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

To appear in SIGCOMM '94

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 fifteen 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 are less persuasively 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 speculations 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. 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, Gusella90, FL91, DJCME92]. Recent work argues convincingly that LAN traffic is much better modeled using statistically *self-similar* processes [LTWW93], 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 arrival processes appear much better modeled using self-similar processes.

For our study we analyze fifteen 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 using 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

Draft

within a single FTP session are clustered into bursts. Both FTPDATA connection and FTPDATA burst arrivals are not 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 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	Drops
Bellcore (BC)	10Oct89	13 days	TCP S/F*	0
DEC (DEC-1)	26Nov91	24 hours	TCP S/F	?
LBL 1-7	See refs.	30 days	TCP S/F	$\leq 15 \cdot 10^{-6}$
UCB (UCB)	31Oct89	24 hours	TCP S/F*	0
LBL PKT-1	17Dec93	2 hours	All TCP	$5 \cdot 10^{-4}$
LBL PKT-2	19Jan94	2 hours	All TCP	$9 \cdot 10^{-4}$
LBL PKT-3	20Jan94	2 hours	All TCP	$2 \cdot 10^{-4}$
LBL PKT-4	21Jan94	1 hour	All pkts.	$7 \cdot 10^{-4}$
LBL PKT-5	28Jan94	1 hour	All pkts.	$5 \cdot 10^{-4}$

Table 1: Summary of Wide-Area Traces

Table 1 summarizes the traces of wide-area traffic used in our study. The first rows represent traces previously analyzed: the BC and UCB traces in [DJCME92][1], the DEC-1 trace in [P93], and the LBL traces in [P94, P93]. The "LBL 1-7" row represents 7 wide-area TCP SYN/FIN traces, each spanning 30 days. The "drops" column gives the fraction of packets dropped by the trace program during each of the traces, always quite low. 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.

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. For NNTP and SMTP, a Poisson model of connection arrivals is questionable. We also show that for individual FTPDATA connections, the arrival process is definitely not Poisson, but, as is discussed later in Section 6, there is noteworthy structure in the arrival process of individual FTPDATA connections within an FTP session.

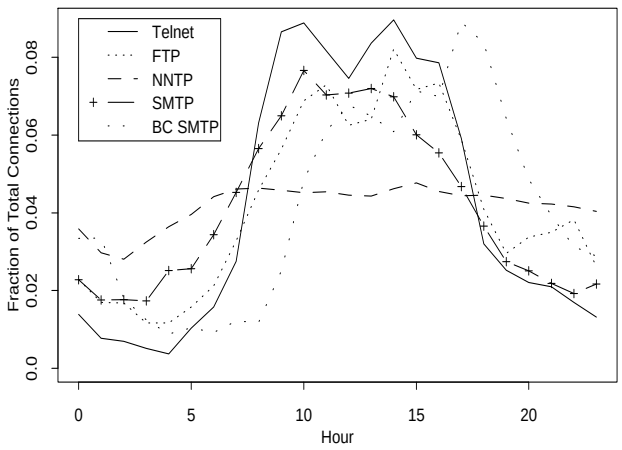

Figure 1: Mean, relative, hourly connection arrival rate for LBL-1 through LBL-4 datasets.

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 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-

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

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

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 period are Poisson with a fixed rate. If the arrivals during the period are truly Poisson, then we would expect 95% of the tested periods to pass the test. Note that we would expect testing ten-minute periods to perhaps be more successful than testing one-hour periods, because using ten-minute periods allows the arrival rate to change six times each hour rather than remaining constant throughout the hour.

We applied our methodology to the LBL-5 and LBL-6 datasets for TELNET, FTP, FTPDATA, SMTP, and NNTP connections. 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 [3]. We also tested arrivals of FTPDATA *bursts* (see Section 6 below) for the LBL-6 and LBL-7 datasets.

