

Wide-Area Traffic: The Failure of Poisson Modeling *

Vern Paxson and Sally Floyd †
Lawrence Berkeley Laboratory and
EECS Division, University of California, Berkeley
1 Cyclotron Road, Berkeley CA 94720
vern@ee.lbl.gov, floyd@ee.lbl.gov

July 18, 1995

Abstract

Network arrivals are often modeled as Poisson processes for analytic simplicity, even though a number of traffic studies have shown that packet interarrivals are not exponentially distributed. We evaluate 24 wide-area traces, investigating a number of wide-area TCP arrival processes (session and connection arrivals, FTP data connection arrivals within FTP sessions, and TELNET packet arrivals) to determine the error introduced by modeling them using Poisson processes. We find that user-initiated TCP session arrivals, such as remote-login and file-transfer, are well-modeled as Poisson processes with fixed hourly rates, but that other connection arrivals deviate considerably from Poisson; that modeling TELNET packet interarrivals as exponential grievously underestimates the burstiness of TELNET traffic, but using the empirical Tcplib [Danzig et al, 1992] interarrivals preserves burstiness over many time scales; and that FTP data connection arrivals within FTP sessions come bunched into “connection bursts,” the largest of which are so large that they completely dominate FTP data traffic. Finally, we offer some results regarding how our findings relate to the possible *self-similarity* of wide-area traffic.

1 Introduction

When modeling network traffic, packet and connection arrivals are often assumed to be Poisson processes because such processes have attractive theoretical properties [FM94]. A number of studies have shown, however, that for both local-area and wide-area network traffic, the distribution of packet interarrivals clearly differs from exponential [JR86, G90, FL91, DJCME92]. Recent work argues convincingly that LAN traffic is much better modeled using statistically *self-similar* processes [LTWW94], which have much different theoretical properties than Poisson processes. For self-similar traffic, there is no natural length for a “burst”; traffic bursts appear on a wide range of time scales. In this paper we show that for wide-area traffic, Poisson processes are valid only for modeling the arrival of user sessions (TELNET connections, FTP control connections); that they fail as accurate models for other WAN arrival processes; and that WAN packet arrival processes appear better modeled using self-similar processes.

* This paper appeared in *IEEE/ACM Transactions on Networking*, 3(3), pp. 226-244, June 1995. A preliminary version of this paper appeared in the Proceedings of SIGCOMM '94.

† This work was supported by the Director, Office of Energy Research, Scientific Computing Staff, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

For our study we analyze 24 traces of wide-area TCP traffic. We consider both previous and new models of aspects of TELNET and FTP traffic, discuss the implications of these models for burstiness at different time scales, and compare the results of the models with the trace data. We show that in some cases commonly-used Poisson models seriously underestimate the burstiness of TCP traffic over a wide range of time scales. (We restrict our study to time scales of 0.1 seconds and larger.)

We first show that for interactive TELNET traffic, *connection* arrivals are well-modeled as Poisson with fixed hourly rates. However, the exponentially-distributed interarrivals commonly used to model *packet* arrivals generated by the user side of a TELNET connection grievously underestimate the burstiness of those connections, and high degrees of multiplexing do not help. Using the empirical Tcplib [DJ91, DJCME92] distribution for TELNET packet interarrivals instead results in packet arrival processes significantly burstier than Poisson arrivals, and in close agreement with traces of actual traffic. From these findings we then construct a model of TELNET traffic parameterized by only the hourly connection arrival rate and show that it accurately reflects the burstiness found in actual TELNET traffic. (We do not model the TELNET response, only the user side.) The success with this model of using Tcplib packet interarrivals confirms the finding in [DJCME92] that the arrival pattern of user-generated TELNET packets has an invariant distribution, independent of network details.

For small machine-generated bulk transfers such as SMTP (email) and NNTP (network news), connection arrivals are not well-modeled as Poisson, which is not surprising since both types of connections are machine-initiated and can be timer-driven. Previous research has discussed how the periodicity of machine-generated IP traffic such as routing updates can result in network-wide traffic synchronization [FJ94], a phenomenon impossible with Poisson models.

For large bulk transfer, exemplified by FTP, the traffic structure is quite different than suggested by Poisson models. As with TELNET connections, user-generated FTP session arrivals are well-modeled as Poisson with fixed hourly rates. However, we find that FTP data connections within a single FTP session (which are initiated whenever the user lists a directory or transfers a file) come clustered in *bursts*. Hereafter we will refer to these data connections as FTPDATA connections, and the corresponding bursts as FTPDATA bursts. Neither FTPDATA-connection nor FTPDATA-burst arrivals are well-modeled as Poisson processes. Furthermore, one of our key findings is that the distribution of the number of bytes in each burst has a very heavy upper tail; a small fraction of the largest bursts carries almost all of the FTPDATA bytes. This implies that

faithful modeling of FTP traffic should concentrate heavily on the characteristics of the largest bursts.

Poisson arrival processes are quite limited in their burstiness, especially when multiplexed to a high degree. Our findings, however, show that wide-area traffic is much burstier than Poisson models predict, over many time scales. This greater burstiness has implications for many aspects of congestion control and traffic performance. We conclude the paper with a discussion of how our burstiness results mesh with self-similar models of network traffic, and then with a look at the general implications of our results.

2 Traces used

Dataset	Date	Duration	What
Bellcore (BC)	10Oct89	13 days	17K TCP conn.
U.C.B. (UCB)	31Oct89	24 hours	38K TCP conn.
coNCert (NC)	04Dec91	24 hours	63K TCP conn.
UK-US (UK)	21Aug91	17 hours	26K TCP conn.
DEC 1-3	See refs.	24 hours × 3	195K TCP conn.
LBL 1-8	See refs.	30 days × 8	3.7M TCP conn.

Table 1: Summary of Wide-Area TCP Connection Traces

Dataset	Date	When	What
LBL PKT-1	Fri 17Dec93	2PM-4PM	1.7M TCP pkts.
LBL PKT-2	Wed 19Jan94	2PM-4PM	2.4M TCP pkts.
LBL PKT-3	Thu 20Jan94	2PM-4PM	1.8M TCP pkts.
LBL PKT-4	Fri 21Jan94	2PM-3PM	1.3M pkts.
LBL PKT-5	Fri 28Jan94	2PM-3PM	1.3M pkts.
DEC WRL-1	Wed 08Mar95	10PM-11PM	3.3M pkts.
DEC WRL-2	Thu 09Mar95	2AM-3AM	3.9M pkts.
DEC WRL-3	Thu 09Mar95	10AM-11AM	4.3M pkts.
DEC WRL-4	Thu 09Mar95	2PM-3PM	5.7M pkts.

Table 2: Summary of Wide-Area Packet Traces

Our study is based on two sets of traces of wide-area network traffic. The first set, shown in Table 1, consisted of TCP SYN/FIN connection start/stop packets. SYN/FIN packets are enough to measure connection start times (and hence connection arrival processes), durations, TCP protocol, participating hosts, and data bytes transferred in each direction. The BC and UCB traces are analyzed in depth in [DJCME92], and also in [P94a], and the UCB trace forms the basis of the connection characteristics used for Teplib [DJ91]. The NC, UK, and DEC traces are analyzed in [P94a], and the LBL traces are analyzed in [P94a, P94b]. The “DEC 1-3” rows represent three wide-area TCP SYN/FIN traces, each spanning 1 day, and the “LBL 1-8” row represents 8 wide-area TCP SYN/FIN traces, each spanning 30 days. The reader is referred to the abovementioned papers for details regarding the characteristics of the traffic in each dataset, including the number of connections and bytes due to each TCP protocol.

These traces are all fairly lengthy, allowing us to assess how traffic varies over the course of a day or longer, and giving us enough TCP connection arrivals to make a statistically sound evaluation of the connection arrival processes. These traces are used in § 3 to evaluate the effectiveness of using Poisson models for TCP connection arrivals. Because SYN/FIN traces allow us to characterize

connection size, we also used these trace in § 6 to investigate the notion of “FTPDATA bursts.”

Because the SYN/FIN traces do not contain information regarding packet arrivals within a connection, to evaluate *packet* arrival processes we acquired nine packet-level traces of wide-area traffic, summarized in Table 2.¹

The “LBL PKT-*n*” rows summarize traces gathered at the Lawrence Berkeley Laboratory’s wide-area Internet gateway. The first three traces captured all TCP packets, and lasted two hours. The final two traces captured all packets and lasted one hour. In the first set of traces, the fraction of dropped packets, where known, was always $\leq 5 \cdot 10^{-6}$. For the second set, it was always ≤ 0.001 .

The “DEC WRL-*n*” rows summarize traces gathered at the primary Internet access point for the Digital Equipment Corporation. The access point is operated by Digital’s Palo Alto research groups, and the traces were supplied by Digital’s Western Research Lab (hence “WRL”). For these traces, the fraction of dropped packets was always ≤ 0.00025 .

The packet traces do not include a large number of TCP connections, unlike the traces in Table 1, so we do not use them for evaluating Poisson models for TCP connection arrivals, nor for the size of FTPDATA bursts (though the traces are used to illustrate the heaviness of the distribution’s upper tail). Instead we use the LBL PKT datasets in § 4 and § 5 to evaluate different models for TELNET packet arrivals, and both the LBL PKT and the DEC WRL datasets in § 7 to investigate the presence of “large-scale correlations” in wide-area network traffic. (We did not include the DEC WRL datasets in our packet-level TELNET evaluation because, due to the use of a firewall proxy server, the DEC TELNET traffic is dominated by a single, heavily-loaded machine.)

To disambiguate between the LBL and DEC SYN/FIN traces and packet traces, we use LBL-*n* and DEC-*n* to refer to SYN/FIN traces, and LBL PKT-*n* and DEC WRL-*n* to refer to packet traces.

3 TCP connection interarrivals

This section examines the connection start times for several TCP protocols. The pattern of connection arrivals is dominated by a 24-hour pattern, as has been widely observed before. We show that for TELNET connection arrivals and for FTP session arrivals, within one-hour intervals the arrival process can be well-modeled by a homogeneous Poisson process; each of these arrivals reflects an individual user starting a new session. Over one hour intervals, no other protocol’s connection arrivals are well-modeled by a Poisson process. Even if we restrict ourselves to ten-minute intervals, only FTP session and TELNET connection arrivals are statistically consistent with Poisson arrivals, though the arrival of SMTP connections and of FTPDATA “bursts” (discussed later in § 6) during ten-minute intervals are not terribly far from what a Poisson process would generate. The arrivals of NNTP, FTPDATA, and WWW (World Wide Web) connections, on the other hand, are decidedly not Poisson processes.

Figure 1 shows the mean hourly connection arrival rate for datasets LBL-1 through LBL-4. For the different protocols, we plot for each hour the fraction of an entire day’s connections of that

¹The BC and UCB traces listed in Table 1 actually include all packets, and are analyzed as such in [DJCME92]. We excluded a packet-level analysis of the BC dataset because of its low traffic rate (on average, about 1 packet/sec over the 11 days), and the UCB dataset because it forms the basis of the Teplib library, against which we compare the packet-level traces.

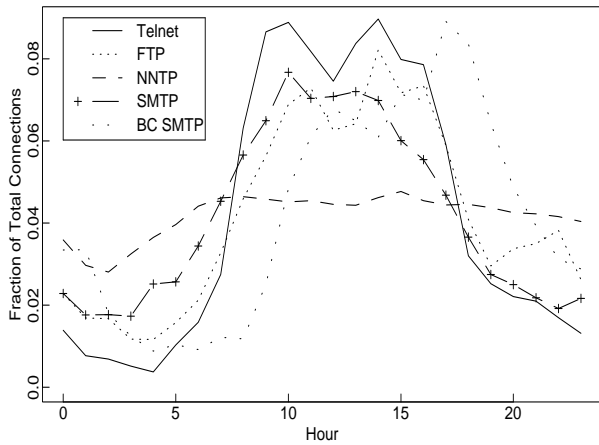


Figure 1: Mean, relative, hourly connection arrival rate for LBL-1 through LBL-4 datasets.

protocol occurring during that hour. (In the figure, FTP refers to FTP sessions.) For example, TELNET connections occur primarily during normal office hours, with a lunch-related dip at noontime; this pattern has been widely observed before. FTP file transfers have a similar hourly profile, but they show substantial renewal in the evening hours, when presumably users take advantage of lower networking delays. The NNTP traffic maintains a fairly constant rate throughout the day, only dipping somewhat in the early morning hours (but the mean size of each connection varies over the course of the day; see [P94a]). The SMTP traffic is interesting because it shows a morning bias for the LBL site (west-coast U.S.) and an afternoon bias for the Bellcore site (east-coast U.S.); perhaps the shift is due to cross-country mail arriving relatively earlier in the Pacific time zone and later in the Atlantic time zone.

Figure 1 shows enough daily variation that we cannot reasonably hope to model connection arrivals using simple homogeneous Poisson processes, which require constant rates. The next simplest model is to postulate that during fixed-length intervals (say, one hour long) the arrival rate is constant and the arrivals within each interval might be well modeled by a homogeneous (fixed-rate) Poisson process. Telephone traffic, for example, is fairly well modeled during one-hour intervals using homogeneous Poisson arrival processes [FL91].

