A Compound Model for TCP Connection Arrivals

Carl J. Nuzman* ‘ Princeton

Iraj Saniee† Bell Labs

Wim Sweldens‡ Bell Labs

Alan Weiss§ Bell Labs

March, 2000

Abstract

We propose a two level model for TCP connection arrivals in local area networks. The first 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 traffic collected a a firewall. We show that traffic synthesized with our model closely characterizes the original data.

1 Introduction

The development of data networking in the 1960s, first as an academic exercise, and subsequently and rapidly as a means of providing a host of new services, presents new challenges and opportunities to the well-established field of teletraffic theory. Apart from a few pioneering investigations, such as Mandelbrot [13], some of the distinguishing features of data traffic from voice telephony were noticed as early as the late 1980s (Folwer and Leland [9], Meier-Hellstern et al. [14]). In their pioneering study of LAN traffic, Willinger et al. [11] presented data and argued for use of alternative models. By showing that sufficiently aggregated data traffic 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 [12], 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 [1, 2] had given a different (M/G/∞) theoretical model of long-range dependent process.

Much of the succeeding investigations (for Web data notably Crovella [3] and Paxson [16]) 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. The potential impacts of self-similarity on the performance of switching systems have been argued not only in favor of the self-similar approach (see Erramilli et al. [6]), but also against it (Elwalid and his coauthors [18] and [4]). In the latter works it was argued that the short-term correlation in data traffic 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 traffic (Reidy and Levi-Vehel [17] and Feldmann et al [8]) and potential performance impacts of multi-fractality (Erramilli et al. [5]).

A key observation made in relation to Web traffic by Paxson [16] was that although the traffic itself exhibited LRD for a variety of applications, the user behavior itself could be modeled as in the standard teletraffic approach using what may be termed loosely as “Poisson models”. As exhaustive as this study was in regard to the variety of protocols and data traces it analyzed, it did not classify the various *layers* present in the data to determine if the same was true at layers below the user.

It is the purpose of this report to model the arrival process at two layers. We used high-resolution packet data, collected at one Bell Labs facility. We aggregated the data into two distinct layers: the “user layer” and the “connection layer”. In contrast to the general methodology followed in much data traffic analysis attempted so far, we look at the “arrival process” for each of the above-mentioned layers. The arrival process has the advantage that it is easily converted to the “count process” used extensively in the study of self-similar and multi-fractal phe-

*cjnuzman@ee.princeton.edu
†iis@bell-labs.com
‡wim@bell-labs.com
§apdoo@bell-labs.com

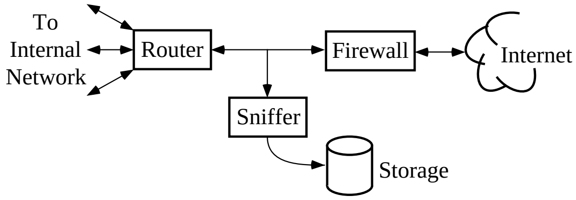

Figure 1: The data collection system

nomenon, while it provides a possible link to the more standard teletraffic theory. Our observations for the "user layer" and "connection layer" largely agree with a previous study by Feldmann [7]. We give an explicit model of the two layers that captures the interarrival characteristics of data traffic, and lends itself easily to simulation. Specifically, we model the arrival process of user sessions as Poisson, and we model user session characteristics and TCP connection timing using heavy-tailed distributions such as Weibull and Pareto. The resulting simulated (synthesized) trace matches many key statistics of the actual data.

Earlier work by Paxson and Floyd [16] and [7] found non-Poisson arrivals of connections. Further, Feldmann observed fat-tailed marginals for TCP connections. Our contribution consists of extending that work to include the correlated structure of interarrival times, and the resulting link with self-similar models based on packet/byte counts.

