Lawrence Berkeley National Laboratory

Recent Work

Title

Empirically-Derived Analytic Models of Wide-Area TCP Connections

Permalink

<https://escholarship.org/uc/item/1pz6j49f>

Journal

IEEE-ACM Transactions on Networking, 2(4)

Author

Paxson, V.

Publication Date

1993-05-01

Lawrence Berkeley Laboratory

UNIVERSITY OF CALIFORNIA

Engineering Division

Submitted for Publication

Empirically-Derived Analytic Models of Wide-Area TCP Connections

V. Paxson

May 1993

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

REFERENCE COPY
Does Not
Circulate
Bldg. 50 Library.

LBL-34086

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-34086
UC-000

**EMPIRICALLY-DERIVED ANALYTIC MODELS OF
WIDE-AREA TCP CONNECTIONS**

**Vern Paxson
vern@ee.lbl.gov**

**Electronic Engineering
Computer Science Division
University of California**

**Engineering Division
Lawrence Berkeley Laboratory
University of California
Berkeley, California 94720**

May 1993

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-AC03-76SF0098.

Empirically-Derived Analytic Models of Wide-Area TCP Connections: Extended Report

Vern Paxson
Lawrence Berkeley Laboratory and
EECS Division, University of California, Berkeley
1 Cyclotron Road
Berkeley, CA 94720
vern@ee.lbl.gov
LBL-34086

May 25, 1993

Abstract

We analyze 2.5 million TCP connections that occurred during 14 wide-area traffic traces. The traces were gathered at five “stub” networks and two internetwork gateways, providing a diverse look at wide-area traffic. We derive analytic models describing the random variables associated with *telnet*, *nntp*, *smtp*, and *ftp* connections, and present a methodology for comparing the effectiveness of the analytic models with empirical models such as *tcplib* [DJ91]. Overall we find that the analytic models provide good descriptions, generally modeling the various distributions as well as empirical models and in some cases better.

1 Introduction

Though wide-area networks have been in use since the early 1970’s, until recently we have known virtually nothing about the characteristics of the individual connections of different protocols. In the last few years a number of papers have appeared giving statistical summaries of traffic on a per-protocol basis [Caceres89, Heimlich90, CW91, EHS92, WLC92], an important first step. The next step in understanding wide-area traffic is to form models for simulating and predicting traffic.

One such model, *tcplib* [DJ91, DJCME92], is now available. *tcplib* is an *empirical* model of wide-area traffic: it models the distribution of the random variables (e.g., bytes transferred, duration) associated with different protocols by using the distributions actually measured for those protocols at an Internet site.

Ideally we would like to have *analytic* traffic models: simple mathematical descriptions rather than empirical distributions. Such models are easier both to convey and to analyze. The key question is whether analytic models can describe the

diverse phenomena found in wide-area traffic as well as empirical models. Our previous work [Paxson91] offered such models but suffered in part from flawed statistical methodology.

In this paper we analyze 14 wide-area traffic traces gathered at seven different sites, five “stub” networks and two internetwork gateways. We derive analytic models describing the random variables associated with *telnet*, *nntp*, *smtp*, and *ftp* connections, and present a methodology for comparing the effectiveness of the analytic models with *tcplib* and with another empirical model constructed from one of the datasets. Table 1 summarizes our main results. Overall we find that the analytic models provide good descriptions, generally modeling the various distributions as well as the empirical models and in some cases better. We develop each of these findings in the remainder of the paper.

In the next section we give an overview of the 14 traffic traces. We describe the gross characteristics of the traces including their traffic mix, and discuss how we filtered the traffic to remove anomalous connections.

The following section presents our statistical methodology. We discuss how we transformed the data and dealt with outliers; our unsuccessful attempts to find “statistically valid” models; the metric we devised for comparing the fit of two different models to a dataset; and our methodology for modeling connection interarrivals, which is more complex than modeling the other random variables associated with a connection.

We then present one section each on modeling *telnet*, *nntp*, *smtp*, and *ftp*. These sections can be read independently if the reader is more interested in one protocol than another, except that the first section describes how to read the plots used in all four sections to compare models. By reading the text accompanying Figures 4, 5, 13, and 30, the reader can if desired skip the remainder of the *telnet* section.

Random variables associated with wide-area network connections can be described as well by analytic models as by empirical models.
When using either type of model, caution must be exercised due to frequent discrepancies in the upper 1% tails.
While in general the analytic models do not match the observed distributions identically in a statistical sense, often a random subsample of hundreds of data points does result in a statistically valid fit, indicating that the analytic models are often <i>close</i> though not <i>exact</i> .
Bulk-transfer traffic (<i>ftpdata</i> , <i>smtp</i> , <i>nntp</i> , and <i>telnet</i> response) is best modeled using log-normal distributions.
Bulk-transfer traffic is not strongly bidirectional; the responses to bulk transfers show little variation relative to the variation in the size of the transfer.
Network traffic varies significantly, both over time and more so from site-to-site, not only in traffic mix but in connection characteristics.
Scaling usually helps significantly in modeling the bytes transferred by <i>nntp</i> , <i>smtp</i> , <i>rlogin</i> , and individual <i>ftpdata</i> connections, but is usually not necessary for adequate fits to <i>telnet</i> connections and full <i>ftp</i> conversations.
Except for <i>nntp</i> , connection interarrivals are well modeled using nonhomogeneous Poisson processes with fixed hourly rates.

Table 1: Major Findings

In the last section we summarize the different analytic models and discuss findings in addition to those listed in Table 1. We also include appendices summarizing how we filtered the data prior to analysis, and exploring the effectiveness of modeling *rlogin* traffic using the *telnet* models.

2 Overview of Network Traffic Traces

To develop and then evaluate our models we acquired a number of traces of wide-area traffic. Our main data were from six month-long traces of all wide-area TCP connections between the Lawrence Berkeley Laboratory (LBL) and the rest of the world. With the help of colleagues we also were able to study traces from Bellcore, the University of California at Berkeley, the University of Southern California, Digital's Western Research Laboratory, the United Kingdom-United States academic network link, and traffic between the coNCert¹ network and the rest of the world. We discuss the general characteristics of each of these datasets in turn and then provide summaries of their TCP traffic.

2.1 The LBL Traces

All off-site communication at LBL funnels through a group of gateways that reside on a network separate from the rest of the Laboratory. The first two datasets were taken using a Sun 3/50 residing on the gateway network, using the *tcpdump*

packet capture tool [JLM89] running the Berkeley Packet Filter [MJ93]. The Sun 3/50 had kernel modifications to gain a clock resolution of 10 msec. These are the traces discussed in [Paxson91]. When we took the last four traces the monitor workstation had been upgraded to a Sun SLC with a consequent improvement of clock resolution to 1 microsecond.

We used a *tcpdump* filter to capture only those TCP packets with SYN, FIN, or RST flags in their headers, greatly reducing the volume and rate of data (but at the cost of no analysis of intra-connection dynamics). From SYN and FIN packets one can derive the connection's TCP protocol, connection duration, number of bytes transferred in each direction (excluding TCP/IP overhead), participating hosts, and starting time. In principle we could derive the same information using RST packets instead of FIN packets, but we found that often the sequence numbers associated with RST packets were erroneous. Since we could not derive reliable byte counts from RST-terminated connections we excluded them from subsequent analysis.

With this packet capture scheme there are two mechanisms by which packets can be lost. The first is that, if a packet arrives at the Ethernet controller and the controller has run out of kernel memory to buffer the packet, it drops the packet and sets a bit indicating that this event occurred. The Ethernet driver subsequently reads the bit and increments a corresponding counter. It is possible that more than one packet will be dropped before the driver is able to read the bit, so the actual number of dropped packets is unknown but at least as large as the driver's counter.

The second packet-drop mechanism occurs when the kernel determines that the packet filter accepts a packet, but has no

¹Communications for North Carolina Education, Research and Technology.

more buffer space for saving the packet (due to the user-level program failing to consume previously accepted packets). In this case the kernel drops the packet and increments a counter. Values reported by this counter thus correspond to exactly the number of acceptable packets (in our case, SYN/FIN/RST packets) dropped.

Dataset	Packets (days)	Start	End	Drops
LBL-1	124M (36)	01Nov90	01Dec90	0 + 0
LBL-2	?	28Feb91	30Mar91	0+?
LBL-3	207M (47)	07Nov91	07Dec91	9 + 24
LBL-4	210M (36)	19Mar92	18Apr92	6 + 233
LBL-5	337M (35)	24Sep92	23Oct92	8 + 1808
LBL-6	447M (31)	24Feb93	26Mar93	3 + 0

Table 2: Summary of LBL Datasets

Table 2 summarizes the LBL datasets. The second column gives the total number of network packets received by the kernel for each dataset, along with the number of days spanned by the entire trace. (The statistics missing for the LBL-2 dataset are due to abnormal termination of the tracing program.) Each dataset was then trimmed to span exactly 30 days, beginning at midnight on a Thursday and ending at midnight on a Saturday (i.e., just after 11:59PM Friday night), except for LBL-6, which begins on a Wednesday and ends on a Friday midnight. The “Drops” column gives the drop count reported by the Ethernet driver followed by the drop count reported by *tcpdump*; this last value represents dropped SYN/FIN/RST packets.²

Finally, since the LBL datasets span 2.5 years at roughly regular intervals, they provide an opportunity to study how a site’s wide-area traffic evolves over time. Such a study is reported in [Paxson93].

2.2 The Additional Traces

As mentioned above, a number of colleagues generously provided access to traffic traces from other sites. The authors of [DJCME92] provided their traces of traffic from Bellcore, U.C. Berkeley, and U.S.C.; Jeffrey Mogul provided traces from DEC-WRL; Wayne Sung provided traces of traffic to/from the coNCert network in North Carolina; and the authors of [WLC92] provided their traces of the UK-US academic network. The first four traces all originate from “stub” sites, while the latter two represent inter-network traffic (though the authors of [WLC92] characterize the UK side

²In the LBL-4 dataset we observed the heated exchange of nearly 400,000 RST packets sent between a lone remote host and three LBL hosts, separated by a geometric mean of 1.3 msec. In LBL-5 we observed the exchange of nearly 120,000 RST packets between a single pair of hosts, virtually all occurring during a 98 second period, separated by a geometric mean of 400 μ sec. LBL-6 did not include any RST bursts.

The LBL-5 RST bursts correspond to enough traffic to consume > 500 kbit/sec. We suspect that the RST bursts are the cause of the relatively large number of dropped SYN/FIN/RST packets in LBL-4 and LBL-5.

of the UK-US traffic as similar to a large stub site since it comprises only a few hosts).

Site	Starting Time	Duration	Drops
Bellcore (BC)	Tue 14:37 10Oct89	13 days	0
UCB (UCB)	Tue 10:30 31Oct89	24 hours	0
USC (USC)	Tue 14:24 22Jan91	26 hours	$\approx 0.6\%$ *
DEC (DEC-1)	Tue 16:46 26Nov91	24 hours	?
DEC (DEC-2)	Wed 17:55 27Nov91	24 hours	?
DEC (DEC-3)	Mon 15:02 02Dec91	24 hours	?
coNCert (NC)	Wed 09:04 04Dec91	24 hours	?
UK-US (UK)	Wed 05:00 21Aug91	17 hours	0

Table 3: Summary of Additional Datasets

The additional datasets are summarized in Table 3. Next to the site name we give in parentheses the abbreviation we will use to identify the dataset. The drop rates for the first three datasets correspond to those listed in [DJCME92]; for the last dataset, to that listed in [WLC92]; and the drop rates for the remaining datasets were unavailable. The USC dataset’s drop rate is marked because we found our copy of the trace plagued throughout by “blackouts” of missing packets, occurring almost exactly a minute apart and each blackout lasting roughly ten seconds.³ Because of these blackouts, we exclude the USC dataset from our interarrival models.

2.3 Filtering of non-WAN traffic

Before proceeding with our analysis we filtered out non-wide-area traffic from the datasets: internal and transit traffic. The details are given in Appendix A. In addition, we removed from the LBL datasets all traffic between LBL and U.C. Berkeley⁴. While traffic with the University forms a significant fraction of LBL’s off-site traffic (20-40% of all connections), it is atypical wide-area traffic due to the close administrative ties and the short, high-speed link between the institutions.

2.4 Traffic Overview

We now turn to characterizing the different datasets in order to gauge their large-scale similarities and differences. Of previous traffic studies, only [FJ70], the related [JS69], and [DJCME92] compare traffic from more than one institution. The first two papers found significant differences between their four traffic sites, which they attribute to the fact that the different sites engaged in different applications and had different hardware. The authors of [DJCME92] found that their three sites (which correspond to the USC and UCB datasets in

³These blackouts do not correspond to network outages; sequence numbers of TCP connections spanning outages show jumps.

⁴Including *nntp*, unlike [Paxson93], which keeps the *nntp* traffic.

Dataset	# Conn	<i>nnip</i>	<i>smtp</i>	<i>ftpdata</i>	<i>ftpctrl</i>	<i>telnet</i>	<i>rlogin</i>	<i>finger</i>	<i>domain</i>	<i>X11</i>	<i>shell</i>	<i>other</i>
LBL-1	146,209	40	26	16	3	4	1	4	4	0.2	0.5	0.5
LBL-2	170,718	34	30	16	3	4	1	5	4	0.2	0.2	0.7
LBL-3	229,835	20	33	17	3	4	1	4	11	0.4	0.3	5
LBL-4	449,357	16	21	15	3	2	1	32	5	0.4	0.2	4
LBL-4*	312,429	23	30	21	4	3	1	3	8	0.5	0.3	5
LBL-5	370,397	14	34	22	5	4	1	6	8	0.9	0.2	5
LBL-6	528,784	11	40	23	6	3	0.8	5	5	0.7	0.4	4
BC	17,225	2	49	30	4	4	2	5	0.1	0.1	0.5	2
UCB	37,624	18	45	18	2	2	0.9	12	0.1	0.02	0.2	0.8
USC	13,097	35	27	14	2	3	1	11	2	0.09	0.3	3
DEC-1	72,821	33	35	11	1	0.08	0.05	0.1	20	0	0.001	0.8
DEC-2	49,050	38	22	8	1	0.04	0.06	0.2	29	0	0.02	1
DEC-3	73,440	26	43	9	1	0.07	0.07	0.2	19	0	0.003	1
NC	62,819	1	42	30	4	5	0.3	5	0.8	0.03	0.3	5
UK	25,669	0.02	42	39	7	4	0.4	0.9	1	0.02	0.02	4

Table 4: Percentage Connection Mixes for All Datasets

Dataset	MB	<i>nnip</i>	<i>smtp</i>	<i>ftpdata</i>	<i>ftpctrl</i>	<i>telnet</i>	<i>rlogin</i>	<i>finger</i>	<i>domain</i>	<i>X11</i>	<i>shell</i>	<i>other</i>
LBL-1	2,852	19	5	65	0.2	6	0.8	0.1	1	3	1	0.1
LBL-2	3,785	14	6	67	0.2	5	1	0.1	0.9	1	3	2
LBL-3	6,710	7	4	67	0.1	4	1	0.1	0.7	3	11	1
LBL-4	11,398	21	4	52	0.1	4	0.9	0.0	0.6	6	10	1
LBL-5	19,269	17	3	57	0.1	3	0.7	0.1	0.4	11	8	1
LBL-6	22,076	22	5	57	0.2	2	0.7	0.1	0.5	8	3	0.8
BC	346	4	8	78	0.3	4	2	0.2	0.1	0.1	2	2
UCB	318	23	16	50	0.3	4	3	0.9	0.0	0.2	0.6	1
USC	362	62	3	18	0.1	2	0.9	0.3	0.3	5	7	2
DEC-1	981	43	17	38	0.2	0.1	0.2	0.0	0.7	0.0	0.0	1
DEC-2	819	54	14	30	0.1	0.0	0.2	0.1	0.6	0.0	0.0	2
DEC-3	1,379	52	16	30	0.1	0.1	0.2	0.1	0.6	0.0	0.0	1
NC	1,553	9	8	68	0.3	5	0.3	0.1	0.3	0.1	0.3	8
UK	625	0.5	11	80	0.4	3	0.5	0.0	0.3	0.1	0.5	4

Table 5: Percentage Byte Mixes for All Datasets

this paper, as well as part of the BC dataset) had quite different mixes of traffic, but that the characteristics of any particular protocol's traffic were very similar (though they did not quantify the degree of similarity).