To evaluate these Poisson models, we developed a simple statistical methodology (Appendix A) for testing whether arrivals during a given one-hour or ten-minute interval are Poisson with a fixed rate. We test two aspects of each protocol’s interarrivals: whether they are consistent with exponentially distributed interarrivals, and whether they are consistent with independent interarrivals. If the arrivals during the interval are truly Poisson, then we would expect 95% of the tested intervals to pass each test. Note that we would also expect testing ten-minute intervals to perhaps be more successful than testing one-hour intervals, because using ten-minute intervals allows the arrival rate to change six times each hour rather than remaining constant throughout the hour.

We applied our methodology to all of the TCP connection traces shown in Table 1. For each trace, we separately tested the trace’s TELNET, FTP, FTPDATA, SMTP, NNTP, and WWW connections. Only two of the traces had significant WWW traffic, but as use of this protocol is rapidly growing, it is worth investigating even given

the limited samples.

FTP here refers to an FTP *session* (i.e., an FTP control connection), while FTPDATA refers to the data-transfer connections spawned by these control connections. Prior to our analysis we removed the periodic “weather-map” FTP traffic discussed in [P94b], to avoid skewing our results. We also tested arrivals of FTPDATA *bursts* (see § 6 below).

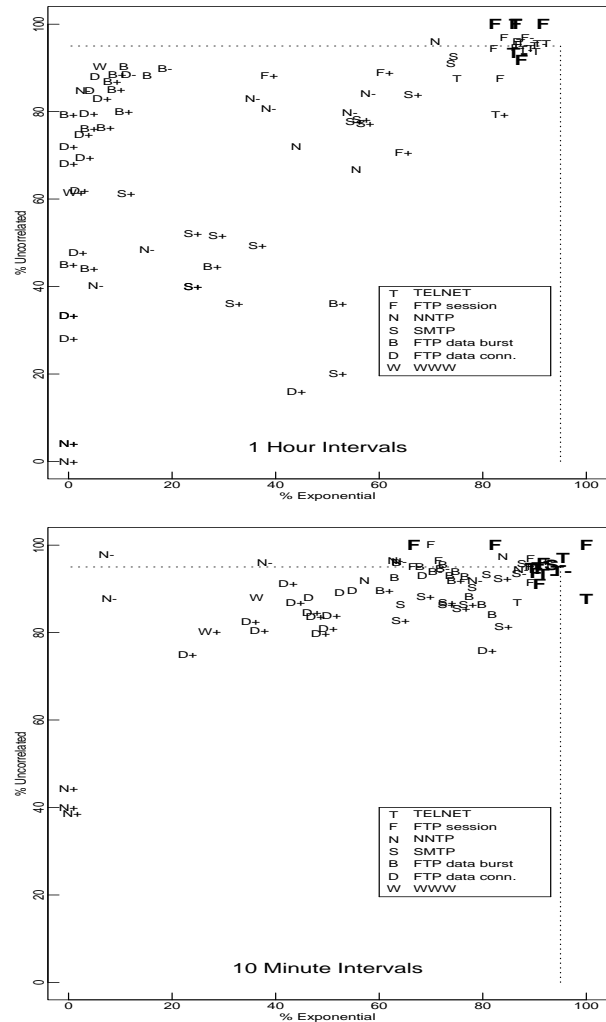


Figure 2: Results of testing for Poisson arrivals.

Figure 2 shows the results of our tests, for both one-hour intervals (top plot) and ten-minute intervals (bottom plot). Along the x -axis we plot the percentage of tested intervals that passed the statistical test for exponentially distributed interarrivals, and along the y -axis the percentage that passed the test for independent interarrivals. The dashed lines correspond to a 95% pass-rate, which we would expect on average if the arrivals were truly Poisson. In general, we expect Poisson arrivals to cluster near the upper right corner of the plots.

Each letter in a plot corresponds to a single trace’s connection arrivals for the given TCP protocol. Letters drawn in large **bold** indicate that the trace’s arrivals are statistically indistinguishable from Poisson arrivals (see Appendix A for details). A + or – after a letter indicates that consecutive interarrival times are consistently either positively or negatively correlated, even if the correlation

itself is not particularly strong (again, see Appendix A).

We see immediately that TELNET connection arrivals and FTP session arrivals are very well modeled as Poisson, both for 1-hour and 10-minute fixed rates. No other protocol’s arrivals are well modeled as Poisson with fixed hourly rates. If we require fixed rates only over 10-minute intervals, then SMTP and FTPDATA burst arrivals are not terribly far from Poisson, though neither is statistically consistent with Poisson arrivals, and consecutive SMTP interarrival times show consistent positive correlation. NNTP, FTPDATA, and WWW arrivals, on the other hand, are clearly not Poisson.

That NNTP and to a lesser extent SMTP arrivals are not Poisson is not too surprising. Because of the flooding mechanism used to propagate network news, NNTP connections can immediately spawn secondary connections as new network news is received from one remote peer and in turn offered to another. NNTP and SMTP connections are also often timer-driven. Finally, SMTP connections are perturbed by mailing list explosions in which one connection immediately follows another, and possibly by timer effects due to using the Domain Name Service to locate MX records [P86].

That FTPDATA connection arrivals are clearly not Poisson can be readily attributed to the fact that “multiple-get” file transfers often result in a rapid succession of FTPDATA connections, one immediately following another [P94a]. Coalescing multiple FTPDATA connections into single *burst* (§ 6) arrivals improves the 10-minute Poisson fit somewhat, but still falls short of statistical consistency.

The finding that TELNET connection arrivals are well-modeled as a Poisson process with fixed hourly rates is at odds with that of [MM85], who found that user interarrival times looked “roughly log-normal”. We believe the discrepancy is due to characterizing the distribution of all of the interarrivals lumped together, rather than postulating separate hourly rates.

Given that TELNET connection arrivals appear Poisson over one-hour intervals, one might imagine that other human-initiated traffic such as RLOGIN and X11 will also fit this model. We find that RLOGIN does and X11 does not. We conjecture that the difference is that during a single X11 *session* (corresponding to running an instance of *xterm*, say) a user initiates multiple X11 connections, while TELNET and RLOGIN sessions are comprised of a single TCP connection. Thus, TELNET connection arrivals correspond to users deciding to *begin* using the network; X11 connection arrivals correspond to users deciding to do something new *during* their use of the network. The former behavior is likely to be close to uncorrelated, memoryless arrivals, since each arrival generally involves a new user. The latter is much more akin to the creation of FTPDATA connections during a single FTP session, since a single user is involved in generating new arrivals. Because X11 connections are created in this way, their arrivals do not have the memoryless property and hence are not Poisson. If we could discern between X11 session arrivals and X11 connection arrivals, then we conjecture we would find the session arrivals to be Poisson.

4 TELNET packet interarrivals

The previous section showed that start times for TELNET connections are well-modeled by Poisson processes. In this section we look at the packet arrival process within a TELNET connection. We restrict our study to packets generated by the TELNET connection originator; this in general is a user typing at a keyboard. We would expect the packets generated by the TELNET connection responder to have a somewhat different arrival process, since they

will usually include both echoes of the user’s keystrokes and larger bursts of bulk-transfer consisting of output generated by the user’s remote commands.

Because the originator packets are initiated by a human, we might hope that the arrival process is to some degree “invariant”; that is, the process may be independent of network dynamics and instead mainly reflect the delays and bursts of activity associated with people typing commands to a computer. Indeed, our empirical results of the interarrival times between packets in a single TELNET connection are consistent with the empirical Tcplib distribution found by previous researchers. Unlike the exponential distribution, the empirical distribution of TELNET packet interarrival times is *heavy-tailed*; we show that using the exponential distribution results in seriously underestimating the burstiness both of TELNET traffic within a single connection and of multiplexed TELNET traffic. Modeling TELNET packet arrivals by a Poisson process, as is generally done, can result in simulations and analyses that significantly underestimate performance measures such as average packet delay.

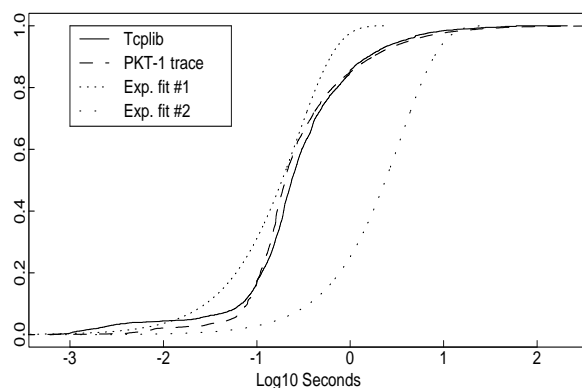


Figure 3: Empirical distributions of packet-interarrivals within TELNET connections.

Figure 3 shows two empirical distributions of the interarrival times of packets within TELNET connections. The solid line shows the distribution used by Tcplib [DJ91, DJCME92]; the dashed line shows the same for the LBL PKT-1 trace. Above 0.1 seconds, the agreement is quite good, especially in the upper tail. That different sites produce the same distribution argues heavily that the distribution is independent of network dynamics and instead reflects human typing dynamics. Below 0.1 seconds the interarrival times probably *are* dominated by network dynamics; but, as stated earlier, in this paper we are not concerned with time scales below 0.1 seconds.

Even ignoring the lower tail, the interarrival distribution is not even close to exponential in shape (note that the x -axis is logarithmically scaled). To dramatize this fact, we have also plotted two logarithmically-scaled exponential distributions. The lefthand one (“fit #1”) has the same geometric mean as the LBL PKT-1 distribution, and the righthand one has the same arithmetic mean.

The exponential fits are very poor. Using the exponential distribution fitted to the same geometric mean will faithfully capture only the distribution of packet interarrivals that are between 200 and 400 msec apart. Shorter interarrivals will be overestimated, and longer interarrivals will be underestimated. For example, the exponential distribution models a full 25% of the interarrivals as being less than 8 msec, and only 2% as being longer than 1 sec, but for the actual data under 2% were less than 8 msec apart, and over 15%

were more than 1 sec apart.

The exponential distribution fitted to the arithmetic mean fares even worse. For example, it predicts nearly 70% of the packets will arrive more than 1 sec apart, when the actual observed distribution is 15% of the packets.

Thus, simple exponential distributions for packet interarrival times, which are necessary for Poisson models of packet arrivals, provide very poor fits to the observed distribution. On the other hand, the main body of the observed distribution fits very well to a Pareto distribution (doubly-exponential; see Appendix B) with shape parameter $\beta \approx 0.9$, and the upper 3% tail to a Pareto distribution with $\beta \approx 0.95$. Interestingly, a Pareto distribution with $\beta < 1$ has infinite mean and variance; a very different beast than an exponential distribution. We will see later that Pareto-distributed interarrivals lead to observable large-scale correlations (Appendix C).

It is not surprising that interactive packet arrivals do not fit a Poisson model, since earlier work looking at many different components of interactive traffic failed to find any statistically significant exponential fits to the observed distributions [FJ70]. This leaves the question: What are the consequences of using Poisson packet arrivals rather than the Tcplib distribution for TELNET traffic?

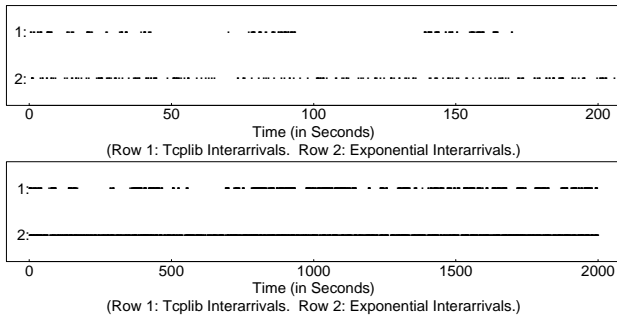


Figure 4: Comparisons between Tcplib and exponential interpacket times.

Figure 4 shows two views of packet arrivals from two simulated TELNET connections, each lasting 2,000 seconds. The first graph shows the first 200 seconds, and the second graph the entire 2,000 seconds. Row 1 for each graph shows a connection using independent, identically-distributed (i.i.d.) interpacket times from the Tcplib distribution, and row 2 shows a connection using i.i.d. interpacket times from an exponential distribution with a mean of 1.1 seconds (to give roughly the same number of packets as the Tcplib distribution). We have plotted a dot for each packet arrival, with the x -axis giving the time of the arrival. In all, there were 1,926 Tcplib interarrivals and 2,204 exponential interarrivals. Over both time scales, the packets from the connection with Tcplib interpacket times are dramatically more clustered.

This difference in burstiness between exponential and heavy-tailed (i.e., Tcplib) interpacket times persists to some extent for multiplexed connections. For example, we ran 10-minute simulations with 100 active TELNET connections, where all connections were active for the entire duration of the simulation. In one simulation each connection used Tcplib interpacket times, and in the other simulation each connection used exponential interpacket times. We found that the multiplexed packet arrival processes with Tcplib interpacket times remained more bursty. For each simulation, consider the number of TELNET packets arriving during successive one-second intervals. For the simulation with individual connec-

tions using Tcplib interpacket times, this aggregate number had a mean of 92 and a variance of 240; for the simulation with exponential interpacket times, the aggregate number had a mean of 92 and a variance of 97. Even a high degree of statistical multiplexing failed to smooth away the difference between the two packet arrival processes.

