LBL-35238
UC-405

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Information and Computing Sciences Division

Presented at the SIGCOMM 1994 Conference, London, England, August 27–September 4, 1994, and to be published in IEEE/ACM Transactions of Networking

Wide-Area Traffic: The Failure of Poisson Modeling

V. Paxson and S. Floyd

August 1994

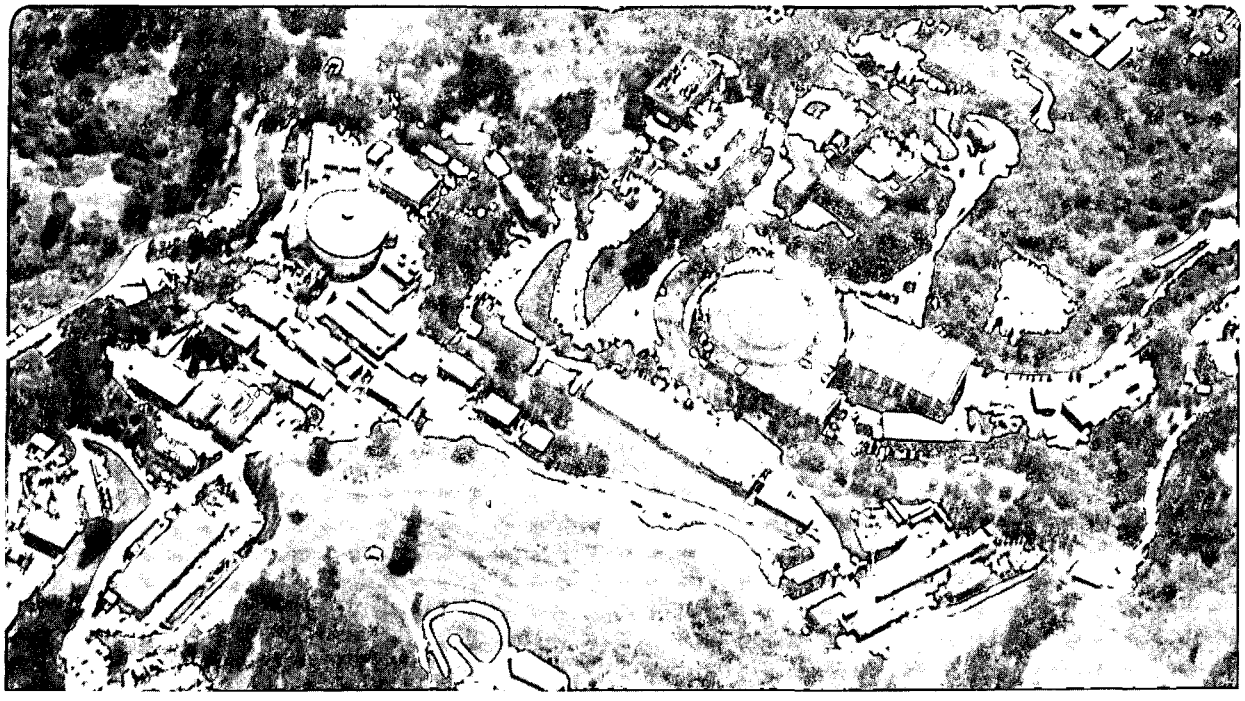

REFERENCE COPY
Does Not
Circulate
Bldg. 50 Library.
LBL-35238
Copy 1

Prepared for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

LBL–35238
UC–405

Wide-Area Traffic: The Failure of Poisson Modeling

Vern Paxson and Sally Floyd
Information and Computing Sciences Division
Lawrence Berkeley Laboratory
University of California, Berkeley, CA 94720

August 1994

This work was supported by the Director, Office of Energy Research,
Office of Information and Computing Sciences Division, of the
U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

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

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 21 wide-area traces, investigating a number of wide-area TCP arrival processes (session and connection arrivals, FTPDATA 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[DJCME92] interarrivals preserves burstiness over many time scales; and that FTPDATA connection arrivals within FTP sessions come bunched into "connection bursts", the largest of which are so large that they completely dominate FTPDATA traffic. Finally, we offer some preliminary 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.

