

Empirically-Derived Analytic Models of Wide-Area TCP Connections*

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

Revised March 1, 1994

Abstract

We analyze 3 million TCP connections that occurred during 15 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. To assess these models we present a quantitative methodology for comparing their effectiveness with that of empirical models such as Tcplib [DJ91]. Our methodology also allows us to determine which random variables show significant variation from site to site, over time, or between stub networks and internetwork gateways. Overall we find that the analytic models provide good descriptions, and generally model the various distributions as well as empirical models.

1 Introduction

In the last few years a number of papers have appeared giving statistical summaries of wide-area traffic on a per-protocol basis [C89, H90, CW91, WLC92, DHS93], an important first step to characterizing WAN traffic. 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 distribu-

tions. Such models are easier both to convey and to analyze. Two key questions are whether analytic models can describe the diverse phenomena found in wide-area traffic as well as empirical models, and whether either type of model faithfully captures the essential characteristics of the traffic.

In this paper we analyze 15 wide-area traffic traces gathered at seven different sites, five “stub” (end-point) 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. Our statistical methodology also allows us to determine which random variables show significant variation from site to site, over time, or between stub networks and internetwork gateways. 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 the findings in the remainder of the paper.

Below, § 2 presents an overview of the traces used in the study, and § 3 gives a discussion of our statistical methodology. § 4 summarizes the models we developed and evaluates their effectiveness. Readers interested in particular protocols will find more detailed summaries and discussions in § 5 through § 8.

A longer, preliminary version of this paper is also available [P93]. In the remainder of this paper we note where, in the interests of brevity, we have relegated details to that report instead.

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

* Appeared in *IEEE/ACM Transactions on Networking*, 2(4), pp. 316-336, August 1994.

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.
Network traffic varies significantly, both over time and more so from site-to-site, not only in traffic mix but in connection characteristics. We believe this variation is the basis for the success of the analytic models; there is enough variation that any model, empirical or analytic, must be a somewhat rough compromise.
The number of data bytes in bulk-transfer traffic (<i>ftpdata</i> , <i>smtp</i> , and <i>nntp</i>) 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.
The ratio between bytes sent by the computer-side of a <i>telnet</i> connection and bytes sent by the user is about 20:1.
Of <i>ftp</i> sessions that are not “failures” (no data transferred), half transfer more than 32 KB, and a sixth transfer more than 500 KB.
The upper tail of <i>ftp</i> “bursts” (§ 8.4) is so large that 2% of the bursts account for 50-80% of <i>all</i> the <i>ftp</i> data bytes.

Table 1: Major Findings

seven 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. We recorded our seven LBL traces using the *tcpdump* packet capture tool [JLM89] running the Berkeley Packet Filter [MJ93]. We used a *tcpdump* filter to capture only those TCP packets with SYN or FIN 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.

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

¹Communications for North Carolina Education, Research and Technology.

Dataset	Packets (days)	Start	End
LBL-1	124M (36)	01Nov90	01Dec90
LBL-2	?	28Feb91	30Mar91
LBL-3	207M (47)	07Nov91	07Dec91
LBL-4	210M (36)	19Mar92	18Apr92
LBL-5	337M (35)	24Sep92	23Oct92
LBL-6	447M (31)	24Feb93	26Mar93
LBL-7	560M (32)	16Sep93	15Oct93

Table 2: Summary of LBL Datasets

by the entire trace.² Each dataset was then trimmed to span exactly 30 days. Very few packets were dropped by the tracing program (always < 15 per million).

Since the LBL datasets span three 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 [P94].

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 traf-

²The statistics missing for the LBL-2 dataset are due to abnormal termination of the tracing program; this termination, however, did not imply any extra-ordinary loss of packets during the 30-day study period.

fic 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” (endpoint) sites, while the latter two represent inter-network traffic (though the authors of [WLC92] characterize the UK side of the UK–US traffic as similar to a large stub site since it comprises only a few hosts).

Site	Starting Time	Duration
Bellcore (BC)	Tue 14:37 10Oct89	13 days
UCB (UCB)	Tue 10:30 31Oct89	24 hours
USC (USC)	Tue 14:24 22Jan91	26 hours
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

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. In general the traces had no packet drops or an unknown number of drops (see [P93] for specifics).

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 [P93]. In addition, we removed from the LBL datasets all traffic between LBL and U.C. Berkeley³. 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 attributed 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 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.

³Including *nntp*, unlike [P94], which keeps the *nntp* traffic.

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.⁴

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). Some of the variation in the mix is due to periodic traffic. For example, the large spike in the LBL-4 *finger* connections, the large jump in *other* connections at LBL-3, the increasing proportion of *ftpctrl* traffic (i.e., the interactive, control side of an *ftp* session), and the large number of *telnet* connections in LBL-7, are all due to periodic traffic. [P94] explores this phenomenon further.

Another factor affecting traffic mix over time (as seen in the LBL datasets) is the large variance of the *nntp* mix, which is due to changes in LBL’s *nntp* peer servers and differences in the rate at which new network news arrives. Again, see [P94] for a discussion.

Regarding the DEC datasets, 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. Finally, the DEC-2 dataset includes part of the Thanksgiving holiday, accounting for the depressed number of 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, explored further in [P94].

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 making up the data shown in Tables 4 and 5.

3 Statistical Methodology

As noted in [P89], one weakness of many network traffic studies to date has been in their use of statistics. Often the studies

⁴The authors also used twenty-minute silences to delimit the end of connections, instead of FIN packets.

Dataset	# Conn	<i>nntp</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-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
LBL-7	606,487	11	34	18	4	15	0.9	4	5	0.9	0.4	6
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>nntp</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
LBL-7	30,910	25	3	51	0.1	2	0.7	0.0	0.4	8	8	1.8
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

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 [P91] 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 shortcomings and to present a general statistical methodology that might serve future work as well.

3.1 Definitions and conventions

We will use a number of terms and concepts taken from statistics. In this section we define the terms and conventions used in the remainder of the paper.

3.1.1 Random variables

For our purposes, we define a *random variable* as a quantity that each time it is measured takes on one of a range of values. Particular values occur with different probabilities. We refer to each separate measurement as an *instance* of the random variable.

An example of a random variable is the number of bytes transferred during an *ftp* session. Another (and closely re-

lated) random variable is the logarithm of this value.

By convention, X represents a generic random variable and x_i the i th instance of X . Unless otherwise stated, we assume there are a total of n instances.

3.1.2 Models

We define a *model* of a random variable as a hypothesized distribution for what values the random variable might take, and with what probability. This definition of model is quite simple, as it assumes that instances of the random variable are independent and identically distributed. Often *correlations* between instances of a random variable are very important, meaning that instances of the random variable might be identically distributed, but not *independent*. For example, previous work has found that large *ftpdata* bursts tend to arrive in clusters [PF94]. While incorporating correlations into models can be very important, doing so is beyond the scope of our study, so we limit ourselves to simple correlation summaries (§ 3.9).

Another important aspect of modeling network connections is modeling the connection *arrival process*. Here we limit ourselves to briefly describing arrival phenomenon such as periodicity. See [PF94] for a more detailed look at the connection arrival processes.

In this paper we distinguish between *empirical* models and *analytic* models. An empirical model, such as Tcplib [DJ91, DJCME92], describes a random variable’s distribution based on the observed distribution of an earlier sample of the variable. For example, the Tcplib models were constructed from the UCB dataset. An analytic model, on the other hand, attempts to capture a distribution in a simple mathematical form. We discuss the advantages of each type of model in § 3.4 below.

For each of the random variables modeled in this paper, we examine three models, one analytic and two empirical. The empirical models were constructed from the UCB and LBL-2 datasets. We refer to these three models as A , U , and L , respectively. As explained above, the U model reflects the behavior of Tcplib.

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 analytic 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, except for UCB and LBL-2, since we used these to construct the U and L empirical models.

3.1.3 Distributions

We define the *distribution* of a random variable as:

$$F(x) = P(X \leq x)$$

That is, $F(x)$ is the probability that an instance of the random variable X takes on a value less than or equal to x .

For our analytic models we draw upon a number of distributions commonly used in statistics. We assume that the reader is familiar with the normal and exponential distributions. Two other distributions we will use are the *extreme* distribution:

$$F(x) = \exp \left[- \exp \left(- \frac{(x - \alpha)}{\beta} \right) \right] \quad (1)$$

and the doubly-exponential *Pareto* distribution:

$$F(x) = 1 - (k/x)^\alpha. \quad (2)$$

The Pareto distribution is noteworthy for having a very heavy upper tail, an important property when considering *self-similarity* in network traffic [LTWW93].

We will be using variants of the normal and extreme distributions called *log-normal* and *log-extreme*; these are discussed in § 3.2 below.

3.1.4 Estimating parameters

All of the distributions mentioned in the previous section are *parameterized* using one or more constants. A normal distribution is parameterized by a mean, \bar{x} , and a standard deviation, σ_x ; an exponential distribution, by a rate λ ; an extreme distribution, by α and β ; and a Pareto distribution, by α and k (k is the lower bound of X).

Sometimes when using an analytic model, the parameterization constants are known in advance. Other times, they must be *estimated* from the same data that we are trying to describe using the analytic model. For a normal distribution, the mean is estimated using:

$$\hat{x} = \sum_{i=1}^n x_i / n \quad (3)$$

and the standard deviation using:

$$\hat{\sigma}_x = \sqrt{\sum_{i=1}^n \frac{(x_i - \hat{x})^2}{(n - 1)}} \quad (4)$$

We often use the \hat{x} and $\hat{\sigma}_x$ estimates to evaluate a distribution even if it is not normal.