One of the natural performance measures for TELNET traffic is average packet delay. It would not be hard to construct simulations, one using Tcplib and the other using exponential interarrivals, where making the mistake of using exponential interarrivals instead of Tcplib significantly underestimates the average queuing delay for TELNET packets.

The above shows that the Tcplib packet interarrival distribution behaves quite differently than a Poisson process, even in the presence of multiplexing. We now show that for measured network traffic, these differences extend far beyond the time scale of individual packets. To look at the difference in burstiness at different time scales, we first extracted all TELNET originator packets, except those consisting of no user data (“pure ack”), from the two-hour LBL PKT-2 trace. These packets belonged to 277 separate TCP connections. Of these connections, 4 were anomalously large and rapid (more than 2^{10} bytes transferred by the originator at sustained data rates exceeding 8 bytes/sec). These are unlikely to correspond to human typing, were clear outliers, and are probably better modeled as bulk transfer connections. Removing the outliers left us with 273 connections.

We then synthesized several two-hour packet traces as follows. For each of the TELNET connections, we synthesized a connection with the same starting time within the two-hour period and the same size (in packets). One of the synthesized traces used the Tcplib empirical distribution for the packet interarrivals within each connection (“TCPLIB”); one used exponential interarrivals with mean 1.1 (“EXP”); and one uniformly distributed each connection’s packet arrivals over the interval between when the connection began and when during the LBL PKT-2 trace the connection terminated (“VAR-EXP”). This last method corresponds to exponential interarrivals with the mean adjusted to reflect the connection’s actual observed packet rate. Thus, for the TCPLIB and EXP schemes, we generated connections with the same starting times and sizes (in packets) as their counterparts in the LBL PKT-2 trace, but perhaps with different durations, while with the VAR-EXP scheme, the generated connections shared starting time, size, and duration.

A valuable tool for assessing burstiness over different time-scales is the “variance-time plot” [LTWW94, GW94], which we describe here by example rather than rigorously. Suppose we have a count process consisting of 72,000 observations, corresponding to a two-hour trace viewed every 0.1 seconds. Each observation gives the number of packet arrivals during that 0.1 second interval. The variance of this count process gives us an indication of how bursty the traffic was when viewed on a time scale of 0.1 seconds.

If however we are interested in the process’s burst-structure on a time scale of 10 seconds, we could construct a “smoothed” version of the process by averaging the first 100 observations to obtain the process’s mean value during the first 10 seconds, the next 100 observations for the next 10 seconds, and so on. In general we can do this sort of smoothing for any aggregation level M , where in this example $M = 100$. The variance of the smoothed process then gives us an indication of how bursty the traffic was when viewed on a 10-second time scale.

A natural question is then: how does the variance change as

we progressively smooth the process? By plotting variance vs. degree of smoothing (M), we can examine how burstiness changes according to the time scale used to view the traffic.

For count processes with rapidly decaying autocorrelation functions, such as Poisson processes, the variance of a process aggregated to level M will be $1/M$ times the variance of the unaggregated process (see § 7.3.1). For processes with more persistent autocorrelation functions, however, the variance will decay more gradually. Given this relationship, we can then construct a variance-time plot by smoothing the process for different values of M and plotting the variance of the smoothed process on the y -axis vs. the aggregation level (M) on the x -axis. We use logarithmic scales because they allow us to immediately assess whether the variance decays as $1/M$ (which will show up on the plot as a straight line with slope -1), or more slowly (a slope more shallow than -1), indicating slowly decaying autocorrelation or possibly non-stationarity; that is, from the plot we can tell a great deal about burstiness at different time scales.

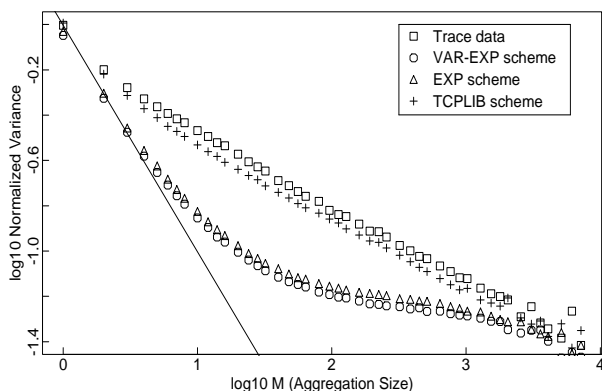


Figure 5: Variance-Time Plot for TELNET packet arrival process. The line from the upper left corner has slope -1 .

Figure 5 shows such a plot for the LBL PKT-2 TELNET trace and for the three schemes discussed above. Here the unaggregated process ($M = 1$) corresponds to 72,000 observations of the number of TELNET originator packets arriving during 0.1-second intervals. The y -axis is the variance of the aggregated process normalized by dividing by the square of the average number of packets per 0.1-second. This normalization allows us to compare the variance of processes with different numbers of arrivals, as the traces consisted of between 82,500 and 86,000 packets.

From the plot it is immediately clear that the variance of the TCPLIB scheme agrees closely with the LBL PKT-2 trace data, while both EXP and VAR-EXP exhibit far less variance, indicating they are much less bursty over a large range of time scales. Thus, the TCPLIB scheme preserves the burstiness present in the measured traffic, while the EXP and VAR-EXP schemes both sacrifice burstiness at larger time scales. At very large time scales ($M = 10^3$), we again get agreement between all of the schemes and the measured traffic, because these time scales are so coarse that we are essentially viewing each connection’s arrivals lumped together as a single observation—differences in the distribution of the arrivals within the connection are lost due to the coarse granularity of our observations.

Figure 6 shows the difference in burstiness between the schemes

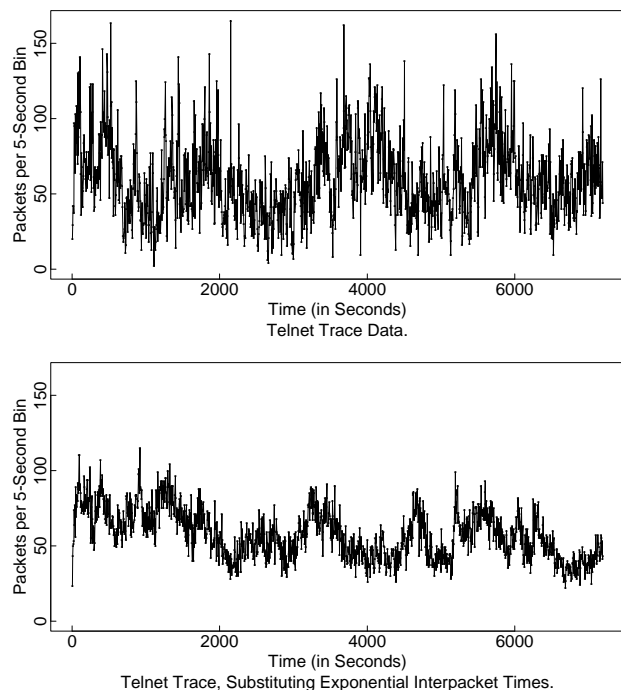


Figure 6: Comparisons of actual and exponential TELNET packet interarrival times.

explicitly. Here we plot the arrival process corresponding to 5-second intervals ($M = 50$) for the LBL PKT-2 trace and for the EXP trace. The x -axis shows the time in seconds, and the y -axis shows the total number of TELNET packets in each 5-second interval. The average number of packets in the two traces are similar; the LBL PKT-2 trace has an average of 59 packets in each 5-second interval, and the fixed-rate exponential trace has an average of 57 packets in each 5-second interval. The variances, however, are quite different. With 5-second bins, the LBL PKT-2 trace has a variance of 672, while the exponential trace has a variance of 260.

Clearly, this difference in the packet-generation rate over 5-second intervals could have consequences for queuing delays in simulations using these two different traces. As the variance-time plot shows, the LBL PKT-2 trace is more bursty over many time intervals, not only over the five-second intervals shown here. The conclusions are that using exponential packet interarrival times for TELNET connections results in substantial underestimations of the burstiness of multiplexed TELNET traffic, but using i.i.d. interarrivals drawn from the Tcplib distribution faithfully reproduces the burst structure.

5 Fully modeling TELNET originator traffic

Section 3 has shown that over 1-hour periods, TELNET connection arrivals are well-modeled as Poisson processes, and § 4 has shown that within a TELNET connection, packet interarrival times can be modeled using the heavy-tailed distribution in Tcplib. The connection size in *bytes* has been previously modeled by a log-extreme distribution [P94a]; the distribution of the connection size in *pack-*

ets is somewhat different, and seems to be better modeled by a log-normal distribution (see below). In this section, we put these three pieces together to construct a complete model of TELNET originator traffic that is parameterized only by the connection arrival rate. Variance-time plots show that this model corresponds well to empirical measurements.

First, we look at the difference in the distributions of originator bytes per connection vs. originator packets. Previous work reports that the number of bytes sent by the originator in a wide-area TELNET connection is well-modeled using a log-extreme distribution with location parameter $\alpha = \log_2 100$ and scale parameter $\beta = \log_2 3.5$ [P94a]. We experimented with using this distribution to produce sizes for an equal number of TELNET connections as appeared in the LBL PKT-2 trace. We found that the distribution consistently generates connection sizes (in bytes) much larger than the connection sizes (in packets) observed in the trace. We attribute this difference to two effects:

- The [P94a] fit was made using month-long traces of TELNET connections, allowing for much longer and larger connections than are present in our two-hour trace;
- The [P94a] fit models connection size in *bytes* and not in *packets*. One generally assumes that each TELNET originator packet conveys one byte of user data, corresponding to a keystroke. Often, however, a packet carries more than one byte, either due to effects of the Nagle algorithm [N84] or because the TELNET connection is operating in “line mode” [B90] or “line-at-a-time mode” [PR83, S94]. For example, the LBL PKT-2 TELNET originator traffic comprised about 85,000 packets carrying 139,000 user data bytes.

Given these difficulties, we attempted to fit the observed TELNET connection sizes (in packets) with another simple analytic distribution. We found that a \log_2 -normal distribution with \log_2 -mean $\bar{x} = \log_2 100$ and \log_2 -standard deviation $\sigma = 2.24$ fit the connection size in packets well visually, considerably better than a log-extreme distribution with parameters fitted to the data. (The exact numerical values of \bar{x} and σ here should not be taken too seriously, as they came from a small sample.) We also found that a log-extreme distribution fit the connection size in bytes better than a log-normal distribution, so our data remains consistent with the models presented in [P94a].

Putting all of this together, we have a complete model for TELNET traffic, FULL-TEL, parameterized only by the TELNET connection arrival rate. FULL-TEL uses Poisson connection arrivals, log-normal connection sizes (in packets), and Tcplib packet interarrivals.

We then used FULL-TEL to generate three synthetic traces of TELNET originator traffic, using a connection arrival rate of 273 connections in 2 hours. Because such traces start off with no traffic and build up to a steady-state corresponding to the connection arrival rate, we trimmed the traces to just their second hour. We then used variance-time plots to compare the traces with the second hour of the LBL PKT-2 TELNET trace.

Figure 7 shows the results of the comparison. In general the agreement is quite good, though the models have slightly higher variance than the trace data for $M > 10^2$. We conclude that FULL-TEL faithfully captures TELNET originator traffic, except to be a bit burstier on time scales above 10 seconds. As a final note, we also tested the model’s fit to the LBL PKT-1 and PKT-3 TELNET traces; the results were similar.

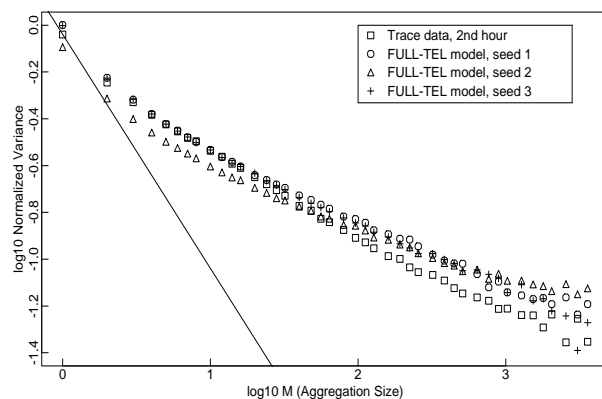


Figure 7: Variance-time plot comparing LBL PKT-2 trace data with the complete TELNET model, FULL-TEL.

6 FTPDATA connection arrivals

This section investigates arrival processes for FTP traffic. Modeling FTP is particularly important because FTPDATA connections currently carry the bulk of the data bytes in wide area networks ([CBP93]). Section 3 showed that while FTP session arrivals can be modeled as Poisson processes, this is not the case for FTPDATA connection arrivals. This section shows that FTPDATA connections within a session are clustered in bursts, and that the distribution of burst sizes in bytes is quite heavy-tailed; half of the FTP traffic volume comes from the largest 0.5% of the FTPDATA bursts. These large bursts are likely to completely dominate FTP traffic dynamics.