For our study we analyze 21 traces of wide-area TCP traffic. We consider both previous and new models of aspects of FTP and TELNET 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 result in serious underestimations of 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. 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 [FJ93], 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 FTPDATA connections within a single FTP session are clustered into *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.

*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.

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 might 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.
U.S.C. (USC)	22Jan91	26 hours	13K 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.
LBL PKT-1	17Dec93	2 hours	1.7M TCP pkts.
LBL PKT-2	19Jan94	2 hours	2.4M TCP pkts.
LBL PKT-3	20Jan94	2 hours	1.8M TCP pkts.
LBL PKT-4	21Jan94	1 hour	1.3M pkts.
LBL PKT-5	28Jan94	1 hour	1.3M pkts.

Table 1: Summary of Wide-Area Traces

Table 1 summarizes the traces of wide-area traffic used in our study. The first set of rows represent traces previously analyzed: the BC, UCB, and USC traces in [DJCME92][1], the NC, UK, and DEC traces in [P93], and the LBL traces in [P93, P94]. The "DEC 1-3" rows represents 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 final five rows reflect new traces we gathered for our study. Each of these traces began at 2PM; the first three 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.

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 connections, on the other hand, are decidedly not Poisson processes.

[1]These traces captured all WAN packets, but our analysis in this paper uses only the TCP SYN/FIN connection start/stop packets.

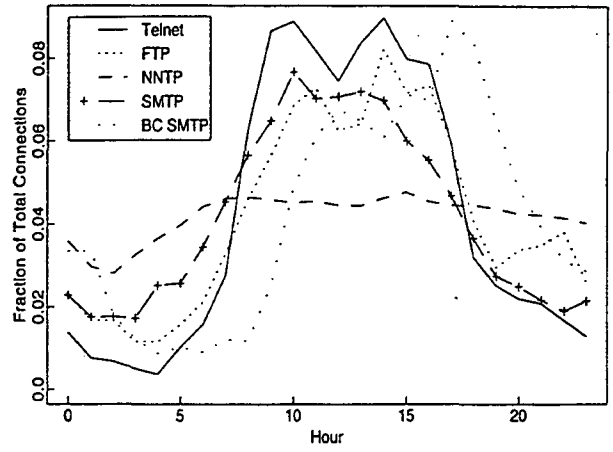

Figure 1: Mean, relative, hourly connection arrival rate for LBL-1 through LBL-4 datasets.

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 protocol occurring during that hour.[2] 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 [P93]). 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

[2]In Figure 1, FTP refers to FTP sessions.

shown in the first half of Table 1. For each trace, we separately tested the trace's TELNET, FTP, FTPDATA, SMTP, NNTP, and WWW (World Wide Web) connections[3]. Here FTP refers to an FTP *session* (i.e., an FTP control connection), while FTPDATA refers to the data-transfer connections spawned by these control connections[4]. We also tested arrivals of FTPDATA *bursts* (see § 6 below).