Table 4 shows the "connection mix" for each of the datasets. The second column gives the total number of connections recorded, and the remaining columns the percentage of the total due to particular TCP protocols. The mixes for BC, UCB, and USC differ from those given in [DJCME92] because the latter reports *conversation* mixes, where multiple related connections have been combined into single conversations. (The authors also used twenty-minute silences to delimit the end of connections, instead of FIN packets.)

From the Table it is immediately clear that traffic mixes for all protocols vary substantially, both from site-to-site and

over time (for LBL). There are also a number of anomalies which merit comment:

- The huge spike in the LBL-4 *finger* connections, the large jump in *other* connections at LBL-3, and the increasing proportion of *ftpctrl* traffic (i.e., the control side of an *ftp* conversation), are all due to the use of background scripts to automate periodic network access. Reference [Paxson93] explores this phenomenon further. LBL-4* shows the LBL-4 connection mix with the periodic *finger* connections removed, as they significantly skew the mix profile.
- The large variance of LBL's *nnip* mix is due to changes in LBL's *nnip* peer servers and differences in the rate at which new news arrives. Again, see [Paxson93] for a

discussion.

- DEC has a “firewall” in place which prohibits traffic other than *nntp*, *smtp*, and *ftp*, and *domain*. The little remaining traffic due to other protocols originated on the outside of the firewall.
- The DEC-2 dataset includes part of the Thanksgiving holiday, accounting for the depressed number of connections.
- As mentioned in [WLC92], the United Kingdom receives its network news from Holland, hence the very low proportion of *nntp* connections.

Table 5 shows the total number of data megabytes transferred (in either direction) for each of the datasets, along with the “byte mix”—the percentage of the total bytes due to each protocol. The LBL datasets show striking growth over time, which we explore further in [Paxson93]. The LBL datasets naturally total more bytes than the others because they span 30-day periods, as opposed to about 1 day for all the other datasets except BC (see Table 3).

We see immediately that, much as with the connection mix, the byte mix also varies considerably both from site-to-site and over time. Some sites (the first three LBL datasets, BC, NC, and UK) are wholly dominated by *ftp* traffic, while others (the last three LBL datasets, UCB, and the DEC datasets) show more of a balance between *nntp* and *ftp* traffic; and USC is dominated by *nntp* traffic. For some sites (UCB, DEC), *smtp* traffic contributes a significant volume, and for others (LBL, USC), traffic due to *X11* and *shell* far outweighs the almost negligible proportion of connections due to those protocols (see Table 4).

We now turn to the development of the statistical methodology that we will use to characterize the individual connections that make up the data shown in Tables 4 and 5.

3 Statistical Methodology

As noted in [Pawlita89], one weakness of many traffic studies to date has been in their use of statistics. Often the studies report only first or perhaps second moments, and distributions are summarized by eye. Frequently they omit discussion of dealing with outliers, and rarely do they report goodness-of-fit methodologies and results. The few cases where goodness-of-fit issues have been discussed are somewhat unsatisfying (the authors of [FJ70] developed their own, apparently never-published goodness-of-fit measure; and in our own previous work [Paxson91] we used the Kolmogorov-Smirnov goodness-of-fit test as a goodness-of-fit *metric*, an inferior choice). We endeavor in this work to address these statistical shortcomings and to present a general statistical methodology that might serve future work as well.

Our initial goal was to develop “statistically valid” analytic models of the characteristics of wide-area network use. By statistically valid we mean models whose distributions for random variables could not be distinguished in a statistical sense from the actual observed distributions of the variables. In this attempt we failed. Most of the models we present do not reflect the underlying data in a statistically valid sense; that is, we cannot say that our analytic distributions do indeed precisely give the distributions of the random variables they purport to model. We discuss our failure in Section 3.8 below, and then in Section 3.9 develop a “metric” for determining which of two statistically invalid models better fits a given dataset. But first we discuss the value of statistically valid analytic models and our methodology for developing them, as these issues remain fundamental to putting our results in perspective.

3.1 Analytic vs. Empirical Models

For our purposes we define an analytic model of a random variable as a mathematical description of that variable’s distribution. Ideally the model has few bound parameters (making it easy to understand) and no free parameters (making it predictive), in which case it fully predicts the distribution of similar random variables derived from datasets other than the ones used to develop the model. But typically the model might include free offset and scale parameters, in which case it predicts the general shape of future distributions but not the exact form. If those parameters are known for a future dataset, then the model becomes fully predictive for that dataset.

In contrast, an empirical model such as *tcplib* describes a random variable’s distribution based on the observed distribution of an earlier sample of the variable. The empirical model includes a great number of bound parameters, one per bin used to characterize the variable’s distribution function; it may be predictive but not easy to understand.

There are a number of advantages of an analytic model compared to an empirical model for the same random variable:

- analytic models are often mathematically tractable, lending themselves to greater understanding;
- analytic models are very concise and thus easily communicated;
- with an analytic model, different datasets can be easily compared by comparing their fitted values for the model’s free parameters.

A key question, though, is whether an analytic model fully captures the essence of the quantity measured by a random variable. An empirical model perfectly models the dataset from which it was derived; the same cannot be said of an analytic model. If the analytic model strays too far from reality, then, while the above advantages remain true, the model no

longer applies to the underlying phenomena of primary interest, and becomes useless (or misleading, if one does not recognize that the model is inaccurate).

The key question then is how to tell that an analytic model accurately reflects reality as represented by a dataset of samples. One approach is to require that the random variable distributions predicted by the model and those actually observed be indiscernable in a statistical sense. To test for such agreement we turn to goodness-of-fit techniques.

3.2 Goodness-of-fit Tests

The random variables we model (amount of data transferred, connection duration, interarrival times, and ratios of these quantities) all come from distributions with essentially unbounded maxima. Furthermore, these distributions are either continuous or, in the case of data transferred, continuous in the non-negative integers. As such the values of the variables do not naturally fall into a finite number of categories, which makes using the well-known chi-squared test less than ideal because it requires somewhat arbitrary choices regarding binning [Knuth81, DS86].

The goodness-of-fit test commonly used with continuous data is the Kolmogorov-Smirnov test. The authors of [DS86], however, recommend the Anderson-Darling (A^2) test [AD54] instead. They state that A^2 is often much more powerful than either Kolmogorov-Smirnov or chi-squared, and that A^2 is particularly good for detecting deviations in the tails of a distribution, often the most important to detect. We followed their recommendation and, in attempting to develop statistically valid models, always used A^2 in assessing goodness-of-fit.

3.3 Logarithmic Transformations

When analyzing data drawn from distributions unbounded in one direction and bounded in the other, often it helps to re-express the data by applying a logarithmic transformation [MT77]. We found that for many of our models logarithmic transformations were required to discern patterns in the large range of values in the data.

For convenience we developed and tested our models using a $\log_2 x$ transformation. Note that, when converting from logarithmic models back to untransformed models, arithmetic means of transformed values become *geometric* means of the untransformed values, and standard deviations become *factors* instead of additive values. For example, a log-normal model with $\bar{x} = 4.0$ and $\sigma = 2.5$ specifies that any observation within a factor of 5.66 ($2^{2.5}$) of 16 ($2^{4.0}$) lies within one standard deviation of the geometric mean. Thus, 2.83 ($= 16/5.66$) and 90.56 ($= 16 \times 5.66$) are the boundaries of values lying within one standard deviation of the geometric mean, which is 16.

3.4 Dealing with Outliers

When applying a logarithmic transformation to non-negative data, one immediately runs into the problem of what to do with data equal to zero. Fortunately for us, in our data such values are rare (and confined to values representing number of data bytes transferred), so we decided to eliminate any connections in which the number of bytes transferred in either direction was zero. We report in Appendix B the number of connections thus eliminated for each dataset; in the worst case they comprised 0.5% of the total connections. An alternative approach would have been to bias our logarithms, by using $\log_2(x+1)$ rather than $\log_2 x$; we rejected this approach as being error-prone when converting to and from the logarithmic models.

Some of our datasets also exhibited values so anomalously large that we removed their associated connections from our study. These outliers were much rarer than those discussed above. Often the values were clearly due to protocol errors (for example, connections in which the sequence numbers indicated $2^{32} - 1$ bytes transferred). We discuss these outliers also in Appendix B.

Finally, we restricted our analysis to datasets with at least 100 connections of interest, to prevent small, anomalous datasets from skewing our results.

3.5 Censored Data

Some of our models describe only a portion of the distribution of a random variable (such as the upper 80% of the distribution). Reference [DS86] discuss modified goodness-of-fit tests (including A^2) to use with such *censored* distributions, in which a known fraction of either tail has been removed from the measurements prior to applying the test. In addition, they describe a method (due to Gupta [Gupta52]) for estimating the mean and variance of such a censored distribution, which can be used to derive estimated parameters of a model from censored data.

3.6 Deriving Model Parameters from Datasets

Often a model has free parameters that must be estimated from a given dataset before testing the model for validity in describing that dataset. For example, a log-normal model may require that the geometric mean and standard deviation be estimated from the dataset. The authors of [DS86] make the important point that estimating free parameters from datasets alters the significance levels corresponding to statistics such as A^2 computed from the fitted model. They then provide both methods to estimate free parameters from datasets, and the required modifications for interpreting the significance of the resulting A^2 (and other) statistics. We followed their approach.

3.7 Model Development vs. Testing

To know if a model is truly predictive, we must test it on data other than that used to develop the model. To this end, we developed all of our models using the first half of the LBL-1 through LBL-4 datasets. We refer to these below as the “test datasets”. We then tested the models against the second half of these LBL datasets along with the entirety of the remaining datasets (including all of LBL-5 and LBL-6).

Below we compare our analytic models with two empirical models: one derived from the UCB dataset, which is essentially the same as the *tcplib* model, and one derived from all of LBL-2. Thus, in keeping with our goal of testing models only on data other than that used to develop them, we do not report results for fits to these datasets. An exception is for our interarrival models, which in general we do not compare to the empirical models (see Section 3.11 below).

3.8 Failure to Find Statistically Valid Models

Using the methodology described above, we attempted to develop models for a number of random variables for TCP connections of various protocols. While we often could find fairly simple analytic models that appeared to the eye to closely match the distributions of the random variables for a given dataset, these models rarely proved valid at a significance level of 5%, or even 1%, when tested against other datasets.⁵

What we found tantalizing, though, is that often, when we subsampled the dataset, we *did* find valid fits to the smaller sample. This pattern held whether the subsamples were constructed randomly or chronologically (for example, testing each day in the LBL datasets separately). We tested whether the pattern was due to daily variations in the model’s parameters by using autocorrelation plots. We found such patterns only in the arrival process and bytes transferred of *nntp*, and bytes transferred by *smtp* connections. We discuss these findings below in Sections 5.2 and 6.2. We did not find any consistent patterns in the LBL *telnet* or *ftp* test datasets, ruling out simple hourly, daily, or weekly patterns in the parameters.

These findings are consistent with our models being *close* to describing the distributions but not statistically *exact*. In such a case it will take a large number of sample points for a goodness-of-fit test to discern a difference between the distributions. When we subsample we present the test with fewer points and the fit is then more likely to be found valid.

Figure 1 illustrates the problem. Here we see the distribution of $\log_2$ of the bytes sent by the *telnet* responder (i.e., not the host that began the connection) for the first half of the LBL-4 dataset. Fitted against the distribution is our responder-bytes model, which uses a normal distribution for

Figure 1: Censored Log-Normal Fit to Upper 80% of LBL-4 TELNET Responder Bytes

the upper 80% of the data (and ignores the lower 20%). The horizontal line indicates the 20th percentile; the goodness-of-fit test applied only to the agreement above this line. While judging visually we might be tempted to call the fit “good”, it fails the A^2 test even at the 1% level.

This sample consisted of 5,448 points. We then subsampled 1,000 points randomly, tested the validity of the model’s fit to the subsample, and repeated the process 100 times. Of these 100 tests, 79 were valid at the 1% level and 55 at the 5% level. Thus we feel confident that the model is *close*, though we know it is not exact.

3.9 Comparing Analytic and Empirical Models

While we must abandon our initial goal of producing statistically valid, “exact” models, we still can produce useful analytic models by building on the work of [DJ91, DJCME92] in the following way. In those papers the authors argue that their empirical models are valuable because the variation in traffic characteristics from site-to-site and over time is fairly small. Therefore the *tcplib* models, which were taken from the UCB dataset, faithfully reproduce the characteristics of wide-area TCP connections. If we can develop analytic models that fit other datasets as well as *tcplib* does, then the analytic models are just as good at reproducing the characteristics of wide-area TCP connections; a network researcher is just as well off using either set of models, and may prefer the analytic descriptions for the advantages discussed in Section 3.1.

The question then remains how to compare an analytic

⁵A significance level of 5% indicates a 5% probability that the A^2 test erroneously declares the analytic model to *not* fit the dataset. A 5% test is more stringent than a 1% test; it errs more often because it demands a closer correspondence between the model and the dataset before declaring a “good fit.” See [Ross87, pp. 205-206] for further discussion of significance levels.

model with an empirical one. Rather than a goodness-of-fit test, we need some sort of goodness-of-fit metric. While under certain conditions one can apply tests like A^2 as metrics [DS86], they are not appropriate metrics for measuring the fit of an empirical model; the tests are designed for comparing a continuous distribution (an analytic one) with an empirical distribution.