For an exponential distribution,

$$\hat{\lambda} = 1 / \hat{x}.$$

For an extreme distribution, α and β can be estimated using an iterative method [DS86], and α and k for a Pareto distribution using a simple least-squares technique [CM80].

When estimating parameters, we will generally drop the “hat” notation (e.g., use \bar{x} instead of \hat{x}). An exception is when discussing the λ^2 discrepancy measure (§ 3.5 below), where the fact that we are only estimating its value becomes an important consideration.

Finally, sometimes when estimating parameters we want to ignore part of the distribution we are using to compute the estimation. For example, below we model the number of bytes sent by a *telnet* responder as a log-normal distribution, but the fit between the model and the dataset is only good for the upper 80% of the distribution (see § 5.3). If we compute \hat{x} from the entire dataset’s distribution, we will spoil the upper-80% fit because of the disagreement between the model and the distribution in the lower 20%. Instead, we *truncate* the distribution by ignoring its lower 20%, and then estimate \hat{x} and $\hat{\sigma}_x$ from the remainder using methods given in [DS86].

3.1.5 The notion of “scaling”

When possible, we would like to avoid having to estimate parameterization constants for our models (discussed in the previous section). For example, consider the problem of modeling the bytes transferred during an *ftp* session. Suppose Model-1 describes the distribution of bytes transferred as log-normal with $\bar{x} = 15$ and $\sigma_x = 4$, but Model-2 simply describes the distribution as log-normal, with a note that the values of \bar{x} and σ_x must be estimated from each dataset to which we want to apply the model. We would prefer Model-1 because it tells us *in advance* what to expect; Model-2 only tells us the general *shape* of what to expect, but not the exact quantities. We can use Model-1 to make quantitative predictions of what we will measure in the future, but Model-2 can only make qualitative predictions.

We will refer to models like Model-1 as *unscaled*, and those like Model-2 as *scaled*. In general, we use the term *scaling* to refer to tailoring a model to fit a dataset by estimating parameters of the model from the dataset. If we find that an unscaled model gives us good fits to many different datasets, then we have reason to believe that the model captures an “invariant” distribution. Such models are particularly powerful because they allow confident prediction of future distributions. Sometimes, however, a scaled model gives us significantly better fits to different datasets than an unscaled version of the model. In this case, the distribution’s general shape (e.g., log-normal) might be invariant, but the particulars of the shape vary. Scaled models are less powerful than unscaled models because they allow less complete predictions, but are still valuable because with them we can explore possible behavior given separate hypotheses as to the values of the model’s parameters.

In § 4 below we discuss how we chose whether to use the unscaled or scaled version of each of our models.

While estimating parameters applies only to analytic models, we can perform an analogous operation on empirical models. Suppose we are modeling a random variable X using an empirical distribution Y . If \bar{x} is significantly different from \bar{y} , or σ_x from σ_y , then there is little hope that the empirical model faithfully describes X ’s distribution. But it may be that if Y were adjusted to have the same mean and standard

deviation as X , then it would also describe X ’s distribution well. This adjustment is easy to make. If we define a new empirical distribution:

$$Y' = \frac{\sigma_x}{\sigma_y}(Y - \bar{y}) + \bar{x} \quad (5)$$

then Y' keeps the same general shape as Y (since we have merely applied a linear transformation to Y) but has the same mean and standard deviation as X . When discussing empirical models, we will use the term *scaling* to refer to the transformation given in Equation 5.

When developing our unscaled analytic models, we picked for each model parameter a round value lying somewhere in the range the parameter exhibited in the LBL test datasets (see [P93] for details regarding the parameter ranges). 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 not particularly exact.

3.1.6 Comparing estimates

Suppose we have two estimated quantities, \hat{a} and \hat{b} , and we want to compare them to see which is smaller. Because the quantities are estimates, we would rather not make the comparison on the basis of testing whether $\hat{a} < \hat{b}$, since perhaps the error in estimating \hat{a} and \hat{b} is large enough that the comparison will be misleading. We can make a more meaningful comparison if we have a quantitative possible-error associated with \hat{a} and \hat{b} (this becomes relevant in § 3.5 below).

A natural measure of possible error in an estimate is a standard deviation. Suppose in addition to \hat{a} and \hat{b} we have standard deviations σ_a and σ_b , which quantify the error in the estimates. Then we can define a comparison operator $<_\sigma$ as follows:

$$\hat{a} <_\sigma \hat{b} \quad \text{iff} \quad \hat{a} + \sigma_a < \hat{b} - \sigma_b$$

Using this comparison operator, we will find a less than b only when the difference between \hat{a} and \hat{b} is greater than what we can account for by the uncertainty in their estimates. Similarly, we define $\hat{a} >_\sigma \hat{b}$ iff $\hat{b} <_\sigma \hat{a}$.

If neither $\hat{a} <_\sigma \hat{b}$ nor $\hat{b} <_\sigma \hat{a}$ then we say that a and b are *unordered*, rather than “equal”, to stress that one of them may in fact be smaller than the other, but we are unable to say so conclusively.

3.2 Logarithmic transformations

In this and the next section we discuss how we transformed the data prior to analysis, including dealing with outliers.

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, which we will sometimes write as $\lg x$.

If the random variable $Y = \log X$ has a normal distribution, then X is said to have a *log-normal* distribution. Similarly, if Y has an extreme distribution, X has a *log-extreme* distribution. (We will often write these distributions as \log_2 -normal and \log_2 -extreme as a reminder that all logarithms in this paper are taken base 2).

With a random variable X that has a large range of values, the computed mean (Equation 3) and standard deviation (Equation 4) are greatly skewed by the largest of the x_i . The mean and standard deviation of the transformed quantity $Y = \log_2 X$ do not have this problem, though, since the logarithmic transformation greatly reduces the range of values. With these types of random variables, it is generally more meaningful to analyze \bar{y} and σ_y than \bar{x} and σ_x .

We can then attach interpretations to the quantities $2^{\bar{y}}$ and 2^{σ_y} . In particular, the *geometric mean* of X is defined as:

$$\text{Geometric mean}(X) = \sqrt[n]{\prod_{i=1}^n x_i} \quad (6)$$

and it is easy to show that this is the same as $2^{\bar{y}}$. This equivalence between the geometric mean and $2^{\bar{y}}$ suggests an analogous definition for the *geometric standard deviation*, which we define as:

$$\text{Geometric std. dev.}(X) = 2^{\sigma_y} \quad (7)$$

If Y is normally distributed, then σ_y characterizes the range of Y . For example, about 68% of the distribution of Y will reside in the range $\bar{y} \pm \sigma_y$. We then can interpret the geometric standard deviation of X as giving an analogous range for X . If X is log-normally distributed, then about 68% of the distribution of X resides in the range $2^{\bar{y} \pm \sigma_y}$. More explicitly:

$$\left(\frac{2^{\bar{y}}}{2^{\sigma_y}}\right) \leq 68\% \text{ of the range of } X \leq (\bar{2} \times 2^{\sigma_y}) \quad (8)$$

In the tables summarizing the different protocols below, when we report quantities such as \bar{x} and σ_x , they reflect $2^{\bar{y}}$ and 2^{σ_y} ; that is, they give values computed using Equations 6 and 7 and not values computed using Equations 3 and 4. As a reminder of this fact, we precede all σ_x values with a “ \times ” symbol, in keeping with their interpretation given in Equation 8. So, for example, a value of $\sigma_x = \times 8$ indicates $\sigma_y = \sigma_{\lg x} = \lg 8 = 3$.

3.3 Dealing with outliers

When applying a logarithmic transformation to non-negative data, one immediately runs into the problem of what to do with data values equal to zero, since after a logarithmic transformation these become $-\infty$. 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. The appendices of [P93] report the number of connections thus eliminated for each dataset; in the worst case they comprised 0.5% of the total connections.

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). Again, see [P93] for a discussion of these outliers.

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

3.4 Empirical vs. analytic models

In this section we look at the relative advantages of empirical and analytic models, which motivates our subsequent pursuit of analytic models.

The main advantage of empirical models is that they are known to fully reflect a portion of Real World behavior. If there are consistent spikes in a distribution or even subtle deviations from “smooth” behavior, the empirical model will capture these nuances if they were present in the dataset from which the model was derived. An analytic model might easily miss these characteristics.

There are, however, several advantages of analytic models compared to empirical models:

- 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 corresponding estimates for the analytic model’s parameters (§ 3.1.4).

While these advantages are certainly attractive, the crucial issue remains whether an analytic model truly 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 interest, and it becomes useless (or, even worse, 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 given dataset. One approach is to require that the distribution predicted by the analytic model and that actually measured from the dataset

be indiscernable in a statistical sense. A large body of literature examines techniques for testing for such statistical exactness (an excellent reference is [DS86]).

In our earlier work we tried to find statistically exact models but failed (see [P93] for details). This failure, however, is not surprising: it is well known in the statistics community that large datasets almost never have statistically exact descriptions [M74].⁵ The next section addresses how to deal with this failure.

3.5 Measuring discrepancy

Even if a model is not statistically exact, we can still attempt to gauge how *close* it is to the distributions it endeavors to describe. To do so, we turn to techniques for measuring *discrepancy*.

One widely-used technique for doing so is based on a modified χ^2 test [M84]. To understand it, we first review the χ^2 test itself.