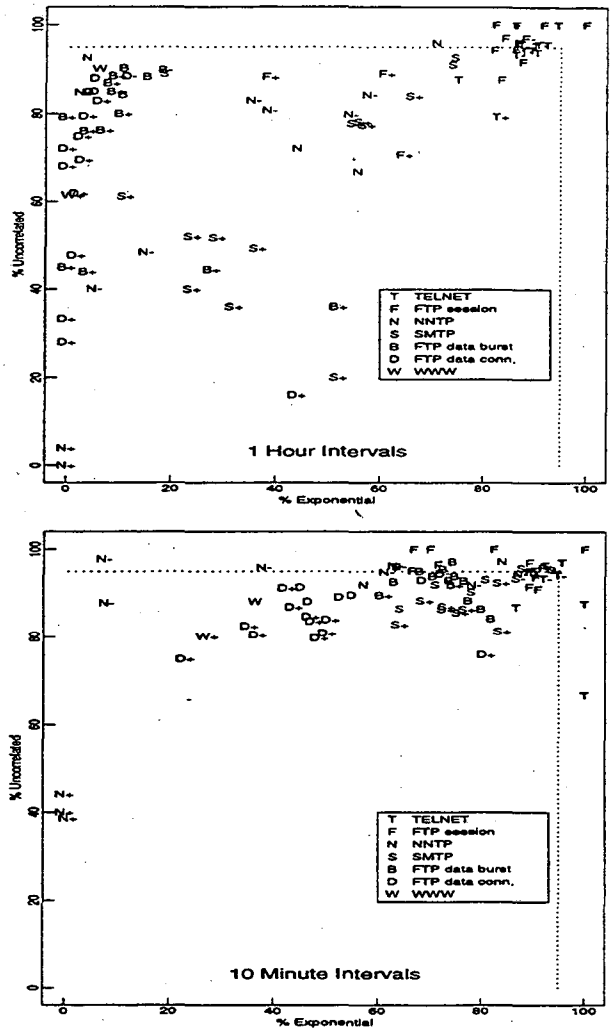

Figure 2: Results of testing for Poisson arrivals.

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 **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 [S94].

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 [P93]. 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 con-

[3] Only two traces had significant WWW traffic.

[4] We first removed the periodic "weather-map" FTP traffic discussed in [P94].

nection originator; this in general is a user typing at a keyboard. Because these 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.

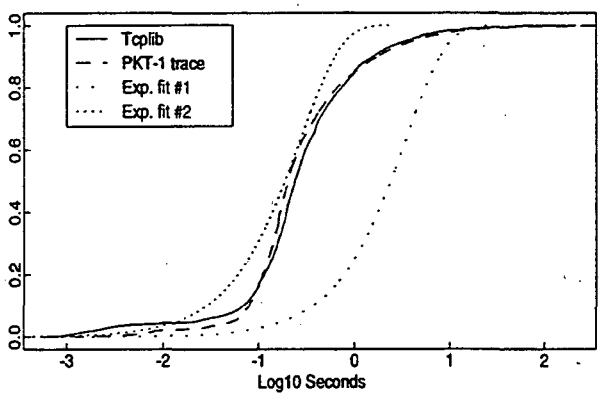

Figure 3: Empirical distributions of packet-interarrivals within TELNET connections.

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 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 PKT-1 distribution, and the righthand one has the same arithmetic mean. The exponential fits are very poor. On the other hand, the main body of the 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.

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?

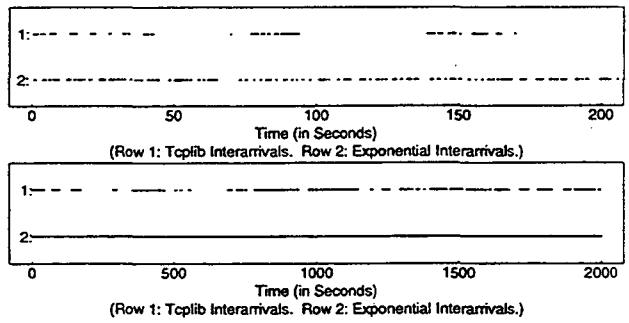

Figure 4: Comparisons between Tcplib and exponential interpacket times.

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. Simulation also shows that the greater burstiness of Tcplib connections persists even with 100 multiplexed TELNET connections [PF94].

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 queueing 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[5] from the two-hour 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

[5]Except for "pure ack" packets, containing no user data.

when during the 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 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. If 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 the previous example $M = 100$.

To construct a variance-time plot, we smooth the process for different values of M and then plot the variance of the smoothed process on the y-axis vs. the aggregation level (M) on the x-axis, using logarithmic scales.

Variance-time plots are useful for gauging burstiness over many different time scales; straight lines on variance-time plots with slopes more shallow than -1 are also suggestive of self-similarity (see § 7 for further discussion).

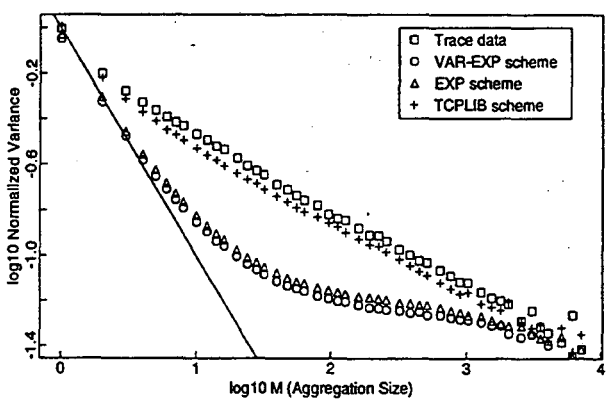

Figure 5: Variance-Time Plot for TELNET packet arrival process. The line from the upper left corner has slope -1.