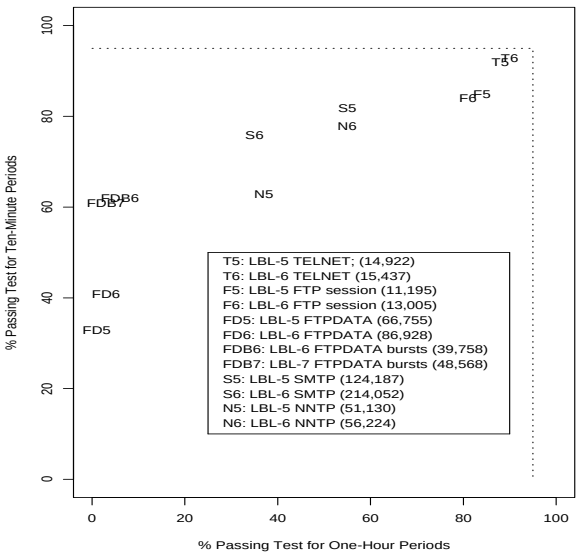

Figure 2: Results of testing for homogeneous Poisson arrivals at the 5% significance level.

Figure 2 shows the results of our tests. Along the x-axis we plot the percentage of tested one-hour intervals that passed the statistical test, and along the y-axis the same for ten-minute intervals. The dashed lines correspond to a 95% pass-rate, indicating arrivals indistinguishable from Poisson. The legend identifies the protocol and dataset corresponding to each of the points, along with the total number of such connections (or bursts) during the dataset.

We see immediately that TELNET connection arrivals are almost exactly Poisson, both for 1-hour and 10-minute fixed rates. FTP session arrivals are also well-modeled as Poisson.

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

On the other hand, from these tests NNTP and SMTP arrivals appear not well-modeled as Poisson. This finding is not conclusive, however, because (as explained in Appendix A) the higher number of NNTP and (especially) SMTP connections may simply be allowing the statistical test more opportunity to detect relatively insignificant deviations from Poisson arrivals.

That NNTP and SMTP arrivals appear not well-modeled as 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 also are 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 [Stevens94].

FTPDATA connection arrivals are clearly not Poisson over one-hour periods and only poorly modeled as such over 10-minute periods. This finding 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-connection FTPDATA bursts (see Section 6) into single arrivals improves the 10-minute Poisson fit somewhat, but the arrivals still fail to fit as well as even NNTP.

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 all interarrivals lumped together, rather than postulating separate hourly rates; for example, the distribution of all LBL-5 TELNET interarrivals looks close to log-normal.

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 arrivals correspond to users deciding to *begin* using the network; X11 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 does not have the memoryless property (and hence is not exponential), since a single user is involved in generating new arrivals. If we could discern between X11 session arrivals and X11 connection arrivals, then we conjecture we would find the session arrivals to be Poisson.

Draft

Draft

4 TELNET packet interarrivals
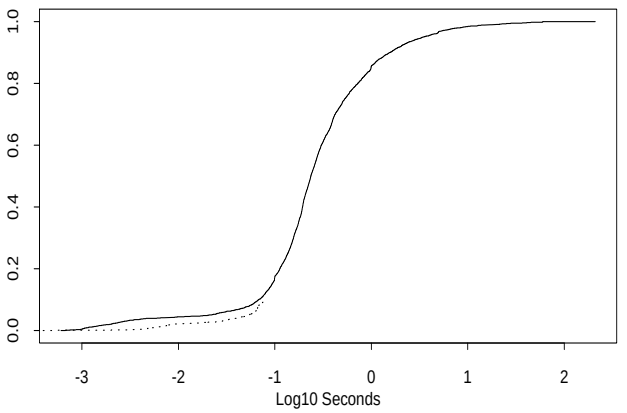
The previous section shows that start times for TELNET connections are well-modeled by a Poisson process. In this section we look at the packet arrival process within a TELNET connection. We restrict our study to packets generated by the TELNET connection originator; this in general is a user typing at a keyboard. 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.