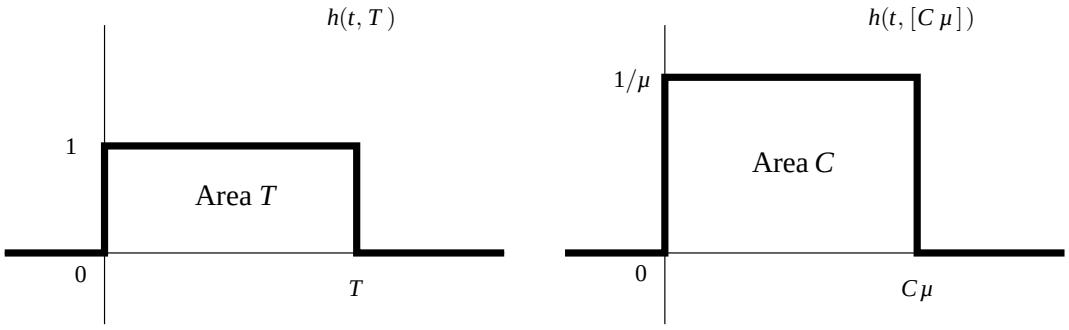

Figure 14: Shot noise impulse responses for the M/G/∞ queue (left) and a generalization (right).

the amount of traf£c generated by the user session and μ_i is the expected time between traf£c events generated by that session, and where the variables C_i and μ_i are independent of one another. Then the impulse response corresponding to the model given in Section 2 is

$$h(t,(C_i,\mu_i)) = \frac{1}{\mu_i} 1_{(0,\mu_i C_i)}(t).$$

As depicted in Figure 14, each user session consists of a rate rectangle of height $1/\mu_i$, width $C_i\mu_i$, and area C_i. Note that if $\mu_i = \mu$ is a determistic constant, this model reduces to the M/G/∞ case.

We are interested in the way in which the distributions of C_i and μ_i affect the autocorrelation $r_\lambda(t)$ of the rate process $\lambda(t)$. In particular, we are interested in how heavy tailed distributions give rise to long range dependence in $\lambda(t)$.

4.2 Mean and autocorrelation of $\lambda(t)$

Suppose that for shot noise given by equation (1), the impulse response h is deterministic, i.e. does not depend on X_i. Then [24] the mean and autocorrelation of $\lambda(t)$ are given by

$$\bar{\lambda} = \mathbf{E}\{\lambda(t)\} = \nu \int_{-\infty}^{\infty} h(t)\,dt$$

and

$$r_\lambda(\tau) = \mathbf{E}\big\{(\lambda(t)-\bar{\lambda})(\lambda(t+\tau)-\bar{\lambda})\big\} = \nu \int_{-\infty}^{\infty} h(t)\,h(t+\tau)\,dt.$$

When $h(t)$ is a rectangular pulse of £xed height $1/\mu$ and area C, the mean is $\bar{\lambda} = \nu C$ and the autocorrelation is

$$r_\lambda(\tau) = \begin{cases} \nu\frac{1}{\mu}\left(C - \frac{|\tau|}{\mu}\right) & |\tau| < \mu C \\ 0 & \text{else} \end{cases} .$$

Suppose that there are n independent deterministic shot noises with parameters ν_i, μ_i, and C_i. Suppose also that $\sum \nu_i = \nu$. Then summing the shot noises together produces an equivalent random shot noise, in which the parameters (μ_i, C_i) are drawn with probability ν_i/ν for each arrival.

Due to the independence of the shot noises, the means and autocorrelations sum. In the limit of large n, we can think of μ and C having a given continous joint distribution, and the mean and autocorrelation can be obtained via integration. Speci£cally,

$$\bar{\lambda} = \nu \mathbf{E}\{C\} = \nu \int_0^\infty p(C) C \, dC$$

and

$$r_\lambda(\tau) = \nu \mathbf{E}\left\{ \frac{1}{\mu}\left(C - \frac{|\tau|}{\mu}\right)^+ \right\} = \nu \int_0^\infty \int_0^\infty p(\mu, C) \frac{1}{\mu}\left(C - \frac{|\tau|}{\mu}\right)^+ d\mu \, dC \quad (2)$$

where $p(C)$ and $p(\mu, C)$ are probability distribution functions. Note that $r_\lambda(0) = \nu \mathbf{E}\{C/\mu\}$, or in case C and μ are independent, $r_\lambda(0) = \nu \mathbf{E}\{C\}\mathbf{E}\{(1/\mu)\}$.