Figure 5 shows such a plot for the 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.[6]

From the plot it is immediately clear that the variance of the TCPLIB scheme agrees closely with the PKT-2 trace data, while

[6]The traces consisted of between 82,500 and 86,000 packets.

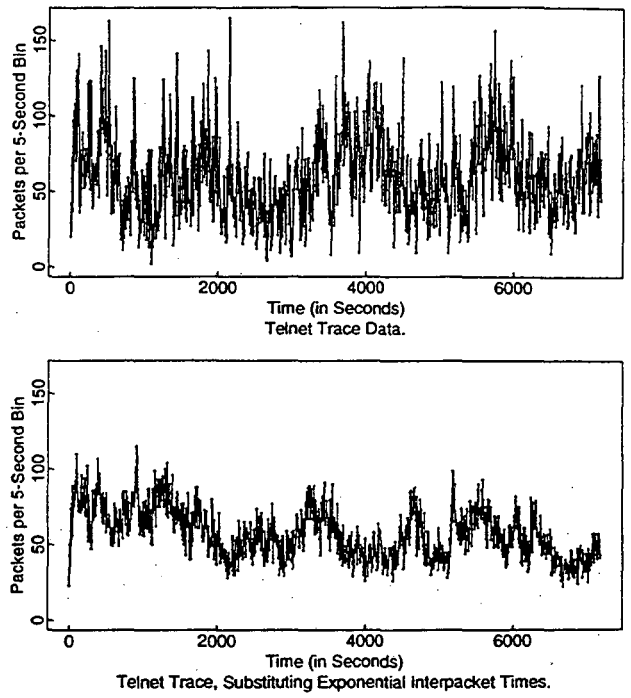

Figure 6: Comparisons of actual and exponential TELNET packet interarrival times.

both EXP and VAR-EXP exhibit far less variance, indicating they are much less bursty over a large range of time scales.

Figure 6 shows this explicitly. Here we plot the arrival process corresponding to 5-second intervals ($M = 50$) for the 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 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 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 queueing delays in simulations using these two different traces. As the variance-time plot shows, the 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 connec-

tion size in *bytes* has been previously modeled by a log-extreme distribution [P93]; the distribution of the connection size in *packets* 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 $\beta = \log_2 3.5$ [P93]. We experimented with using this distribution to produce sizes for an equal number of TELNET connections as appeared in the 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 [P93] 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 [P93] 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 or because the TELNET connection is operating in "line mode" [S94]. For example, the 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[7], considerably better than a log-extreme distribution with parameters fitted to the data.[8]

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 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.[9]

[7] The exact numerical values of $\bar{x}$ and σ should not be taken too seriously, as they came from a small sample.

[8] We also found that a log-extreme distribution fit the connection size in bytes better than a log-normal distribution.

[9] We also tested the model's fit to the PKT-1 and PKT-3 TELNET traces; the results were similar.

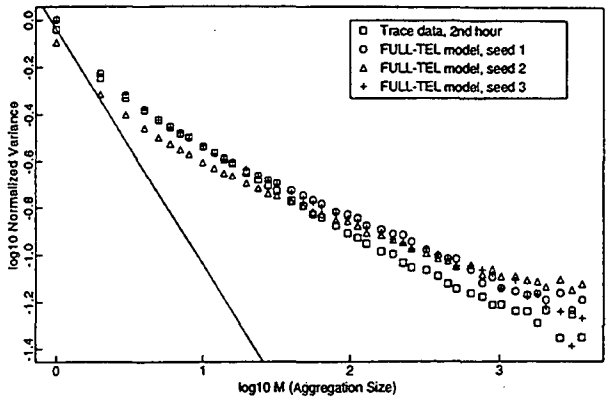

Figure 7: Variance-time plot comparing PKT-2 trace data with the complete TELNET model, FULL-TEL.