We chose as our metric a measure of “bin” frequencies, similar to a chi-squared test. A chi-squared test computes:

$$X^2 = \sum_{i=1}^M \frac{(N_i - np_i)^2}{np_i}$$

where M is the number of bins, p_i the fraction of all observations predicted to fall into the i th bin, n the total number of observations, and N_i the number of observations actually falling into the i th bin. We make one important change, though. If a chi-squared test is used to compare non-identical distributions, then the resulting X^2 increases with n , making it difficult to compare X^2 values when testing a distribution against different-sized datasets to see which it more closely matches. If two distributions are different, then for large values of n , $N_i/(np_i)$ will approach some fixed factor ρ_i , and the squared term in the X^2 computation approaches $(\rho_i - 1)(np_i)$. We then see that the metric:

$$K^2 = \sum_{i=1}^M \frac{(N_i - np_i)^2}{(np_i)^2}$$

remains invariant with increasing n . If the bins have equal width, then we have:

$$K^2 = \sum_{i=1}^M (\rho_i - 1)^2$$

which allows us to compute μ , the “average deviation” in each bin:

$$\mu = \sqrt{K^2/M} \quad (1)$$

We interpret μ as follows: the value of K^2 we observed is consistent with what we would observe if in *each* bin the proportion of observations deviates from the predicted proportion by μ , i.e., $|N_i - np_i|/(np_i) = \mu$. While in general the deviation will vary from bin to bin, we can use μ to summarize the “average” deviation.

We are faced with several problems when using this metric:

- Similar to the problems using chi-squared tests mentioned in Section 3.2 above, we are forced to make a somewhat arbitrary choice as to how many bins to use. We chose to use ten equal-sized bins, so as to measure the deviation from the predicted distribution within each 10th percentile.⁶

⁶In one case below we use nine bins, to accommodate censored data.

- The metric does not inform us of deviations in the distribution tails, often the most important type of deviation. We address this shortcoming in the next section.

- The metric does not inform of us interesting, localized spikes or clumps. Within a single bin we may miss considerable departure from a model; the danger is particularly acute when testing analytic models, since their continuous nature does not usually allow for clumping. Empirical models, on the other hand, may exactly predict the clumping.

We do not believe this problem to be major because in our studying of the LBL test datasets to form our models we rarely encountered consistent clumping (we make mention below of those occasions when we did). We also note that if clumping exists and is not accompanied by nearby sparseness, then the clump will “pull” more values into the bin than a model without the clump would predict, which will raise the μ value. So a major clump may be detected as an overall poor fit by the model.

- An empirical model does not always allow us to create equal-sized bins. It may be that the model has a single-valued spike straddling a bin boundary (for an exaggerated example, suppose that the lower 20% of an empirical distribution are all equal and we want to create bins 10% wide). We deal with this case by placing the entire spike in the lower bin and adjusting the bin widths accordingly. If the spike is substantial and not aberrant, then this procedure will aid the fit of the empirical model more than that of an analytic model.
- Since an empirical model has bounds on the range of values it allows for, the tested dataset may have values not corresponding to any bin. We removed such values from the dataset prior to computing its fit to the model. We did, however, include these values in the summary of deviation in the tails (see Section 3.10 below).

We use the μ metric to gauge how closely the distributions of different models match that of a particular observed distribution. We deem the model distribution with the lowest μ value as corresponding to the best-fitting model for the observed distribution. In general we tested each dataset against three model distributions: one produced by our analytic model, one produced using the empirical distributions found in the UCB dataset, and one drawn from the LBL-2 dataset. As mentioned in Section 2.4, the distributions in *tcplib* come from the UCB dataset, with some minor differences in the data reduction. Thus, how well the UCB dataset fits the other datasets should closely match the fit of *tcplib* to those datasets. If the analytic models fit the datasets as well or better, then we argue that the analytic models provide as good or better an overall model. Finally, to guard against the possibility that the UCB dataset is atypical and that better

empirical models might exist, we also constructed and tested an empirical model consisting of the entire LBL-2 dataset.

We developed and settled on this metric prior to observing the values it gave for the different models. We tested two versions of each model. In the first version all parameters were fixed; none were derived from the dataset being tested. When developing our analytic models we picked for each free parameter a round value lying somewhere in the range the parameter exhibited in the LBL test datasets. We chose round values as reminders that there is in general considerable range in the possible values of the parameters, and that our choice was therefore somewhat arbitrary (nearby choices would work just as well).

In the second version of each model we derived the model's free parameters from the dataset being tested. For empirical models we applied a linear transformation to the empirical distribution so that its mean and standard deviation matched that computed for the tested dataset. We refer to this second type of model as *scaled*.

3.10 A Metric for Deviation in the Tails

We summarize each model's fit to the extreme tails as follows. Suppose we test the model against n datasets. For the i th dataset, let x_i be the number of observations predicted to lie in the tail, and y_i be the number actually found to do so. Define:

$$\xi = \frac{1}{n} \sum_{i=1}^n \log \frac{x_i}{y_i} \quad (2)$$

ξ then gives the mean of the natural logarithm of the proportion by which the model overestimates the population of the tail. Positive values of ξ indicate that the model overestimates the tail, either consistently or in a few cases grossly. Similarly, negative values indicate the model underestimates the tail.

With this definition, an underestimate by a factor of two ($x_i/y_i = 1/2$) is just as bad as an overestimate by the same factor ($x_i/y_i = 2$), though if the two occur in different datasets they will cancel out one another. Values of ξ close to 0.0 indicate that either the model consistently does well in modeling the tail, or overestimates for some datasets and underestimates for others. In the latter case there probably is great diversity in the distribution's tail across the different datasets, and the model's estimate of the tail is a good compromise.

One problem arises when using this definition of ξ : if y_i is 0 then ξ becomes undefined. We address this problem by replacing x_i/y_i with 100 in these cases.

In comparing models we summarize how well each model does in the 10% and 1% tails. For models describing bytes transferred, we only summarize the upper tails, as in these cases disagreement in the lower tails is a matter of predicting a few bytes too many (or few) in small connections, while

disagreement in the upper tails can result in large connections that are megabytes too big or small. For other models we summarize both the upper and lower tails.

3.11 Modeling Interarrivals

The final aspect of our methodology is how we model connection interarrivals. Our hope was to successfully model interarrivals as Poisson processes, as these have many attractive properties and a natural interpretation (uncorrelated, memoryless arrivals).

We cannot hope for much success, though, if we simply model the interarrival distribution directly: we expect that the arrival process will vary over the course of each day, since computer users tend to work during daylight hours, take lunch breaks, and so on; we do not expect a *homogeneous* Poisson arrival process. Instead we first look at the relative rate of connection arrivals over the course of a day in order to develop a nonhomogeneous Poisson model.

Figure 2 shows the mean, normalized, hourly connection rate for the test datasets. For each hour we plot the fraction of the entire day's connections that occurred during that hour. We see, for example, that *telnet* connections are particularly prominent during the 8AM-6PM working 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 hums along at a fairly constant rate, only dipping somewhat in the early morning hours (but the mean size of each connection varies over the course of the day; see Section 5.2). The *smtp* traffic is interesting because it shows more of a morning bias than either *telnet* or *nntp*. To explore this bias we have also plotted the hourly rates for the BC dataset's *smtp* connections. Here we see a significant afternoon bias. As LBL lies on the west coast of the United States and Bellcore to the east, three time zones away, we can interpret this difference as being due to cross-country mail: mail sent by east-coast users arrives early in the day for west-coast users, and mail sent by west-coast users late in the day for east-coast users.

We can then use this data to attempt to model interarrivals as Poisson processes. First we compress datasets consisting of more than one day into a single "superday" by grouping together all connections beginning during each hour of the day. For example, all connections arriving between 9:00AM and 9:59AM are placed in one 9AM "superhour", regardless of during which day the connection arrived. The hope is that the daily variations are considerably less than the hourly variations, which is true in general except for weekends, during which much less traffic is generated. But because weekends have many fewer arrivals, the effect of aggregating them with weekday connections of the same hour is small.

Next we predict the number of connections occurring dur-

Figure 2: Mean Daily Variation in the Test Dataset Connection Rate

ing each hour by multiplying that hour's fraction as given in Figure 2 by the total number of connections during the superday. Call this quantity a_h , for the number of arrivals during hour h . If we have a_h arrivals from a Poisson process during a single hour, then we expect the mean interarrival time in seconds to be

$$\bar{m}_h = 3600\text{sec}/a_h \quad (3)$$

and if we divide the interarrival times by $\bar{m}_h$, then they should be exponentially distributed with a mean of 1. Now that each hour's interarrivals have been normalized to the same mean, we test the distribution of all of the superday's normalized arrivals together against that predicted by an exponential model with mean 1.

We can also test a "scaled" version of this model which does not rely on the rates given by Figure 2. Instead of computing $\bar{m}_h$ as given in Equation 3, we simply compute each super-hour's interarrival mean directly and divide by that value, guaranteeing a resulting mean of 1.

tcplib does not presently include empirical models for interarrivals, probably because creating such empirical models requires a fair amount of transformation to the raw interarrival times. We therefore do not compare the performance of the analytic interarrival model against that of empirical models, but instead compare the scaled version of the model against the unscaled. If we find that for both versions μ is quite low, then the analytic model is successful and the rates given by

Figure 2 are widely applicable. If μ is only low for the scaled model, then the arrivals are indeed from a nonhomogeneous Poisson process, but with rates different from those given in Figure 2. If μ is high for both versions, then the arrivals are not from a Poisson process with a fixed hourly rate. If μ were to be high for the scaled model but low for the unscaled model, then we would be left with a puzzle, but fortunately this never happened.

Note that we do not model the arrival of a site's inbound and outbound connections separately, though the two might well have different hourly rates; nor do we model the correlations between inbound and outbound arrivals. We leave these important refinements to future work.

4 TELNET

We now turn to analyzing the characteristics of individual protocols and developing models to describe them. We begin with *telnet*.⁷

⁷Appendix C presents a similar overview for *rlogin* traffic, along with results of modeling it with the *telnet* models developed in this section.

Dataset	# Conn	# Rej	$\bar{x}_{orig}$	σ_{orig}	max_{orig}	$\bar{x}_{resp}$	σ_{resp}	max_{resp}	$\bar{x}_{dur}$	σ_{dur}	max_{dur}
LBL-1	5,734	9	199B	$\times 4.4$	207KB	4.2KB	$\times 7.9$	1.9MB	266 s	$\times 6.8$	90.5 h
LBL-2	7,582	12	199B	$\times 4.6$	282KB	4.3KB	$\times 7.5$	3.2MB	237 s	$\times 6.8$	78.2 h
LBL-3	9,607	23	214B	$\times 4.7$	537KB	4.1KB	$\times 7.6$	5.5MB	226 s	$\times 6.9$	167.9 h
LBL-4	10,897	58	237B	$\times 4.3$	613KB	5.3KB	$\times 7.4$	86.6MB	271 s	$\times 6.8$	270.0 h
LBL-5	14,922	81	237B	$\times 3.9$	215KB	5.2KB	$\times 6.8$	19.3MB	248 s	$\times 7.1$	386.8 h
LBL-6	17,425	52	147B	$\times 7.3$	777KB	3.8KB	$\times 8.7$	14.0MB	256 s	$\times 6.9$	102.9 h
LBL-6*	15,437	52	242B	$\times 4.5$	777KB	5.7KB	$\times 7.3$	14.0MB	270 s	$\times 7.7$	102.9 h
BC	744	2	145B	$\times 4.1$	9.7KB	2.9KB	$\times 8.7$	0.6MB	193 s	$\times 6.4$	8.1 h
UCB	655	4	155B	$\times 4.7$	27KB	2.5KB	$\times 9.1$	0.7MB	166 s	$\times 6.9$	7.9 h
USC	405	0	184B	$\times 4.3$	12KB	4.1KB	$\times 7.2$	0.6MB	168 s	$\times 6.5$	5.5 h
NC	3,023	34	112B	$\times 3.9$	146KB	2.6KB	$\times 10.6$	3.4MB	106 s	$\times 7.4$	6.8 h
UK	962	35	143B	$\times 3.6$	30KB	2.5KB	$\times 9.3$	0.7MB	175 s	$\times 5.2$	7.2 h

Table 6: Summary of TELNET Connections

4.1 Overview of TELNET Connections

Table 6 summarizes some basic statistics of the datasets' *telnet* connections. The Table is read as follows.

The second column gives the number of "valid" connections recorded for the dataset and the third column the number of "rejected" connections; Appendix B details the rejected connections. As discussed in [Paxson93], the LBL-6 *telnet* traffic included 1,988 connections due to periodic traffic. LBL-6* summarizes the LBL-6 traffic with these connections removed. For the remainder of this section we use LBL-6* instead of LBL-6.

The 4th through 6th columns summarize the number of data bytes transmitted by the originator (the user end of the remote-terminal connection). The values given are the geometric mean, the geometric standard deviation, and the maximum. As noted in Section 3.3, except for interarrival times we applied logarithmic transformations to the data prior to analysis. This transformation is also important for summary statistics such as those presented in this Table, because arithmetic means and standard deviations are quickly dominated by upper-tail outliers; compare the figures given in this paper with those of our previous work [Paxson91]. The latter tend to be much larger.

The 7th through 9th columns give the same summary for the number of bytes transmitted by the responder (remote computer), and the 10th through 12th columns the same for the duration of the connections, with 's' used to indicate seconds and 'h' for hours.

We note that the geometric mean duration of *telnet* connections ranges from 2 to 4 minutes, while Jackson and Stubbs [JS69] reported average connection lengths for local logins of 17 to 34 minutes, and [Bryan67] gives a local-login median of 20 minutes and a mean of 45-50 minutes. Jackson and Stubbs infer that connection time "may be considerably reduced by providing a high-speed channel from the computer to the user", so we might suspect the difference between their measurements and the *telnet* data is due to the higher commu-

nication speeds of today's computers. More recently (1985), Marshall and Morgan found that local-area remote logins had an average duration of 45 minutes [MM85], and non-network logins had an average duration of 150 minutes. Thus the distance between the user and the computer appears inversely correlated with the login duration. Since bandwidth usually decreases with distance, we appear to be seeing Jackson and Stubbs' effect but rescaled to reflect today's range of communication speeds.

The LBL *telnet* connections were on average substantially longer and consisted of more bytes than those at other sites. We would expect slightly longer average durations for LBL connections since the datasets span several weeks, giving an opportunity to detect long-lived connections that would be missed by the short spans of the other datasets (except for BC, which spans 13 days and has the next highest average). But this effect is small: if we eliminate from LBL-1 all connections spanning more than one day (i.e., crossing midnight), then $\bar{x}_{orig}$ drops to 197B, $\bar{x}_{resp}$ drops by 53B, and $\bar{x}_{dur}$ drops to 260 s. Given the difference in these parameters even after this adjustment, we are forced to conclude that, at least with regard to mean bytes transferred and duration, the LBL *telnet* traffic is significantly different from that at other sites.

We also note a definite trend over the LBL datasets towards increasing values of $\bar{x}_{orig}$, and a similar though less convincing trend in $\bar{x}_{resp}$, too, indicating that *telnet* connections are growing larger with time. Connection durations, on the other hand, are not growing longer, suggesting that higher network bandwidths are enabling users to engage in more work during each session.