Suppose we have observed n instances of a random variable Y which we want to model using another model distribution Z . We partition the distribution Z into N bins. Each bin has a probability p_i associated with it, which is the proportion of the distribution Z falling into the i th bin. Let Y_i be the number of observations of Y that actually fell into the i th bin. Then one computes the statistic:

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

The χ^2 test for a statistically-exact fit involves testing where the X^2 statistic falls in the range of a corresponding χ^2 distribution (the exact details of the comparison are secondary to understanding the remainder of this section).

The χ^2 discrepancy measure is then simply X^2/n . This is essentially the measure used in our preliminary work [P93]. There is, however, a problem with using X^2/n as a discrepancy measure [PJ90]. In general, the optimal value of N (the number of bins) to use when computing X^2 varies with the size of Y and its standard deviation σ_y (see § 3.6 below). The X^2/n discrepancy measure, however, cannot be used to compare discrepancies for different values of N .

Pederson and Johnson [PJ90] describe a related discrepancy measure, λ^2 , which corrects the X^2/n measure so that λ^2 can be used to compare discrepancies for different values of N . They also give a way to compute σ_λ , the standard deviation associated with estimating $\hat{\lambda}^2$ for a particular dataset Y and model distribution Z . Knowing σ_λ lets us use the $<_\sigma$ operator (§ 3.1.6) for comparing estimated discrepancies. Given two models, Z_1 and Z_2 , we can then state in a meaningful way whether Z_1 is a better description of Y than Z_2 (if we find that

⁵As well as giving a general discussion of this problem, [M74] also analyzes an experiment in which 26,306 throws of 12 dice failed a χ^2 test for an exact fit to the predicted binomial distribution.

$\lambda^2_{Z_1} <_\sigma \lambda^2_{Z_2}$), or that Z_2 is better than Z_1 ($\lambda^2_{Z_2} <_\sigma \lambda^2_{Z_1}$), or that the two models are *unordered*, indicating that they are roughly equal.

We now summarize how to compute λ^2 for assessing the discrepancy between a random variable Y and a model distribution Z . First, let $E_i = np_i$ be the expected count for the i th bin, and $D_i = Y_i - E_i$ be the discrepancy in the i th bin. Then define:

$$K = \sum_{i=1}^N \frac{D_i}{E_i}. \quad (9)$$

(Note that K is quite similar to X^2 except the numerator of the summation is not squared.) We then define:

$$\hat{\lambda}^2 = \frac{X^2 - K - \text{df}}{n - 1},$$

where “df” is the number of *degrees-of-freedom* in computing X^2 and K . For our purposes, $\text{df} = N - 1 - \text{Est}$, where “Est” is the number of parameters estimated from Y (0 for unscaled models, 1 or 2 for scaled models; see § 3.1.4-3.1.5 above).

The *variance* associated with this estimate of λ^2 is given by:

$$\hat{v}(\hat{\lambda}^2) = [2\text{df} + 4n\hat{\lambda}^2 + 4n\hat{\lambda}^4 + 4T]/n^2,$$

where:

$$T = \sum_{i=1}^N [\hat{E}_i - 2D_i/E_i + \frac{5}{2}D_i^2 + \frac{3}{2}(D_i + E_i)] / E_i^2$$

([PJ90] states that this expression for T is not quite exact when the parameters of Z are estimated from Y , but they found in practice the correction makes little difference.)

The standard deviation associated with estimating λ^2 is then:

$$\sigma_\lambda = \sqrt{\hat{v}(\hat{\lambda}^2)}.$$

3.6 Considerations when measuring discrepancy

There are a number of considerations when evaluating models using a discrepancy measure such as λ^2 : how to pick the number of bins to use; how to capture significant discrepancies in the distribution’s tails; what to do about significant “spikes” that fall into a single bin; and how to deal with dataset values not falling into any of the bins. We address each of these in turn.

When computing λ^2 , we are forced to make a choice as to how many bins to use (i.e., the value of N in Equation 9). If N is too small then we will be measuring discrepancies only on a gross scale, and similarly if N is too large then we will be sensitive to quite small discrepancies which perhaps are of no real interest. Fortunately, statistics provides some guidance. Scott has shown that to minimize the mean-square

error in approximating a distribution using fixed-sized bins, the bin-width should be:

$$w = 3.49\hat{\sigma}_x n^{-1/3} \quad (10)$$

where $\hat{\sigma}_x$ is the estimated standard deviation of the distribution (Equation 4), n is the number of instances in the distribution, and w the bin-width [S79]. Given the range of the distribution, it is straightforward to compute the value of N to use such that each bin has width w . Fortunately, the value of w is not strict; any value close to w provides a satisfactory estimate of the distribution, so we rounded N to the nearest multiple of 5.

One important point we found, though, was to compute σ_x after first applying the logarithmic transformation to X (§ 3.2). That is, we used for σ_x the geometric standard-deviation (Equation 7) and not the value given in Equation 4. In general, our untransformed data had such large ranges that using the untransformed values of X in Equation 10 resulted in a very large number of bins.

One other point is that when using Equation 10 to compute N for an empirical model, we had to decide between computing w using $\hat{\sigma}_x$ and n from the dataset being modeled, or from the empirical model. We chose whichever had the smaller n .

The number of bins tested ranged from 5 to 240; the average was around $N = 35$.

A related consideration with using the λ^2 discrepancy measure is that when using constant-sized bins the measure does not give any particular emphasis to the distribution’s tails, which sometimes are the most important aspect of a distribution. We address this consideration in § 3.8 below.

A third consideration is that the discrepancy measure does not inform us of 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 consideration 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).

A final consideration is that since an empirical model has bounds on the range of values it allows for, the tested dataset may have values outside the range of any bin. We removed such values from the dataset prior to computing its fit to the model. We did, however, include these values when evaluating the model’s tail discrepancy (§ 3.8).

3.7 Testing for significant differences

The λ^2 discrepancy measure and the $<_\sigma$ operator allow us to determine whether, given two datasets and a model, the model is significantly better at modeling one of the datasets than the other. They also allow us to determine, given one dataset and

two models, which of the models (if either) is significantly better at modeling the dataset.

In this section we build on these techniques to develop methods for comparing datasets and models in a general sense. The first method allows us to determine with high confidence whether a given model is better at describing one set of datasets than it is at describing a second set. The second method allows us to determine with high confidence whether one model is better than another model at describing a set of datasets.

3.7.1 Method I

In this section we describe Method I, a method for determining whether a given model is significantly better at modeling one set of datasets than at modeling another.

Suppose we have a model Y and two sets of datasets, D_1 and D_2 , with m and n datasets in each set, respectively. For example, Y might be a model of the number of bytes transferred during an *smtp* connection, D_1 the set of LBL datasets, and D_2 the set of non-LBL datasets. We would like to determine whether model Y does substantially better at modeling D_1 than at modeling D_2 , because we suspect that *smtp* connections may show considerable variation between network sites. If our suspicion is correct, and if model Y was derived from some of the LBL datasets, then we would expect to find that Y is significantly better at modeling D_1 than D_2 .

For each dataset in D_1 and D_2 we can compute the corresponding value of λ_Y^2 , the discrepancy between the model and the dataset. It is not obvious, though, how to take this collection of discrepancies and reduce it to a simple statement that Y does or does not appear to model D_1 significantly better than D_2 .

The approach we take is as follows. We assume the null hypothesis that Y performs equally well when modeling D_1 as when modeling D_2 .⁶ We compare each dataset d_1 in D_1 against each dataset d_2 in D_2 . Let l be the number of comparisons for which we found $d_1 <_\sigma d_2$, g the number of times $d_1 >_\sigma d_2$, and u the number of times two datasets were unordered. There are m datasets in D_1 and n in D_2 , so we have $l + g + u = mn$. We now restrict ourselves to just the $t = l + g$ comparisons that were not unordered. If the null hypothesis is correct, then the probability that in t tests we would find l instances for which $d_1 <_\sigma d_2$ is:

$$P(\text{In } t \text{ trials, } d_1 <_\sigma d_2 \text{ occurs } l \text{ times}) = \binom{t}{l} 2^{-t}$$

We can easily generalize the above to find the probability that $d_1 <_\sigma d_2$ occurs at least l times, rather than exactly l times. For Method I, then, we pick a value of k such that, given the null hypothesis, the probability that $d_1 <_\sigma d_2$ occurs

⁶Here “equally well” means that given datasets d_1 from D_1 and d_2 from D_2 , if Y models one better than the other (in a $<_\sigma$ sense), then the probability that it models d_1 better than d_2 is 1/2.

at least k times is $\leq 5\%$. We then test whether indeed $d_1 <_{\sigma} d_2$ occurs k or more times. If so, then with 95% confidence we declare that the null hypothesis is incorrect and that Y does in fact model D_1 better than D_2 . If, however, $d_1 <_{\sigma} d_2$ occurs fewer than k times, then the test is inconclusive, and we refrain from ruling out the null hypothesis that Y is equally good (or bad) at modeling both D_1 and D_2 .

One concern when applying Method I is: what if u , the number of times d_1 and d_2 were unordered, is large relative to $l+g$? For example, if $u = 95$, $l = 5$, and $g = 0$, then Method I will declare that Y models D_1 better than D_2 , even though it might be more reasonable to say that it basically models the two equally well. For our purposes, this did not turn out to be a problem: when Method I (or Method II; see below) declared a significant difference, we always had $u < \frac{4}{3}(l + g)$.