4.3 Long-range dependence

Heavy-tailed random distributions of μ and C can lead to long-range dependence in the autocorrelation of $\lambda(t)$. Cox [4] analyzed this relationship for the $M/G/\infty$ queue, i.e. for $\mu = 1$ and random C, showing that $x^{-\alpha}$ behavior in the tail distribution of C leads to $t^{1-\alpha}$ behavior in the autocorrelation of the queue size. To be more precise, a *regularly varying function U with parameter p* is de£ned as a function which satis£es

$$\lim_{t\to\infty} \frac{U(xt)}{U(t)} = x^p$$

for all $x > 0$. If $p = 0$, U is referred to as a *slowly varying function*. It turns out that each regularly varying function can be written in the form $U(t) = t^p L(t)$, where L is slowly varying. Examples of slowly varying functions include functions with £nite, non-zero limits and powers of $\log t$.

Cox's result states that if the complementary distribution function $\bar{F}_C(x)$ of C is regularly varying with £nite mean, i.e. if

$$\bar{F}_C(x) = x^{-\alpha}L(x)$$

for any $\alpha > 1$ and slowly-varying function $L(x)$, then the autocorrelation function of the queue size process is

$$r_\lambda(\tau) = \nu \int_\tau^\infty p(C)(C-\tau)\,dC = \nu \int_\tau^\infty \bar{F}_C(C)\,dC \equiv \tau^{1-\alpha}L^*(\tau)$$

where L^* is a slowly varying function (see Theorem A.1). If $1 < \alpha < 2$, then $r_\lambda(\tau)$ is unsummable, and hence the rate process is long-range dependent.

In the case that both C and μ are random, it turns out that power-law tails in either C or μ can lead to power-law decay in r_λ. As an illustrative example, consider the case in which C is Pareto(α,1), μ is Pareto(β,μ_0), and C and μ are mutually independent. After evaluating (2), the autocorrelation turns out to be

$$r_\lambda(t) = \frac{\nu}{\mu_0}\left(\frac{\alpha\beta}{(\alpha-1)(\beta+1)} - \frac{\beta}{\beta+2}\frac{t}{\mu_0}\right)$$

for $|t| \le \mu_0$ and

$$r_\lambda(t) = \frac{\nu\beta}{\mu_0(\alpha-1)(\beta+2-\alpha)}\left(\frac{t}{\mu_0}\right)^{1-\alpha} + \frac{\nu\alpha\beta}{\mu_0(\beta+1)(\beta+2)(\alpha-\beta-2)}\left(\frac{t}{\mu_0}\right)^{-\beta-1}$$

for $|t| \ge \mu_0$. If $\alpha - 1 < \beta + 1$, then the autocorrelation decays with exponent $\alpha - 1$. If the inequality is reversed, the autocorrelation decays with exponent $\beta + 1$. As in the M/G/∞ queue, long-range dependence is induced whenever $1 < \alpha < 2$. The following theorem shows that these observations can be generalized to distributions with regularly varying tail distributions. To simplify the statement of the theorem, we de£ne the function

$$G(t) = \int_t^\infty \frac{1}{u} dF_\mu(u). \tag{3}$$

When F_μ is regularly varying with parameter $-\beta$, Lemma A.2 shows that G is regularly varying with parameter $-\beta - 1$. Also, we denote the Mellin transform of a function f by

$$\hat{f}(z) = \int_0^\infty t^{-z} f(t)\,dt/t$$

and we denote the relationship

$$\lim_{t\to\infty} \frac{f(t)}{g(t)} = 1$$

by $f(t) \sim g(t)$.

Theorem 4.1 *Suppose that C and μ are independent random variables with tail distributions $\bar{F}_C$ and $\bar{F}_\mu$, and that $\mathbf{E}\{1/\mu\} < \infty$.*

I. Suppose that $\bar{F}_C$ varies regularly with parameter $-\alpha < -1$, and that $\bar{F}_\mu$ is bounded by a regularly varying function with parameter $-\beta < -1$, where $\alpha - 1 < \beta + 1$. Then

$$r_\lambda(t) \sim t\bar{F}_C(t)\bar{G}(1-\alpha),$$

and $r_\lambda(t)$ is regularly varying with parameter $1-\alpha$.

II. Suppose that $\bar{F}_\mu$ varies regularly with parameter $-\beta < -1$ and that $\bar{F}_C$ is bounded by a regularly varying function with parameter $-\alpha < -1$, where $\beta + 1 < \alpha - 1$. Then

$$r_\lambda(t) \sim G(t)\bar{F}_C(-\beta-2),$$

and $r_\lambda(t)$ is regularly varying with parameter $-\beta - 1$.

A proof of this theorem is provided in Appendix A. Although the theorem as stated does not cover the case $\alpha - 1 = \beta + 1$, the appendix gives asymptotic bounds. These bounds show that if $r_\lambda(t)$ varies regularly in this case, then as expected it varies regularly with exponent $1-\alpha$.