2 TCP Connection Arrival Process

2.1 Data Set

Our data was collected by a system at Bell Laboratories in Murray Hill, New Jersey by Don Sun, Dong Lin, and William Cleveland. As pictured in Figure 1, the math and computer science research center at Bell Labs is connected to the Internet through a firewall. The firewall connects to the math center via a router; the data was collected from the link between this router and the firewall. A packet sniffer recorded all the packet headers. Our data comes from a 27-hour period beginning the evening of November 17, 1998.

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 traffic. The majority of TCP packets are HTTP.

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

3 Model-based Measurements

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.

3.1 Definition 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 definition 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 defined 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.

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. The grouping of packets into connections ad user sessions is sketched in Figure 2.

This method, of course, has the problem that different users sharing the same machine cannot be distinguished, and this can lead to extremely long users session.

In choosing a particular value of τ, we wanted to use a time that was on the order of Web browsing speed. In a typical Web browsing behavior, a human user clicks on a link. The Web browser opens up several simultaneous connections to retrieve the information. Then, when the information is presented, the user may take a few seconds

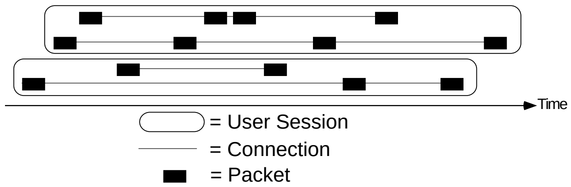

Figure 2: Sketch of user sessions, TCP connections, and packets.

or minutes to digest the information before locating the next desired link. We wanted τ to be large enough to keep such a sequence of clicks together, but not so large as to group essentially unrelated Web uses together. If τ is very large, each user generates only a single user session, and many more users are considered to be active at the same time. Note that if the pool of potential users were not large, then the arrival rate might depend on the number of active users, and a Poisson model would not be appropriate.

3.2 User Session Arrivals Are Poisson

User sessions are supposed to resemble phone calls, at least in the process of initiation (arrival process). We tested this supposition several ways; the bottom line is that overall user session arrivals look like a time-varying Poisson process, with very slow (diurnal) variation in arrival rate.

Figure 3 compares the quantiles of an hour's worth of user session interarrivals with quantiles taken from an exponential distribution. The linear nature of the plot indicates that this distribution provides a good fit.

The interarrivals were uncorrelated, as depicted in Figure 4, and appeared independent. We arrived at the latter conclusion by examining the correlation of powers (squares, cubes) of interarrival times; the results were so similar to Figure 4 as to not be worth showing. Note, too, that the autocorrelation pictured in Figure 4 lies almost entirely between two dotted horizontal lines that represent the 95% confidence interval for an autocorrelation of zero.

3.3 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 define C as the number of TCP connections that are part of that user session. If

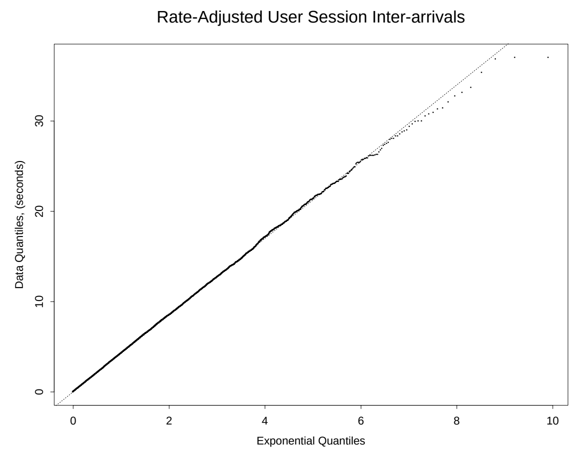

Figure 3: Fit of TCP user session interarrivals to an exponential distribution.