In this paper, we do not attempt to model FTPDATA packet arrivals within a connection. Unlike TELNET connections, where the originator packet arrival process is largely determined by the packet generation pattern at the source, the packet arrival process for an FTPDATA connection is largely determined by network factors such as the available bandwidth, congestion, and details of the transport-protocol congestion control algorithms. Previous studies have found that FTPDATA packet interarrivals are far from exponential [DJCME92]; this is not surprising, since the above network factors lead to a process quite different from memoryless arrivals.

To begin, § 3 showed that FTPDATA connection arrivals are not well-modeled as Poisson. Each FTP session spawns a number of FTPDATA connections; one key question is how these connections are distributed within the duration of the FTP session.

We computed the distribution of spacing between FTPDATA connections spawned by the same FTP session for six datasets: LBL-1, LBL-5, LBL-6, LBL-7, DEC-1, and UCB. Here, “spacing” refers to the amount of time between the end of one FTPDATA connection within a session and the beginning of the next. Figure 8 plots the results. In each case the upper tail of the distribution is much heavier than exponential (the x -axis is logarithmic), and is better approximated using a log-normal or log-logistic distribution. Furthermore, all of the distributions show inflection points at spacings between 2 and 6 seconds, indicating bimodality. We conjecture that spacings shorter than these points reflect sequential FTPDATA connections due to multiple transfers (the FTP “mget” command) or a user issuing a “list directory command” very shortly followed by a “get.” Such closely-spaced connections might well be interpreted as corresponding to a single “burst” of file-transfer activity. We somewhat arbitrarily chose a spacing of ≤ 4 seconds (the dashed

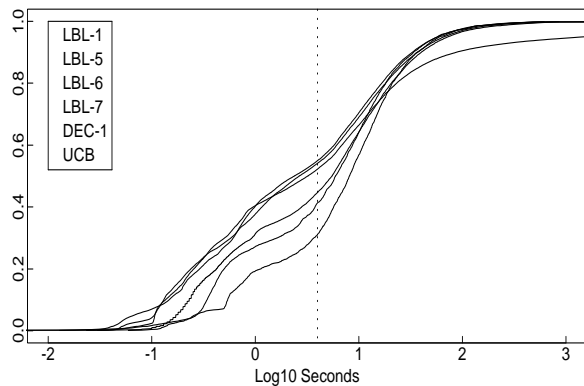


Figure 8: FTPDATA Intra-session Connection Spacing.

vertical line) as defining connections belonging to the same *burst*, and we note that such spacings are not inordinately larger than the 1-2 second spacings that can occur internal to a single FTPDATA connection due to TCP retransmission timeouts. Here, “somewhat arbitrarily” means that, for example, using a cutoff spacing of 2 seconds instead (which actually slightly better delimits the two modes of activity) results in virtually identical results as those discussed in the remainder of this section.

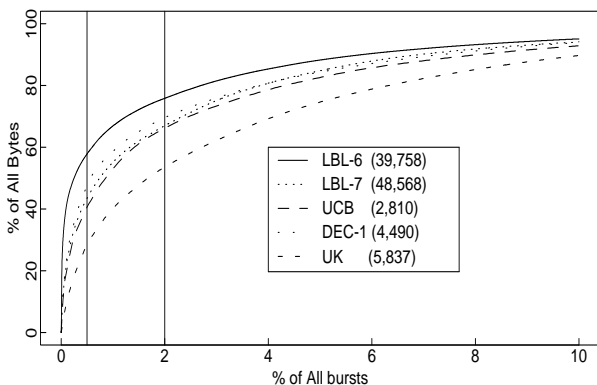


Figure 9: Percentage of all FTPDATA bytes due to largest 10% FTPDATA bursts.

With this definition of a burst of FTPDATA connections, we analyzed the same datasets to measure the distribution of the number of bytes transferred during a single connection burst. The distribution proves to be remarkably heavy-tailed. Figure 9 shows the percentage of all FTPDATA bytes (*y*-axis) due to the largest 10% of the FTPDATA bursts (*x*-axis). The numbers in parentheses in the legend give the total number of FTPDATA bursts occurring during each trace. The first vertical line marks the upper 0.5% of the FTPDATA bursts, and the line to its right, the upper 2%.

The key point to draw from this figure is that the upper 0.5% tail of the FTPDATA bursts holds *between 30% and 60% of all of the data bytes*. Thus, at any given time FTP traffic will most likely be *completely dominated by a single or small handful of bursts*. Note that this phenomenon is present in *all* of the connection datasets we studied. The dataset with the least heavy tail is UK (shown in the figure), which still held 30% of the data bytes in the upper 0.5% tail

and 55% in the 2% tail. The NC dataset lies about halfway between UK and the others in the figure, and the remainder lie within the bounds of the others shown in the figure.

This finding means that for many aspects of network behavior, modeling small FTP sessions or bursts is irrelevant; all that matters is the behavior of a few huge bursts. The sizes and durations of these bursts will vary considerably from one time to another; but they *will* be present. We also note that our finding that the size of an FTPDATA burst has a heavy tail matches a survey conducted by Irlam [I93] of the sizes of files in 1,000 file systems comprising 12 million files and 250 GB of data: 1.9% of the files accounted for 71% of the bytes, and 0.5% accounted for 54% of the bytes.

We performed fitting of the upper tail of the distribution of data bytes per FTPDATA burst and found that the upper 5% tail fits well to a Pareto distribution with $0.9 \leq \beta \leq 1.4$ [P94a]. As the Pareto distribution is heavy-tailed (see Appendix B), this agrees with our findings in Figure 9. In contrast, the upper 0.5% tail of an exponential distribution always holds about 3% of the entire mass of the distribution, regardless of the distribution’s mean.

Figures 10 and 11 graphically illustrate the dominance of the upper FTPDATA-burst tail. The four plots in Figure 10 show the FTPDATA traffic rate in bytes/minute for the LBL PKT-1, PKT-2, PKT-3, and PKT-5 datasets, and in Figure 11 the same is shown for the DEC WRL datasets. The shaded areas represent traffic contributed by the largest 2% of the bursts, and the black areas the largest 0.5%. The numbers in parentheses give the number of bursts and FTPDATA connections comprising the 2% burst upper-tail. (For example, the upper 2% tail of the PKT-1 bursts was made up of 7 bursts consisting of a total of 19 FTPDATA connections, while for WRL-2 this tail was made up of 16 bursts and 1,796 connections.) We see that sometimes bursts contain many separate connections and sometimes not. Indeed, the distribution of the number of connections per burst is well-modeled as a Pareto distribution. For example, a single burst in the LBL-7 dataset was made up of 979 separate FTPDATA connections.

For PKT-1 (364 bursts) and PKT-3 (552 bursts), the upper 2% and 0.5% tails hold around 50% and 15% of all the traffic; for PKT-2 (483 bursts) and PKT-5 (238 bursts), 85% and 60%. The large degree of difference between PKT-1/PKT-3 and PKT-2/PKT-5 illustrates how volatile the upper-tail behavior is; a trace comprising 400 bursts (and substantially more FTPDATA connections) might well be completely dominated by 2 of the bursts, or it might not, since 2 is a very small sample of the upper-tail behavior. Thus we are left in the difficult position of knowing that upper-tail behavior dominates traffic, but with such small numbers of bursts that we cannot reliably use large-number laws to predict what we are likely to see during any given trace. Furthermore, the PKT-2 and PKT-5 bursts were not geographically anomalous, either: the largest PKT-2 burst was to a government site in Colorado, and the largest PKT-5 burst was to a commercial site in Washington state. These sites are about 1,500 km and 1,000 km distant from LBL, respectively.

For the DEC datasets, the difference in the size of the burst tails is not so pronounced: in WRL-1 (971 bursts), WRL-3 (2,161 bursts), and WRL-4 (2,100 bursts) the 2% and 0.5% tails hold 54-70% and 33-42% of all the traffic, while for WRL-2 (788 bursts) they hold 45% and 18%. The lesser degree of difference between the datasets is what we would expect: since the datasets have considerably more bursts than their LBL counterparts, large-number laws become more reliable in predicting the size of the tails.

We would also like to know whether the arrivals of the upper-tail

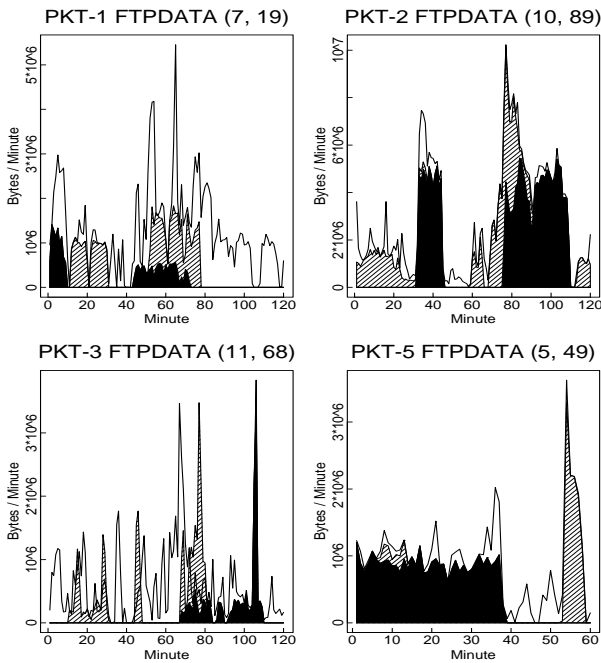


Figure 10: Proportion of LBL PKT FTPDATA traffic due to largest 2% (shaded) and 0.5% (black) connection bursts.

bursts can be modeled as a Poisson process, as that would provide a first step toward predicting their effect on network traffic. We analyzed the 199 upper-0.5%-tail LBL-6 bursts, first removing effects due to daily variation in traffic rates by looking at interarrivals in terms of number of intervening bursts instead of seconds. We found that the dataset failed the statistical test (Appendix A) for exponential interarrivals at all significance levels. Thus, caution must be used if approximating large-burst arrivals using a Poisson process; further analysis is needed to model the burst-clustering.

7 Large-scale correlations and possible connections to self-similarity

We have argued in the previous sections that on any time-scale smaller than user-session arrivals, modeling wide-area TCP traffic using Poisson processes fails to faithfully capture the traffic’s dynamics. Recent work [LTWW94] shows that local-area Ethernet traffic (and perhaps wide-area TCP traffic) is much better modeled as a *self-similar* process, which displays substantially more burstiness over a wide range of time scales than do Poisson processes.

In this section we discuss the degree of “large-scale correlation” present in the LBL PKT traces of TELNET traffic, and the LBL PKT and DEC WRL traces of FTPDATA traffic and aggregate wide-area traffic. We also consider the evidence for whether such correlation is well modeled using self-similar processes. We begin with a discussion of the concepts of “large-scale correlation,” “long-range dependence,” and “self-similarity.” We next give an overview of two existing methods for generating truly self-similar traffic, along with a new method for producing “pseudo-self-similar” traffic. We then discuss how the traffic models developed in this paper might match these methods. We finish with a preliminary assessment of the possible self-similarity of general wide-area traffic. We find the

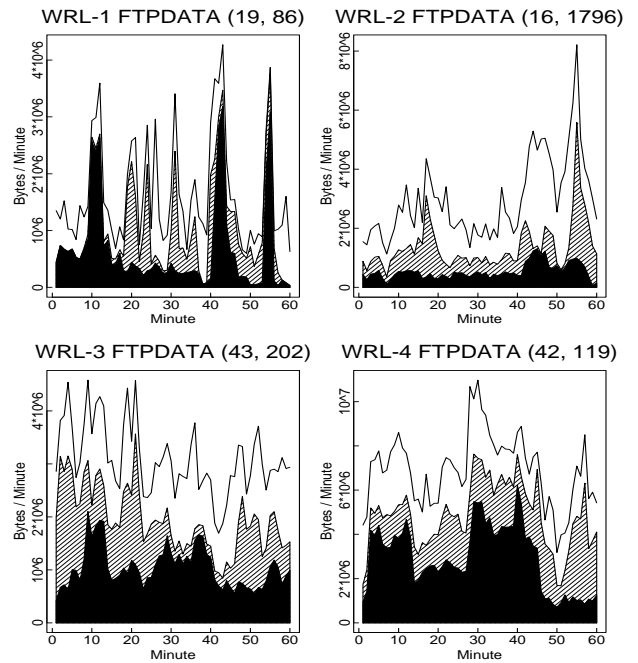


Figure 11: Proportion of DEC WRL FTPDATA traffic due to largest 2% (shaded) and 0.5% (black) connection bursts.

evidence inconclusive, though the traffic clearly exhibits large-scale correlations inconsistent with Poisson processes.

7.1 Definitions

We use the term “large-scale correlation” as an informal way of describing correlations that persist across large time scales. For example, the lower right plot in Figure 10 shows a 40-minute burst of highly correlated traffic.

A related, more precise notion of sustained correlation is that of “long-range dependence.” A stationary process is *long-range dependent* if its autocorrelation function $r(k)$ is nonsummable (i.e., $\sum_k r(k) = \infty$) [C84]. Thus, the definition of long-range dependence applies only to infinite time series.