BiPareto distributions have regularly varying tail distributions, and hence rate and count processes based on biPareto C and μ will have regularly varying autocorrelation functions. In practice, it is important to keep in mind that observations have a limited temporal range. If the regularly varying behavior of C and μ is only in the extreme tail, then the regular variation induced in the autocorrelation may not be apparent at time scales of practical interest.

The second-order temporal structure of a point process is speci£ed by its co-occurence function. In the case of a doubly stochastic Poisson process, this function is proportional to the autocorrelation of the rate function [18]. For count processes derived from the point process, the autocorrelation is closely related to the co-occurence function and typically has the same asymptotic decay. Hence long-range dependent rate functions $\lambda(t)$ lead to long-range dependent arrival count processes. Under suf£cient levels of aggregation, any such count process eventually approaches the usual fractional Brownian motion model.

4.4 Other properties

As Theorem 4.1 shows, the shot noise rate function we have proposed leads to such a long-range dependent process whenever C is heavy-tailed with parameter $1 < \alpha < 2$. However, the model retains a great deal of ¤exibility in the distributions of μ and C, even after the parameter α is £xed. These distributions determine the remaining properties of the rate function, including the marginal distribution, short-term correlation properties, and the amplitude of the slowly decaying autocorrelation tail. For example, a continuous distribution on μ leads to a continuous marginal distribution for the rate function. If very small values of μ are allowed, the variance of the rate function increases with $\mathbf{E}\{1/\mu\}$, and the short-term autocorrelation may also change shape.

While the tail densities of C and μ determine the Hurst parameter of the rate process, its coef£cient of variation is controlled by the arrival rate ν of the Poisson process of user session arrivals. Speci£cally,

$$\frac{(r_\lambda(0))^{1/2}}{\bar{\lambda}} = \nu^{-1/2}\frac{\mathbf{E}\{(C/\mu)\}^{1/2}}{\mathbf{E}\{C\}}$$

Increasing ν corresponds exactly to increasing the level of spatial aggregation, as dicussed in Section 3.1. As the population generating rate events increases, the relative ¤uctuations in total rate become smaller.

The properties of the interarrival processes corresponding to point processes modulated by $\lambda(t)$ are also determined by the distributions of μ and C. Although the count processes obtained from these point processes inherit properties such as long range dependence directly from the rate process, the relationship between the statistics of the rate process and those of the inter-arrival sequence is much more complex. For example, the second order statistics of $\lambda(t)$ are not suf£cient to determine even the marginal distribution of the interarrival process [5].

5 Conclusions

The compound, shot noise approach appears to be a useful and ¤exible modeling method. The distributions of parameters such as μ and C can be estimated based on empirical measurements, as in this work, or based on system knowledge and experience. Once the distributions are £xed, the model automatically reproduces complex point process behavior, including long-range dependence counts and bursty interarrivals. In [27], a complexity/¤exibility tradeoff is noted between

Cox's model and Kurtz's model [16]. Generalizing Cox's model along the lines that we have described provides an intermediate point in this tradeoff spectrum.

One £nal observation relates to Section 3.1. The smoothing effect observed through large-scale synthesis, as well as that predicted by the analytical model, has a signi£cant impact on traf£c engineering issues. For LAN data traf£c, variability around the mean is substantial and good traf£c engineering needs to take account of these unusually large ¤uctuations. However, for very large aggregates, such as those in wide area traf£c, as in the core of networks, the variability is considerably reduced. This reduction substantially simpli£es the traf£c engineering problem for WANs and implies that high utilizations are feasible *regardless* of the short or long term temporal correlations in the traf£c.

There are many areas for further study. Initial investigation suggests that the relationship between TCP connections and aggregated IP packets may be analogous to that between user sessions and TCP connections, and that it may be possible to simulate the packet-level arrivals using a bi-level model based directly on user sessions, or using a three-level model including user sessions, TCP connections, and packets. The timing of packets within a TCP connection is clearly affected in a complicated way by closed-loop features of the TCP protocol such as congestion avoidance. However, it is possible that a simple timing model such as that described in Section 2.6 may be suf£cient in order to reproduce empirical behavior in suf£ciently aggregated packet traf£c. Detailed queueing analysis of the present two-level model would test the applicability of our results to buffer statistics, and it would be desirable to have greater analytic understanding of properties of this family of models.

A Regularly Varying Functions

The standard reference on regularly varying functions is [2]. Useful introductions to regularly varying functions and their applications in probablity also include relevant chapters of [12] and [23]. In this section, we give a proof of Theorem 4.1, preceeded by some supporting results.

De£nition A.1 *A measurable function $U : \mathbb{R}^+ \to \mathbb{R}^+$ is regularly varying with parameter p if for each $x > 0$,*

$$\lim_{t\to\infty} \frac{U(tx)}{U(t)} = x^p.$$

Regularly varying functions behave in some sense like t^p with regards to integration, as in the following theorem taken from [12].