3.7.2 Method II

In this section we describe Method II, a method for determining whether one of two models is significantly better than the other at modeling a given set of datasets. It is quite similar to Method I.

Suppose we have a set D of n datasets, and two models, Y and Z . For example, D might be all of the datasets listed in Tables 2 and 3, Y might be an analytic model for bytes transferred in an *smtp* connection, and Z an empirical model of the same random variable. We would like to determine whether either of the models is significantly better than the other for describing D .

For each dataset in D we can compute two discrepancy values, λ_Y^2 and λ_Z^2 . We want to take this collection of discrepancies and reduce it to a statement that Y or Z (or neither) appears significantly better than the other at modeling the datasets in D .

We take a similar approach to that described for Method I. We assume the null hypothesis that Y and Z perform equally well when modeling D . Now, however, for each of the n datasets in D we compare the corresponding values of λ_Y^2 and λ_Z^2 . Let l be the number of times $\lambda_Y^2 <_{\sigma} \lambda_Z^2$, g the number of times $\lambda_Y^2 >_{\sigma} \lambda_Z^2$, and u the number of times the comparison is unordered. We consider just the comparisons that were not unordered. Let $t = l + u$. If the null hypothesis is correct, we have:

$$P(\text{In } t \text{ trials, } \lambda_Y^2 <_{\sigma} \lambda_Z^2 \text{ occurs } l \text{ times}) = \binom{t}{l} 2^{-t}$$

We again generalize the above to find the probability that $\lambda_Y^2 <_{\sigma} \lambda_Z^2$ occurs at least l times, and for Method II pick a corresponding value of k such that, given the null hypothesis, the probability that $\lambda_Y^2 <_{\sigma} \lambda_Z^2$ occurs at least k times is $\leq 5\%$. As before, this test, if successful, lets us state with 95% confidence that the null hypothesis is incorrect and that Y does in fact model D better than Z does. If the test is inconclusive, however, we refrain from ruling out the null hypothesis that Y and Z are equally good (or bad) at modeling D .

3.7.3 A note on Methods I and II

One important point regarding Method I and Method II is that they are fairly conservative. Their use of both the $<_{\sigma}$ comparison and 95% confidence levels assures that differences they uncover are very likely to be significant. It is quite possible for the tests to fail to declare two unequal models as being different, but it is not likely that they will erroneously declare two equal models as being different.

3.8 Evaluating deviation in the tails

Often a distribution’s behavior in its lower or upper tail can be crucially important. For example, as mentioned in Table 1 and developed in § 8.4 below, for the distribution of bytes in *ftpdata* “bursts”, the upper 2% tail is so heavy that large-though-rare bursts will often completely dominate *ftp* traffic. As discussed in § 3.6 above, the \bar{X} discrepancy measure does not give any special weight to agreement with a distribution’s tails. In this section we present a simple way of qualitatively evaluating how well a model captures a distribution’s tails.

Suppose we test the model against n datasets. For the i th dataset, let a_i be the number of instances predicted to lie in a given tail, and b_i be the number actually found to do so. Define:⁷

$$\xi_i = \log_2 \frac{a_i}{b_i}$$

Positive values of ξ_i indicate that the model *overestimates* the tail, and negative values that it *underestimates* the tail.

With this definition, an underestimate by a factor of two ($a_i/b_i = 1/2$, $\xi_i = -1$) is considered just as bad as an overestimate by the same factor ($a_i/b_i = 2$, $\xi_i = 1$).

A value of $\xi_i = 0$ indicates that the model perfectly captures the i th dataset’s tail. With this in mind, we then compute:

$$\sigma_{\xi} = \sqrt{\sum_{i=1}^n \xi_i^2 / n} \quad (11)$$

σ_{ξ} then represents the standard deviation, from a mean of 0, of the model’s accuracy in the tail. If we find $\sigma_{\xi} < 1$ then the model typically predicts the tail population within a factor of two; we deem this “acceptable”. If $1 \leq \sigma_{\xi} < \log_2 5$, then the model typically errs in predicting the tail by a factor between 2 and 5; we deem this “bad”. Similarly, if $\sigma_{\xi} \geq 5$, then we evaluate the model’s tail behavior as “very bad”.

These definitions of “acceptable”, “bad”, and “very bad” appear quite generous; after all, one might wonder how a model could possibly misrepresent a dataset’s tail by more than a factor of 5. Yet it turns out that for the extreme 1% tails, a fair number of our models are evaluated as “bad” or “very bad”; see § 4 below.

⁷One problem arises when using this definition of ξ_i : if b_i is 0 then ξ_i becomes undefined. We address this problem by replacing b_i with 0.5.

One final note in evaluating tails. For models describing bytes transferred, we only evaluate the upper tails, as in these cases disagreement in the lower tails is of little consequence, 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. See [P93] for more details concerning the models’ tail-behavior.

3.9 Evaluating correlation

As noted in § 3.1.2, it is often important to model the correlations between instances of a random variable. While doing so is beyond the scope of this paper, in § 4 we present a simple summary of each random variable’s correlation, computed as follows.

The *autocorrelation function* (which we abbreviate as $\gamma_x(l)$, but do not further define here) of a random variable X measures the degree to which instances of X are correlated. For a given value of l (called the “lag”), $\hat{\gamma}_x(l)$ is a number ranging from -1 to 1. A value of $\gamma_x(l)$ close to 1 indicates that if x_k , the k th instance of X , is higher (lower) than \bar{x} , then x_{k+l} will also tend to be higher (lower). A value of $\gamma_x(l)$ close to -1 indicates that x_k and x_{k+l} are anti-correlated; if x_k is high then x_{k+l} will tend to be low, and vice versa. A value of $\gamma_x(l)$ close to 0 indicates that x_k and x_{k+l} are not linearly correlated; to first order, knowing the value of x_k does not help in predicting the value of x_{k+l} .

With many random variables, $\gamma_x(1)$ is a particularly significant value of γ_x , because if a random variable is correlated, often the correlation is greatest at a lag of 1. For example, a high value of $\gamma_x(1)$ indicates that successive instances of X tend to have comparable values.

For our summary of correlation in Table 6 below, we computed $\gamma_x(1)$ for each random variable and dataset. We looked at how often $|\gamma_x(1)| > 2/\sqrt{n}$ (where n is the number of instances in the dataset), because if X is uncorrelated then $|\hat{\gamma}_x(l)|$ will exceed this value only 5% of the time. If we found $|\gamma_x(1)| > 2/\sqrt{n}$ occurred for more than half of the datasets, then we considered the random variable to be correlated. (This happened for all but one of the random variables.) We then looked at the magnitude of the mean lag-1 autocorrelation, $|\hat{\gamma}_x(1)|$. If it was ≤ 0.1 , we considered the random variable to be “weakly correlated”, otherwise “significantly correlated”. Finally, we looked at the range of values of $\gamma_x(1)$. If they were significantly positive (measured by a modified Method II test), we considered the correlation to be positive. Otherwise (since they were never significantly negative), we considered the correlation to be “undirected”.

4 Summary and Evaluation of Models

Using the methodology described in § 3, we constructed analytic and empirical models of random variables associated

with wide-area *telnet*, *nntp*, *smtp*, and *ftp* TCP connections. As discussed in § 3.1.2, we generically refer to the analytic model for a random variable as A , the UCB-derived empirical model (reflecting Tcplib; [DJ91]) as U , and the empirical model derived from LBL-2 as L .

Table 6 summarizes the random variables and the corresponding analytic models. The “Variable” column lists the random variable being modeled and the “Abbr.” columns the label and short name we will use subsequently to identify the variable. The “§” column lists the section in the paper that develops the model in detail.

The “model” column lists the analytic models used to describe the random variable’s distribution. Almost all first apply a \log_2 transformation to the data, as described in § 3.2. One model is \log_2 -extreme, where the “extreme” distribution is defined by Equation 1; one is exponential; one is Pareto (Equation 2); and the remainder are \log_2 -normal.

Five of the models have restrictions. The *telnet* responder bytes model describes only the upper 80% of the responses. The *telnet* E and F 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 *telnet* connections. The second model describes this ratio for those connections in the upper 10% of all responses. The *smtp* H model uses parts of two different \log_2 -normal distributions in its description of the bytes transferred by the *smtp* connection originator. The lower 80% of the originator distribution is modeled using the lower 80% of the first \log_2 -normal distribution; similarly, the upper 20% is modeled using the upper 20% of the second \log_2 -normal distribution. Finally, as developed in § 8.4 below, for *ftpdata* connection bursts it is crucial to accurately model the upper tail, so the *ftp* K model describes only the upper 5% of the distribution of bytes in *ftp* bursts.

The “Parameters” column gives the parameters we used for the unscaled (§ 3.1.5) version of the model, after applying \log_2 transformations. Parameters listed using “ \approx ” instead of “=” correspond to models that performed significantly better when scaled than when unscaled (see below). For these models, the parameters given in Table 6 should only be used with considerable caution.

The final column in Table 6, “Corr.,” gives a simple summary of the correlations present in the datasets, using the methodology given in § 3.9 above. A “+” sign indicates weak positive correlation, “++” indicates significant positive correlation, “±” indicates weak undirected correlation, and “no” indicates that the random variable does not appear to be correlated. Two values are given for the F_{burst} variable; the first reflects the correlation of all of the instances of the variable, and the second, the correlation of just the top 5% tail.