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 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. 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 vertical line)

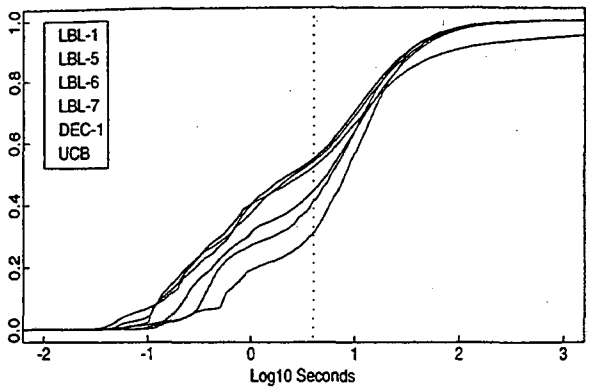

Figure 8: FTPDATA Intra-session Connection Spacing.

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.

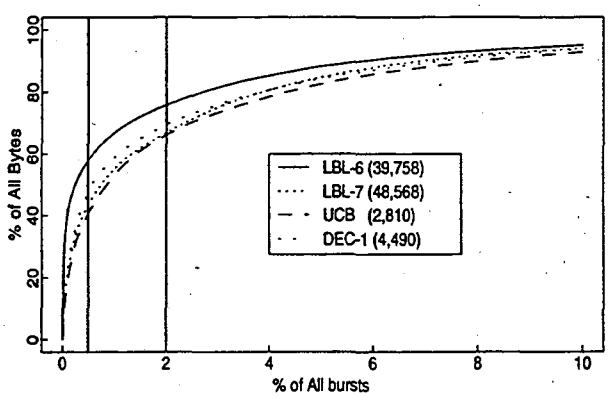

Figure 9: Percentage of all FTPDATA bytes due to largest 10% FTPDATA bursts.

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 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 40% 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*. 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.[10]

We did simple fitting of the upper tail of the distribution of data bytes per FTPDATA burst and found that for all six datasets, the upper 5% tail is fit well to a Pareto distribution with $0.9 \leq \beta \leq 1.1$. As the Pareto distribution is heavy-tailed (see Appendix B), this agrees with our findings in Figure 9.[11]

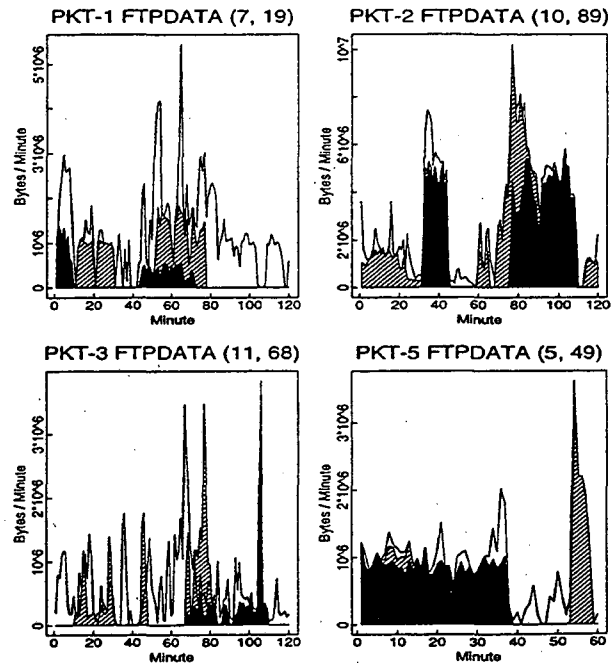

Figure 10: Proportion of FTPDATA traffic due to largest 2% (shaded) and 0.5% (black) connection bursts.

Figure 10 graphically illustrates the dominance of the upper FTPDATA-burst tail. The four plots show the FTPDATA traffic rate in bytes/minute for the PKT-1, PKT-2, PKT-3, and PKT-5 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.) 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.[12]

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

[10]Our finding that the size of an FTP 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.

[11]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.

[12]For example, one of the bursts in the LBL-7 dataset was made up of 979 separate FTPDATA connections.

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.