Theorem A.1 *If Z varies regularly with exponent γ and the integral*

$$Z_p^*(t) = \int_t^\infty u^p Z(u)\,du$$

exists, then

$$\frac{t^{p+1}Z(t)}{Z_p^*(t)} \to \lambda \tag{4}$$

where $\lambda = -(p+\gamma+1) \geq 0$. Conversely, if (4) holds with $\lambda > 0$, then Z and Z^* vary regularly with exponents $\gamma = -\lambda - p - 1$ and $-\lambda$, respectively.

Lemma A.2 *Suppose that $\bar{F}_\mu$ is regularly varying with parameter $-\beta < -1$, and that $\mathbf{E}\{1/\mu\}$ is £nite. Then the function $G(t)$ de£ned in (3) is £nite and monotonically decreasing for all $t \geq 0$, is regularly varying with parameter $-\beta - 1$, and satis£es*

$$G(t) \sim \frac{\beta}{\beta+1} t^{-1}\bar{F}_\mu(t).$$

Likewise, if $\bar{F}_\mu$ is bounded by a regularly varying function with parameter $-\beta$, then G is bounded by a regularly varying function with parameter $-\beta - 1$.

Proof:
By assumption, $G(0) = \mathbf{E}\{1/\mu\}$ is £nite, and G is monotonically decreasing since F_μ is a distribution function. For $t > 0$, integration by parts establishes that

$$G(t) = \int_t^\infty \frac{1}{u}\,dF_\mu(u) = \int_t^\infty \left[\bar{F}_\mu(t) - \bar{F}_\mu(u)\right] \frac{du}{u^2}.$$

From Theorem A.1 we know that

$$\int_t^\infty \bar{F}_\mu(u)u^{-2}\,du \sim (\beta+1)^{-1}t^{-1}\bar{F}_\mu(t),$$

Then

$$G(t) \sim t^{-1}\bar{F}_\mu(t) - (\beta+1)^{-1}t^{-1}\bar{F}_\mu(t) = \frac{\beta}{\beta+1}t^{-1}\bar{F}_\mu(t)$$

so that G is regular varying with parameter $-\beta-1$. Since $G(t) \le \bar{F}_\mu(t)/t$, a regularly varying bound on $\bar{F}_\mu$ becomes a bound on G, with parameter reduced by one. □

The Mellin convolution of two functions k and g is de£ned to be

$$k * g(x) = \int_0^\infty k(x/t)g(t)\,dt/t = \int_0^\infty g(x/t)k(t)\,dt/t.$$

Under appropriate conditions on k and g, the convolution $k * g$ inherits regular variation from g [2, Theorem 4.1.6].

Theorem A.3 (Arandelović (1976)) *Suppose that a measurable function f is regularly varying with parameter ρ. Suppose that the Mellin transform $\hat{k}$ of k converges at least in the strip $\sigma \le \mathrm{Re}\,(z) \le \tau$, where $-\infty < \sigma < \rho < \tau < \infty$, and that $f(x)x^{-\sigma}$ is bounded on every interval $(0,a]$. Then*

$$k * f(x) \;\sim\; \hat{k}(\rho)\, f(x)$$

(and $k * f$ is regularly varying with parameter ρ.)

We are interested in the case in which both functions are at least bounded by regular variation.

Lemma A.4 *Suppose that f and g are positive, bounded functions on $[0,\infty)$. Suppose that f is regularly varying with parameter $-\rho$, and that g is bounded by a regularly varying function with parameter $-\sigma$, where $\sigma > \rho > 0$. Then the Mellin convolution $f * g$ satis£es*

$$f * g(x) \;\sim\; \hat{g}(-\rho)\, f(x),$$

and is regularly varying with parameter $-\rho$. If $\rho = \sigma$, then then for any $T, \varepsilon > 0$, and any postive, bounded function h regularly varying with parameter $\varepsilon - \rho$, the convolution obeys the bounds

$$(1-\varepsilon)\, f(x) \int_0^T g(t)t^\rho\,dt/t \;\le\; f * g(x) \;\le\; h(x)$$

for suf£ciently large x.

Proof:

The case $\sigma > \rho$ is a direct application of Theorem A.3. Since g is bounded by a constant and by a regularly varying function with parameter $-\sigma$, its Mellin

transform exists at least in the strip $-\sigma < \mathrm{Re}\,(z) < 0$. (See [2, Section 1.5.6]; The transform may or may not converge at $-\sigma$.) Also $f(x)x^{\sigma}$ is bounded on £nite intervals, and we may conclude that $f * g(x) \sim \hat{g}(-\rho) f(x)$.