The simplest models with long-range dependence are *self-similar* processes, which are characterized by hyperbocally-decaying autocorrelation functions. Self-similar and asymptotically self-similar processes are particularly attractive models because the long-range dependence can be characterized by a single parameter, the Hurst parameter (which can be estimated using Whittle’s procedure [GW94, LTWW94]).

In the following sections, we look at ways in which long-range dependence in general, and self-similarity in particular, might arise in wide-area network traffic. An important point to bear in mind is that, even if the finite arrival process derived from a particular packet trace does not appear self-similar, if it exhibits large-scale correlations suggestive of long-range dependence then that process is almost certainly better approximated using a self-similar process than using Poisson processes. Thus, we believe that self-similar modeling is a promising successor to Poisson modeling. It may not be exactly right, but given our current understanding of networking phenomena, it appears in any case a good approximation.

7.2 Producing self-similar traffic

There are several methods for producing self-similar traffic that could account for self-similarity in wide-area TCP traffic. As discussed in [LTWW94], self-similar traffic can be produced by multiplexing ON/OFF sources that have a fixed rate in the ON periods and ON/OFF period lengths that are *heavy-tailed* (see Appendix B).

A second method for generating self-similar traffic that could fit TCP traffic is an $M/G/\infty$ queue model, where customers arrive according to a Poisson process and have service times drawn from a heavy-tailed distribution with infinite variance [C84, LTWW94]. In this model, X_t is the number of customers in the system at time t . The count process $\{X_t\}_{t=0,1,2,\dots}$ is asymptotically self-similar (see Appendix D for further discussion). The $M/G/\infty$ model implies that multiplexing constant-rate connections that have Poisson connection arrivals and a heavy-tailed distribution for connection lifetimes would result in self-similar traffic.

We investigated an additional method of producing arrival processes that appear to some extent self-similar. This method involves constructing arrivals using i.i.d. Pareto interarrivals with $\beta \approx 1$, and then considering the corresponding count process (the number of arrivals in consecutive intervals). The goal behind the method is to explore how a simple model of TELNET traffic might lead to self-similarity. We refer to this method as “pseudo-self-similar” because while the traffic it generates has large-scale correlations and the “visual self-similarity” property [LTWW94] over many time scales, we show in Appendix C that the traffic generated is not actually long-range dependent (and thus not self-similar).

7.3 Relating the methods to traffic models

7.3.1 TELNET

As explained in [LTWW94], straight lines on variance-time plots with slopes more shallow than -1 , such as that for the PKT-2 TELNET trace in Figure 5, are suggestive of self-similarity. In general, the slope of an arrival process’s variance-time plot is a function of the process’s autocorrelation function [C84], and a long-range dependent process will exhibit *slowly-decaying variances* on such a plot. That is, the variance-time plot will decline in a more shallow fashion than with slope -1 , though not necessarily in a straight line. An important point is that such slow decline can also occur due to the presence of non-stationarity.

In addition to looking at variance-time plots of the TELNET traffic, we also used Whittle’s procedure [GW94, LTWW94] and Beran’s goodness-of-fit test [B92a] to gauge the agreement between the traffic and the simplest type of self-similar process, *fractional Gaussian noise* [B92b]. All of the results are consistent with self-similarity on scales of tens of seconds or more.

We postulate that two different mechanisms contribute to the apparent self-similarity of TELNET traffic. On smaller time scales, apparent self-similarity might arise from the fact that within individual TELNET connections, packet interarrivals are well modeled as i.i.d. Pareto (§ 4). Thus, individual TELNET connections match the i.i.d. Pareto method of generating pseudo-self-similar traffic that appears self-similar over a range of time scales (Appendix C). On larger time scales, we note that our source model of TELNET connections presented in § 5 in some respects matches the $M/G/\infty$ model described in the previous section. TELNET connection sizes in packets have a *long-tailed* [WT92] distribution, in that the tail function of a log-normal distribution decreases more slowly than

any exponential function. While we show in Appendix E that the $M/G/\infty$ queue with log-normal service times does *not* result in long-range dependent or self-similar traffic, the difference in tail weight between a log-normal distribution and a Pareto distribution may be small enough that over the time scales of interest (seconds to minutes) the traffic still appears self-similar.

Put together, these models of TELNET traffic suggest why the traffic might appear self-similar (or at least long-range dependent) over many time scales. While individually the models fall short of proving self-similarity, it could be the case that the combination of i.i.d. Pareto interpacket times and the $M/G/\infty$ effect due to multiplexing makes TELNET traffic truly self-similar. At a minimum, these models explain why the traffic exhibits large-scale correlations. Further work is needed for a definitive statement regarding actual self-similarity.

7.3.2 FTP

Like the model of TELNET traffic discussed in the previous section, our model of FTP traffic also fits in some respects to the $M/G/\infty$ model of Poisson arrivals with heavy-tailed lifetimes. The distribution of bytes per FTPDATA *burst* is heavy-tailed (§ 6), and FTP *sessions* have Poisson arrivals (§ 3). Over larger time scales the packet arrival process within an FTPDATA burst can be plausibly approximated as constant-rate. If we approximated FTPDATA burst arrivals as Poisson (a bit of a stretch, as shown in § 3 above), and assumed that each FTPDATA burst received the same average rate, then multiplexed FTP traffic would fit the $M/G/\infty$ model above, and should be self-similar.

It turns out, though, that variance-time plots, Whittle’s procedure, and goodness-of-fit tests of our FTP traces all suggest that our FTPDATA traces are not well-modeled as fractional Gaussian noise, although the heavy-tailed distribution of FTPDATA bursts clearly leads to large-scale correlations. The sole exception to this finding is the DEC WRL-3 trace, for which the tests are consistent with self-similarity at time scales of 1 second or greater.

One reason the FTP traces might not be well-modeled as fractional Gaussian noise is that the traces exhibit extremely high burstiness, including lengthy periods during which there is no FTP traffic. These “lulls” mean that the marginal distribution function of the arrival process has a large peak at zero arrivals. Since fractional Gaussian noise is a form of Gaussian process, its marginal distribution is normal, and cannot accommodate such a peak. It is still possible that FTP traffic is well-modeled using different self-similar processes; or that it instead is not well-modeled as self-similar. In this paper we do not try to resolve this issue, but limit our discussion to the interplay between mechanisms affecting FTP traffic dynamics and large-scale correlations in the traffic.

Unlike TELNET traffic, where the timing of packets generated at the source is reasonably close to the timing of the same packets transmitted on the network, the timing of FTPDATA packets transmitted on the network is intimately related to the dynamics of TCP’s congestion control algorithms. The following paragraphs discuss several ways that, due in part to the effects of TCP, multiplexed FTP traffic differs from the $M/G/\infty$ model of self-similar traffic with constant-rate connections. While these factors could account for our FTP traces not appearing statistically self-similar, they do not imply the absence of long-range dependence.

Unlike the $M/G/\infty$ model, which best fits an environment where all connections have the same fixed constant rate, different FTP connections have quite different average rates, and within a single

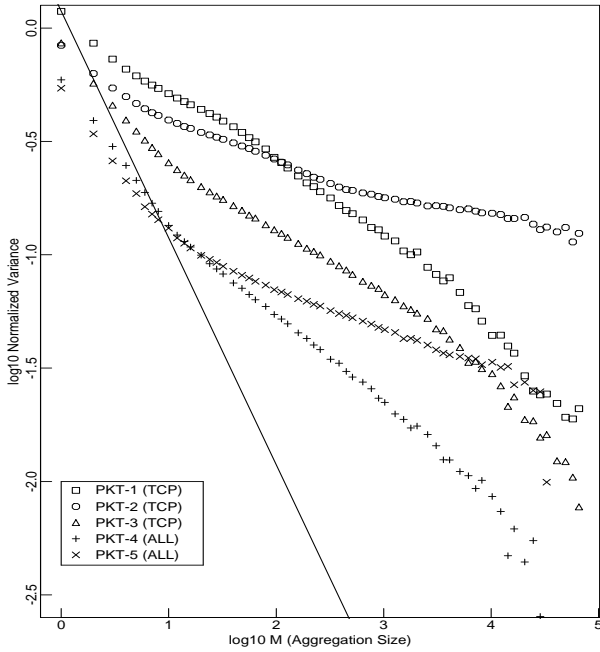


Figure 12: Variance-time plot for all TCP / all link-level packet arrivals in the LBL PKT datasets.

FTP connection the average rate varies over time. TCP’s congestion control algorithms increase the TCP congestion window to probe for additional bandwidth, and reduce the congestion window again in response to congestion (packet drops). TCP’s window flow control has several separate effects on the traffic pattern for an individual FTP connection. First, over intervals less than a roundtrip time the FTP connection does not have a constant rate; each packet is sent only after the TCP source receives an acknowledgement for an earlier packet. Second, if there is congestion in the network, then an FTP connection does not have a constant rate even over longer time intervals; the average rate over a roundtrip time varies as the TCP congestion control window varies. Third, whether or not there is congestion in the network, different FTP connections will have different average rates, depending on such factors as the TCP window and packet sizes, the connection’s roundtrip time, and the congestion encountered in the network. These factors give rise to serious discrepancies between our trace data and the $M/G/\infty$ model.

One way to incorporate the effect of limited bandwidth into the $M/G/\infty$ model would be to explore a model of an $M/G/k$ queue instead of an $M/G/\infty$ queue. In an $M/G/k$ queue, because there are only k servers, the actual arrival times of individuals at a server would occasionally have to be delayed until there was available capacity. While this limited capacity would have the effect of reducing the fit of the multiplexed traffic to a self-similar model, it does not eliminate the underlying large-scale correlations in the $M/G/\infty$ model. However, the $M/G/k$ model as applied to FTP connections assumes that all active connections have the same constant rate, and this is not the case in actual FTP traffic.

Another discrepancy between the $M/G/\infty$ model and our link traces concerns the effect of FTP traffic competing with other families of traffic on a congested link. The four main classes of traffic in our link traces were TCP, Mbone (primarily multicast UDP audio traffic), Domain Name System requests and replies (UDP-based),

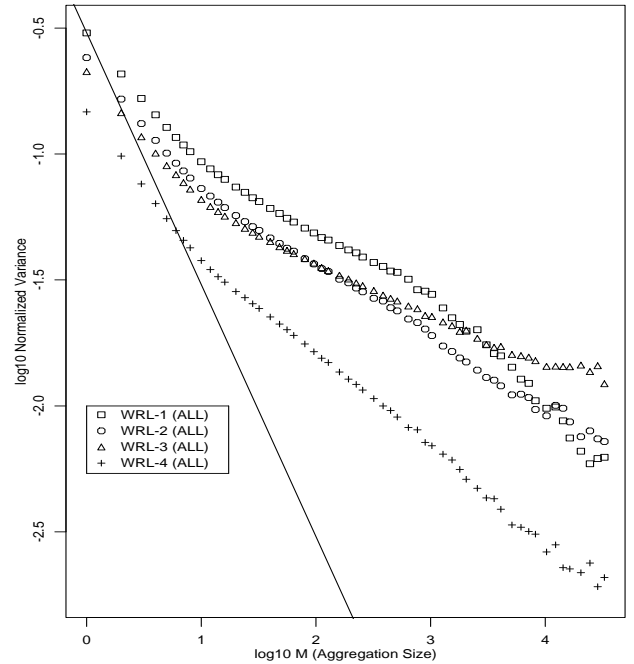


Figure 13: Variance-time plot for all link-level packet arrivals in the DEC WRL datasets.

and DECnet. Unlike TCP, the UDP protocol does not incorporate congestion-avoidance mechanisms. Therefore, when TCP-based FTP traffic is competing for bandwidth with Mbone UDP sources, only the FTP traffic will adjust to fit the available bandwidth. The UDP traffic will continue unimpeded. The effect of this interaction on the overall structure of FTP traffic remains an open question.

7.4 Large-scale correlations in general wide-area traffic

We finish with a preliminary look at whether wide-area traffic multiplexed over different protocols appears self-similar. Figure 12 shows variance-time plots for all of the LBL PKT traces listed in Table 2. Here, the unaggregated process ($M = 1$) corresponds to observing the packets arriving during each 0.01 second interval.

Recall that the first three LBL PKT traces captured all TCP packets for two hours, and the last two captured all wide-area packets appearing on the gateway Ethernet for one hour. The first three traces consist of between 1.7 and 2.4 million packets, and the last two traces each have around 1.3 million packets. The corresponding rates of packets/hour are above those of the “low hours” in [LTWW94], so we would hope to find that the traces exhibit exact self-similarity.

We see in Figure 12 that PKT-4 and PKT-5, the full link-level traces, both yield straight lines with shallow slope, consistent with asymptotic self-similarity for $M \geq 10$ (0.1 second). For the TCP traces, PKT-1 is concave down for small and large M , inconsistent with exact self-similarity, PKT-2 appears consistent with asymptotic self-similarity for $M \geq 10^3$ (10 seconds), and PKT-3 has a straight section between $M = 10$ and $M = 10^3$, but not before or after, also inconsistent with exact self-similarity.