We see that successive *telnet* connections are weakly correlated, perhaps due to small hour-to-hour variations in the characteristics of *telnet* connections. Successive bulk-transfer

Proto.	Variable	Abbr.		§	Model	Parameters	Corr.
<i>telnet</i>	originator bytes	A	T_{orig}	5.2	log ₂ -extreme (Eqn 1; § 3.2)	$\alpha \approx \log_2 100; \beta \approx \log_2 3.5$	+
	responder bytes	B	T_{resp}	5.3	log ₂ -normal, 80-100%	$\bar{x} = \log_2 4500; \sigma_x = \log_2 7.2$	+
	duration secs.	C	T_{dur}	5.4	log ₂ -normal	$\bar{x} = \log_2 240; \sigma_x = \log_2 7.8$	+
	resp. / orig.	D	T_{ratio}	5.5	log ₂ -normal	$\bar{x} = \log_2 21; \sigma_x = \log_2 3.6$	+
	resp. / dur.	E	T_{rd1}	5.6	exponential, 0-90% resp.	$\lambda \approx 1/30$	+
	resp. / dur.	F	T_{rd2}	5.6	log ₂ -normal, 90-100% resp.	$\bar{x} = 5.3; \sigma_x = 1.5$	±
<i>nntp</i>	originator bytes	G	N_{orig}	6.2	log ₂ -normal	$\bar{x} \approx 11.5; \sigma_x \approx 3$	++
<i>smtp</i>	originator bytes	H	S_{orig}	7.2	log ₂ -normal + 300B, 0-80%;	$\bar{x} \approx 10; \sigma_x \approx \log_2 2.75$	++
					log ₂ -normal + 300B, 80-100%	$\bar{x} \approx 8.5; \sigma_x \approx \log_2 3$	
<i>ftp</i>	connection bytes	I	F_{conn}	8.2	log ₂ -normal	$\bar{x} \approx \log_2 3000; \sigma_x \approx 4$	++
	session bytes	J	F_{sess}	8.3	log ₂ -normal	$\bar{x} = 15; \sigma_x = 4$	no
	burst bytes	K	F_{burst}	8.4	Pareto (Eqn 2), 95-100%	$\alpha \approx 1; k \approx 10^{5.5}$	++ , ±

Table 6: Summary of Analytic Models of Connection Characteristics

Model		Scaling			Variation			Ordering	Range mean λ^2	1% Tail analytic dev.
		A	U	L	site	wide	time			
A	T_{orig}	✓	✓	✓	✓	✓	✓	$U < \{A, L\}$	0.05-0.09	$u: ++$
B	T_{resp}		✓		✓	✓	✓	$\{A, U\} < L$	0.04-0.07	$u: +$
C	T_{dur}		✓		✓	✓		none	0.04-0.09	$l: +, u: \pm\pm$
D	T_{ratio}				✓	✓	✓	$U < \{A, L\}$	0.08-0.19	
E	T_{rd1}	✓			✓	✓	✓	$\{U, L\} < A$	0.02-0.10	
F	T_{rd2}			✓	✓	✓		none	0.06-0.10	$l: \pm, u: \pm$
G	N_{orig}	✓			*		✓	$\{U, L\} < A$	0.36-2.00	$u: ++$
H	S_{orig}	✓	✓	✓	✓			$U < \{A, L\}$	0.15-0.34	
I	F_{conn}	✓	✓		†	†		$U < A$	0.18-0.27	$u: +$
J	F_{sess}				†	†		none	0.09-0.14	$u: +$
K	F_{burst}	✓	✓	✓	†			$L < \{A, U\}$	0.16-0.47	$u: +^\dagger$

Table 7: Evaluation of Models

connections, on the other hand, all tend to be significantly correlated, except that the size of complete *ftp* sessions is not correlated. These correlations are not hard to explain: *nntp* connections will tend to be correlated as new network news first arrives in an inbound connection and then is propagated soon afterward in an outbound connection. *smtp* connections can be correlated due to mailing lists expansions generating similar connections one after another, or inbound connections spawning outbound connections due to mail forwarding. *Ftp*-data connection sizes will tend to be correlated because directory listings are substantially smaller than file transfers, so back-to-back directory listings will contribute to positive correlation. This effect will also contribute to correlation in the size of *ftp* bursts, if directory listings occur more than 4 seconds apart (§ 8.4) and thus constitute separate bursts. That the size of *ftp* sessions does not appear correlated suggests that *ftp* sessions are statistically independent, in line with the finding in [PF94] that *ftp* session arrivals appear well-modeled as Poisson processes.

Table 7 summarizes our evaluation of the different models we constructed. A check (✓) in one of the “Scaling” columns indicates that we found (using Method II, § 3.7.2) that the scaled version (§ 3.1.5) of the corresponding model performed significantly better than the unscaled version. When this was the case, for the remainder of our evaluation we used the scaled version of the model; otherwise, we used the unscaled version. As noted in § 3.1.5, a successful unscaled model suggests that the random variable being modeled is in some sense “invariant”, and we can make strong predictions about future behavior using such a model.

The “Variation” columns summarize three tests for significant variation among the performance of the models. A check in the “site” column means that the *A* and *L* models did significantly better (as tested using Method I; § 3.7.1) in modeling the LBL datasets than in modeling the non-LBL datasets. From this we infer that the associated random variable shows significant site-to-site variation, as models derived from one site reflect that site significantly better than other sites.

A check in the “wide” column means that the A and L models did significantly better (using Method I) in modeling the “stub” sites than in modeling the internetwork gateway sites (the UK and NC datasets). From this we infer that there is significant variation in the corresponding random variable when modeling “very wide” traffic as opposed to “less wide” traffic, perhaps due to a richer degree of connection multiplexing.

A check in the “time” column means that the A and L models did significantly better (Method I) modeling the second fifteen days of the LBL-1, LBL-3, and LBL-4 datasets, than in modeling the complete LBL-5, LBL-6, and LBL-7 datasets. From this we infer that the corresponding random variable changes significantly over time, since the A model was derived from the first fifteen days of the first four LBL datasets (the “test datasets”; § 3.1.2) and the L model from the LBL-2 dataset, one of the early LBL datasets.

We have marked the G variable’s site variation with a “*” because while the Method I test did not indicate a significant difference between the LBL sites and the non-LBL sites, we believe this is simply because there is so much variation among the LBL datasets themselves (as indicated in part by the “time” checkmark) that it exceeds the considerable site-to-site variation. See Figure 4 below for an illustration of the large site-to-site variation.

We marked the I , J , and K variable’s “site” and “wide” entries with “†”s because of the following curious phenomenon: we found that the A and L models did significantly *worse* (as indicated by a Method I test) modeling the LBL datasets than modeling the non-LBL datasets. We do not have a firm explanation for this behavior. Evidently the LBL datasets are substantially “noisier” than the non-LBL datasets, perhaps because there are unusually high variations or spikes in the size of files transferred by scientists, or perhaps because the 30-day length of the LBL dataset allows more opportunity for rare phenomena to skew the distribution.

The “ordering” column gives the results of using Method II to compare the effectiveness of the A , U , and L models. An entry like “ $U < \{A, L\}$ ” indicates that the Method II test found that the U model performed significantly less well than both the A and L models, but that the A and L models are unordered. An entry of “none” indicates that all three models were unordered. We note that except for variable B , the A model does as well as better than both of the empirical models, indicating that analytic models can perform as well as empirical models for describing wide-area connection characteristics.

While the “ordering” column gives an indication as to the relative performance of the three models, the “Range” column gives an indication as to absolute performance. Here we list the range of $\bar{\lambda}^2$, the average value of λ^2 for each of the three models. For example, for variable A , the best model (which must have been either A or L) had an average λ^2 value of 0.05, and the worst (U), 0.09. Because λ^2 has the property that it can be meaningfully compared for different

models, the variation in the “Range” column tells us a good deal about how each model performs. For example, we see that of the four protocols, *telnet* is consistently modeled the most successfully, even by the worst-performing of the three models. *Nntp*, on the other hand, is the most poorly modeled, not at all surprising given the irregular distributions of *nntp* originator bytes (see § 6.2 and Figure 4 below). *Smtplib* is not well-modeled either, though considerably better than *nntp*. Finally, *ftp* connection and burst sizes are modeled about as well as *smtplib* originator sizes, but *ftp* session sizes appear as well modeled as *telnet* connections.

The final column in Table 7 summarizes the deviation of the A model when describing the lower and upper 1% tails of the random variable’s distribution. Here we have use the methodology discussed in § 3.8 above. A blank entry indicates that the A model’s description of the upper (and lower, if appropriate) 1% tail is “acceptable”. A “+” indicates that the description is “bad” and that the A model consistently *overestimates* the tail. A “++” indicates that the description is “very bad” and also consistently an overestimation. “±” and “±±” are analogous to “+” and “++” except the model sometimes overestimates the tail and sometimes underestimates it. Finally, the K variable’s tail is marked with a “‡” as a reminder that the F_{burst} model only models the upper 5% of *ftpdata* connection bursts, so the upper 1% tail of this model reflects only on the upper 0.05% tail of the entire distribution.

While we have only reported the tail of the A model, the U and L models have similar sorts of problems. Of the three models, the A model is not, overall, especially good or bad in describing the tails.

We also tested the 10% tails using the same methodology. We found that all of the 10% tails were “acceptable”.

The next four sections develop in greater detail the models summarized in Tables 6 and 7 above. One aspect of wide-area TCP connections we do not discuss in this paper is the connection arrival process (other than to note the presence of periodic patterns). Instead, we refer the reader to [PF94].