We would also like to know whether the arrivals of the upper-tail 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 PKT-1 through PKT-5 traces of TELNET traffic, FTPDATA traffic, and general 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 two new methods of generating "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 aggregate wide-area traffic. We find the 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 long 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 hyperbolically-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 Methods for generating 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 Poisson start times, 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/∞ 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 $\{X_t\}_{t=0,1,2,\ldots}$ is asymptotically self-similar (see [PF94] for further discussion). The M/G/∞ model implies that multiplexing constant-rate connections with Poisson connection arrivals and a heavy-tailed distribution for connection lifetimes would result in self-similar traffic.

We have found two additional methods of generating arrival processes that possibly exhibit self-similarity. We refer to these methods as generating "pseudo-self-similar" traffic, because we have not shown in any solid way that they are truly self-similar processes. Both methods are very fast. The first is Fourier-transform based (see [PF94] for details). Traffic generated using this method passes Beran's goodness-of-fit test for fractional Gaussian noise [B92a]. The second method is to construct arrivals using i.i.d. Pareto interarrivals with $\beta \approx 1$, and to consider the corresponding count process (the number of arrivals in consecutive intervals). In [PF94] we develop some intuition behind why this method might generate asymptotically self-similar traffic.

7.3 Relating the methods to traffic models

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[13].

In addition to looking at variance-time plots of the TELNET traffic, we also applied Whittle's procedure [GW94, LTWW94] and Beran's goodness-of-fit test [B92a]. All of the results are consistent with self-similarity on scales of tens of seconds or more. One way of explaining such findings of self-similarity is to note that

[13] That is, the variance-time plot declines 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.

our model of TELNET connections presented in § 5 in some respects matches the M/G/∞ model described in the previous section. For example, 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 (although the log-normal distribution is not heavy-tailed [PF94]). Thus, the M/G/∞ model can give some intuition to the suggested self-similarity of our TELNET traces and models.

Another source of possible TELNET self-similarity arises 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. Since the aggregation of multiple self-similar traffic sources remains self-similar, this would lead to aggregate TELNET traffic appearing self-similar.

Our model of FTP traffic also fits in some respects to the M/G/∞ model of Poisson arrivals with heavy-tailed lifetimes. The distribution of bytes per FTPDATA *burst* is heavy-tailed, and FTP *sessions* have Poisson arrivals. 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 the aggregate FTP traffic would fit the M/G/∞ 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 FTP traces are not self-similar, although the heavy-tailed distribution of FTPDATA bursts clearly leads to large-scale correlations. The following paragraphs discuss several ways that aggregate FTP traffic differs from the M/G/∞ model of self-similar traffic described earlier. While these factors could account for our lack of finding self-similarity in our FTP traces, they do not imply the absence of long-range dependence.

First, even in the absence of congestion, different FTP connections can have quite different average rates; the average rate for a particular connection depends on such factors as the TCP window and the roundtrip time. This could be a major discrepancy between our trace data and the M/G/∞ model. Of particular relevance would be the average rates of the biggest FTP bursts.

A second factor concerns the effect of bandwidth limitations on multiplexed FTP traffic. The simplest way to incorporate the limited bandwidth on a congested link would be to assume a limited capacity in the M/G/∞ model for generating self-similar traffic, where the actual arrival times of individuals would occasionally have to be delayed until there was available capacity. This would transform the M/G/∞ queue into an M/G/k queue. While this limited capacity would have the effect of reducing the fit of the aggregate traffic to a self-similar model, it does not eliminate the underlying large-scale correlations.

A third factor concerns the effect of FTP traffic competing with other families of traffic on a congested link. The three main classes of traffic in our link traces are TCP, Mbone (primarily multicast UDP audio traffic) and DECnet. Because the UDP protocol does not incorporate congestion-avoidance mechanisms, when FTP traffic is competing for bandwidth with 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 aggregate wide-area traffic