In contrast, use of Whittle’s procedure and goodness-of-fit tests suggest that the link-level PKT-4 trace and the TCP PKT-1 and

PKT-3 traces are consistent with fractional Gaussian processes, while the link-level PKT-5 trace and the TCP PKT-2 trace are not. As Figure 10 shows, the FTP traffic in the PKT-5 and PKT-2 traces is heavily dominated by a few large FTPDATA bursts. Thus, while large-scale correlations are clearly present in these traces, it might be difficult to characterize the correlations over the entire trace with a single Hurst parameter.

Figure 13 shows the same sort of variance-time plot for the DEC WRL datasets listed in Table 2. The least active of the WRL datasets exceeds the most active in [LTWW94], so we would again expect to find exact self-similarity. The variance-time plots for WRL-2 and WRL-4 are encouraging in this regard, lying in essentially straight lines for time scales of 0.1 seconds and higher. WRL-3 lies in a straight line at time scales of 1 second and higher, while WRL-1 does so only at 10 seconds and higher. But of these datasets, Whittle’s procedure and Beran’s goodness-of-fit test indicate that only WRL-3 is consistent with fractional Gaussian noise (at time scales of 1 second and greater). The others, while clearly exhibiting large-scale correlations, do not appear to be well-modeled by a simple self-similar process. This could be due to distorting effects of short-range dependence, better fits to other self-similar models such as *fractional ARIMA* processes [B92b], or the presence of non-stationarity. WRL-3 was also the only dataset whose FTP traffic appears consistent with fractional Gaussian noise, though we have not assessed whether this coincidence is significant. Clearly, further work is required to fully understand the correlational structure of wide-area traffic.

We end with a comment regarding the balance between link-level modeling and protocol-specific modeling. One approach to investigating self-similarity is to model multiplexed link traffic as self-similar, without attempting to model individual connections. This approach could have many uses in simulations and in analysis. For example, self-similar traffic could be used instead of Poisson traffic to model cross-traffic, or self-similar traffic could be used in simulations investigating link-sharing between two different classes of traffic.

However, for many simulations, the simulator needs to model individual sources. In particular, it is only from modeling of individual sources, and a direct implementation of TCP’s congestion control algorithms, that a simulation can take into account the effects of the TCP algorithms in different environments. TCP’s congestion control algorithms contribute long-term oscillations to the traffic pattern for a particular connection, as the TCP congestion window changes over the lifetime of the connection. In addition, TCP’s window flow control contributes a shorter-term periodicity to the traffic pattern, as each packet is transmitted in response to an acknowledgement returned for an earlier packet [FJ92]. It is particularly important to take into account these effects in simulations investigating changes to either TCP, the gateway scheduling algorithms, or the network’s packet-dropping algorithms.

8 Implications

This paper’s findings are summarized in the Introduction. In this section we conclude with a look at the implications of our results.

Several researchers have previously discussed the implications of long-range dependence (burstiness across different time scales) in network traffic. Modeling TCP traffic using Poisson or other models that do not accurately reflect the long-range dependence in actual traffic will result in simulations and analyses that significantly

underestimate performance measures such as average packet delay or maximum queue size.

[FL91] examines the burstiness of data traffic over a wide range of time scales, and discusses the impact of this burstiness for network congestion. Their conclusions are that congested periods can be quite long, with losses that are heavily concentrated; that, in contrast to Poisson traffic models, linear increases in buffer size do not result in large decreases in packet drop rates; and that a slight increase in the number of active connections can result in a large increase in the packet loss rate. They suggest that, because the level of busy period traffic is not predictable, it would be difficult to efficiently size networks to reduce congestion adequately. They observe that, in contrast to Poisson models, in reality “traffic ‘spikes’ (which cause actual losses) ride on longer-term ‘ripples’, that in turn ride on still longer-term ‘swells’.” They suggest that a filtered variable can be used to detect the low-frequency component of congestion, giving some warning before packet losses become significant.

[LTWW94] discusses some additional implications of long-range dependence of packet traffic. These include an explanation of the inadequacy of many commonly-used notions of burstiness, and the somewhat counter-intuitive observation that the modeling of individual connections can gain insight from an understanding of the fundamental characteristics of multiplexed traffic. In this paper, observations of the characteristics of multiplexed traffic motivated our revisitation of models for individual connections; indeed, we originally set out to challenge the notion that wide-area traffic might be self-similar, and have come full circle.

[GW94] examines the long-range dependence of variable-bit-rate (VBR) video traffic. Their empirical measurements of VBR traffic show strong low-frequency components, and they propose source models for video traffic that display the same long-range dependence. Given the likelihood that VBR traffic will soon comprise a significant fraction of Mbone traffic, we soon will have wide-area traffic of which a substantial portion is perforce self-similar, simply due to the source characteristics of its individual connections.

There are some additional respects in which the burstiness and long-range dependence of aggregate traffic can affect traffic performance. Consider a link with priority scheduling between classes of traffic, where the higher-priority class has no enforced bandwidth limitations (other than the link bandwidth itself). In such a partition, interactive traffic such as TELNET might be given priority over bulk-data traffic such as FTP. If the higher-priority class has long-range dependence and a high degree of variability over long time scales, then the bursts from the higher-priority traffic could starve the lower-priority traffic for long periods of time.

A second impact of the long-range dependence of packet traffic concerns classes with admissions control procedures that are based on measurements of recent traffic, rather than on policed traffic parameters of individual connections [CSZ92]. As has been shown by numerous researchers, such admissions control procedures could lead to a much more effective use of the available bandwidth [YKTH93]. Nevertheless, if the measured class has high burstiness consisting of both a high variance and significant long-range dependence, then an admissions control procedure that considers only recent traffic could be easily misled following a long period of fairly low traffic rates. (This is similar to a situation in California geology some decades ago. Because there hadn’t been a large earthquake for a long time, people began to believe it unlikely that there would be another one.)

In summary: we should abandon Poisson-based modeling of

wide-area traffic for all but user session arrivals. For TELNET traffic, we offer a faithful model of originator traffic parameterized by only the hourly connection arrival rate. Modeling the TELNET responder remains to be done. For FTP traffic, we have shown that modeling should concentrate heavily on the extreme upper tail of the largest bursts. A wide-area link might have only one or two such bursts an hour, but they tend to strongly dominate that hour’s FTP traffic. Finally, our look at multiplexed TCP and all-protocol traffic suggests that anyone interested in accurate modeling of wide-area traffic should begin by studying self-similarity.

9 Acknowledgments

We gratefully acknowledge helpful discussions with Johan Bengtsson, Kim Claffy, Domenico Ferrari, Mark Garrett, Van Jacobson, Steve McCanne, Allyn Romanow, and Bin Yu, as well as the comments of the referees.

We particularly want to thank Walter Willinger for many fruitful discussions, and also for making available Jan Beran’s S programs to compute Whittle’s estimator and Beran’s goodness-of-fit test.

We would also very much like to thank Jeff Mogul for both the DEC SYN/FIN datasets and the DEC WRL packet datasets; Peter Danzig and his coauthors for the UCB and BC datasets; Ian Wake-man and Jon Crowcroft for the UK dataset; Wayne Sung for the NC dataset; Ramón Cáceres and Sugih Jamin, who between them made available all of the non-LBL SYN/FIN datasets; and Kate Lance, Robert Elz, Geoff Martin, Tony Nicholson, and Douglas Ray, for their efforts in capturing additional packet traces.

The LBL traces were gathered with the help of Craig Leres and Steve McCanne. The Bellcore traces were gathered by D. V. Wilson.

A Methodology for testing for Poisson arrivals

To test whether a trace of connection arrivals corresponds to a non-homogeneous Poisson process, we first pick an interval length I over which we hypothesize that the arrival rate does not change. If the trace spans a total of T time units, we divide the entire trace into $N = T/I$ intervals each of length I . We then separately test each interval to see whether the arrivals during the interval are consistent with arrivals from a Poisson process with rate fixed so that the expected number of arrivals is the same as the number actually observed. Thus, we reduce the problem of testing for nonhomogeneous Poisson arrivals to that of testing a number of intervals for homogeneous Poisson arrivals.

Poisson arrivals have two key characteristics: the interarrival times are both exponentially distributed, and independent. We discuss testing for each in turn.

For each interval, we test the interarrivals for an exponential distribution using the Anderson-Darling (A^2) test, recommended by Stephens in [DS86] because it is generally much more powerful than either of the better-known Kolmogorov-Smirnov or χ^2 tests. A^2 is also particularly good for detecting deviations in the tails of a distribution. A^2 is an *empirical distribution test*; it looks at the entire observed distribution, rather than reducing the distribution into bins as is required by χ^2 .

We associate a *significance level* with each A^2 test. For example, a test with a significance level of 5% will correctly confirm the null

hypothesis (if it is correct) with probability 0.95; with probability 0.05, the test will erroneously declare the hypothesis false. Thus, the significance level indicates the proportion of “false negatives” (in general it is difficult to assess the corresponding percentage of “false positives”). We can use significance-level testing as follows. Suppose we test N intervals for exponential interarrivals and K of them pass the A^2 test at the 5% significance level. If the null hypothesis is correct, then the probability of K successes in N trials will be given by a binomial distribution with parameter $p = .95$. If we find that the probability of observing K successes was less than 5%, then we conclude with 95% confidence that the arrival process is inconsistent with exponential interarrivals.

There are two important details for correctly applying and interpreting the A^2 test. The first is that estimating the parameters of our model from the data to be tested alters the significance levels of the A^2 test (this applies to our null hypothesis above, in which we derive the mean of the exponential fit from the data rather than knowing it *a priori*). The second is that the number of data points tested also alters the significance levels. In general, the more points tested, the more likely the test will detect an incorrect null hypothesis. [DS86] gives procedures for incorporating both of these considerations into A^2 tests.

We also need to test the interarrivals for independence. One indication of independence is an absence of significant autocorrelation among the interarrivals. Autocorrelation can be significant in two different ways: it can be too strong in magnitude, or it can be too frequently positive or negative. We address each of these in turn.

Given a time series of n samples from an uncorrelated white-noise process, the probability that the magnitude of the autocorrelation at any lag will exceed $1.96/\sqrt{n}$ is 5%. Thus we can test for independence by observing how often this occurs and using a binomial test similar to the one outlined above. (Because for many non-Poisson processes autocorrelation among interarrivals peaks at lag one, to keep our test tractable we restrict it to just the lag one autocorrelation.)

We also apply one further test for independent interarrivals. If the interarrivals are truly independent, then we would expect their autocorrelation to be negative with probability 0.5 and positive with probability 0.5. For Poisson arrivals, then, the number of positive lag one autocorrelation values should be binomially distributed with parameter $p = 0.5$. Given this assumption, we assess the probability of at least the observed number of positive values occurring. If its probability is too low ($< 2.5\%$) then we conclude that the interarrivals are significantly positively correlated. Similarly, if the observed number of negative values has probability $< 2.5\%$, then the interarrivals are significantly negatively correlated.

B Pareto distributions

In this paper the Pareto distribution plays a role both in TELNET packet interarrivals and in the size of FTPDATA bursts. This appendix discusses the Pareto distribution and its occurrence in the physical world.

The classical Pareto distribution with shape parameter β and location parameter a has the cumulative distribution function [HK80]:

$$F(x) = P[X \leq x] = 1 - (a/x)^\beta, \quad a, \beta \geq 0, \quad x \geq a,$$

with the corresponding probability density function:

$$f(x) = \beta a^\beta x^{-\beta-1}.$$

If $\beta \leq 2$, then the distribution has infinite variance, and if $\beta \leq 1$, then it has infinite mean.

The Pareto distribution (also referred to as the power-law distribution, the double-exponential distribution, and the hyperbolic distribution) has been used to model distributions of incomes exceeding a minimum value, and sizes of asteroids, islands, cities and extinction events [K93, M63]. Leland and Ott also found that a Pareto distribution with $1.05 < \beta < 1.25$ is a good model for the amount of CPU time consumed by an arbitrary process [LO86].

In communications, heavy-tailed distributions have been used to model telephone call holding times [DMRW94] and frame sizes for variable-bit-rate video [GW94]. The discrete Pareto (Zipf) distribution [A83, p.95]:

$$P[x = n] = 1/((n + 1)(n + 2)) \quad \text{for } n \geq 0.$$

arises in connection with platoon lengths for cars at different speeds traveling on an infinite road with no passing [A83, p.95] [F66, p.40], a model suggestively analogous to computer network traffic.

Following [LTWW94], we define a distribution as *heavy-tailed* if:

$$P[X \geq x] \sim cx^{-\beta}, \quad \text{as } x \rightarrow \infty, \quad \beta \geq 0. \quad (1)$$

By this, we mean that for some β and some constant c , the ratio $P[X \geq x]/(cx^{-\beta})$ tends to 1 as $x \rightarrow \infty$. This definition includes the Pareto and Weibull distributions [DMRW94].