Finally, we note that the data provide support for the observation in [DJCME92] that "interactive applications can generate 10 times more data in one direction than the other," and actually suggest the factor is around 20:1. The observation that the computer end of a terminal session generates an order-of-magnitude more data than the user end can be found as far back as reference [JS69], though [Bryan67] found the ratio to

be 2.85:1 on a line-by-line basis (the author also states, however, that the studied system was substantially different from a general-purpose, on-line, time-shared system). Marshall and Morgan found ratios as high as 35:1 for teletypewriters in technical use, with half that being a representative average, and as low as 3:1 for teletypewriters used for word processing [MM85].

In Section 4.5 below we present a model for this ratio.

4.2 TELNET Originator Bytes

With the bulk transfer protocols we examine in subsequent sections, we usually are only interested in modeling the number of bytes transferred and the connection interarrival process. With interactive applications, on the other hand, we not only are interested in the bytes transferred in both directions but also the connection duration and the relationships between these variables.

We begin by modeling the number of bytes sent by the originator of a *telnet* connection (typically a human typing at a keyboard). The best fit we found to the LBL *telnet* test datasets came using the *extreme* distribution:

$$F(x) = \exp \left[-\exp \left(-\frac{(x - \alpha)}{\beta} \right) \right] \quad (4)$$

Reference [DS86] gives a procedure for estimating α and β for a given dataset. For our originator-bytes model, x in Equation 4 is $\log_2$ of the number of bytes transmitted by the connection originator.

Figure 3: TELNET Originator-Bytes Model for LBL-2: Log-Extreme Distribution

Figure 3 shows the distribution for the first half of the

LBL-2 dataset, along with the fitted model. We see apparently good agreement except in the tails, but when tested with A^2 the fit fails to be valid; the same holds for the other LBL half-datasets.

For the four test datasets, α varied from 6.55 to 6.93; we chose $\alpha = \log_2 100 \approx 6.64$. β varied from 1.74 to 1.92. For our fixed model we chose $\beta = \log_2 3.5 \approx 1.81$.

Figure 4: Empirical vs. Analytic Models for TELNET Originator Bytes

Figure 4 shows the computed values of μ for this analytic model plotted against both the UCB and LBL-2 empirical models, where μ is defined as in Equation 1. The X axis gives the value of μ corresponding to one of the empirical models, and the Y axis the value corresponding to the analytic model.

We read the plot as follows. Each point on the plot is labeled with the name of the corresponding dataset. "L1" through "L6" represent the LBL datasets and "D1" through "D3" the DEC datasets (not present in this particular plot). Labels written in lower case (e.g., "usc") reflect μ values for *unscaled* models; that is, using the raw UCB or LBL-2 data for an empirical model, and the fixed version (no fitted parameters) for the analytic model. Labels in upper case ("usc") reflect the *scaled* models.

We plot the text label at the point corresponding to comparing the UCB empirical model, on the X axis, with the analytic model, on the Y axis. We then draw a line from that point to the corresponding point comparing the LBL-2 empirical model with the analytic model. This line is always horizontal because the two comparisons share the same μ value, for the analytic model, on the Y axis.

Thus for each dataset four different points are plotted: the unscaled analytic model vs. the UCB empirical model (e.g.,

“usc”); the scaled version of the same (“usc”); the unscaled analytic model vs. the LBL-2 empirical model (the line drawn from “usc”); and the scaled version of the same (line drawn from “usc”).

For example, the lowest pair of points indicate that the UCB empirical model had $\mu \approx 0.3$; the LBL-2 model, $\mu \approx 0.1$; and the analytic model, $\mu = 0.1$. Since the line drawn from “usc” goes to the left, the LBL-2 empirical model provided a better fit to the unscaled USC dataset than did the UCB empirical model. In general, if the lines head to the left of the labels then the LBL-2 empirical model surpasses the UCB model; and vice versa if the lines go to the right.

The diagonal line indicates where $\mu_{\text{analy}} = \mu_{\text{emp}}$, i.e., where the analytic and empirical models yield the same closeness-of-fit metric. Points below and to the right of this line indicate datasets for which the analytic model fitted better than the empirical model; points above and to the left, where the empirical model fitted better. For example, from this plot we see that the unscaled analytic fit to the LBL-5 dataset was much better than that of the UCB empirical model (“15”) but about the same as that of the LBL-2 empirical model.

We see in this plot that the LBL-2 empirical model almost always does better than the UCB empirical model, and that the analytic model performs comparably. The points tend to lie either just above the diagonal, indicating a slightly better empirical fit, or a bit further away from and below the diagonal, indicating a better analytic fit. For this model, scaling sometimes results in a big improvement (NC, LBL-4), no improvement (BC), or an improved empirical model but worsened analytic model (USC, UK). Thus in this case it makes sense to scale the empirical models when predicting traffic, but not the analytic model.

Closer observation reveals that for every dataset except UK and NC (two extreme cases), the analytic model fits the dataset better than the UCB model, while the LBL-2 model fits best in every case except for a few points very close to the line of equality. Thus we can order the models: the LBL-2 empirical model is better than the analytic model, which in turn is better than the UCB model.

The overall fit of the model to the datasets does not tell the entire story, however. As is generally the case with bytes-transferred models, for *telnet* originator bytes the models’ fits to the upper tail are much more important than fits to the lower tail. Figure 5 summarizes the upper-tail fits. The plot is labeled with “a” for the unscaled analytic model, “u” for the unscaled empirical UCB model, and “l” for the unscaled empirical LBL-2 model. The upper-case versions of these letters correspond to the tails for the scaled versions of these models. The X axis gives the ξ value for the upper 10% tail, and the Y axis the value for the upper 1% tail, where ξ is computed as given in Equation 2. In this plot we see that scaling had little effect on fitting the upper tails, as all of the uppercase letters are near their lowercase counterparts.

A letter close to the origin, such as “u”, indicates excellent

Figure 5: Tail Summary for TELNET Originator Bytes

modeling in both the upper 10% tail and the upper 1% tail. That all of the models are clustered around the Y-axis indicates they all model the upper 10% tail well. But we see that both the analytic model and the LBL-2 model have high values of ξ for the upper 1% tail. As explained in Section 3.10 above, this indicates that those models *overestimate* the distribution in the 1% tail. That is, they tend to predict more values in the 1% tail than were actually present in the datasets. As the axes are scaled logarithmically, the deviations shown are quite large. Indeed, the unscaled analytic model overestimates the upper 1% tail for every single dataset, and for all except the LBL and NC datasets not a single observation actually resided in the predicted tail. Thus a value of $\xi \approx 2$ corresponds to exceptionally poor tail fitting.

While the UCB empirical model does poorly versus the other models in fitting the datasets over the entire distribution of originator bytes, it is the obvious champion when it comes to fitting the upper 1% tail. Thus predicting *telnet* originator bytes leaves us in a quandary: we must decide which is more important to us, the overall fit to the distribution, in which case LBL-2 or the analytic model is recommended, or the upper 1% tail, in which case UCB is recommended. If fitting just the upper 10% tail well is adequate, then either the analytic model or LBL-2 is recommended.

In the interest of conserving space, for the remaining models we relegate their outlier summaries to Appendix D.

4.3 TELNET Responder Bytes

We next turn to modeling the bytes transferred by the *telnet* responder. Figure 6 shows a log-normal fit to the upper 80%

Figure 6: TELNET Responder-Bytes Model for LBL-2: Log-Normal Fit to Upper 80%

of the responder bytes in the LBL-1 test dataset. This fit is excellent; it passes the A^2 test at the 25% significance level (compare with Figure 1, which shows the same fit for the LBL-4 dataset and fails A^2 even at 1% significance). We see, however, that the lower 20% (below the horizontal line, corresponding to less than 1 KB transferred) is not smoothly distributed, making it unlikely we might find a simple analytic model encompassing it. We speculate that this roughness is due to the varying sizes of log-in dialogs and message-of-the-day greetings. Fortunately the lower tail is the least important part of this distribution.

We found in the test datasets that the log-mean ($\bar{x}$) varied from 12.0 to 12.4, generally closer to 12.0, and we chose for our fixed model $\bar{x} = \log_2 4500 \approx 12.1$. σ_x varied from 2.79 to 2.89; we chose $\sigma_x = \log_2 7.2 \approx 2.85$.

For this one model we evaluated the metric μ using 9 bins, from 0.2 to 1.0, instead of 10 bins (0.1 to 1.0), because the analytic model only fits the upper 80% of the data and it did not seem worthwhile to develop a separate model for the lower 20%. Figure 7 summarizes the fits. Except for NC, the analytic model uniformly performs well, with μ always ≤ 0.2 . The LBL-2 model also fares quite well, while the UCB model is not as good except for UK and BC. Scaling these models does not always improve things (USC in particular) but in general helps. Figure 8 explains the terrible performance fitting NC: the distribution suffers from two large clusterings, one between 240 and 265 bytes, and the other between 400 and 425 bytes. The first consumes 13% of all the connections, the second 5%. A single host originated virtually all of the connections in the first cluster, but to a number of different

Figure 7: Empirical vs. Analytic Models for TELNET Responder Bytes

Figure 8: Distribution of NC TELNET Responder Bytes

Figure 9: Empirical vs. Analytic Models for TELNET Duration

hosts, and two other hosts originated almost all of the connections in the second cluster, primarily to two remote hosts. We were unable to find obvious patterns in the interarrivals (see Figure 18 below for an example of clear one-minute patterns in connection arrivals); therefore, unlike many of the spikes discussed in [Paxson93], the connections were probably not generated by background scripts. Perhaps they correspond to cracking attempts, or more benign searches. Overall they remain puzzling.

Figure 28 in Appendix D shows the performance of the models with regard to the upper tails. Each model except for unscaled UCB does well in the upper 10% tail. All of the models overestimate the upper 1% tail somewhat; the unscaled UCB model surprisingly doing the best. On the basis of these plots we would prefer the empirical models if the upper 1% tail is important to us; otherwise either the analytic model or LBL-2 is preferable. LBL-2 provides the best overall model.

4.4 TELNET Duration

We model *telnet* connection durations using a simple log-normal distribution. For the test datasets we found $\bar{x}$ ranging from 7.67 to 8.03 and chose $\bar{x} = \log_2 240 \approx 7.91$. σ_x ranged from 2.83 to 3.02; we chose $\sigma_x = \log_2 7.8 \approx 2.96$.

Figure 9 shows the fits for the duration models. In general the models are fairly good, with the metric μ falling between 0.1 and 0.3. NC again proves troublesome, though not so when scaled. No model emerges a clear winner, and, while the analytic model appears to do worst, it is not considerably

Figure 10: Empirical vs. Analytic Models for TELNET Resp./Orig. Ratio

worse than the other two.

Figure 28 in Appendix D summarizes the tail performance of the models. In general the models do well in the upper 10% tail, though the unscaled analytic and LBL-2 models overestimate somewhat. In the upper 1% tail these same models do quite poorly, while the UCB models are excellent in both tails. Because the UCB model did well in the general fitting shown in Figure 9, its good performance here makes it the model of choice for *telnet* duration.

4.5 TELNET Responder/Originator Ratio

If we wish to use these models to generate or predict *telnet* traffic, then we also need models giving the relationships between the various distributions. In particular, we would like to know how many responder bytes to expect given a particular number of originator bytes, and how long a connection will last given how many bytes it transfers.

We model the ratio between the number of responder bytes and originator bytes using a simple log-normal distribution. For the test datasets we found $\bar{x}$ ranged from 4.17 to 4.46, tending toward the high end, and σ_x from 1.77 to 1.89, also tending to the larger value. For the fixed model we chose $\bar{x} = \log_2 21 \approx 4.39$ and $\sigma_x = \log_2 3.6 \approx 1.85$.

Figure 10 shows the performance of each model. Other than the unscaled UK and NC datasets, the analytic model does quite well, with $\mu \leq 0.2$ except for the scaled NC, with $\mu = 0.25$. In general the LBL-2 empirical model does a little better than the analytic model, and almost always better than UCB. Scaling improves some fits considerably and has only

marginal effect on others. The overall success of the unscaled analytic model gives solid evidence that the ratio between the bytes generated by the computer in a remote login session and those generated by the user is about 20:1, since the fixed model uses a ratio of 21:1.

For the responder/originator ratio we are interested in agreement in both the upper and lower tails, as disagreement in either could result in skewed predictions when the number of originator bytes is large. Figure 28 in Appendix D shows the performance for the upper and lower tails. All of the models do fairly well for the upper tails except for the unscaled UCB model, which underestimates both upper tails. The analytic model does best. With the LBL-2 model, scaling trades off better performance in the 10% tail for worse in the 1% tail.

In the lower tail for both the analytic and LBL-2 models scaling helps the 10% tail but worsens the 1% tail, indicating that the 1% tail is distributed differently than the other 99%. The UCB model does well in the lower tails, though. All in all we are left with no clear best model, and none of the models is really bad.

One might wonder whether the responder/originator ratio's distribution itself varies according to the number of bytes transferred; for example, perhaps when many originator bytes are transferred, the ratio tends to be low, because relatively speaking not so many responder bytes are transferred. For the test datasets we found that the correlation coefficient between $\log_2$ of the originator bytes and $\log_2$ of the responder/originator ratio varied from 0.07 to 0.10, indicating at most a mild positive correlation.

When using the responder/originator ratio to generate *telnet* traffic, a subtle point arises: one can either derive the originator bytes and the ratio, and multiply to obtain the responder bytes, or one can proceed in the opposite fashion, generating the responder bytes and the ratio, and dividing to obtain the originator bytes. While these two approaches appear equivalent, they are not, and the former (deriving the responder bytes from the originator) is preferable. The difference arises because while both the responder bytes and the ratio are log-normal distributed, the originator bytes are *extreme* distributed. Multiplying the originator byte's extreme distribution by the ratio's log-normal distribution yields a distribution close to log-normal; but dividing the responder byte's log-normal distribution by the ratio's log-normal distribution yields *exactly* a log-normal distribution (since the difference of two normal distributions is a normal distribution), and not an extreme distribution. Alternatively, we can think of the originator bytes as having a somewhat skewed log-normal distribution. Multiplying this distribution by another log-normal distribution smears out the deviations, and the result is close to log-normal; but chances are dividing two log-normal distributions will never reproduce the skewed distribution.

Thus, to generate traffic we should begin by generating the number of originator bytes and the responder/originator ratio, and then multiply to derive the responder bytes. This ap-

Figure 11: Responder/Duration Distributions for LBL-1: Exponential Fits

proach is not ideal, however, because it ignores the responder-bytes model we outlined above, which is more successful than the originator-bytes model.

4.6 TELNET Responder/Duration Ratio

Just as we want a way to relate the originator bytes sent with the responder bytes, we also would like to relate these random variables to the connection duration. We investigated analytic models for three different ratios: originator bytes to duration, responder bytes to duration, and total bytes to duration. We found the best fits came using the responder/duration model.

For most connections the responder/duration ratio was well modeled by an exponential distribution, but “large” connections—those whose responder bytes were in the upper 10% of all connections—had a different distribution. For these, the ratio was fairly well modeled by a log-normal distribution.