5 TELNET

We now turn to analyzing the characteristics of individual protocols and developing models to describe them. We begin with *telnet*.⁸

5.1 Overview of TELNET connections

Table 8 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 (§ 3.3); [P93] details the rejected

⁸[P93] presents a similar overview for *rlogin* traffic, along with results of modeling it using 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	15,437	52	242B	$\times 4.5$	777KB	5.7KB	$\times 7.3$	14.0MB	270 s	$\times 7.7$	102.9 h
LBL-7	17,998	106	235B	$\times 4.2$	651KB	5.7KB	$\times 6.9$	3.10MB	252 s	$\times 7.5$	172.8 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 8: Summary of TELNET Connections

connections. As discussed in [P94], the LBL-6 and LBL-7 *telnet* traffic included a large number of connections due to periodic traffic. We removed those connections prior to our analysis, and they do not appear in Table 8.

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 (Equation 6), the geometric standard deviation (Equation 7), and the maximum. As noted in § 3.2, we applied \log_2 transformations to the data prior to analysis.

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 summarize 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 [B67] gives a local-login median of 20 minutes and a mean of 45-50 minutes. Jackson and Stubbs inferred that connection time “may be considerably reduced by providing a high-speed channel from the computer to the user”, so we appear to be seeing Jackson and Stubbs’ effect 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, even if we adjust for the fact that the LBL datasets span more days and hence give an opportunity to detect long-lived connections missed by the other datasets. We conclude that, at least with regard to mean bytes transferred and duration, the LBL *telnet* traffic is significantly different from that at other sites (and this is what is shown in Table 7 above).

We also note an apparent trend over the LBL datasets towards increasing values of \bar{x}_{orig} and \bar{x}_{resp} , 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 (again reflected in Table 7).

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. 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 § 5.5 below we present a model for this ratio.

5.2 TELNET originator bytes (A)

With the bulk transfer protocols examined in subsequent sections, we usually are only interested in modeling the number of bytes transferred. 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 (generally a human typing at a keyboard). The best fit we found to the LBL *telnet* test datasets came using the \log_2 -extreme distribution (§ 3.2).

Figure 1 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, where, as indicated in Table 7, the upper 1% tail is grossly overestimated.

One important point regarding this model is that it is easy to assume that the number of *bytes* generated by the *telnet* originator equates to the number of *packets* generated. As mentioned in [PF94], this equation is erroneous. Often an originator packet holds more than a single keystroke. [PF94] finds that the number of *packets* generated by the originator appears better modeled using a log-normal distribution; as we did not have any intra-connection information, we were unable to test this finding.




Figure 1: TELNET Originator-Bytes Model for LBL-2: Log₂-Extreme Distribution

5.3 TELNET responder bytes (B)

We next model the bytes transferred by the *telnet* responder. As shown in Figure 2, the upper 80% of the distribution is well-modeled using a log₂-normal distribution, but 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. No doubt this roughness is due to the varying sizes of log-in dialogs and message-of-the-day greetings. Fortunately, exactly modeling the lower tail is of little importance, so we limit ourselves to modeling just the upper 80% (in doing so, we apply data *censoring*; see § 3.1.4).

5.4 TELNET duration (C)

We model *telnet* connection durations using a simple log₂-normal distribution. An alternative is to simply not model the duration at all, for the following reason: Paxson and Floyd show that by using Poisson arrivals, log-normal originator connection sizes (in packets), and the Tcplib packet interarrival distribution, one can synthesize *telnet* traffic which reproduces the same burstiness at different time scales as observed in actual traffic traces [PF94]. The duration of their *telnet* connections is simply the sum of the packet interarrivals. That their model is successful without using a separate distribution for the connection duration suggests that to simulate *telnet* traffic with realistic durations, it is sufficient to use the Tcplib interarrival distribution.

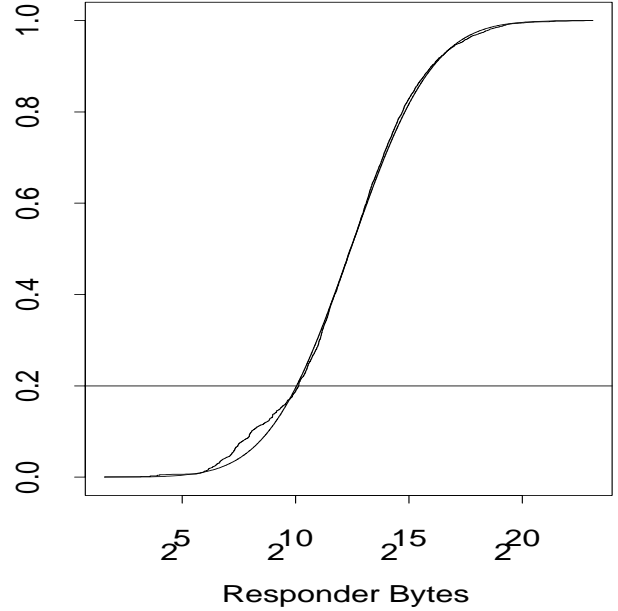


Figure 2: Censored Log₂-Normal Fit to Upper 80% of LBL-4 TELNET Responder Bytes

5.5 TELNET responder/originator ratio (D)

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 log₂-normal distribution. 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 unscaled model uses a ratio of 21:1.

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 log₂-normal distributed, the originator bytes are log₂-extreme distributed. Multiplying the originator byte's log-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




Figure 3: Responder/Duration Distributions for LBL-1: Exponential Fits

of two normal distributions is a normal distribution), and not a log-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.

5.6 TELNET responder/duration ratio (E and F)

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 (though see the last paragraph of this section).

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 3 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. 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 transferred) occur over relatively short durations: while \bar{x}_{resp} in these large connections is 45 times that of the smaller (lower 90%) connections, \bar{x}_{dur} is only 16 times that of the smaller connections. This phenomenon was also observed by the authors of [SC93], 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 upper 10% of the responders we compared considerably fewer datasets. Our requirement that each dataset include at least 100 measurements (§ 3.3) ruled out any dataset with fewer than 1,000 *telnet* connections, leaving just the LBL and NC datasets. The fit remains good, though.

The use of two separate models for the responder/duration ratio is not wholly satisfying, but was the best we could find. One other somewhat successful model for relating *telnet* connection size and duration was the ratio between the duration and the originator size. We found that when this ratio was ≥ 0.5 (about 80% of the connections), then it was well modeled using a Pareto distribution (Equation 2). But when the ratio was < 0.5 , we found no simple yet accurate analytic description of the marginal distribution. Unlike when modeling *telnet* responder bytes, for the duration/originator ratio both the lower and upper parts of the distribution are important; we cannot simply model the upper 80% of the connections and ignore the remainder. So we chose to use models E and F above instead.

6 NNTP

6.1 Overview of NNTP connections

Table 9 summarizes *nnntp* connections. As *nnntp* is non-interactive, the connection duration is not of much interest and has been omitted. [P93] 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: (1) a server contacts a peer and is informed that the peer presently cannot talk to the server; (2) the server offers the peer news articles but the peer already has the articles; (3) 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,

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
LBL-7	70,842	143	7 %	41.7KB	$\times 7.1$	10.8MB	1032B	$\times 2.2$	2.6MB
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 9: Summary of NNTP Connections

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 originating peer then replies with “QUIT” followed by a carriage-return and a line-feed, it will have sent 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. 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-7 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 is most likely also a reflection of the global growth in USENET *nnntp* traffic, which increases in volume about 75%/year [P94].




Figure 4: Distribution of NNTP Originator Bytes

6.2 NNTP originator bytes (G)

Figure 4 shows the distributions of bytes sent by the originator in non-failure *nnntp* connections at LBL, DEC, and coNCert. The distributions show a huge degree of variance (recall that the X axis is scaled logarithmically).

Given the great variation in originator bytes transferred, we decided to simply use a log₂-normal model to describe the connections, with the caveat that we do not expect the model to perform well (we also do not expect empirical models to do well). Indeed, as shown in Table 7, none of the models do well.




Figure 5: Daily Variation in \log_2 -mean of LBL NNTP Originator Bytes

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 5 shows the hourly \bar{x}_{orig} for LBL-1 and LBL-4 non-failure *nntp* connections. 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.⁹

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 late-night hours to take advantage of minimal loads (LBL-1).

6.3 NNTP responder bytes

As seen in Table 9 above, there is in general much less variation in the bytes sent by an *nntp* responder than by the originator. For the majority of the datasets, the responder sent fewer than 1500 bytes in 80% or more of the connections. Thus we decided not to model *nntp* responder bytes, as in general the datasets do not show interesting variations.

⁹The test datasets also showed a weekly pattern, with LBL-1 and LBL-4 (and to a lesser extent LBL-2) having minimal \bar{x}_{orig} during weekends, while LBL-3 had a *maximum* \bar{x}_{orig} on Saturdays.




Figure 6: One-Minute Variation in DEC-2 NNTP Arrivals

6.4 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.

6.5 NNTP arrival patterns

Nntp arrivals have a definite one-minute periodicity about them.¹⁰ Figure 6 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 past. All of the *nntp* datasets show this pattern to varying degrees except for LBL-3; LBL-4 shows two distinct spikes. Sometimes the spike is quite sharp. With the other datasets, it is broad, like in Figure 6. In general, periodicity such as this can lead to global synchronization of network processes; see [FJ93].

7 SMTP

7.1 Overview of SMTP connections

Table 10 summarizes the *smtp* connections. Again, [P93] summarizes the reasons for removing the connections marked

¹⁰We also found three-, five-, fifteen-, and twenty-minute patterns in various datasets.