A more general definition of *heavy-tailed* defines a distribution as heavy-tailed if the conditional mean exceedance (CME_x) of the random variable X is an increasing function of x [HK80], where

$$\text{CME}_x = E[X - x | X \geq x].$$

Using this second definition of heavy-tailed, consider a random variable X that represents a waiting time. For waiting times with a light-tailed distribution such as the uniform distribution, the conditional mean exceedance is a decreasing function of x . For such a light-tailed distribution, the longer you have waited, the sooner you are likely to be done. For waiting times with a medium-tailed distribution such as the (memoryless) exponential distribution, the expected future waiting time is independent of the waiting time so far. In contrast, for waiting times with a heavy-tailed distribution, the longer you have waited, the longer is your expected future waiting time. For the Pareto distribution with $\beta > 1$ (that is, with finite mean), the conditional mean exceedance is a linear function of x [A83, p.70]:

$$\text{CME}_x = x/(\beta - 1).$$

The Pareto distribution is scale-invariant, in that the probability that the wait is at least $2x$ seconds is a fixed fraction of the probability that the wait is at least x seconds, for any $x \geq a$.

A related result shows that the Pareto distribution is the only distribution that is “invariant under truncation from below” [M83, p.383] [A83, p.81]. That is, for the classical Pareto distribution, for $y \geq x_0$,

$$P[X > y | X > x_0] = P[(x_0/a)X > y]. \quad (2)$$

Hence the conditional distribution is also a Pareto distribution, with the same shape parameter β and new location parameter $a' = x_0$. We make use of this property in the next section.

Finally, we note that Mandelbrot argues that because the asymptotic behavior of Pareto distributions with $\beta \leq 2$ is unchanged for a wide variety of filters (including aggregation, maximums, and the weighted mixture of distributions), and because this is true of no other distribution, this invariance could in some respects explain

the widespread observance of Pareto distributions in the social sciences [M63] [M83, p.344].

C Pareto interpacket times

In this section we give some intuition for the observed long-range dependence of traces of TELNET traffic. Recall that the main body of the distribution of TELNET interpacket times fits a Pareto distribution with shape parameter 0.9, while the upper 3% tail fits a Pareto distribution with shape parameter 0.95. In this section we consider packets generated by a single connection using i.i.d. Pareto interpacket times, for a Pareto distribution with shape parameter β and location parameter a . We then consider the associated count process $X = \{X_i\}_{i=0,1,2,\dots}$, where X_i is the number of packets arriving during the i th time interval, each time interval being a bin of width b . We give an intuitive explanation for the observed long-range dependence of the count process by looking at the properties of the point process of packet arrivals, concentrating on the interpacket times. We show that while this process is not truly long-range dependent, when observed over a finite time scale it exhibits properties we associate with self-similar processes. In particular, we show that aggregating the process by increasing b does not change the dominant features of the process.

Let $\{X_i^{(b)}\}$ denote the count process associated with counting arrivals using bins of size b . We are interested in the behavior of $\{X_i^{(b)}\}$ for different sizes of b .

Rather than analyzing relationships between the precise values of different bins, we simplify the problem by just looking at whether, for a given i , $X_i^{(b)} = 0$ or $X_i^{(b)} > 0$. We refer to the former as an *empty* bin and the latter as an *occupied* bin. Further, for $j \geq i$, we call $X_{i,\dots,j}^{(b)}$ a *burst* of occupied bins if for all k , $i \leq k \leq j$, bin k is occupied. Similarly, $X_{i,\dots,j}^{(b)}$ is a *lull* if all the corresponding bins are empty. Sample paths of X are made up of alternating bursts and lulls.

We are interested in the relative predominance of bursts vs. lulls, as we change the bin size b and the Pareto shape parameter β .

Suppose bin i is occupied and bin $i - 1$ is empty. Then bin i begins a burst. Associated with each bin is a set of Pareto interarrival times, beginning with I_n , the arrival that first fell into the bin. For bin i , we know that $I_n > b$ because the previous bin is unoccupied. Consider now the subsequent interarrivals $I_{n+1} \dots I_{n+t}$ contributing to the burst of consecutive occupied bins. Clearly each of these interarrivals must be $< 2b$, as otherwise they will skip a bin and end the burst. Furthermore, any interarrival in the range $b < I < 2b$ has the potential of skipping a bin, depending on where we are positioned in the current bin prior to the arrival. Thus, any interarrival $I > 2b$ definitely will end the burst, and $I > b$ possibly will end the burst.

Since the interarrivals are independent, we have a situation similar to that of a geometric random variable: for any given interarrival, it will with probability p_t terminate the burst, and with probability $1 - p_t$ continue the burst. Here p_t is a function of exactly where we are in the current bin, but is bounded as follows:

$$\left(\frac{a}{2b}\right)^\beta \leq p_t \leq \left(\frac{a}{b}\right)^\beta, \quad (3)$$

where a and β are the Pareto location and shape parameters, and b is the bin width.

We can then bound the expected length of a burst using the expected value of the geometric random variables that correspond to

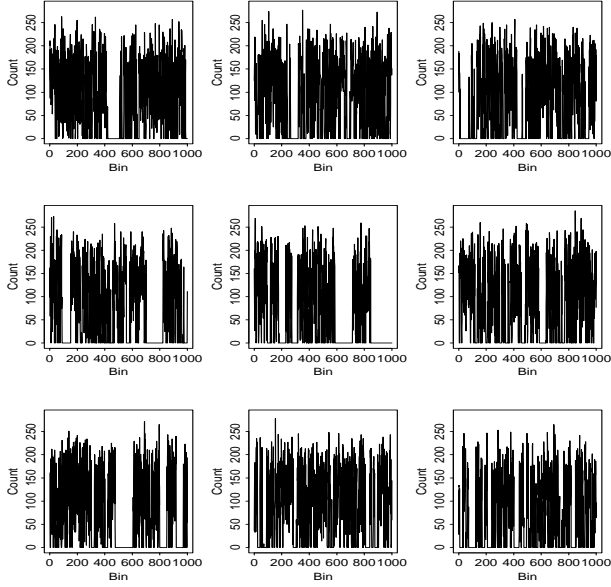


Figure 14: Count process for i.i.d. Pareto interarrivals, bin size $b = 10^3$ ($\beta = 1$, $a = 1$), 9 different seeds.

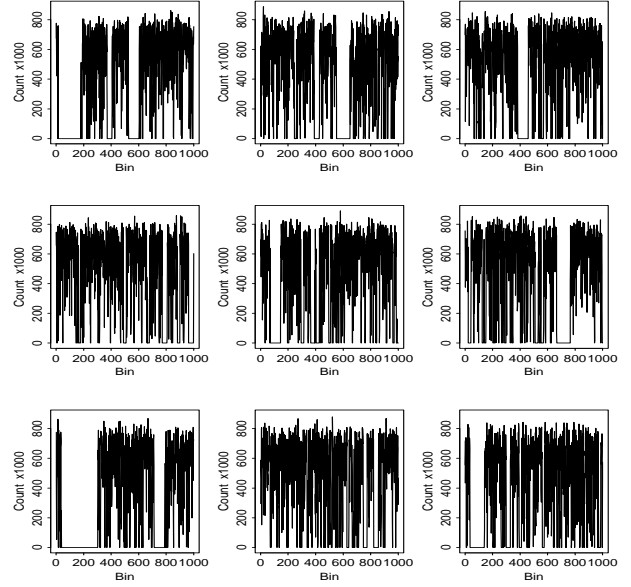


Figure 15: Count process for i.i.d. Pareto interarrivals, bin size $b = 10^7$ ($\beta = 1$, $a = 1$), 9 different seeds.

the lower and upper bounds in Equation 3. Let B be the expected number of bins spanned by a burst. It can be shown that:

$$B \approx \begin{cases} b/a, & \text{if } \beta = 2, b \gg a, \\ \log(b/a), & \text{if } \beta = 1, b \gg a, \text{ and} \\ \in [1, \sqrt{2}] & \text{if } \beta = \frac{1}{2}, \end{cases}$$

where $b \gg a$ holds if $b - a \approx b$.

Thus, for $\beta = 2$, as we “widen” our view by choosing b larger and larger, we will observe longer and longer bursts; for $\beta = 1$, the bursts grow longer with increasing bin size, but only very slowly; and for $\beta = \frac{1}{2}$, the bursts have a constant length regardless of the size of the bins (!).

Consider now the length of the lulls separating bursts. Let L be the length of a lull, and L_b be the number of bins (of size b) spanned by the lull. Each lull is due to a single interarrival that is possibly greater than $2b$ and definitely greater than b . Due to the Pareto distribution’s invariance to truncation from below (Equation 2), this means that the distribution of L will be stochastically bounded between $\mathcal{P}(b, \beta)$ and $\mathcal{P}(2b, \beta)$, where $\mathcal{P}(a, \beta)$ denotes the Pareto distribution with parameters a and β .

From this observation, it follows that:

$$1 - \left(\frac{2}{k}\right)^\beta \leq P[L_b \leq k] \leq 1 - \left(\frac{1}{k}\right)^\beta.$$

Thus, the distribution of L_b is *invariant with respect to b* . That is, regardless of the time scale over which we view the count process, the lulls between bursts will “look” the same.

We now can summarize the behavior of the count process for varying values of β :

- For $\beta = 2$, the number of bins spanned by the bursts grows linearly with b , while bins spanned by the lulls remains constant, so aggregation fairly quickly smoothes out the main variations of the count process.

- For $\beta = \frac{1}{2}$, the burst lengths are constant across all time scales, as are the lull lengths: *the process appears self-similar over all time scales.*
- For $\beta = 1$, the burst lengths (in bins) grow only very slowly (logarithmically). This means that over a large time scale, the predominance of bursts vs. lulls remains virtually unchanged: *the process appears self-similar over many time scales.*

Figures 14 and 15 illustrate the “visual self-similarity” [LTWW94] of this process. Each figure plots 1,000 observations of the count process corresponding to i.i.d. Pareto interpacket times for $\beta = 1$ and $a = 1$. Nine different random seeds were used in generating each figure. The first figure corresponds to using a bin-width of $b = 10^3$, while the second figure uses $b = 10^7$. To the eye, the two sets of arrivals exhibit the same general activity in terms of alternations of bursts and lulls and the fairly regular ceiling of activity, though the occupied bins of the $b = 10^7$ arrivals appear to have a higher mean than those of the $b = 10^3$ arrivals. As predicted by the analysis above, the average number of bins in a burst for $b = 10^7$ is somewhat higher than for $b = 10^3$ (a factor of 2.6), while the average lull size is virtually the same (a factor of 1.2). Overall, the sustained variation even when the process is aggregated by a factor of 10^4 is striking.

In general, the process associated with $\beta = 1$ is similar to that of a single TELNET connection’s traffic, which we model using i.i.d. Pareto interpacket times with $\beta = 0.95$ for the upper tail of the distribution. Thus this model explains in part why TELNET traffic appears self-similar.

We finish with an explanation of why the count processes associated with $\beta = 1$ and $\beta = \frac{1}{2}$ are not, in fact, self-similar, even though the balance they exhibit between bursts and lulls suggests they might be. We have shown that the lull length L is stochastically bounded between two Pareto distributions with the same shape parameter β . But for $\beta \leq 1$, the mean of a Pareto-distributed random variable is infinite. The expected burst size, on the other hand, is

finite. Using these facts, and viewing the count process's bursts and lulls as an alternating renewal process, it follows that, for $\beta \leq 1$, once the process reaches steady-state, each bin is empty with probability 1 (regardless of the value of b). The autocorrelation function of the process is thus 0 everywhere, and hence summable, so the process is not long-range dependent (and so cannot be self-similar).

Even though the count processes are not strictly self-similar, an important point remains that, when viewed over a finite time scale (i.e., before settling into steady-state), the count process associated with i.i.d. Pareto interarrivals (with $\beta \leq 1$) appears in many ways like a self-similar process. Assuming that this likeness persists when the process is multiplexed, this finding gives an understanding as to why observed TELNET traffic appears self-similar. The fact that the count process is not truly long-range dependent does *not* imply that TELNET traffic is not truly self-similar. It may be that TELNET traffic *is* truly self-similar but the simplifying assumptions in our argument (i.i.d. arrivals; no multiplexing) fail to faithfully model the traffic properties necessary for true self-similarity.

This argument also shows that it is possible for a process which is not long-range dependent to appear to be so over many time scales. This illustrates some of the dangers of arguing for true self-similarity (or, more generally, long-range dependence) based on (necessarily finite) measurements alone, without a corresponding model from which to argue for self-similarity analytically.

At the same time, the question of whether a particular (infinite) model based on a finite process is long-range dependent is only one of the questions we are exploring. Equally important is whether or not long-range dependent models in general are useful as parsimonious approximations to particular finite processes arising in network traffic. Finally, we should not underestimate the value of the fundamental insights and shifts in focus that come from considering questions of self-similarity and long-range dependence.