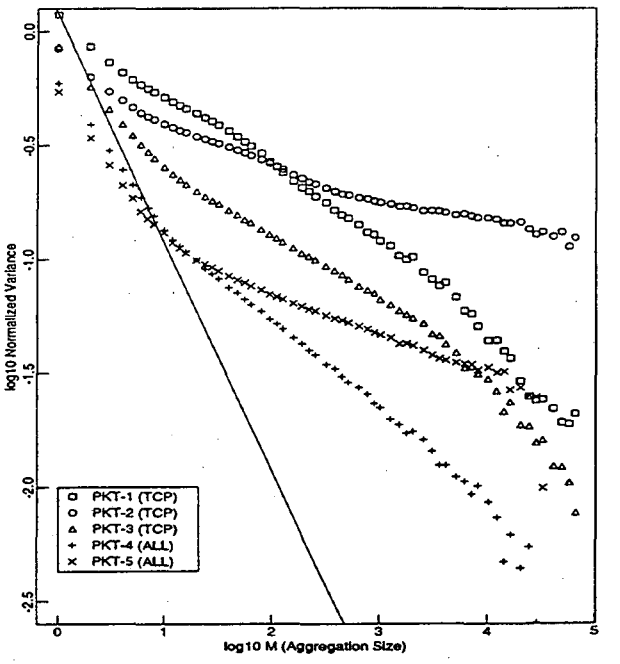

Figure 11: Variance-time plot for all TCP / all link-level packet arrivals.

We finish with a preliminary look at whether wide-area traffic aggregated over different protocols appears self-similar. Figure 11 shows variance-time plots for all of the "PKT" traces listed in Table 1. Here, the unaggregated process ($M = 1$) corresponds to observing the packets arriving during each 0.01 second interval.

Recall that the first three 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 11 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 self-similar 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 FTP 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. 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 aggregate link traffic as self-similar, without attempting to model individual connections. This approach could have many uses in simulations and in analysis. For example, aggregate self-similar traffic could be used instead of Poisson traffic to model cross-traffic, or aggregate 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. For example, in simulations that investigate the effects of different transport protocols or different gateway scheduling algorithms on network traffic, the simulator requires source models; the traffic patterns on the link will depend on the transport protocols and scheduling algorithms that are used in the simulations, as well as on the pattern of traffic generated by the source.

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 aggregate traffic. In this paper observations of the characteristics of aggregate 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] has examined 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 large 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 TCP 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 enforced traffic parameters of individual connections. As has been show by numerous researchers, such admissions control procedures could lead to a much more effective use of the available bandwidth. 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 mislead following a long period of fairly low traffic rates.[14]

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 busy 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 aggregate 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, Allyn Romanow, and Bin Yu, as well as the comments of the anonymous reviewers.

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

We would also like to thank Peter Danzig and his coauthors for the UCB, USC, and BC datasets; Jeff Mogul for the DEC datasets; Ian Wakeman and Jon Crowcroft for the UK dataset; Wayne Sung for the NC dataset; and especially Ramón Cáceres and Sugih Jamin, who between them made available all of the non-LBL datasets.

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

[14] 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.

A Methodology for testing for Poisson arrivals

To test whether a trace of connection arrivals corresponds to a nonhomogeneous 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.

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]:

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

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].

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 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 for some constant c,

$$P[X \geq x] \sim cx^{-\beta}, \quad \text{as } x \to \infty, \; \beta \geq 0.$$

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, and for x such that $f(x) > 0$, 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 $y \geq x_0$,

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

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

References

[A83] B. Arnold., *Pareto Distributions*, International Cooperative 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.

[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.

[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 excalibur.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, 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.

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

[FJ93] S. Floyd and V. Jacobson, "The Synchronization of Periodic Routing Messages," *SIGCOMM '93*, September 1993, pp. 33-44. Also to appear in *IEEE/ACM Transactions on Networking*.

[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", to appear in *SIGCOMM '94*.

[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.

[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*, 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.

[P93] V. Paxson, "Empirically-Derived Analytic Models of Wide-Area TCP Connections: Extended Report", technical report LBL-34086, Lawrence Berkeley Laboratory, May, 1993. Available from ftp.ee.lbl.gov as *papers/WAN-TCP-models.prelim.1.ps.Z* and *papers/WAN-TCP-models.prelim.2.ps.Z*.

[P94] V. Paxson, "Growth Trends in Wide-Area TCP Connections", to appear in *IEEE Network*, July, 1994.

[PF94] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling (Extended Version)", technical report LBL-35238, Lawrence Berkeley Laboratory, August, 1994. Available as *papers/WAN-poisson-extended.ps.Z* from ftp.ee.lbl.gov.

[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.

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
TECHNICAL INFORMATION DEPARTMENT
BERKELEY, CALIFORNIA 94720

AAV339
LBL Libraries