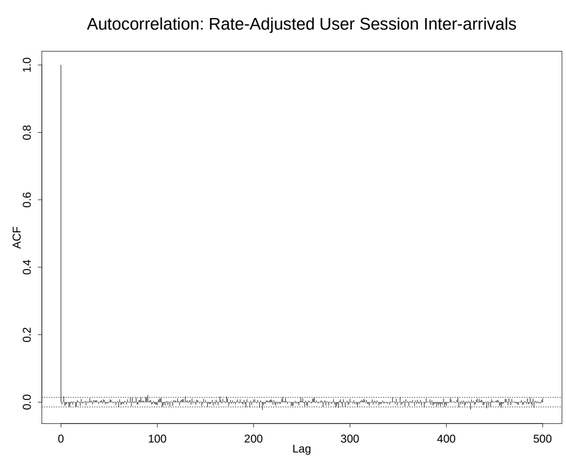

Figure 4: Autocorrelation of the TCP user session interarrivals sequence.

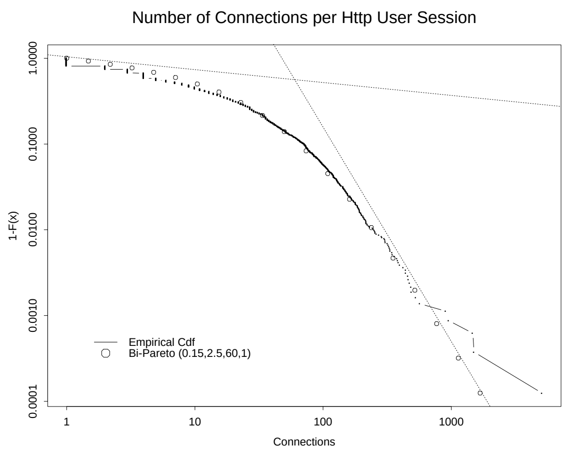

Figure 5: Empirical tail distribution and biPareto fit of C, the number of connections per user session.

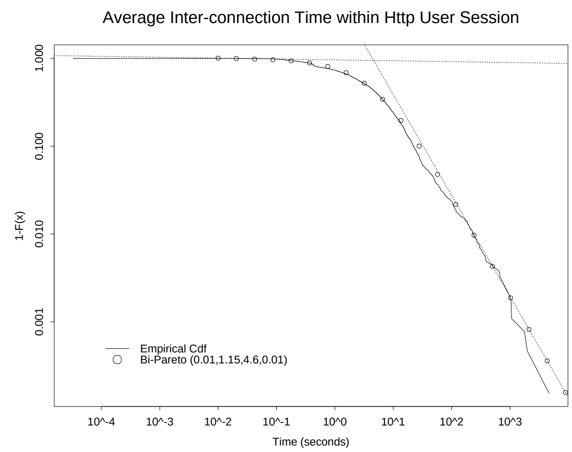

Figure 6: Empirical tail distribution and biPareto fit of μ, the mean connection interarrival time within a user session.

$C > 1$ we define two statistics. T is the total interarrival time; it is the time from the first connection arrival to the last connection arrival. The *average connection interarrival time* is defined 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 3.6.

3.4 Model Statistics for Connection Interarrivals

Figure 5 shows the tail distribution function of C, the number of connections generated by a user session. Figure 6 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 6, there are two nearly linear regimes, with a smooth transition between them. In Figure 5, there appear to be two regimes, although the regime on the left is never linear.

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 flexibility needed to model empirical distributions such as that in Figure 6 effectively, we defined the *biPareto* distribution, which has a tail distribution function given by

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

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 $\beta > 0$. The Pareto and ParetoII distributions are special cases of the biPareto distribution [15].

The circles in Figures 5 and 6 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 fit between the measured data and the biPareto distributions is excellent.

3.5 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 coefficient between μ and C was measured to be only -0.013 in our data set. A further illustration of the the relationship between μ and C is given in Figure 7, which depicts the estimated mean of μ conditioned on various ranges of C. The user sessions are grouped according to the number of connections they contain; for example, the sessions with

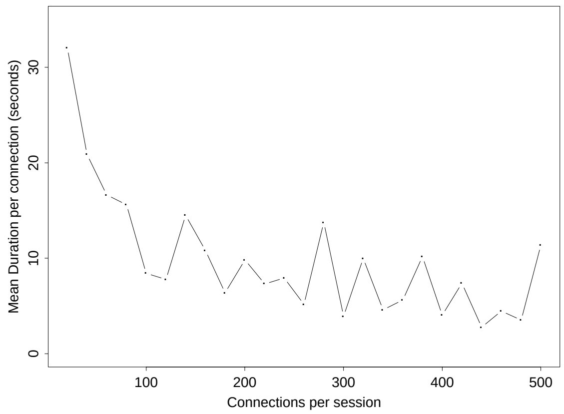

Figure 7: Conditional mean of μ based on ranges of C.