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
LBL-7	205,668	6,306	1.9KB	$\times 2.9$	8.0MB	314B	$\times 1.4$	16.6KB
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 10: Summary of SMTP Connections

as rejects. Based on the values for \max_{orig} it is clear that *smtp* is sometimes used to transfer quite large files.

There is quite a bit of variation in \bar{x}_{orig} (and just about none in \bar{x}_{resp}). In [WLC92] the authors note that the UK *smtp* data show a substantially higher (arithmetic) \bar{x}_{orig} than for the LBL-1 and LBL-2 datasets reported in [P91]. 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 their hypothesis. The DEC traffic has similar σ_{orig} values, and Mogul also states that an “FTP-by-mail” facility is responsible for about 150 rather lengthy *smtp* messages at DEC-WRL each day [M92].

Another explanation is that perhaps the 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 bytes¹¹. One would expect the greater number of hops to be correlated with “wider” wide-area traffic, presumably a property of the NC and UK traffic, as these sites are at inter-network gateways.

We see a definite trend in the LBL data indicating larger and larger mail messages. As discussed in [P94], LBL’s wide-area traffic did become “wider” during the three year period spanned by the LBL datasets, in agreement with the “hops overhead” explanation.

7.2 SMTP originator bytes (H)

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

¹¹A check of one of the author’s mail folders revealed an average Received header length of more than 100 bytes.




Figure 7: Bimodal Log₂-Normal Fit to LBL-3 SMTP Originator 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 bimodal, not surprisingly given that *smtp* is sometimes




Figure 8: Daily Variation in \log_2 -mean of LBL SMTP Originator Bytes

used to transfer files. We model the distribution using two \log_2 -normal distributions, one (called f here) for the lower 80% of the data, and one for the remaining 20% (g). Figure 7 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).

As was the case for *nntp*, for *smtp* we found that the originator bytes distribution is not stationary. Figure 8 shows the hourly \bar{x}_{orig} for LBL-1 and LBL-4 *smtp* connections. Unlike *nntp*, which suffered from inconsistent variations, here the pattern is more stable: connection sizes peak during off-hours, and reach minima during peak working hours. We conjecture that uses of *smtp* to transfer files tend to happen off-hours and cause this pattern.

7.3 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% of all connections the responder sent between 300 and 400 bytes, and in every dataset more than 97% of the connections sent between 100 and 1000 bytes. While reference [DJCME92] finds that *smtp* connections are bidirectional, this finding must be interpreted with the rather fixed nature of the *smtp* responder in mind.

8 FTP

8.1 Overview of FTP connections

Table 11 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 (a *put* command). 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, [P93] gives details regarding the connections we rejected.

A considerable portion of the UCB, LBL-5, LBL-6, and LBL-7 *ftp* connections were due to periodic traffic, as discussed in [P94]. As with periodic *telnet* connections, we eliminated these prior to analysis.

There clearly is quite a range in \bar{x}_{bytes} , and the uniformly large values of σ_{bytes} shows that in general file sizes vary widely.

Table 12 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 (see [P93] for details). We refer to such a collection of an *ftpctrl* connection and its associated *ftpdata* connections as an *ftp session*.

The fourth through seventh columns in Table 12 summarize the number of *ftpdata* connections that occurred during each *ftp session*. The “0 xfer” column lists the percentage of all *ftp sessions* that did not have *any* associated *ftpdata* connections, presumably due to failed attempts to provide log-in information. These numbers are somewhat lower than the 42.9% reported in [DHS93], but still surprisingly high.

The \bar{x}_{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 (both result in an *ftpdata* connection), 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.

The \bar{x}_{bytes} and σ_{bytes} columns show the geometric mean and standard deviation for the total number of bytes transferred via *ftpdata* connections during each *ftp session* (for those connections with at least one *ftpdata* transfer). We note that these means are 5-10 times greater than those for individual *ftpdata* connections, an increase larger than that due simply to the multiplying effect of \bar{x}_{xfers} . We suspect that this dispar-

Dataset	# Conn	# Rej	Get	\bar{x}_{bytes}	σ_{bytes}	maXbytes
LBL-1	23,555	287	80 %	2.3KB	$\times 15.3$	54MB
LBL-2	27,917	335	92 %	2.4KB	$\times 17.4$	124MB
LBL-3	39,552	349	91 %	3.3KB	$\times 17.7$	62MB
LBL-4	65,860	335	86 %	3.8KB	$\times 14.7$	67MB
LBL-5	66,411	344	80 %	4.5KB	$\times 16.0$	177MB
LBL-6	86,464	464	91 %	2.1KB	$\times 14.9$	292MB
LBL-7	105,821	468	94 %	2.8KB	$\times 15.2$	223MB
BC	5,199	58	97 %	2.5KB	$\times 12.6$	16MB
UCB	4,529	77	96 %	1.0KB	$\times 13.5$	22MB
USC	1,870	29	93 %	1.3KB	$\times 14.5$	5MB
DEC-1	7,970	6	100 %	2.2KB	$\times 16.5$	5MB
DEC-2	4,013	13	100 %	1.3KB	$\times 17.1$	7MB
DEC-3	6,775	25	99 %	1.9KB	$\times 16.7$	13MB
NC	19,076	183	98 %	1.8KB	$\times 19.0$	44MB
UK	10,018	58	97 %	3.4KB	$\times 14.2$	7MB

Table 11: Summary of FTP Data Connections

Dataset	# Conn	# Rej	0 xfer	\bar{x}_{xfers}	σ_{xfers}	maXxfers	\bar{x}_{bytes}	σ_{bytes}
LBL-1	3,757	51	19 %	3.3	$\times 2.9$	1,006	28KB	$\times 15.2$
LBL-2	5,312	72	25 %	3.2	$\times 2.8$	388	27KB	$\times 17.0$
LBL-3	6,916	90	21 %	3.1	$\times 2.9$	612	30KB	$\times 18.4$
LBL-4	7,941	189	17 %	3.3	$\times 3.0$	1,951	33KB	$\times 17.6$
LBL-5	9,968	1,227	26 %	3.0	$\times 3.0$	975	31KB	$\times 16.7$
LBL-6	12,470	535	24 %	3.1	$\times 2.9$	2,996	31KB	$\times 16.8$
LBL-7	17,556	319	27 %	3.0	$\times 2.9$	1,666	34KB	$\times 16.8$
BC	669	19	32 %	3.3	$\times 2.7$	426	13KB	$\times 14.2$
UCB	756	19	26 %	3.9	$\times 2.6$	350	12KB	$\times 14.9$
USC	272	6	22 %	3.8	$\times 2.8$	133	20KB	$\times 14.5$
DEC-1	727	8	26 %	5.4	$\times 3.2$	961	36KB	$\times 15.6$
DEC-2	491	8	13 %	5.0	$\times 3.0$	106	36KB	$\times 17.8$
DEC-3	811	17	25 %	4.8	$\times 2.9$	232	36KB	$\times 15.3$
NC	2,500	59	31 %	5.0	$\times 2.9$	392	26KB	$\times 18.6$
UK	1,733	35	24 %	3.4	$\times 3.0$	368	22KB	$\times 16.0$

Table 12: Summary of FTP Control Connections

ity is due to a typical *ftp* session including at least one true file transfer. As files will tend to be significantly larger than directory listings, the mean number of transferred bytes during an *ftp* session 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.

One important point regarding the *ftpdata* connections occurring during an *ftp* session is that they frequently arrive in *bursts* [PF94]. An *ftpdata* connection burst is defined as one or more *ftpdata* connections belonging to the same session that are spaced less than 4 seconds apart. That is, each connection in the burst begins less than 4 seconds after the previous connection ended. A burst can be due to a “multiple-get” transfer,

or to a user doing a remote directory listing and shortly after it completes, fetching a file. A key finding in [PF94] regarding bursts is that the distribution of the number of bytes in a burst has an extremely heavy tail: just a handful of the largest bursts carry the majority of *all* of the *ftp* data bytes.

Table 13 summarizes the *ftpdata* bursts. We see from the second column that each dataset had roughly half as many bursts as *ftpdata* connections (Table 11). The third through fifth columns summarize the number of bytes transferred per burst of *ftpdata* connections. The values of \bar{x}_{bytes} are surprising—they are lower than the corresponding values for *ftpdata* connections! This appears paradoxical, because each *ftpdata* connection burst is made up of at least one *ftpdata* connection, so we would expect the bursts on average to be

Dataset	# Bursts	\bar{x}_{bytes}	σ_{bytes}	maxbytes	2% Tail	0.5% Tail	α	k
LBL-1	13,055	2.0KB	$\times 17.6$	102MB	70 %	49 %	1.00	375KB
LBL-2	16,111	1.8KB	$\times 19.2$	124MB	70 %	44 %	1.01	445KB
LBL-3	22,388	2.3KB	$\times 19.0$	96MB	68 %	41 %	1.03	624KB
LBL-4	27,084	2.5KB	$\times 18.8$	83MB	68 %	41 %	1.05	717KB
LBL-5	30,358	2.8KB	$\times 19.1$	754MB	79 %	61 %	0.93	620KB
LBL-6	39,740	2.7KB	$\times 16.7$	465MB	77 %	58 %	1.06	595KB
LBL-7	48,542	2.8KB	$\times 17.1$	200MB	67 %	44 %	1.16	748KB
BC	2,077	1.4KB	$\times 13.9$	16MB	64 %	41 %	1.21	290KB
UCB	2,804	0.9KB	$\times 13.2$	12MB	66 %	42 %	1.11	148KB
USC	830	1.1KB	$\times 19.1$	5MB	53 %	31 %	1.37	315KB
DEC-1	4,487	1.3KB	$\times 13.9$	32MB	70 %	48 %	1.06	248KB
DEC-2	2,743	1.3KB	$\times 13.5$	12MB	66 %	37 %	1.18	374KB
DEC-3	4,276	1.3KB	$\times 13.2$	15MB	67 %	45 %	1.09	244KB
NC	13,086	1.3KB	$\times 17.7$	44MB	57 %	37 %	1.34	308KB
UK	5,837	1.9KB	$\times 14.6$	8MB	54 %	29 %	1.40	375KB

Table 13: Summary of FTP Bursts

at least as large as the individual connections. The key to understanding this discrepancy is that *large ftpdata connections tend to arrive together in single bursts*. This means that the upper tail of the distribution of the number of bytes per burst is heavier than the corresponding upper tail for *ftpdata* connections; there are fewer big bursts, but those few are very large. Because x_{bytes} is a *geometric* mean (Equation 6), and the geometric mean is relatively insensitive to outliers, we find x_{bytes} becomes *lower* when we shift distribution weight higher into the upper tail.