For the upper bound in the case $\sigma = \rho$, de£ne for any $\delta > 0$ the function $k_\delta = \max\{g, \delta h\}$. The ratio h/g is regularly varying with parameter $\varepsilon > 0$ and hence increases without limit, so that $k_\delta(x) = \delta h(x)$ for suf£ciently large x. Applying the £rst part of the lemma, we have

$$f * g(x) \;\le\; f * k_\delta(x) \;\sim\; k_\delta(x)\,\hat{f}(\varepsilon - \rho) \;\sim\; \delta \hat{f}(\varepsilon - \rho) h(x),$$

and we may take $\delta < 1/\hat{f}(\varepsilon - \rho)$.

To obtain the lower bound, let g_T be the truncation of g to the interval $[0, T]$. Then g_T can be upperbounded by a regularly varying function with arbitrary parameter, and applying the £rst part of the theorem, we have

$$f * g(x) \geq f * g_T(x) \;\sim f(x)\,\hat{g}_T(-\rho)$$

as desired. □

Finally, we may prove Theorem 4.1.

Proof of Theorem 4.1 Applying integration by parts to (2), we have

$$\begin{aligned}
r_\lambda(t) &= \nu \int_0^\infty \int_0^\infty \frac{1}{\mu} \left(C - \frac{t}{\mu}\right)^+ dF_C(C)\, dF_\mu(\mu) \\
&= \nu \int_0^\infty \frac{1}{\mu} \int_{t/\mu}^\infty \left(C - \frac{t}{\mu}\right) dF_C(C)\, dF_\mu(\mu) \\
&= \nu \int_0^\infty \frac{1}{\mu} \int_{t/\mu}^\infty \bar{F}_C(C) dC\, dF_\mu(\mu) \\
&= \nu \int_0^\infty \bar{F}_C(C) dC \int_{t/C}^\infty \frac{1}{\mu} dF_\mu(\mu) \\
&= \nu H * G(t)
\end{aligned}$$

where $H(t) \equiv t\bar{F}_C(t)$ and G is de£ned by (3).

Since H decays asymptotically to zero, and because $H(t) \le t$, H is bounded on $[0, \infty)$. If $\bar{F}_C$ is (resp. is bounded by) a regularly varying function with parameter $-\alpha$, then H is (resp. is bounded by) a regularly varying function with parameter $1 - \alpha$, and Lemma A.2 similarly relates G and $\bar{F}_\mu$. Hence if $\alpha < \beta + 2$ we may apply Lemma A.4 with $\rho = \alpha - 1$, $\sigma = \beta + 1$ to obtain the desired results. Otherwise if $\beta < \alpha - 2$, we take $\rho = \beta + 1$ and $\sigma = \alpha - 1$. □

If H and G have the same coef£cient of variation $1-\alpha = -\beta - 1$, Lemma A.4 provides the asymptotic lower bounds

$$(1-\varepsilon)xF_C(x)\int_0^T G(t)t^{\beta+1}\,dt/t \text{ and } (1-\varepsilon)G(x)\int_0^T F_C(t)t^{\alpha}\,dt/t$$

for $r_\lambda(x)$, and shows that $r_\lambda(t)$ is asymptotically upper bounded by any positive, regularly varying function with parameter strictly greater than $1-\alpha$.

B BiPareto distribution

B.1 Pareto Distribution

The Pareto distribution has cumulative distribution function

$$\text{Pareto}(\alpha,k): \quad F(x) = 1-\left(\frac{x}{k}\right)^{-\alpha}, \quad x \geq k.$$

The decay exponent α must be strictly positive, and the ν-th order moment will exist if and only if $\alpha > \nu$. In particular, the mean and variance are de£ned when $\alpha > 1$ and $\alpha > 2$, respectively. The mean is $k\alpha/(\alpha-1)$ and the variance is given by

$$k^2\frac{\alpha}{(\alpha-1)^2(\alpha-2)}.$$

The k corresponds simply to scaling: multiplying a Pareto(k,α) random variable by a results in a Pareto(ak,α) variable. We restrict k to be £nite and strictly positive.

The probability density function of the Pareto distribution is

$$f(x) = \alpha k^{\alpha}x^{-\alpha-1}, \quad x > k.$$

B.2 Bi-Pareto Distribution

The bi-Pareto distribution has cumulative distribution function

$$\text{BiPareto}(\alpha,\beta,c,k): \quad F(x) = 1-\left(\frac{x}{k}\right)^{-\alpha}\left(\frac{x/k+c}{1+c}\right)^{\alpha-\beta} \quad x \geq k.$$

The distribution looks like a Pareto(α,k) distribution for small values of x and like a Pareto$(\beta,k(1+c)^{1-\alpha/\beta})$ distribution in the tail. The parameter c marks the boundary between the two types of behavior.