Figure 11 shows the responder/duration ratio for both the lower 90% of the LBL-1 connections (in terms of responder bytes) and the upper 10%. The distribution on the left is for the lower 90%; though it is hard to tell due to scaling, an exponential with the same mean has been drawn and lies squarely on top of it. This fit is very good; it passes A^2 at the 5% level. To the right we show the distribution of the upper 10%, plotted with an exponential with the same mean. We see that the distribution is qualitatively different, and the corresponding exponential not a good fit.

We find the bimodality shown in this figure a bit puzzling. It says that very large connections (in terms of bytes trans-

Figure 12: Empirical vs. Analytic Models for TELNET Responder/Duration Ratio

ferred) occur over relatively short durations: while the geometric mean of the responder bytes in these large connections is 45 times that of the smaller (lower 90%) connections, the geometric mean of their durations is only 16 times that of the smaller connections. This phenomenon was also observed by the authors of [SC92], who found that “users transmitting large amounts of data over a link tend to transmit that data within 15 minutes.” We do not have a good explanation for this phenomenon.

For the lower-90% model, the test datasets gave $\bar{x}$ ranging from 27 to 33 for the responder/duration ratio; we chose $\bar{x} = 30$. For the upper-10% model, $\bar{x}$ ranged from 5.19 to 5.41 and σ_x from 1.38 to 1.61; we chose $\bar{x} = 5.3$ and $\sigma_x = 1.5$.

Figure 12 shows the fit of the models for the lower 90% of the responders. The analytic fit is good, with $\mu \leq 0.3$ and often $\mu \leq 0.2$; in general it fits better than either empirical model.

For the upper 10% of the responders we compared considerably fewer datasets. Our requirement that each dataset include at least 100 measurements ruled out any dataset with fewer than 1,000 *telnet* connections, leaving just the LBL and NC datasets. The fit remains good, though: the analytic model does well, with $\mu \leq 0.3$ except for the unscaled NC dataset (where $\mu \approx 0.6$ for all three models), quite a bit better than the UCB model and about equal to the LBL-2 model.

Figure 28 in Appendix D summarizes the upper and lower tail distributions for the fit to the lower 90% of the responders. In the upper tails the analytic model does best, only mildly underestimating the upper 1% tail; only the scaled LBL-2 model is roughly comparable. In the lower tails the scaled analytic

Figure 13: Interarrivals for TELNET

model does very well, with the unscaled version overestimating the 1% tail somewhat. Again the empirical models do considerably worse and the scaled UCB model is completely inadequate, though the unscaled model is acceptable.

For the models of the upper 10% of the responders, every model underestimates the upper 1% tail somewhat, with the analytic models and the scaled LBL-2 model about the same at $\xi \approx -0.5$. The unscaled empirical models fare poorly in the 10% tail, too, considerably underestimating it, while analytic models and the scaled empirical models match the 10% tail well. In the lower tails the unscaled models do fairly well with the 10% tail, and the scaled models do quite well. Except for the scaled UCB and analytic models, though, the lower 1% tail is considerably underestimated.

4.7 TELNET Interarrivals

We now turn to modeling *telnet* interarrivals, using the methodology discussed in Section 3.11 above. Figure 13 compares the μ values for the unscaled and scaled arrival models. As explained in Section 3.11, instead of comparing the analytic model to the empirical models, we compare the analytic model’s scaled version with its unscaled version. We plot μ for the scaled analytic model on the Y axis vs. μ for the unscaled model on the X axis. Also, as mentioned in Section 2.2, we omit the USC dataset from our interarrival models because of the trace’s periodic blackouts.

As expected, the scaled model in general does uniformly better, but we note that even for the unscaled model, $\mu \leq 0.25$, which, when compared to the fits of other models above, we see is quite good. The arrivals are thus well-

Dataset	# Conn	# Rej	% Failures	$\bar{x}_{orig}$	σ_{orig}	max_{orig}	$\bar{x}_{resp}$	σ_{resp}	max_{resp}
LBL-1	57,898	2	38 %	2.0KB	$\times 9.2$	4.2MB	305B	$\times 2.0$	923KB
LBL-2	57,997	1	36 %	2.4KB	$\times 7.8$	1.1MB	328B	$\times 2.1$	584KB
LBL-3	46,167	6	19 %	2.4KB	$\times 6.2$	1.9MB	384B	$\times 1.9$	128KB
LBL-4	73,179	39	2 %	6.0KB	$\times 8.5$	5.6MB	398B	$\times 2.2$	1.4MB
LBL-5	50,969	161	8 %	14.5KB	$\times 8.5$	16.5MB	633B	$\times 2.9$	9.5MB
LBL-6	55,176	1048	8 %	28.4KB	$\times 6.8$	15.7MB	888B	$\times 2.2$	1.3MB
BC	345	116	25 %	15.5KB	$\times 6.2$	2.4MB	1005B	$\times 3.0$	81KB
UCB	6,899	0	1 %	2.1KB	$\times 7.2$	720KB	307B	$\times 2.0$	1.7MB
USC	4,615	15	4 %	11.5KB	$\times 10.3$	3.6MB	709B	$\times 2.3$	74KB
DEC-1	23,864	5	2 %	1.1KB	$\times 11.6$	5.8MB	264B	$\times 2.2$	75KB
DEC-2	18,819	88	3 %	1.3KB	$\times 11.7$	26MB	292B	$\times 2.4$	356KB
DEC-3	19,244	7	7 %	2.2KB	$\times 14.1$	18MB	339B	$\times 2.7$	223KB
NC	904	206	9 %	12.9KB	$\times 12.3$	12MB	1182B	$\times 4.5$	3.2MB

Table 7: Summary of NNTP Connections

modeled as a non-homogeneous Poisson process with hourly rates given by Figure 2. This finding is at odds with that of [MM85], who found that “user interarrival times look roughly lognormal”. Perhaps the discrepancy is due to the authors characterizing all interarrivals lumped together, rather than postulating separate hourly rates.

Figure 30 in Appendix D summarizes the tail distributions for the scaled and unscaled arrival models. (See the text in Appendix D for an explanation of the symbols in the figure.) Note the range shown in the figure: even the worst fits have $|\xi| \leq 0.25$. Thus both the unscaled and scaled models do quite well, and the scaled model does exceptionally well.

5 NNTP

5.1 Overview of NNTP Connections

Table 7 summarizes *nnntp* connections. As *nnntp* is non-interactive, the connection duration is not of much interest and has been omitted. Appendix B discusses the connections we rejected due to protocol errors.

We expect *nnntp* connections to show considerable variation because they can come in at least three modes: a server contacts a peer and is informed that the peer presently cannot talk to the server; the server offers the peer news articles but the peer already has the articles; the server offers articles and the peer does not have the articles. Each of these modes will result in significantly different distributions of the bytes transferred during the connection. Furthermore, the second and third modes are somewhat indistinct, since the remote peer may have some but not all of the offered articles.

The first mode is easy to detect. If upon initially being contacted a responder peer is unable to communicate with the originating peer, it sends a message with response code 400 (“service discontinued”) as per [RFC977]. When the origi-

nating peer then replies with “QUIT” followed by a carriage-return and a line-feed, it will send a total of 6 bytes during the connection. Indeed, we find large spikes of 6 originator bytes in the *nnntp* datasets, as did the authors of [DJCME92]. Thus we can recognize a connection in which the originating host sent 6 bytes as a “failure”.

Not surprisingly, the failure rate varies greatly from site to site and from time to time, since it is often due to transient phenomena such as full disks. These failure rates are given in the “% Failures” column. Note that even over a period of 7 days, the DEC failure rate moved from 2% to 7%. To compute the remaining statistics in the Table, we first removed all failure connections from the datasets.

Not only can the failure rate vary significantly, but so can the bytes transferred during non-failure connections. For example, as can be seen by the large increase in $\bar{x}_{orig}$ between LBL-3 and LBL-4, the LBL *nnntp* server became much more effective in propagating news over a five month period. LBL-5 and LBL-6 continue the impressive growth in $\bar{x}_{orig}$. A similar effect can be seen between DEC-1 and DEC-3, only a week apart. Such changes can be due in part to circumstances wholly outside of the local site. Whether the articles a server attempts to propagate to its peers are accepted depends on whether those peers already have the articles; a subtle change in the *nnntp* peer topology can swing a server’s position from one of holding mostly “stale” news to holding mostly “fresh” news. The steadily increasing $\bar{x}_{orig}$ value for the last four LBL datasets, though, is most likely simply a reflection of the global growth in USENET *nnntp* traffic, which increases in volume about 75%/year (see [Paxson93]).

5.2 NNTP Originator Bytes

Figure 14 shows the distributions of bytes sent by the originator in non-failure *nnntp* connections at LBL, DEC, and coNCert. The distributions show a large degree of variance

Figure 14: Distribution of NNTP Originator Bytes

(recall that the X axis is scaled logarithmically), suggesting that scaling is vital when modeling *nntp* traffic.

Given the great variation in originator bytes transferred, we decided to simply use a log-normal model to describe the connections, with the caveat that we do not expect the model to perform well in either scaled or unscaled forms (but we also do not expect empirical models to do well, either).

For the LBL test datasets, the log-mean ranged from 10.8 to 12.5, and the log-standard deviation from 2.6 to 3.3. For our fixed model we chose $\bar{x} = 11.5$ and $\sigma_x = 3$, respectively. We note that $\bar{x} \approx \log_2 3$ KB appears a bit higher than the median *nntp* article size of around 2 KB reported in [Adams92]. This difference probably means that, when an *nntp* server has a “fresh” article, it tends to have more than one.

Figure 15 shows the performance of the various models. As expected, none of the models does well due to the great variations from dataset to dataset, though scaling the models helps somewhat. In general, the analytic model performs acceptably only when scaled, and the empirical models only on the earlier LBL datasets and the scaled BC datasets. But the analytic model does as well as than either empirical model, indicating that the log-normal approximation is no worse than the inherent variation in the distributions.

Figure 29 in Appendix D shows the tail performance of the models. Again we show only the upper tail because with bulk transfer the lower tail is not of much interest. The unscaled

models do quite badly, not surprisingly, given their poor overall performance. The upper 10% tail is only safe when scaling the models, and the upper 1% tail only when scaling empirical models. From this figure we conclude that the *nntp* models must always be scaled, and even then the otherwise somewhat successful analytic model is problematic due to grossly overestimating the upper 1% tail.

One final important point regarding modeling *nntp* originator bytes is that the distribution is not stationary but changes over the course of a day. Figure 16 shows the hourly $\bar{x}_{\text{orig}}$ for LBL-1 and LBL-4 non-failure *nntp* connections (this plot was made by constructing “superhours” as discussed in Section 3.11). We see considerable but not consistent variation. The peak-to-peak differences for both datasets is about a factor of 3.4; but LBL-1’s connections tended to be largest in the middle of the night, with secondary peaks during “prime-time” work hours. LBL-4’s connections peaked during working hours and were lowest at precisely the time when LBL-1’s were highest. The test datasets also showed a weekly pattern, with LBL-1 and LBL-4 (and to a lesser extent LBL-2) having minimal $\bar{x}_{\text{orig}}$ during weekends (with a peak-to-peak variation of about a factor of 3), while LBL-3 had a *maximum* $\bar{x}_{\text{orig}}$ on Saturdays (peak-to-peak about a factor of 2).

The variation in the daily pattern may be due to the influence of key *nntp* gateways either propagating news as soon as it comes in (consistent with the LBL-4 case) or waiting till the

Figure 15: Empirical vs. Analytic Models for NNTP Originator Bytes

Figure 17: Interarrivals for NNTP

Figure 16: Daily Variation in $\log_2$ -mean of LBL NNTP Originator Bytes

Figure 18: One-Minute Variation in DEC-2 NNTP Arrivals

late-night hours to take advantage of minimal loads (LBL-1). The weekly variation is more difficult to explain. We expect that most news articles are written during the week, and it seems unlikely that *nntp* gateways would queue a significant number of articles till the weekend. So the strong LBL-3 Saturday peak remains a puzzle.

5.3 NNTP Responder Bytes

As seen in Table 7 above, there is in general much less variation in the bytes sent by an *nntp* responder than by the originator. For all of the datasets except LBL-5, BC and NC, the responder sent fewer than 1500 bytes in 82% or more of the connections. For LBL-5 this value was 77%, for BC 65%, and for NC 63%. Thus we decided not to model *nntp* responder bytes, as in general the datasets do not show interesting variations. We did compute the correlation coefficients between $\log_2$ of the originator and responder bytes and found a range from 0.37 (NC and UCB) to about 0.8 (USC, BC). In general we would expect to find positive correlation since the more articles offered by the originator to the responder, the more replies the responder must generate.

5.4 NNTP Duration

Since *nntp* is a bulk-transfer protocol and not interactive, we do not model connection durations, because these are presumably dominated by networking latencies and not a fundamental aspect of the *nntp* protocols. Similarly, below we do not model *smtp* or *ftp* durations.

5.5 NNTP Interarrivals

Figure 17 shows the scaled and unscaled interarrival models.⁸ The results appear puzzling. The UCB, LBL-3, LBL-6, and to some degree NC arrivals all appear well-modeled as Poisson processes with hourly rates corresponding to those in Figure 2. The other datasets, including the remaining LBL datasets, are poorly fitted both when scaled and when unscaled, indicating that they are *not* Poisson processes.

The poor fit to the BC data is in part due to its very low *nntp* arrival rate: less than two connections per hour on average. The other poorly-fitted datasets turn out to have interesting periodic behavior. In particular, *nntp* arrivals have a definite one-minute periodicity about them. Figure 18 shows the number of DEC-2 *nntp* connections that arrived during each second (i.e., ignoring minutes and larger units of the arrival time). Clearly, arrivals tended to show up at about 19 seconds past the minute, though some tended to arrive about 7 seconds

⁸The actual points for DEC-3 and LBL-4, and for LBL-3, LBL-6 and UCB, overlapped on this plot and became hard to distinguish, so we have added some horizontal bias away from the diagonal. These points all actually lie on the diagonal.

past. All of the *nntp* datasets show this pattern to varying degrees except for LBL-3; LBL-4 shows two distinct spikes. With NC, UCB, and LBL-6, the spike is quite sharp. With the other datasets, it is broad, like in Figure 18. The sharp spikes mean that only a relatively small fraction of the interarrivals are skewed, evidently enough to preserve sufficient approximation to a Poisson process.

We also investigated one-hour variation. We found a three-minute pattern in LBL-1, five-minute patterns in LBL-5 and UCB, fifteen-minute patterns in LBL-6 and BC, and a less-strong twenty-minute pattern in DEC-2.

Figure 30 in Appendix D summarizes the tails corresponding to the scaled and unscaled models. Both models underestimate the lower 10% tail but overestimate the lower 1% tail; the upper 10% tail is well modeled but the upper 1% considerably underestimated. As in Figure 15 above, we see little difference between the scaled and unscaled models.

6 SMTP

6.1 Overview of SMTP Connections

Table 8 summarizes the *smtp* connections. Again, Appendix B summarizes the reasons for removing the connections marked as rejects. Based on the values for $\max_{\text{orig}}$ it is clear that *smtp* is sometimes used to transfer quite large files, though that is not its main purpose.