1 to 20 connections form the first group and the sessions with 21 to 40 connections form the second group. The mean of the average connection interarrival time is plotted for each group. If the two parameters were independent, this mean value should be constant, which is certainly not the case. However, the mean of μ appears roughly independent of C for sufficiently large values of C.

In the mathematical model, we define C and μ to be mutually independent random variables. This simplification 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 define a few different condition distributions for μ conditioned on ranges of C.

3.6 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 defined 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 [10] and because its coefficient of variation can be easily adjusted using the shape parameter, c.

The Weibull distribution function is

$$F(x) = 1 - \exp\left(-(x/b)^c\right), \quad x > 0, \quad b, c > 0,$$

where b and c are known as the scale and shape parameters, respectively. A Weibull(b,c) random variable $Y = X^{1/c}$ can be obtained by taking the c-th root of an exponential random variable X with mean b. When $c = 1$, the result is an exponential distribution, while for $c < 1$, the Weibull distribution is more variable than the exponential, in the sense that extremely small or large values are more likely. The HTTP connection interarrivals were well-modeled by a Weibull distribution with shape parameter $c = 0.75$. The coefficient of variation, or ratio of variance to squared mean, of the interarrivals was about twice that of an exponential distribution.

The coefficient of variation of a distribution is the ratio of the variance to the square of the mean; for positive distributions, it is a normalized measure of variability. The coefficient of variation of an exponential distribution is unity. For a Weibull(b,c) distribution, the coefficient of variation $V(c)$ is given by

$$V(c) = \frac{\Gamma(1+2/c)}{\Gamma(1+1/c)^2} - 1.$$

When $c < 1$, $V(c)$ is greater than one, and the Weibull interarrivals are more variable than exponential interarrivals.

In the synthesis described below, we used $c = 0.48$ so that the coefficient of variation matched the measured *average* coefficient of variation for within-session connection interarrival times, which was 5.6 in our data set.

3.7 Truncated Distributions

Some of the measurements in the data set appear to come from a *truncated* biPareto distribution. By a truncated distribution, we mean the distribution that arises when all realizations of a random variable which exceed some threshold M are thrown out. Let $F(x)$ be any cdf, $M > 0$ be a

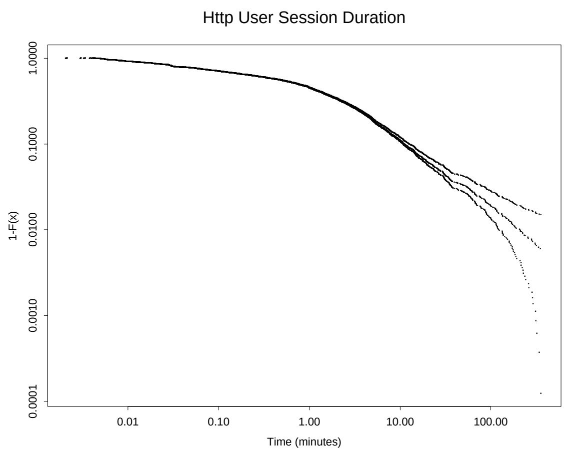

Figure 8: Empirical tail distribution of HTTP user session duration. Bottom: original estimate. Middle: assuming a truncation probability of 0.006. Top: assuming a truncation probability of 0.012.

truncation threshold, and $F(M)$ be the probability that the original random variable does not exceed the threshold. Then the distribution function $F_M(x)$ of the truncated distribution is

$$F_M(x) = \begin{cases} \frac{F(x)}{F(M)} & x < M \\ 1 & x \geq M \end{cases}$$

From each distribution $F(x)$, truncation produces a family of distributions $F_M(x)$ parameterized by M. How can we identify the $F(x)$ which gives rise to $F_M(x)$? If $F(M)$ is known, $F(x)$ can be recovered for $x < M$ as $F(x) = F(M)F_M(x)$. If $F(M)$ is not known, we can search for a scaling of $F_M(x)$ which produces a plausible $F(x)$.