The probability density function of the biPareto distribution is

$$f(x) = k^{\beta}(1+c)^{\beta-\alpha}x^{-\alpha-1}(x+kc)^{\alpha-\beta-1}(\beta x+\alpha kc), \quad x > k.$$

We make the obvious restrictions $\beta > 0$, $c \geq 0$, and $k > 0$. The initial exponent must satisfy $\alpha \geq -\beta/c$ and hence can be negative in general. In the applications found the present study, we always had $0 < \alpha < \beta$.

Because the tail is like that of a Pareto distribution with decay exponent β, the ν-th order moments occur if and only if $\beta > \nu$.

The k parameter is related to scaling just as for the Pareto distribution. On the log-log plot, changing k simply shifts the graph to the left or right.

To examine special cases of the distribution, is is useful to rewrite the cdf as

$$F(x) = 1 - k^{\beta}(1+c)^{\beta-\alpha}x^{-\alpha}(x+kc)^{\alpha-\beta}, \quad x \geq k.$$

Then the following special cases and distributional limits become clear.

- biPareto$(\alpha,\beta,0,k)$ = Pareto(β,k)
- $\lim_{c\to\infty}$ biPareto(α,β,c,k) = Pareto(α,k)
- biPareto(α,α,c,k) = Pareto(α,k)
- $\lim_{k\to 0}$ biPareto$(0,\beta,b/k,k)$ = ParetoII(β,b)

The ParetoII(β,b) distribution has the cdf

$$F(x) = 1 - (x/b+1)^{-\beta} \quad x \geq 0.$$

B.3 Moments

For £nite values of k and c, the existence of moments is determined by the tail-decay exponent β. In particular, the r-th moment $\mathbf{E}\{X^r\}$ is £nite for every value of r strictly less than β.

For $r < \beta$, the r-th moment of a biPareto distribution is

$$\mathbf{E}\{X^r\} = k^r\left(1 + rc^{r-\beta}(1+c)^{\beta-\alpha}B_{c/(1+c)}(\beta-r, r-\alpha)\right)$$

where the incomplete beta function is de£ned as

$$B_x(a,b) = \int_0^x u^{a-1}(1-u)^{b-1}\,du$$

for $0 \leq x < 1$ and $a > 0$. The above expression can be derived using the formula

$$\mathbf{E}\{X^r\} = r\int_0^\infty x^{r-1}\,(1-F(x))\,dx,$$

valid for any positive random variable such that $x^r(1-F(x)) \to 0$ as $x \to \infty$.

The expression for the mean can be seen to simplify in several special cases. When $\alpha = 0$, we get $\mu = k(\beta+c)(\beta-1)$. Taking the limit as $k \to 0$ with $c = b/k$, we obtain $kb/(\beta-1)$, the mean of the ParetoII(β, b) distribution. Some other special cases correspond to the Pareto distribution, as explained above.

B.4 Maximum Likelihood Estimator

The maximum likelihood estimator based on N independent samples from a bi-Pareto distribution £nds the four parameters which maximize the joint density function

$$f(\mathbf{x}) = \prod_{i=1}^{N} f(x_i)$$

or equivalently the logarithm of this function,

$$\ln f(\mathbf{x}) = \sum_{i=1}^{N} \ln f(x_i)$$

This equation can't be solved in closed form, but is amenable to numerical optimization. It is convenient for this development to replace the parameters k and c by $u = \ln k$ and $\gamma = \ln c$. Taking partial derivatives with respect to each parameter, we have

$$\frac{\partial}{\partial u} f(x_i) = \beta + \frac{(\alpha-\beta-1)e^{\gamma+u}}{x_i + e^{\gamma+u}} + \frac{\alpha e^{\gamma+u}}{\beta x_i + \alpha e^{\gamma+u}} \tag{5}$$

$$\frac{\partial}{\partial \gamma} \ln f(x_i) = \frac{(\beta-\alpha)e^{\gamma}}{1+e^{\gamma}} + \frac{(\alpha-\beta-1)e^{\gamma+u}}{x_i + e^{\gamma+u}} + \frac{\alpha e^{\gamma+u}}{\beta x_i + \alpha e^{\gamma+u}} \tag{6}$$

$$\frac{\partial}{\partial \alpha} \ln f(x_i) = -\ln(1+e^{\gamma}) + \ln x_i + \ln(x_i + e^{\gamma+u}) + \frac{e^{\gamma+u}}{\beta x_i + \alpha e^{\gamma+u}} \tag{7}$$

$$\frac{\partial}{\partial \beta} \ln f(x_i) = u + \ln(1+e^{\gamma}) - \ln(x_i + e^{\gamma+u}) + \frac{x_i}{\beta x_i + \alpha e^{\gamma+u}} \tag{8}$$