There is quite a bit of variation in $\bar{x}_{\text{orig}}$ (and just about none in $\bar{x}_{\text{resp}}$). In [WLC92] the authors note that the UK *smtp* data show a substantially higher (arithmetic) $\bar{x}_{\text{orig}}$ than for the LBL-1 and LBL-2 datasets reported in [Paxson91]. They attribute this difference to the fact that since the U.K. academic network (JANET) was not at that time fully connected to the Internet, U.K. users were more likely to use *smtp* to transfer files. The large UK σ_{orig} variance supports this hypothesis. The DEC traffic has similar σ_{orig} values, and Mogul also states in [Mogul92] that an “FTP-by-mail” facility is responsible for about 150 rather lengthy *smtp* messages at DEC-WRL each day. It is less clear whether this theory explains the large NC message sizes.

Another explanation is that perhaps the DEC, NC, and UK traffic tends to make more *smtp* “hops”, each of which adds a Received header to the mail message [RFC822], pushing up the average number of bytes⁹. One would expect the greater number of hops to be correlated with “wider” wide-area traffic, presumably a property of the NC and especially the UK traffic, as these sites are at inter-network gateways. But this explanation does not address why the DEC traffic might tend to make more hops, unless due to the structure of DEC’s internal mail gateways.

⁹A check of one of the author’s mail folders revealed an average Received header length of more than 100 bytes.

Dataset	# Conn	# Rej	$\bar{x}_{orig}$	σ_{orig}	max_{orig}	$\bar{x}_{resp}$	σ_{resp}	max_{resp}
LBL-1	38,481	286	1.4KB	$\times 2.8$	2.1MB	331B	$\times 1.2$	1.9KB
LBL-2	51,240	572	1.5KB	$\times 2.9$	7.2MB	334B	$\times 1.2$	6.5KB
LBL-3	75,418	333	1.6KB	$\times 2.6$	1.6MB	334B	$\times 1.2$	2.9KB
LBL-4	92,694	1583	1.7KB	$\times 3.0$	1.2MB	335B	$\times 1.3$	2,980KB
LBL-5	123,741	446	1.7KB	$\times 2.9$	2.4MB	320B	$\times 1.3$	8.0KB
LBL-6	207,485	6,567	1.9KB	$\times 3.0$	37.0MB	321B	$\times 1.3$	9.5KB
BC	8,428	121	1.3KB	$\times 2.8$	1.1MB	324B	$\times 1.3$	10.2KB
UCB	16,929	61	1.3KB	$\times 3.0$	0.5MB	334B	$\times 1.3$	2.0KB
USC	3,498	3	1.4KB	$\times 2.3$	0.1MB	337B	$\times 1.2$	1.6KB
DEC-1	25,160	19	2.0KB	$\times 3.1$	2.5MB	340B	$\times 1.2$	4.7KB
DEC-2	10,777	5	2.1KB	$\times 3.5$	4.9MB	341B	$\times 1.2$	4.7KB
DEC-3	31,631	70	2.0KB	$\times 3.2$	5.1MB	338B	$\times 1.2$	3.5KB
NC	26,161	511	1.9KB	$\times 2.9$	1.8MB	340B	$\times 1.4$	10.6KB
UK	10,729	129	1.9KB	$\times 3.3$	4.6MB	319B	$\times 1.3$	6.0KB

Table 8: Summary of SMTP Connections

We see a definite trend in the LBL data indicating larger and larger mail messages. As discussed in [Paxson93], LBL’s wide-area traffic did become “wider” during the 29 month period spanned by the LBL datasets, in agreement with the “hops overhead” explanation.

6.2 SMTP Originator Bytes

When modeling the number of bytes sent by the *smtp* originator, we found that nearly all connections transferred more than 300 bytes, while the connections transferring fewer bytes showed sporadic distributions. We hypothesize that the first 300 bytes of these connections constitute a more-or-less fixed overhead, and that connections with fewer total originator bytes correspond to “failures”: either invalid email addresses or busy remote machines unable to accept mail at the moment. In constructing our models we therefore removed any connections of ≤ 300 bytes (anywhere from 0.6% to 2.3% of all connections) and subtracted 300 bytes from the remaining connections.

We found the distribution of *smtp* originator bytes to be bi-modal, not surprisingly given that *smtp* is also used to transfer files. We model the distribution using two log-normal distributions, one (called f here) for the lower 80% of the data, and one for the remaining 20% (g). Figure 19 shows this model’s fit to the LBL-3 test data after removing failures and subtracting 300 bytes; the horizontal line indicates the dividing line between using distribution f (below the line) and g (above).

For our fixed model, we found the mean of distribution f to range from 9.90 to 10.16, and chose $\bar{x} = 10$; the standard deviation ranged from 1.42 to 1.52, and we chose $\sigma_x = \log_2 2.75 \approx 1.46$. For g , the upper distribution, the means ranged from 8.43 to 9.06, and the standard deviations from 2.55 to 3.52. We chose $\bar{x} = 8.5$ (since three of the test datasets had means quite close to 8.5) and $\sigma_x = 3$. Note that, while the distribution g has a lower mean than distribution

Figure 19: Log-Normal Fit to LBL-3 SMTP Originator Bytes

f , we use only g ’s upper 20% tail, which is larger due to g ’s significantly higher standard deviation.

Figure 20 shows that this analytic model is highly successful compared to the empirical models. In virtually every case it performs as well or better than the empirical models, usually better. We also see that scaling consistently improves the performance of all three models, often substantially, indicating that there is considerable site-to-site variation in bytes transferred.

Figure 29 in Appendix D shows the tail distribution for this model. The models all do well, in general slightly underestimating the upper tails, except for the unscaled UCB model,

Figure 20: Empirical vs. Analytic Models for SMTP Originator Bytes

which is more severe in its underestimation.

As was the case for *nntp*, for *smtp* we found that the originator bytes distribution is not stationary. Figure 21 shows the hourly $\bar{x}_{\text{orig}}$ for LBL-1 and LBL-4 *smtp* connections after removing those less than 300 bytes and subtracting 300 bytes from the remainder. Unlike *nntp*, which suffered from inconsistent variations, here the pattern is more stable: connection sizes peak during off-hours, the evening and early morning, and reach minima during peak working hours. We conjecture that uses of *smtp* to transfer files and not messages typed in by users tend to happen off-hours and cause this pattern. Of the four test datasets, LBL-4 shows the greatest peak-to-peak variation, about a factor of 3.3, comparable to the *nntp* variation. The other datasets are closer to a factor of 2. Unlike *nntp*, we did not detect a noteworthy weekly pattern.

6.3 SMTP Responder Bytes

We did not model the distribution of the responder bytes in *smtp* connections, as the responder's role shows little variation. For the LBL test datasets, in about 75% (73% to 79%) of all connections the responder sent between 300 and 400 bytes, and more than 99% of the connections sent between 100 and 1000 bytes. Of all the datasets, LBL-5 had the lowest proportion of connections sending between 100 and 1000 bytes, still a very high 98%. We also found that the coefficient of correlation between $\log_2$ of the originator bytes and $\log_2$ of the responder bytes for the LBL datasets varied from .035 to .246; thus we found little interesting behavior to model in the responses. While reference [DJCME92] finds that *smtp*

Figure 21: Daily Variation in $\log_2$ -mean of LBL SMTP Originator Bytes

Figure 22: Interarrivals for SMTP

connections are strongly bidirectional, this finding must be interpreted with the rather fixed nature of the *smtp* responder in mind.

6.4 SMTP Interarrivals

Figure 22 shows the unscaled and scaled fits to the *smtp* interarrivals.¹⁰ Both models do extremely well for almost all of the datasets, indicating these arrivals are well described by the pattern shown in Figure 2. We are somewhat puzzled that the BC interarrivals fared so well with the unscaled model, given the roughly three-hour shift between BC's arrival activity and LBL's as shown in Figure 2; evidently there is enough similar overlap during the busy 11AM-4PM times to bring the overall distributions into fairly close agreement. We do not have an explanation for the poor fit to DEC-2's interarrivals, though the traffic, which spanned the Thanksgiving holiday, is certainly atypical in one sense: only 6% of the DEC-2 connections originated from a DEC host, while for DEC-1 and DEC-3 the figure is 40-46%.

Figure 30 in Appendix D summarizes the tail behavior for the interarrival models. Both the scaled and unscaled models do quite well in the lower and upper tails, only underestimating the upper tail somewhat.

7 FTP

7.1 Overview of FTP Connections

Table 9 summarizes *ftpdata* connections. Each connection is unidirectional, with sometimes data flowing from the connection originator to the responder (corresponding to an *ftp get* command) and sometimes in the other direction. The "Get" column shows the percentage of connections that were get commands; the remainder were put commands. The next three columns show the (geometric) mean, standard deviation, and maximum for the number of bytes transferred. As before, Appendix B gives details regarding the connections we rejected.

Two rows are given for each of LBL-5 and LBL-6. As discussed in [Paxson93], a considerable portion of LBL traffic, particularly in LBL-5 and LBL-6, was generated by background scripts fetching weather maps from a remote anonymous ftp site. LBL-5* and LBL-6* show the *ftpdata* statistics with this traffic removed.¹¹

Two rows are given for the UCB dataset. The first includes 2,315 connections of 74 bytes each, all but three of which

¹⁰Here again we have added some horizontal space between the points in the lower-left corner to aid legibility. The LBL points all have unscaled $\mu < 0.05$; the excursion up to $\mu = 0.1$ is an artifact of the plot.

¹¹Weather-map/ftp traffic comprised 3.2% of LBL-3 connections and 8.5% of LBL-4 connections. Excluding this traffic lowered LBL-3's mean to 3.2KB and raised LBL-4's mean to 4.0KB. The corresponding standard deviations rose to 18.0 and fell to 14.0, respectively.

were between the same two hosts, and 95% of which came between 30 and 45 seconds apart. The UCB* row summarizes the UCB data with this anomaly removed.

In testing our models below we used LBL-5*, LBL-6*, and UCB*.

There clearly is quite a range in $\bar{x}_{\text{bytes}}$, even day-to-day as shown in the DEC data (though the low-point there, DEC-2, includes Thanksgiving, and might therefore be uncharacteristic). We might be tempted to declare a trend towards increasing file sizes with time in the LBL datasets, save for the LBL-6 dataset, which shows a sharp drop. We do not know whether LBL-6 was atypical, or whether the mean file size simply fluctuates a great deal. The uniformly large values of σ_{bytes} shows that in general file sizes vary widely.

Finally, computing the coefficient of variation (i.e., $\sigma / \bar{x}$) for the *ftpdata* interarrival times gives values from 2.4 to 8.0, significantly higher than for the other protocols. If the arrivals came from a homogeneous Poisson process, then for $\sigma / \bar{x}$ we would get 1, and if they were perfectly periodic, 0. These high values show that the traffic is quite bursty. This result is not surprising, as a "multiple-get" file transfer results in a rapid succession of *ftpdata* connections, sometimes quite large (see below).

Table 10 summarizes the *ftpctrl* connections. We have not shown statistics for bytes transferred and duration of the *ftpctrl* connections themselves since the primary use of *ftpctrl* connections is to spawn *ftpdata* connections, either for file transfer or to list remote directories. Instead, we grouped with each *ftpctrl* connection its associated *ftpdata* connections. We considered an *ftpdata* connection to belong to a *ftpctrl* connection if it occurred during the span of the *ftpctrl* connection and was between the same two hosts.

The starred LBL rows summarize the LBL datasets with the weather-map traffic removed. Unlike with *ftpdata*, here we include LBL-3 and LBL-4 in the filtering, since weather-map traffic had a substantial influence on their *ftpctrl* connections. Again, we use the starred datasets in our analysis below.

The "Orphans" column lists the percentage of *ftpdata* connections for which we could not associate an *ftpctrl* connection. High percentages of orphans were often due to *ftpctrl* connections that were terminated by RST packets instead of FIN packets, which, as explained in Section 2.1 above, were not included in our analysis. The authors of [EHS92] reported about 3% of *ftpctrl* connections were terminated by RST packets. As seen in the Table there appears to be considerable variation in this value.

The "# Overlap" column lists the number of overlapping *ftpctrl* connections between two hosts. For our analysis we merged such overlaps into a single conversation. The large number of LBL-4 overlapping connections is almost all due to overlapping connections to one of the weather-map sites, as can be seen by the appreciably lower value for LBL-4*. In LBL-4 we observed up to five overlapping connections (typically four), all virtually identical in bytes transferred and

Dataset	# Conn	# Rej	Get	$\bar{x}_{bytes}$	σ_{bytes}	maxbytes
LBL-1	23,555	287	80 %	2.3KB	$\times 15.3$	54.0MB
LBL-2	27,917	335	92 %	2.4KB	$\times 17.4$	124.3MB
LBL-3	39,552	349	91 %	3.3KB	$\times 17.7$	61.6MB
LBL-4	65,860	335	86 %	3.8KB	$\times 14.7$	67.2MB
LBL-5	82,025	344	83 %	5.1KB	$\times 16.0$	176.5MB
LBL-5*	66,411	344	80 %	4.5KB	$\times 16.0$	176.5MB
LBL-6	123,773	464	89 %	3.8KB	$\times 16.2$	291.6MB
LBL-6*	86,464	464	91 %	2.1KB	$\times 14.9$	291.6MB
BC	5,199	58	97 %	2.5KB	$\times 12.6$	16.1MB
UCB	6,844	77	98 %	0.4KB	$\times 11.5$	22.1MB
UCB*	4,529	77	96 %	1.0KB	$\times 13.5$	22.1MB
USC	1,870	29	93 %	1.3KB	$\times 14.5$	4.7MB
DEC-1	7,970	6	100 %	2.2KB	$\times 16.5$	4.9MB
DEC-2	4,013	13	100 %	1.3KB	$\times 17.1$	7.1MB
DEC-3	6,775	25	99 %	1.9KB	$\times 16.7$	12.8MB
NC	19,076	183	98 %	1.8KB	$\times 19.0$	43.8MB
UK	10,018	58	97 %	3.4KB	$\times 14.2$	6.8MB

Table 9: Summary of FTP Data Connections

Dataset	# Conn	# Rej	Orphans	# Overlap	0 xfer	$\bar{x}_{xfers}$	σ_{xfers}	maxxfers	$\bar{x}_{bytes}$	σ_{bytes}
LBL-1	3,757	51	5 %	57	19 %	3.3	$\times 2.9$	1,006	28KB	$\times 15.2$
LBL-2	5,312	72	5 %	49	25 %	3.2	$\times 2.8$	388	27KB	$\times 17.0$
LBL-3	7,920	93	5 %	135	19 %	2.7	$\times 2.8$	612	24KB	$\times 17.8$
LBL-3*	6,916	90	5 %	135	21 %	3.1	$\times 2.9$	612	30KB	$\times 18.4$
LBL-4	11,587	191	6 %	1,012	15 %	2.8	$\times 2.7$	1,951	24KB	$\times 17.5$
LBL-4*	7,941	189	7 %	112	17 %	3.3	$\times 3.0$	1,951	33KB	$\times 17.6$
LBL-5	18,501	1,227	5 %	160	15 %	2.2	$\times 2.5$	975	28KB	$\times 12.3$
LBL-5*	9,968	1,227	6 %	108	26 %	3.0	$\times 3.0$	975	31KB	$\times 16.7$
LBL-6	31,734	535	7 %	212	21 %	2.2	$\times 2.5$	2,996	30KB	$\times 12.5$
LBL-6*	12,470	535	10 %	196	24 %	3.1	$\times 2.9$	2,996	31KB	$\times 16.8$
BC	669	19	40 %	2	32 %	3.3	$\times 2.7$	426	13KB	$\times 14.2$
UCB	756	19	15 %	7	26 %	3.9	$\times 2.6$	350	12KB	$\times 14.9$
USC	272	6	26 %	2	22 %	3.8	$\times 2.8$	133	20KB	$\times 14.5$
DEC-1	727	8	6 %	18	26 %	5.4	$\times 3.2$	961	36KB	$\times 15.6$
DEC-2	491	8	6 %	14	13 %	5.0	$\times 3.0$	106	36KB	$\times 17.8$
DEC-3	811	17	14 %	18	25 %	4.8	$\times 2.9$	232	36KB	$\times 15.3$
NC	2,500	59	7 %	49	31 %	5.0	$\times 2.9$	392	26KB	$\times 18.6$
UK	1,733	35	5 %	133	24 %	3.4	$\times 3.0$	368	22KB	$\times 16.0$

Table 10: Summary of FTP Control Connections

duration, repeating every half hour for days on end. Evidently a number of weather-map scripts were run in the background on the same host and managed to synchronize. The large number of overlapping UK connections, on the other hand, is due to the high frequency of connections between pairs of popular hosts, such as one vendor's main Internet site in the U.K. connecting to the anonymous ftp archives of Washington University in Missouri. The authors of [WLC92] noted the Missouri site as the single most popular U.S. ftp site (and, indeed, Missouri was the most popular state in general for

U.K.-U.S. traffic).