D The M/G/ ∞ model for generating self-similar traffic

This section briefly discusses the M/G/ ∞ model for generating self-similar traffic [CI80, p.136] [C84, p.67]. The M/G/ ∞ queue model considers customers that arrive at an infinite-server queue according to a Poisson process with rate ρ . In the count process $\{X_t\}_{t=0,1,2,\dots}$ produced by the M/G/ ∞ queue model, X_t gives the number of customers in the system at time t . From [CI80, p.139], for customers with a service time with distribution function F , the autocorrelation function $r(k)$ for the count process is as follows:

$$r(k) = \text{cov}\{X(t), X(t+k)\} = \rho \int_k^\infty (1 - F(x)) dx. \quad (4)$$

D.1 The M/G/ ∞ model and the Pareto distribution

Consider customers with independent service times (or lifetimes) drawn from the Pareto distribution with location parameter a and shape parameter β , for $1 < \beta < 2$. From Equation 4, the autocorrelation function $r(k)$ is as follows:

$$r(k) = \rho \int_k^\infty \left(\frac{a}{x}\right)^\beta dx.$$

$$= \frac{\rho a^\beta}{\beta - 1} k^{(1-\beta)}.$$

Following [BSTW94], the process $\{X_t\}_{t=0,1,2,\dots}$ is *asymptotically self-similar* if

$$r(k) \sim k^{-D} L(k) \text{ as } k \rightarrow \infty, \quad (5)$$

for $0 < D < 1$ and L a slowly-varying function.² Thus, for $a \geq 0$ and $1 < \beta < 2$, the count process of the M/G/ ∞ model with Pareto lifetimes is asymptotically self-similar, and therefore long-range dependent.

From [BSTW94], the process $\{X_t\}_{t=0,1,2,\dots}$ is *exactly self-similar* only if

$$r(k) = 1/2 \left((k+1)^{2H} - 2k^{2H} + (k-1)^{2H} \right)$$

for $1/2 < H < 1$ [BSTW94] [C84, p.59]. In this case the process $\{X_t\}$ and the aggregated process $\{X_t^{(m)}\}$ have the same autocorrelation function. From this result, for Pareto service times and an arbitrary arrival rate ρ , the count process of the M/G/ ∞ model is not exactly self-similar.

From [CI80, p.138], $\{X_t\}$ has a Poisson marginal distribution with mean $\rho\mu$, where μ is the expected service time. For the M/G/ ∞ model with Pareto service times, the expected service time is $\beta a / (\beta - 1)$, for $\beta > 1$. Thus, in this case $\{X_t\}$ has a Poisson marginal distribution with mean $\rho\beta a / (\beta - 1)$.

E Log-normal distributions

From [WT92], the log-normal distribution is called *sub-exponential* because, along with the Pareto and Weibull distributions, the tail function is subexponential (i.e., decreases slower than any exponential function). In that paper, the Pareto, log-normal, and Weibull distributions are all defined as *long-tailed*. In this section we show that the log-normal distribution is not heavy-tailed, according to the definition given in Equation 1.

We use the estimate of the upper tail function for a standard normal random variable N as

$$P[N \geq y] \sim \frac{1}{\sqrt{2\pi y}} e^{-y^2/2}$$

[F50, p.175]. Thus for X a log-normal random variable with scale parameter 1 and shape parameter 1,

$$P[X \geq x] \sim \frac{1}{\sqrt{2\pi \log x}} e^{-\log^2 x/2}. \quad (6)$$

Thus, for some constant c ,

$$P[X \geq x] \sim c \frac{e^{-\log^2 x/2}}{\log x}.$$

So X is only heavy-tailed if for some constant c_1 and some $\beta \geq 0$,

$$x^\beta \sim c_1 \log x e^{\log^2 x/2}.$$

But we can show that for any n ,

$$\log x e^{\log^2 x/2} > x^n$$

²For a slowly-varying function L , $\lim_{t \rightarrow \infty} L(tx)/L(t) = 1$ for all $x > 0$. Constants and logarithms are examples of slowly-varying functions.

for x sufficiently large. (This follows because $\log x > n$, therefore $\log^2 x > n \log x$, and therefore $e^{\log^2 x} > x^n$.) So the log-normal distribution is not heavy-tailed. Note that the log-normal distribution is not heavy-tailed even if we expand our definition of heavy-tailed to include slowly-varying functions, as in Equation 5.

E.1 The M/G/ ∞ model and the log-normal distribution

We consider the M/G/ ∞ model for service times with distribution function F . It is already known (Appendix D) that if F is a Pareto distribution, then the count process from the M/G/ ∞ model is asymptotically self-similar, and therefore long-range dependent. In this section we show that if the lifetimes have a log-normal distribution, then the count process from the M/G/ ∞ model is not long-range dependent.

From Equations 4 and 6, we have:

$$\begin{aligned} r(k) &\sim \rho \int_k^\infty \log^{-1} x \frac{1}{(2\pi)^{1/2}} e^{-\log^2 x / 2} dx \\ &\sim \frac{\rho}{(2\pi)^{1/2}} \int_k^\infty \frac{1}{\log x x^{(\log x)/2}} dx \end{aligned}$$

The count process from the M/G/ ∞ model with log-normal lifetimes is long-range dependent only if $\sum_{k=K}^\infty r(k)$ is infinite. For large K ,

$$\begin{aligned} \sum_{k=K}^\infty r(k) &\sim \sum_{k=K}^\infty \frac{\rho}{(2\pi)^{1/2}} \int_k^\infty \frac{1}{\log x x^{(\log x)/2}} dx \\ &\sim \frac{\rho}{(2\pi)^{1/2}} \sum_{k=K}^\infty \sum_{x=k}^\infty \frac{1}{\log x x^{(\log x)/2}} \\ &\sim \frac{\rho}{(2\pi)^{1/2}} \sum_{x=K}^\infty \frac{(x-K+1)}{\log x x^{(\log x)/2}}. \end{aligned}$$

Because $\sum_{x=1}^\infty 1/x^2$ is finite and

$$\frac{(x-K+1)}{\log x x^{(\log x)/2}} \leq \frac{x}{x^{(\log x)/2}} \leq \frac{1}{x^2}$$

for x sufficiently large, then $\sum_{k=K}^\infty r(k)$ is finite, and the count process of the M/G/ ∞ model with log-normal lifetimes is not long-range dependent.

This analysis shows that, *in the limit*, the behavior of the M/G/ ∞ queue completely changes if the service times are log-normal and not Pareto. An important open question, however, is over what sort of finite time scales are these differences actually significant?

References

[A83] B. Arnold., *Pareto Distributions*, International Co-operative Publishing House, Maryland, 1983.

[B92a] J. Beran, "A Goodness-of-fit Test for Time Series with Long Range Dependence," *Journal of the Royal Statistical Society B*, 54(3), pp. 749-760, 1992.

[B92b] J. Beran, "Statistical Methods for Data with Long-Range Dependence", with discussion, *Statistical Science*, 7(4), pp. 404-427, 1992.

[BSTW94] J. Beran, R. Sherman, M. Taqqu, and W. Willinger, "Variable-Bit-Rate Video Traffic and Long-Range Dependence," to appear in *IEEE Transactions on Communications*.

[B90] D. Borman, "Telnet Linemode Option," RFC 1184, Network Information Center, SRI International, Menlo Park, CA, October, 1990.

[CBP93] K. Claffy, H.-W. Braun and G. Polyzos, "Long-term traffic aspects of the NSFNET," *Proceedings of INET '93*, SDSC Report GA-A21238, San Diego Supercomputer Center, 1993.

[CSZ92] D. Clark, S. Shenker, and L. Zhang, "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism," *Proceedings of SIGCOMM '92*, pp. 14-26, August, 1992.

[CI80] D. Cox and V. Isham, *Point Processes*, Chapman and Hall, 1980.

[C84] D. R. Cox, "Long-Range Dependence: A Review," in *Statistics: An Appraisal*, Proceedings 50th Anniversary Conference, Iowa State Statistical Library, H. A. David and H. T. David, editors, Iowa State University Press, pp. 55-74, 1984.

[DS86] R. B. D'Agostino and M. A. Stephens, editors, *Goodness-of-Fit Techniques*, Marcel Dekker, Inc., 1986.

[DJ91] P. Danzig and S. Jamin, "tcplib: A Library of TCP Internetwork Traffic Characteristics," *Report CS-SYS-91-01*, Computer Science Department, University of Southern California, 1991. Available via FTP to catarina.usc.edu as `pub/jamin/tcplib/tcplib.tar.Z`.

[DJCME92] P. Danzig, S. Jamin, R. Cáceres, D. Mitzel, and D. Estrin, "An Empirical Workload Model for Driving Wide-area TCP/IP Network Simulations," *Internetworking: Research and Experience*, 3(1), pp. 1-26, March, 1992.

[DMRW94] D. Duffy, A. McIntosh, M. Rosenstein, and W. Willinger, "Statistical Analysis of CCSN/SS7 Traffic Data from Working CCS Subnetworks," *IEEE JSAC*, 12(3), pp. 544-551, April, 1994.

[F50] W. Feller, *An Introduction to Probability Theory and Its Applications, Volume I*, John Wiley and Sons, 1950.

[F66] W. Feller, *An Introduction to Probability Theory and Its Applications, Volume II*, John Wiley and Sons, 1966.

[FJ92] S. Floyd and V. Jacobson, "On Traffic Phase Effects in Packet-Switched Gateways," *Internetworking: Research and Experience*, 3(3), pp. 115-156, September, 1992.

[FJ94] S. Floyd and V. Jacobson, "The Synchronization of Periodic Routing Messages," *IEEE/ACM Transactions on Networking*, 2(2), pp. 122-136, April 1994.

[FL91] H. Fowler and W. Leland, "Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management," *IEEE JSAC*, 9(7), pp. 1139-1149, September, 1991.

[FM94] V. Frost and B. Melamed, "Traffic Modeling for Telecommunications Networks," *IEEE Communications Magazine*, 32(3), pp. 70-80, March, 1994.

- [FJ70] E. Fuchs and P. E. Jackson, "Estimates of Distributions of Random Variables for Certain Computer Communications Traffic Models," *Communications of the ACM*, 13(12), pp. 752-757, December, 1970.
- [GW94] M. Garrett and W. Willinger, "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic," *Proceedings of SIGCOMM '94*, pp. 269-280, September, 1994.
- [G90] R. Gusella, "A Measurement Study of Diskless Workstation Traffic on an Ethernet," *IEEE Transactions on Communications*, 38(9), pp. 1557-1568, September, 1990.
- [HK80] T. Hettmansperger, and M. Keenan, "Tailweight, Statistical Inference, and Families of Distributions - A Brief Survey," in *Statistical Distributions in Scientific Work*, V.1, G. P. Patil et al (eds), Kluwer Boston, pp. 161-172, 1980.
- [I93] G. Irlam (gordoni@netcom.com), "ufs'93 [Updated file size survey results]," USENET newsgroup comp.os.research, message 2ddp3b\$jn5@darkstar.UCSC.EDU, Nov. 29, 1993.
- [JR86] R. Jain and S. Routhier, "Packet Trains — Measurements and a New Model for Computer Network Traffic," *IEEE JSAC*, 4(6), pp. 986-995, September, 1986.
- [K93] S. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*, Oxford University Press, 1993.
- [LO86] W. Leland and T. Ott, "Load-balancing Heuristics and Process Behavior," *PERFORMANCE '86 and ACM SIGMETRICS 1986 Joint conference on Computer Performance Modelling, Measurement and Evaluation*, North Carolina State University, pp. 54-69, May 1986.
- [LTWW94] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)," *IEEE/ACM Transactions on Networking*, 2(1), pp. 1-15, February 1994.
- [M63] B. Mandelbrot, "New Methods in Statistical Economics," *Journal of Political Economy*, 71(5), pp. 421-440, October, 1963.
- [M83] B. Mandelbrot, *The Fractal Geometry of Nature*, Freeman, New York, 1983.
- [MM85] W. Marshall and S. Morgan, "Statistics of Mixed Data Traffic on a Local Area Network," *Computer Networks and ISDN Systems* 10(3,4), pp. 185-194, 1985.
- [N84] J. Nagle, "Congestion Control in IP/TCP Internetworks," RFC 896, Network Information Center, SRI International, Menlo Park, CA, January, 1984.
- [P86] C. Partridge, "Mail Routing and the Domain System," RFC 974, Network Information Center, SRI International, Menlo Park, CA, January, 1986.
- [P94a] V. Paxson, "Empirically-Derived Analytic Models of Wide-Area TCP Connections," *IEEE/ACM Transactions on Networking*, 2(4), pp. 316-336, August, 1994.
- [P94b] V. Paxson, "Growth Trends in Wide-Area TCP Connections," *IEEE Network*, 8(4), pp. 8-17, July/August, 1994.
- [PR83] J. Postel and J. Reynolds, "Telnet Protocol Specification," RFC 854, Network Information Center, SRI International, Menlo Park, CA, May, 1983.
- [S94] W. Richard Stevens, *TCP/IP Illustrated, Volume 1: The Protocols*, Addison-Wesley, 1994.
- [WT92] E. Willekens and J. Teugels, "Asymptotic expansions for waiting time probabilities in an M/G/1 queue with long-tailed service time," *Queueing Systems* 10, pp. 295-312, 1992.
- [YKTH93] D. Yates, J. Kurose, D. Towsley, and M. Hluchyj, "On per-session end-to-end delay distributions and the call admission problem for real-time applications with QOS requirements," *Proceedings of SIGCOMM '93*, pp. 2-12, September, 1993.