If we exclude the limiting cases $c = 0$ and $c = \infty$, the dimensionality of the problem can be reduced from a four-dimensional optimization to three, because the ML estimate for k is simply $\hat{k} = \min_i x_i$. To see this, note that

$$\left(\frac{\partial}{\partial u} - \frac{\partial}{\partial \gamma}\right) \ln f(x_i) = \beta + \frac{(\alpha - \beta)e^{\gamma}}{1 + e^{\gamma}} = \frac{\beta + \alpha e^{\gamma}}{1 + e^{\gamma}}, \tag{9}$$

independent of x_i. Suppose that D denotes the domain over which the parameters are de£ned and on which the probability function is non-zero, namely $\alpha \geq 0$, $\beta > 0$, $\gamma \in \mathbb{R}$, and $u \leq \min_i \ln(x_i)$. For any set of parameters in the interior of D, (9) is strictly positive. Hence the maximum of the probability function must occur on the boundary of D in either the γ or u dimensions. If the maximum point occurs at a £nite value of $\hat{\gamma}$, the partial derivative with respect to γ must be zero, the partial derivative with respect to u must be positive, and so $\hat{u}$ must be at its upper limit, $\min_i \ln(x_i)$.

Hence, the £rst step in maximum likelihood estimation of a N iid biPareto samples is to take $\hat{k} = \min_i x_i$, or equivalently $\hat{u} = \min_i \ln x_i$. Secondly, use numerical optimization techniques to £nd £nd the values of α, β, and γ which maximize

$$\ln f(\alpha, \beta, \gamma, \hat{u}, \mathbf{x}) = \ln f(\alpha, \beta, \gamma, 0, \mathbf{x}/\hat{k}) - N \ln \hat{k}.$$

From the righthand expression, it follows that the second step can be performed by forming the normalized observation vector $\mathbf{y} = \mathbf{x}/\hat{k}$, and £nding the MLE estimate of the normalized vector over the family of distributions with $k = 1$.

For the biPareto probability function with $k = 1$, the £rst- and second-order partial derivatives are fairly simple to compute, and can be used in the numerical optimization.

$$\begin{aligned}
\frac{\partial}{\partial \alpha} \ln f(x_i) &= -\ln(1 + e^{\gamma}) + \ln x_i + \ln(x_i + e^{\gamma}) + \frac{e^{\gamma}}{\beta x_i + \alpha e^{\gamma}} \\
\frac{\partial}{\partial \beta} \ln f(x_i) &= \ln(1 + e^{\gamma}) - \ln(x_i + e^{\gamma}) + \frac{x_i}{\beta x_i + \alpha e^{\gamma}} \\
\frac{\partial}{\partial \gamma} \ln f(x_i) &= \frac{(\beta - \alpha)e^{\gamma}}{1 + e^{\gamma}} + \frac{(\alpha - \beta - 1)e^{\gamma}}{x_i + e^{\gamma}} + \frac{\alpha e^{\gamma}}{\beta x_i + \alpha e^{\gamma}} \\
\frac{\partial}{\partial \alpha}\frac{\partial}{\partial \alpha} \ln f(x_i) &= -\frac{e^{2\gamma}}{(\beta x_i + \alpha e^{\gamma})^2} \\
\frac{\partial}{\partial \alpha}\frac{\partial}{\partial \beta} \ln f(x_i) &= -\frac{x_i e^{\gamma}}{(\beta x_i + \alpha e^{\gamma})^2}
\end{aligned}$$

$$
\begin{aligned}
\frac{\partial}{\partial\alpha}\frac{\partial}{\partial\gamma}\ln f(x_i) &= \frac{e^\gamma}{x_i+e^\gamma} - \frac{e^\gamma}{1+e^\gamma} + \frac{\beta x_i e^\gamma}{(\beta x_i+\alpha e^\gamma)^2} \\
\frac{\partial}{\partial\beta}\frac{\partial}{\partial\beta}\ln f(x_i) &= -\frac{x_i^2}{(\beta x_i+\alpha e^\gamma)^2} \\
\frac{\partial}{\partial\beta}\frac{\partial}{\partial\gamma}\ln f(x_i) &= \frac{e^\gamma}{1+e^\gamma} - \frac{e^\gamma}{x_i+e^\gamma} + \frac{\alpha x_i e^\gamma}{(\beta x_i+\alpha e^\gamma)^2} \\
\frac{\partial}{\partial\gamma}\frac{\partial}{\partial\gamma}\ln f(x_i) &= \frac{(\beta-\alpha)e^\gamma}{(1+e^\gamma)^2} + \frac{(\alpha-\beta-1)x_i e^\gamma}{(x_i+e^\gamma)^2} + \frac{\alpha\beta x_i e^\gamma}{(\beta x_i+\alpha e^\gamma)^2}
\end{aligned}
$$

References

[1] Paul Barford and Mark Crovella. Generating representative web workloads for network and server performance evaluation. In *ACM SIGMETRICS '98*, pages 151–160, June 1998.