The next four columns in Table 10 show statistics regarding the number of *ftpdata* connections that occurred during each *ftpctrl* connection. The "0 xfer" column lists the percentage of all *ftpctrl* connections that did not have any associated *ftpdata* connection. These numbers are lower than the 44% reported in [EHS92] because the authors of that paper were able to distinguish between file transfers and remote directory listings; we consider any *ftpdata* connection to be a "file transfer". Presumably a large proportion of the "0 xfer" connections are

The $\bar{x}_{\text{xfers}}$ and σ_{xfers} columns give the geometric mean and standard deviation for the number of files transferred, given that at least one file was transferred. That the mean is substantially higher than one is not surprising since we classify remote directory listings as file transfers, and probably the most common use of *ftp* is to connect to a remote archive site, do several listings to find the file or files of interest, and then transfer those files. We did find the values for $\max_{\text{xfers}}$ surprising, as we would expect that large sets of related files would be grouped together into a single archive file, unless the archive would be too large. This latter hypothesis turns out to be the case: the 1,951 files transferred at one time during LBL-4 together totaled 93MB, certainly too much to conveniently pack into an archive. Similarly, the 961 transferred during DEC-1 totaled 50MB.

The $\bar{x}_{\text{bytes}}$ and σ_{bytes} columns show the geometric mean and standard deviation for the total number of bytes transferred via *ftpdata* connections during each *ftpctrl* connection (for those connections with at least one *ftpdata* transfer). We note that these means are 5-10 times greater than those for *ftpdata*, an increase larger than that due simply to the multiplying effect of $\bar{x}_{\text{fers}}$. We suspect that this disparity is due to a typical *ftpctrl* connection including at least one true file transfer. As files will tend to be significantly larger than directory listings, the mean number of transferred bytes will approach the mean file size, and not be held down, as are the *ftpdata* connection summaries, by a large number of smaller directory listings. The σ_{bytes} values are quite large, again showing a wide range in transfer sizes.

We model the bytes transferred during an *ftpdata* connection using a log-normal distribution. Figure 23 shows this model fitted to the first half of the LBL-4 dataset. While the model appears to match the overall shape, a number of clumps and spikes make the actual distribution rather irregular. For example, LBL-4 has a spike of 1,269 connections, each transferring 1,856 bytes. For the most part, unfortunately, these spikes do not occur in predictable locations, making it difficult to incorporate them into our analytic model. Such unpredictability also impairs the ability of empirical models to fit other datasets. One spike stands out, however, being present in all the DEC datasets, the NC dataset, LBL-4, and LBL-5 (but not LBL-6). This spike occurs at 524,288 bytes (i.e., exactly 2^{19} bytes), a size often used when splitting a large distribution archive into manageable pieces.

Using the LBL test datasets, we found a range for the log-mean of 11.27 to 11.88, and chose $\bar{x} = \log_2 3000 \approx 11.55$ for our fixed model. The log-standard deviation varied from 3.83 to 4.21: we chose $\sigma_x = 4$.

A scatter plot comparing Analytic Fit (Y-axis) against Empirical Fit (X-axis). Both axes range from 0.0 to 0.8 with major ticks every 0.2. A diagonal line representing the identity line (y=x) is drawn from (0,0) to (0.8,0.8). Data points are labeled with model names: d2-, D2-, D3-, usc-, d3-, nc, L2-, L6-, d1-, L3-, L6-, L5c-, L4-, uk, L1-, BC, and 15. Most points are clustered between 0.2 and 0.4 on both axes, with some points like '15' extending towards 0.8 on the X-axis.

26

pirical models for the bytes transferred during an *ftpdata* connection. For the most part, scaling considerably improves the fit, as we might expect given the wide range in $\bar{x}_{\text{bytes}}$ shown in Table 9. The LBL-2 model almost always fits better than the UCB model, even though the UCB model has had its anomalous spike at 74 bytes removed. In general the analytic model does well, though it suffers somewhat when fitted with the DEC datasets, which are very noisy.

For modeling *ftpdata* data bytes we are not particularly concerned with the degree of fit in the lower tail. The upper tail, on the other hand, is of particular importance because large file transfers can consume a tremendous amount of network resources. Figure 29 in Appendix D summarizes the tail performance of the various models. While most of the models fit the upper 10% tail fairly well (with the exception of the unscaled UCB model), all except the scaled LBL-2 model fare poorly in the upper 1% tail, each overestimating the tail except, again, for unscaled UCB. Thus all of these models must be used with particular care concerning their predictions for large *ftpdata* connections. The overestimation of the analytic model might be understood at least in part by the tendency to split huge files into several pieces (see discussion of the 2^{19} byte spike, above). In this case we would expect such large files to be fetched together as a group, and hope that models of the total bytes transferred during an entire *ftp* conversation might more accurately predict the upper tail. Unfortunately those models actually prove worse; see below.

7.3 FTP Conversation Bytes

Perhaps more important than an *ftpdata* bytes-transferred model is a model for the total number of *ftpdata* bytes transferred due to an *ftpctrl* connection. Such a model gives an indication of the total impact of each file transfer conversation. Figure 25 shows this distribution for the LBL-1 test dataset, again fitted to a log-normal model. In this case the fit is visually quite satisfying, and indeed an A^2 test indicates this fit is valid at the 1% level. Unfortunately this fit is also the best of those to the LBL test datasets; the others fail validity.

For the LBL test datasets, the log-mean ranged from 14.85 to 15.20, and the log-standard deviation from 3.82 to 4.18. For the fixed analytic model we took $\bar{x} = 15$ and $\sigma_x = 4$.

Figure 26 summarizes the performance of the models. The overall variance of the analytic model is fairly low; ignoring the BC datasets, the fits all fall in or quite close to the range $0.2 \leq \mu \leq 0.3$. Scaling has only a minor effect on the caliber of the fits (except for BC), and in some cases worsens them. Since scaling is beneficial for modeling *ftpdata* connection bytes but not *ftp* conversation bytes, we conjecture that the “mix” between short directory listings and larger file transfers varies considerably from site-to-site. Such variation would mean that scaling would aid *ftpdata* bytes considerably more than *ftp* conversation bytes, since the latter are dominated by the actual files transferred and are relatively unaffected by

Figure 25: Log-Normal Fit to Bytes in LBL-1 FTP Conversations

Figure 26: Empirical vs. Analytic Models for FTP Conversation Bytes

directory listings.

The analytic model performs noticeably better than the UCB empirical model, but not as well as LBL-2, which overall does quite well. We note that the authors of [DJCME92] reported that 80% of *ftp* conversations transfer less than 10 KB. But once we remove the 20-30% of conversations that did not transfer any data, half of the remainder transfer more than 32 KB, and a sixth transfer more than 500 KB. Thus if a file transfer conversation is not a "failure", it should not be assumed small.

As with *ftpdata* connections, we again are most concerned with the behavior of the models in the upper tails, summarized in Figure 29 of Appendix D. Each model except unscaled UCB does well in the upper 10% tail, but the analytic model greatly overestimates the upper 1% tail, even with scaling, and only scaled UCB performs well in both regards.

# Items	Range	Range _{LBL}
0	13-32 %	17-26 %
1	10-24 %	20-24 %
2	8-15 %	12-15 %
3	5-11 %	8-10 %
4	6-8 %	6-7 %
5	4-6 %	4-5 %
6	3-5 %	3-4 %
7-10	7-13 %	7-8 %
11-20	5-14 %	5-7 %
21-99	3-9 %	3-5 %
100+	0-1 %	0-1 %

Table 11: FTP Data Items Per FTP Conversation

7.4 Data Items Per FTP Conversation

While *ftpdata* connection bytes and total *ftp* conversation bytes are closely related, we were unable to produce a good model of the number of *ftpdata* connections in each *ftp* conversation, both because the distribution varies considerably from dataset to dataset and because of the heavy upper tail in the distributions. Table 11 lists the range over all of the datasets for the distribution. For example, for one dataset 8% of all conversations transferred 2 items, while for another 15% did. Site-to-site variation is considerable. Furthermore, there is too much mass in the upper tail to accommodate a geometric distribution, our most likely candidate for modeling.

The third column lists the same ranges for just the six LBL datasets, with the weather-related *ftp* conversations removed. Here the variation is substantially less, indicating that the mix at a particular site is fairly stable over time. Looking at the Table one might wonder whether the LBL data is "holding down" the lower-end of the ranges in the second column for 4 or more items. This turns out not to be the case; removing the LBL data from the tabulation does not change any of the

ranges for 4 or more items except to narrow the 11-20 range from 5-14% to 6-14%. Thus the LBL data is not particularly atypical.

7.5 FTP Interarrivals

Figure 27 shows the fits of the scaled and unscaled arrival models for *ftp* conversations. Overall both models perform quite well, with the maximum μ for the unscaled model about 0.3 and for the scaled model 0.2. We found periodicity in the DEC datasets, with arrivals peaking on the hour and the half hour, too great an interval to much affect individual interarrivals.

The deviation in the tails is similarly quite low, as shown in Figure 30 in Appendix D, with almost no distortion in the 10% tails except for the lower tail of the unscaled model, and only moderate distortion in the 1% tails.

Figure 27: Interarrivals for FTP

8 Summary

We have presented a number of analytic models for describing the characteristics of *telnet*, *nnip*, *smtp*, and *ftp* connections, drawn from wide-area traces collected from seven different sites, comprising more than 2.5 million connections. While these models are rarely exact in a statistical sense, we developed a methodology for comparing their effectiveness to that of other models, and found that in general they capture the essence of the connections as well or better than the empirical *tcplib* library. We also compared the models to an empirical model derived from a one-month trace of traffic at the

Protocol	Variable	Model	Parameters	Abs.	Rel.	Tails
<i>telnet</i>	orig. bytes	lg-extreme	$\alpha = \log_2 100; \beta = \log_2 3.5$	0	0	u: over
	resp. bytes	lg-norm, upper 80%	$\bar{x} = \log_2 4500; \sigma_x = \log_2 7.2$	+	0	u: over
	duration secs.	lg-norm	$\bar{x} = \log_2 240; \sigma_x = \log_2 7.8$	0	0	u: Over
	resp. / orig.	lg-norm	$\bar{x} = \log_2 21; \sigma_x = \log_2 3.6$	+	0	u: good/okay; l: good/okay
	resp. / dur.	exp., 0-90% resp.	$\bar{x} = 30$	0	+	both good
	resp. / dur.	lg-norm, 90-100% resp.	$\bar{x} = 5.3; \sigma_x = 1.5;$	0	0	u: okay; l: OVER/good
<i>nntp</i>	orig. bytes	lg-norm	$\bar{x} = 11.5; \sigma_x = 3;$	-/0	0	u: Over
<i>smtp</i>	orig. bytes	lg-norm + 300B, lower 80%;	$\bar{x} = 10; \sigma_x = \log_2 2.75;$	0/+	+	u: good
		lg-norm + 300B, upper 20%	$\bar{x} = 8.5; \sigma_x = \log_2 3;$			
<i>ftp</i>	conn. bytes	lg-norm	$\bar{x} = \log_2 3000; \sigma_x = 4;$	0/+	0	u: over
	conv. bytes	lg-norm	$\bar{x} = 15; \sigma_x = 4;$	0	0	u: over

Table 12: Summary of Analytic Models of Connection Characteristics

Protocol	Abs. Fit	Scaling Helpful?	Lower Tail	Upper Tail
<i>telnet</i>	+	Sometimes	good	good
<i>nntp</i>	-	No	1% over, 10% under	1% under
<i>smtp</i>	+	Sometimes	over	under
<i>ftp</i>	+	Sometimes	okay/over	okay/good

Table 13: Summary of Analytic Interarrival Models

Lawrence Berkeley Laboratory, which we found in general to be slightly better at modeling the traffic than the analytic models.

Table 12 summarizes the models characterizing the different protocols' individual connections. The "Variable" column lists the random variable being modeled, where "orig." stands for the bytes sent by the connection originator, "resp." for those sent by the responder, "conn. bytes" the total number of bytes transferred during the connection (for *ftpdata*), and "conv. bytes" the total bytes transferred during a *conversation* (an entire *ftp* session). For *telnet*, we also modeled ratios between some of these variables, to capture their interdependence.

The "model" column lists the models used. Almost all first apply a $\log_2$ transformation to the data. One model is log-extreme, where the extreme distribution is defined by Equation 4; one is exponential; and the remainder are log-normal. Four of the models have restrictions. The *telnet* responder bytes model describes only the upper 80% of the responses. The *telnet* "resp. / dur." models describe the ratio of the responder bytes to the connection's duration. The first such model does so for those connections whose number of responder bytes fell into the lower 90% of all connections. The second model describes this ratio for those connections in the upper 10% of all responses. Finally, the *smtp* originator model uses parts of two different log-normal distributions in

its description. The lower 80% of the originator distribution is modeled using the lower 80% of the first log-normal distribution; similarly, the upper 20% is modeled using the upper 20% of the second log-normal distribution.