The next two columns explore the tail-weight further. The “2% Tail” column gives the percentage of all *ftp* burst data bytes due to the 2% largest bursts, and similarly for the “0.5% Tail” column. We see that the 2% upper tail in all cases accounts for more than 50% of all of the data bytes! Thus *ftp* traffic is heavily dominated by a few rare but huge bursts. As stated in [PF94], this finding means that modeling *ftp* traffic should concentrate heavily on the upper tails of *ftpdata* bursts (as is done in § 8.4 below). Note that *ftpdata* connections alone do *not* have nearly as heavy a tail. For example, in the DEC-1 dataset the upper 2% of the *ftpdata* connections holds about 25% of the data bytes, vs. 70% for the bursts. It is the fact that large *ftpdata* connections tend to arrive together that leads to the very heavy *ftp* burst tail.

In § 8.4 below we model the upper 5% tail of the *ftpdata* burst distribution using a Pareto distribution (Equation 2). The Pareto distribution has an extremely heavy tail, heavier than that of any of the other distributions discussed in § 3.1.3 or their logarithmic versions (§ 3.2). To this end, the final columns in Table 13 gives the estimated values of α and k (corresponding to Equation 2) for each dataset. That the values of α are smaller for the LBL datasets than for the others, and the values for k larger, indicates that the LBL *ftp* bursts are significantly different than those found in the other datasets.




Figure 9: Log₂-Normal Fit to LBL-4 FTP Data Bytes

In particular, the LBL bursts have heavier tails. Probably this difference is due to the prevalence of LBL users exchanging large scientific datasets.

8.2 FTP connection bytes (I)

We model the bytes transferred during an *ftpdata* connection using a log₂-normal distribution. Figure 9 shows this model fitted to the first half of the LBL-4 dataset. While the model




Figure 10: Log₂-Normal Fit to Bytes in LBL-1 FTP Sessions

appears to match the overall shape, a number of clumps and spikes make the actual distribution 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, LBL-5, and LBL-7. This spike occurs at 524,288 bytes ($= 2^{19}$), a size often used when splitting a large distribution archive into manageable pieces.

8.3 FTP session bytes (J)

Figure 10 shows an example of the distribution of the total number of bytes transferred during *ftp* sessions, for the LBL-1 test dataset, again fitted to a log₂-normal model. In this case the fit is visually fairly satisfying.

The authors of [DJCME92] reported that 80% of *ftp* sessions transfer less than 10 KB. But once we remove the 20-30% of sessions that do 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 session is not a “failure”, it should not be assumed small.

8.4 FTP bursts (K)

As mentioned in § 8.1 above, modeling *ftpdata* bursts is particularly important. Not only does the distribution of bytes per burst have an extremely large tail, but because the inter-

connection spacing within a burst is (by definition) < 4 seconds, from a link-level (or queueing) viewpoint, there is little difference between a burst of *ftpdata* connections and a single large connection transferring the same total number of bytes.

Because the upper tail of this distribution is so dominant, we decided to concentrate on modeling the size of the largest 5% of the bursts. We found that upper-tail burst size is well-modeled using a Pareto distribution (doubly-exponential; Equation 2). As discussed further in [PF94], that *ftp* burst sizes are Pareto-distributed suggests a mechanism by which *ftp* traffic might contribute to the presence of *self-similarity* ([LTWW93]) in wide-area network traffic.

9 Summary

We have presented a number of analytic models for describing the characteristics of *telnet*, *nntp*, *smtp*, and *ftp* connections, drawn from wide-area traces collected from seven different sites, comprising more than 3 million connections. While these models are inexact in a statistical sense, we developed a methodology for comparing their effectiveness to that of other models. We found that in general the analytic models reflect the connection characteristics as well as or better than two empirical models, one corresponding to the Tcplib library [DJ91] and one corresponding to a one-month trace of traffic at the Lawrence Berkeley Laboratory. We also found that wide-area connection characteristics exhibit significant variation from site to site and over time.

The essence of the argument presented in this paper is that while wide-area traffic cannot be modeled exactly in a statistical sense, we can usually construct simple analytic models that are a good approximation. Furthermore, these analytic models are as accurate as empirical models, meaning we can reap the benefits of using analytic models without losing accuracy in the process. We believe the methodology presented in this paper will prove beneficial for developing future analytic models and for gauging their effectiveness.

10 Acknowledgments

This work would not have been possible without the support, guidance, 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 Jon Crowcroft, Domenico Ferrari, Sally Floyd, Van Jacobson, Kimmo Raatikainen, and the referees for 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

- [B67] G. E. Bryan, "JOSS: 20,000 Hours At A Console", Proc. of the Fall 1967 AFIPS Conference, Vol. 31.
- [C89] R. Cáceres, "Measurements of Wide Area Internet Traffic", *Report UCB/CSD 89/550*, Computer Science Division, University of California, Berkeley, California, 1989.
- [CM80] W. Cook and D. Mumme, "Estimation of Pareto Parameters by Numerical Methods", in *Statistical Distributions for Scientific Work*, edited by Charles Taillie et al, D. Reidel Pub. Co., Boston, 1980, pp. 127-132.
- [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.
- [DHS93] P. Danzig, R. Hall, and M. Schwartz, "A Case for Caching File Objects Inside Internetworks", Proceedings of SIGCOMM '93, San Francisco, California, September, 1993.
- [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", *Internet-working: Research and Experience*, 3 (1), pp. 1-26, 1992.
- [FJ93] S. Floyd and V. Jacobson, "The Synchronization of Periodic Routing Messages," Proceedings of SIGCOMM '93, September, 1993, pp. 33-44. To appear in *Transactions on Networking*, 1994.
- [FJ70] E. Fuchs and P. E. Jackson, "Estimates of Distributions of Random Variables for Certain Computer Communications Traffic Models", *Communications of the ACM*, 13(12), pp. 752-757, December, 1970.
- [H90] 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.
- [LTWW93] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the Self-Similar Nature of Ethernet Traffic", Proceedings of SIGCOMM '93, September, 1993, pp. 183-193.
- [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.
- [M74] P. Martin-Löf, "The Notion of Redundancy and Its Use as a Quantitative Measure of the Discrepancy between a Statistical Hypothesis and a Set of Observational Data", *Scandinavian Journal of Statistics*, 1(1), pp. 3-18, 1974.
- [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.
- [M92] J. C. Mogul, "Observing TCP Dynamics in Real Networks", Proceedings of SIGCOMM '92, Baltimore, Maryland, August, 1992.
- [M84] D. Moore, "Measures of Lack of Fit from Tests of Chi-Squared Type", *Journal of Statistical Planning and Inference*, 10(2), pp. 151-166, 1984.
- [MT77] F. Mosteller and J. W. Tukey, "Data Analysis and Regression", Addison Wesley, 1977.
- [P89] 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.
- [P94] V. Paxson, "Growth Trends in Wide-Area TCP Connections", to appear in *IEEE Network*, 1994.

- [PF94] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling", submitted to SIGCOMM '94.
- [P93] V. Paxson, "Empirically-Derived Analytic Models of Wide-Area TCP Connections: Extended Report", *Report LBL-34086*, Lawrence Berkeley Laboratory, May, 1993. Available as papers/WAN-TCP-models.prelim.1.ps.Z and papers/WAN-TCP-models.prelim.2.ps.Z via anonymous ftp to ftp.ee.lbl.gov.
- [P91] V. Paxson, "Measurements and Models of Wide Area TCP Conversations", *Report LBL-30840*, Lawrence Berkeley Laboratory, 1991.
- [PJ90] S. Pederson and M. Johnson, "Estimating Model Discrepancy", *Technometrics*, 32(3), August, 1990, pp. 305-314.
- [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.
- [SC93] A. Schmidt and R. Campbell, "Internet Protocol Traffic Analysis with Applications for ATM Switch Design", *Computer Communication Review*, 23(2), April, 1993, pp. 39-52.
- [S79] D. Scott, "On optimal and data-based histograms", *Biometrika*, 66(3), 1979, pp. 605-610.
- [WLC92] I. Wakeman, D. Lewis, and J. Crowcroft, "Traffic Analysis of Trans-Atlantic Traffic", Proceedings of INET'92, Kyoto, Japan, 1992.