[2] N. H. Bingham, C. M. Goldie, and J. L. Tuegels. *Regular Variation.* Cambridge University Press, Cambridge, 1987.

[3] Jin Cao, William S. Cleveland, Dong Lin, and Don X. Sun. On the nonstationarity of internet traf£c. In *Proc. ACM Sigmetrics*, pages 102–112, 2001.

[4] D. Cox. Long-range dependence: A review. In H. A. David and H. T. David, editors, *Statistics: An Appraisal*, pages 55–74. Iowa State University Press, 1984.

[5] D. Cox and V. Isham. *Point Processes*. Chapman and Hall, 1980.

[6] M.E. Crovella and A. Bestavros. Self-similarity in world wide web traf£c: evidence and possible causes. *Proc. 1996 ACM Sigmetrics conference*, May 1996.

[7] A. Elwalid, D.P. Heyman, T.V. Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for ATM multiplexers with applications to video teleconferencing. *IEEE J. Select. Ar. Comm.*, 13:1004–1016, 1995.

[8] Ashok Erramilli, Onuttom Narayan, Arnold Neidhardt, and Iraj Saniee. Performance impacts of multi-scaling in wide area TCP/IP traf£c. In *IEEE Infocom 2000*, March 2000.

[9] Ashok Erramilli, Onuttom Narayan, and Walter Willinger. Experimental queueing analysis with long-range dependent packet traf£c. *IEEE/ACM Trans. on Networking*, 4:209–223, 1996.

[10] Anja Feldmann. Characteristics of TCP connection arrivals. In *Self-Similar Network Traf£c and Performance Evaluation*. Wiley, New York, 2000.

[11] Anja Feldmann, Anna C. Gilbert, Polly Huang, and Walter Willinger. Dynamics of IP traf£c: A study of the role of variability and the impact of control. In *Proceedings of the ACM/SIGCOMM*, 1999.

[12] William Feller. *An Introduction to Probability Theory and Its Applications*, volume II. Wiley, New York, 2 edition, 1971.

[13] H. Fowler and W. Leland. Local area network traf£c characteristics with implications for broadband network congestion management. *IEEE J. Select. Ar. Comm.*, 9(7):1139–1149, 1991.

[14] E. N. Gilbert and H. O. Pollak. Amplitude distribution of shot noise. *Bell Systems Technical Journal*, 39:333–350, 1960.

[15] Norman L. Johnson and Samuel Kotz. *Distributions in Statistics: Continuous Univariate Distributions -1*. Wiley Interscience Publication, 1975.

[16] Thomas G. Kurtz. Limit theorems for workload input models. In *Stocastic Networks: Theory and Applications*. Clarendon Press, Oxford, 1996.

[17] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traf£c (extended version). *IEEE/ACM Transactions on Networking*, 2(1), February 1994.

[18] Steven B. Lowen and Malvin C. Teich. Doubly stochastic point process driven by fractal shot noise. *Physical Review A*, 43(8):4192–4214, April 1990.

[19] Steven B. Lowen and Malvin C. Teich. Power-law shot noise. *IEEE Transactions on Information Theory*, 36(6):1302–1318, November 1990.

[20] B. Mandelbrot. Self-similar error clusters in communication systems and the concept of conditional stationarity. *IEEE Trans. Communication Technology*, pages 71–90, 1965.

[21] C. Meier-Hellstern, P. Wirth, Y. Yan, and D. A. Hoe¤in. Traf£c models for ISDN data users: of£ce automation application. In A. Jansen and V. B. Iverson, editors, *ITC-13, Copenhagen*, pages 167–172, 1991.

[22] Vern Paxson and Sally Floyd. Wide-area traf£c: The failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.

[23] Sidney I. Resnick. *Extreme Values, Regular Variation, and Point Processes*. Springer-Verlag, New York, 1986.

[24] S. O. Rice. Mathematical analysis of random noise. *Bell System Technical Journal*, 23:1–51, 1945. (Reprinted in *Selected Papers on Noise and Stochastic Processes*, N. Wax, Ed. New York: Dover, 1954, pp. 133-294.).

[25] B.K. Ryu and A. Elwalid. The importance of long-range dependence of VBR video traf£c in ATM traf£c engineering: myths and realities. In *Proc. ACM/SIGCOMM*, pages 3–14, 1996.

[26] R. Sherman, M. S. Taqqu, and W. Willinger. Proof of a fundamental result in self-similar traf£c modeling. *ACM SIGCOM, Computer Communications Review*, 1998.

[27] Walter Willinger, Vern Paxson, and Murad Taqqu. Self-similarity and heavy tails: Structural modeling of network traf£c. In *A Practical Guide to Heavy Tails : Statistical Techniques and Applications*, pages 27–53. Birkhäuser, Boston, 1998.