This latter technique is illustrated in Figure 8. The lower line in the plot depicts the empirical distribution of the duration D of user sessions from the HTTP data set, in seconds. The duration is defined to be the time from the beginning of the first connection to the end of the last connection within a user session. Note that on this log-log scale, the distribution turns sharply off toward negative infinity, indicating that the distribution may be truncated. An estimate of the truncation threshold M is given by the maximum value observed, which is 22011 seconds, or about 6 hours. If we assume the truncation probability is $1-F(M) = 0.006$ and plot $1-F(x) = 1-F(M)F_M(x)$, we obtain the middle line of the plot, which looks like a good candidate for a biPareto model. If we assume a truncation probability of $1-F(M) = 0.012$, we get the top line; we have gone "too far" if we want to use a biPareto model.

4 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 coefficient 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 3.6. 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 is depicted in Figure 9. A straight line would indicate a perfect fit. In fact, the plot is curved "downward", indicating that the marginal of the simulated interarrivals is slightly more extreme than that of the data.

The autocorrelation of the actual and simulated interarrivals are shown in Figure 10. Again the match is very good, though not exact. Because of the long-range dependence in the data, the empirical autocorrelation varies significantly from realization to realization.

A measure of the accuracy of the model across different scales is depicted in Figure 11. In this figure, 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 figure, we have simply measured the coefficient 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" is a technique which directly matches the marginal and autocorrelation of the observed interarrival sequence, as described in the full version of

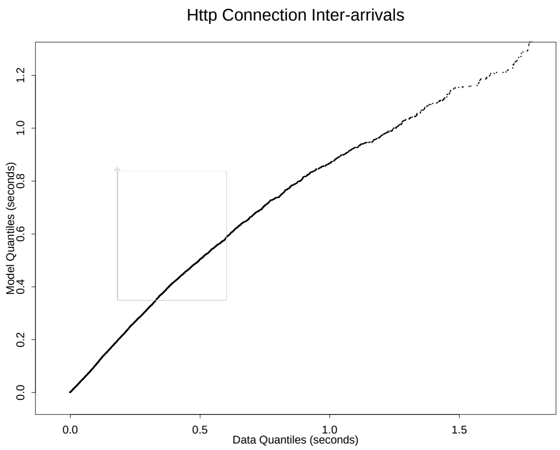

Figure 9: Fit of synthesized TCP connection interarrivals to data.

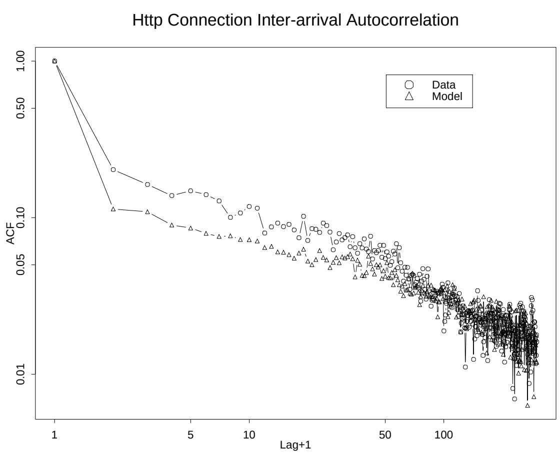

Figure 10: Empirical autocorrelation function of TCP connection interarrival sequence: from data and from model.