The "Parameters" column gives the parameters we used for the fixed (i.e., unscaled) version of the model. The "Abs." column summarizes the quality of the model's fit in absolute terms: how well we assess the model as describing the random variable's distribution. A "0" indicates the model describes it adequately, a "+" that the model describes it well, and a "-" that it does poorly. When two values are given, the first is for the unscaled version of the model and the second for the scaled version. When only one value is given, scaling did not significantly improve the model's fit. The Table shows that we assess the models as being at least adequate for every random variable except when modeling *nntp* originator bytes; in that case the scaled version of the model is required for an adequate description.

The "Rel." column compares the model's performance to that of the two empirical models, one constructed from the UCB dataset and corresponding to *tcplib*, and one constructed from the LBL-2 dataset. A "0" indicates that the analytic model performs about equally and a "+" that it performs better. In all cases we found the analytic model does overall at least about as well as the empirical models, though for some of the "0" entries it did somewhat better than the UCB model and

somewhat worse than the LBL-2 model.

The final column summarizes the analytic model's performance in modeling the tails. A "u:" entry gives the fit to the upper tail and an "l:" entry to the lower tail. A value of "over" indicates the model substantially overestimates the 1% tail; "Over" that it also somewhat overestimates the 10% tail; and "OVER" that it grievously overestimates both the 1% and 10% tail. Similarly for "under". For models that do well describing their tails, we chose subjective evaluations of "okay" and "good". Some models have two evaluations reported for one of their tails; in this case the first is for the unscaled model and the second for the scaled model.

Table 13 summarizes the interarrival models for each protocol. Here "Abs. Fit" summarizes the absolute fit of the model using the same notation as before.

Since we did not compare the analytic interarrival models to empirical ones (due to difficulties in constructing such empirical models), we omit the "relative" fit. The "Scaling Helpful?" column indicates whether scaling substantially improved the model. The lone "No" entry, for *nntp*, reflects our finding that the *nntp* connection arrival process is not Poisson. The other "Sometimes" entries indicate that for many of the datasets scaling was not needed to produce good fits, and the arrivals can be modeled as a non-homogeneous Poisson process with hourly rates given by Figure 2. For those datasets requiring scaling, the arrivals can also be better modeled as a non-homogeneous Poisson process with different hourly rates than those in Figure 2.

The last two columns summarize the arrival models fit in the lower and upper tails. When two values are given, such as for *ftp*, then the first is for the unscaled model and the second for the scaled model.

Table 1, at the beginning of the paper, states our major conclusions. Here we summarize our additional findings:

- The ratio between bytes sent by a user in a remote-login session and those sent back by the remote computer is about 1:20.
- Of *ftp* conversations that are not "failures" (no data transferred), half transfer more than 32 KB, and a sixth transfer more than 500 KB.
- *smtp* and *nntp* connections show variations in size over the course of the day, with the largest *smtp* connections coming during evening and early morning hours, while the peaks of *nntp* varied considerably.
- *rlogin* traffic can be described by models for *telnet* traffic (see Appendix C), but requires scaling for acceptable fits, and even then does not in general fit as well as *telnet* traffic.
- We believe the site-to-site and month-to-month variations in network traffic characteristics are in part responsible for the success of the analytic models: the inter-site

differences are large enough that analytic models tend to be just as good a compromise among the varying datasets as empirical models.

The essence of the argument presented in this paper is that while wide-area traffic cannot be easily modeled exactly, if we can abide some inexactness then we can reap the benefits of using analytic models instead of empirical ones, without any relative loss of accuracy. We believe the approach discussed in this paper will prove beneficial for developing future analytic models and for gauging their effectiveness.

9 Acknowledgments

This work would not have been possible without the support and patience of Van Jacobson and Domenico Ferrari. I am also much indebted to 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 to Ramón Cáceres and Sugih Jamin, who between them made all of the non-LBL datasets available.

The LBL traces were gathered with the help of Craig Leres and Steve McCanne. Craig was also helpful in understanding *nntp*-related phenomena.

The Bellcore traces were gathered by D. V. Wilson.

I also want to thank Terry Speed and particularly Sally Floyd for valuable discussions on both modeling the data and presenting the results; and Domenico Ferrari, Sally Floyd, Van Jacobson, and Jon Crowcroft for their many helpful comments on earlier drafts of this paper.

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.

References

- [Adams92] R. Adams, *USENET Readership Summary Report for Oct 91*, Internet Society News, 1(1), Winter, 1992.
- [AD54] T. W. Anderson and D. A. Darling, *Asymptotic theory of certain goodness-of-fit criteria based on stochastic processes*, Ann. Math. Statist. 23, pp. 193-212, 1954.
- [Bryan67] G. E. Bryan, "JOSS: 20,000 Hours At A Console", Proc. of the Fall 1967 AFIPS Conference, Vol. 31.
- [Cáceres89] R. Cáceres, "Measurements of Wide Area Internet Traffic", Report UCB/CSD 89/550, Computer Science Division, University of California, Berkeley, California, 1989.

- [CW91] J. Crowcroft and I. Wakeman, "Traffic Analysis of some UK-US Academic Network Data", Proceedings of INET'91, Copenhagen, June, 1991.
- [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.
- [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.
- [EHS92] D. Ewing, R. Hall, and M. Schwartz, "A Measurement Study of Internet File Transfer Traffic", *Report CU-CS-571-92*, Department of Computer Science, University of Colorado, Boulder, Colorado, 1992.
- [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.
- [Gupta52] A. K. Gupta, *Estimation of the mean and standard deviation of a normal population from a censored sample*, Biometrika 39, pp. 266-273, 1952.
- [Heimlich90] S. Heimlich, "Traffic Characterization of the NSFNET National Backbone", Proceedings of the 1990 Winter USENIX Conference, Washington, D.C.
- [JS69] P. E. Jackson and C. D. Stubbs, "A study of multiaccess computer communications", Proc. of the Spring 1969 AFIPS Conference, Vol. 34.
- [JLM89] V. Jacobson, C. Leres, and S. McCanne, *tcpdump*, available via anonymous ftp to ftp.ee.lbl.gov, June, 1989.
- [Knuth81] D. Knuth, "Seminumerical Algorithms", Second Edition, Addison-Wesley, 1981.
- [MM85] W. T. Marshall and S. P. Morgan, *Statistics of Mixed Data Traffic on a Local Area Network*, Computer Networks and ISDN Systems 10(3,4), pp. 185-194, 1985.
- [Mogul92] J. C. Mogul, "Observing TCP Dynamics in Real Networks", Proceedings of SIGCOMM '92, Baltimore, Maryland, August 1992.
- [MT77] F. Mosteller and J. W. Tukey, "Data Analysis and Regression", Addison Wesley, 1977.
- [MJ93] S. McCanne and V. Jacobson, "The BSD Packet Filter: A New Architecture for User-level Packet Capture", Proceedings of the 1993 Winter USENIX Conference, San Diego, CA.
- [Pawlita89] P. Pawlita, "Two Decades of Data Traffic Measurements: A Survey of Published Results, Experiences and Applicability", *Teletraffic Science for New Cost-Effective Systems, Networks and Services*, ITC-12, M. Bonatti (Editor), Elsevier Science Publishers B.V. (North-Holland), 1989.
- [Paxson91] V. Paxson, "Measurements and Models of Wide Area TCP Conversations", *Report LBL-30840*, Lawrence Berkeley Laboratory, Berkeley, California, 1991.
- [Paxson93] V. Paxson, "Growth Trends in Wide-Area TCP Connections", in submission to *IEEE Network*. Available as *WAN-TCP-growth-trends.ps.Z* via anonymous ftp to ftp.ee.lbl.gov.
- [RFC822] D. Crocker, "Standard for the Format of ARPA Internet Text Messages", RFC 822, Network Information Center, SRI International, Menlo Park, CA, 1982.
- [RFC977] B. Kantor and P. Lapsley, "Network News Transfer Protocol", RFC 977, Network Information Center, SRI International, Menlo Park, CA, 1986.
- [Ross87] S. Ross, "Introduction to Probability and Statistics for Engineers and Scientists", John Wiley & Sons, 1987.
- [SC92] A. Schmidt and R. Campbell, "Internet Protocol Traffic Analysis with Applications for ATM Switch Design", *Report No. UIUCDCS-R-92-1735*, Department of Computer Science, University of Illinois at Urbana-Champaign, May, 1992.
- [WLC92] I. Wakeman, D. Lewis, and J. Crowcroft, "Traffic Analysis of Trans-Atlantic Traffic", Proceedings of INET'92, Kyoto, Japan, 1992.

A Details of filtering non-WAN traffic

We filtered the traffic datasets as followed:

- As mentioned in Section 2.3, we removed all connections between LBL and U.C. Berkeley.
The LBL datasets did not include any internal or transit traffic.
- The BC dataset consisted of about 65% internal and transit traffic, which we removed, keeping only traffic between a Bellcore host and an external host.
- Similarly, 10% of the UCB dataset was transit or internal traffic, and another 2% was traffic with LBL. We removed these connections.
- The USC dataset had no transit or internal traffic.
- The DEC traffic consisted of about 3% transit traffic, virtually all (e.g., 99.8% of the DEC-1 transit traffic) of which was *smtp*. We left this traffic in the dataset. We did not identify any internal traffic in the DEC datasets.
- The NC dataset had no transit traffic. We removed internal traffic, comprising 2% of the connections.
- About 2.5% of the UK traffic was internal, which we removed. Another 3.2% of the traffic was transit (source and destination both outside of the United Kingdom); as the UK traffic represents a truly wide-area link, we felt it appropriate to include this traffic in our analysis.

B Outliers removed from datasets

As noted in Section 3.4, we removed from our analysis connections that did not transmit at least 1 byte in each direction (or, for *ftp* data transfers, 1 byte total). In addition, we removed connections that were clearly the result of protocol errors. These latter connections were identified by their high purported data volume, transmitted over time scales that required unlikely data rates:

- For *telnet*, the additional removals were an LBL-4 connection of 4GB at 600KB/s, four LBL-5 connections of 4GB at rates from 480KB/s to 6.2MB/s, and three NC connections, two of more than 100MB at rates greater than 1.2MB/s, and one of 33MB at 42KB/s. It is possible (though improbable) that this last was a legitimate connection.

The very large value given in Table 6 for LBL-4's $\max_{\text{resp}}$ appears to be due to a legitimate connection; it lasted just over 18 hours, for an average data rate of 1.3KB/s. We removed it as an extreme outlier.

Each DEC dataset had fewer than 100 *telnet* connections due to the DEC firewall, so we did not include the DEC traffic in our *telnet* study.

- For *nnntp*, we removed four LBL-5 connections of 3.8GB or more at rates of 5.9MB/s or higher.
Two connections were removed from DEC-3 since they purported to have transferred in excess of a 1GB at rates of 1MB/s or higher, along with two USC connections of 3GB or more at rates of 40KB/s or higher.
We omitted the UK dataset because it contained only four connections.
- For *smtp*, we removed one DEC-3 connection as it purported to have transferred 2GB in 20 seconds.
- For *ftpdata*, we discarded three LBL-4 connections due to each involving 500MB or more at rates exceeding 400KB/s.
- For *rlogin*, we rejected one USC connection, purporting 4GB transferred in 40 seconds, as a protocol error. We removed the DEC and UK connections because in each case there were fewer than 100 connections.

C Modeling *rlogin* traffic using *telnet* models

Table 14 summarizes the *rlogin* traffic in the same manner as Table 6 does for *telnet*.

We expect to find the *rlogin* characteristics similar to those of *telnet* connections, since the two protocols are very similar in purpose, though *rlogin* will usually involve a pair of Unix hosts, which would eliminate one source of variance in the characteristics. Indeed, the LBL *rlogin* traffic shown in Table 14 is quite similar to that of the other *rlogin* datasets, suggesting that the difference in *telnet* traffic shown in Table 6 is due to LBL *telnet* connections tending to be with different types of computers than the other *telnet* datasets.¹²

We tested all of the analytic and empirical *telnet* models discussed in Section 4 on the *rlogin* datasets as well. In general we found that scaling improved fitting considerably and was required to achieve adequate fits. From Tables 6 and 14 we see that *rlogin* connections have smaller σ_{orig} 's, larger $\bar{x}_{\text{resp}}$'s, and shorter $\bar{x}_{\text{dur}}$'s, so it is not surprising that the *telnet* models must be scaled. Performance in the tails also often improved when the models were scaled. For the ratio models, however, as with *telnet*, scaling sometimes improved the *rlogin* fits and sometimes did not.

Even after scaling, though, the *rlogin* fits tend not to be as good as those for *telnet*. The arrival fits, however, were

¹²Perhaps because the scientific community still favors mainframe and VMS machines

Figure 28: Tail Summaries for TELNET Models

Dataset	# Conn	# Rej	$\bar{x}_{orig}$	σ_{orig}	max_{orig}	$\bar{x}_{resp}$	σ_{resp}	max_{resp}	$\bar{x}_{dur}$	σ_{dur}	max_{dur}
LBL-1	1,436	2	170B	$\times 3.2$	49KB	3.3KB	$\times 6.6$	0.4MB	251 s	$\times 6.5$	89.3 h
LBL-2	1,617	5	199B	$\times 3.8$	70KB	4.9KB	$\times 7.5$	2.1MB	287 s	$\times 7.0$	143.9 h
LBL-3	3,009	5	211B	$\times 3.6$	44KB	5.8KB	$\times 6.0$	1.4MB	308 s	$\times 6.5$	185.4 h
LBL-4	3,305	13	240B	$\times 3.8$	51KB	6.1KB	$\times 6.1$	17.3MB	306 s	$\times 7.2$	109.3 h
LBL-5	4,303	29	202B	$\times 3.4$	37KB	6.0KB	$\times 6.9$	2.4MB	302 s	$\times 7.3$	170.4 h
LBL-6	4,424	9	201B	$\times 3.7$	246KB	6.3KB	$\times 7.1$	5.7MB	303 s	$\times 8.6$	257.3 h
BC	307	1	177B	$\times 3.3$	7.9KB	5.8KB	$\times 7.3$	0.4MB	462 s	$\times 6.8$	36.8 h
UCB	340	2	185B	$\times 4.3$	23KB	3.0KB	$\times 10.3$	0.7MB	225 s	$\times 6.8$	6.1 h
USC	144	0	208B	$\times 3.8$	9.0KB	4.6KB	$\times 9.6$	0.4MB	308 s	$\times 6.9$	21.2 h
NC	201	0	134B	$\times 3.2$	5.8KB	3.7KB	$\times 8.1$	0.5MB	283 s	$\times 6.1$	4.3 h

Table 14: Summary of RLOGIN Connections

quite good for the LBL datasets, not requiring scaling, and acceptable for BC, UCB, and NC, with scaling required for good fits.

model's fit in the upper tail; and "s" and "S" the same for the scaled model.

D Tail Summaries

Figure 28 shows tail summaries for the various *telnet* models. The text associated with Figure 5 in the main body of the text explains how to read the summaries.

Figure 29 shows similar tail summaries for the *nnntp*, *smtp*, and *ftp* models.

Figure 30 shows tail summaries for the various connection arrival models. A lower-case "u" indicates the fit for the unscaled model in the lower tail; upper-case "U" the same

Figure 29: Tail Summaries for NNTP, SMTP, and FTP Models

Figure 30: Tail Summaries for Arrival Models

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