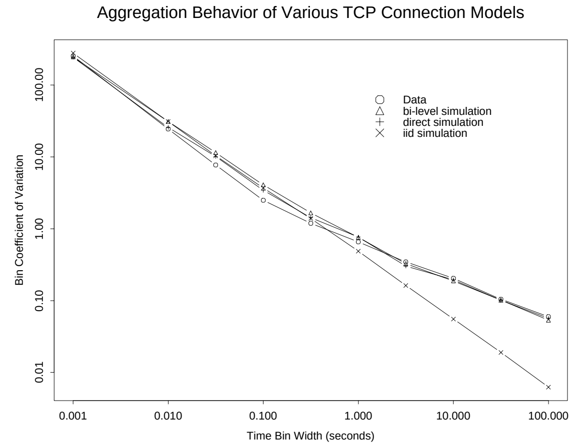

Figure 11: TCP connection count process coefficient of variation as a function of bin length, for data and various models.

this paper. 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 the data in terms of aggregated coefficient of variation, while the i.i.d. simulation departs significantly from the data as the time scale increases. It is interesting to note a dip in the coefficient of variation of the data on the order of 100 ms, which is not reproduced by any of the three models.

We have, in Figures 12 and 13, 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 efficient; 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.

5 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. It may be possible to simulate the packet-level arrivals us-

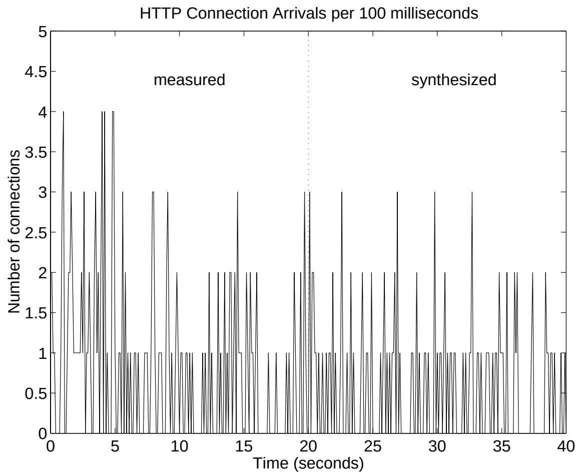

Figure 12: HHTP Connection arrivals in 100 ms intervals

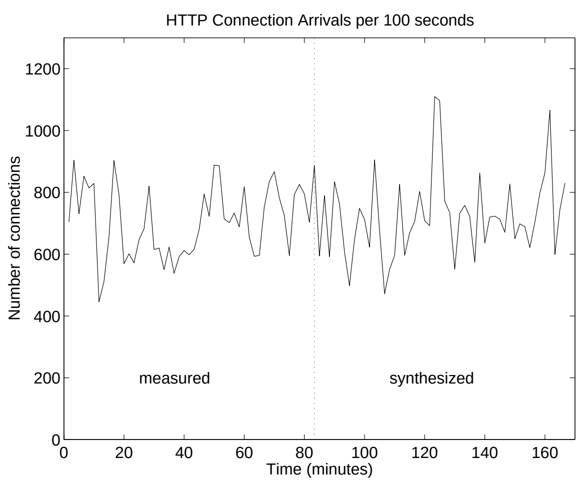

Figure 13: HHTP Connection arrivals in 100 second intervals

ing 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 3.6 may be sufficient in order to reproduce empirical behavior in sufficiently aggregated packet traffic. Also, detailed queueing analysis of the present two-level model would test the applicability of our results to buffer statistics.

References

[1] 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.

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

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

[4] 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.

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

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

[7] Anja Feldmann. Characteristics of TCP connection arrivals. 1998. Unpublished.

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

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

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

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

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

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

[14] C. Meier-Hellstern, P. Wirth, Y. Yan, and D. A. Hoeflin. Traffic models for ISDN data users: office automation application. In A. Jansen and V. B. Iverson, editors, *ITC-13, Copenhagen*, pages 167–172, 1991.

[15] Carl J. Nuzman. Notes on the biPareto distribution. Unpublished, January 2000.

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

[17] R.H. Reidy and J. Levy Vehel. TCP traffic is multifractal: a numerical study. *INRIA Research Report*, (3129), 1997. unpublished.

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