

Lognormal and Pareto Distributions in the Internet

Allen B. Downey*

*Olin College of Engineering
Olin Way
Needham, MA 02492*

Abstract

Numerous studies have reported long-tailed distributions for various network metrics, including file sizes, transfer times, and burst lengths. We review techniques for identifying long-tailed distributions based on a sample, propose a new technique, and apply these methods to datasets used in previous reports. We find that the evidence for long tails is inconsistent, and that lognormal and other non-long-tailed models are usually sufficient to characterize network metrics. We discuss the implications of this result for current explanations of self-similarity in network traffic.

Key words: File sizes, interarrival times, transfer times, transfer bursts, long-tailed distribution, self-similarity

1 Introduction

Researchers have reported traffic patterns in the Internet that show characteristics of self-similarity (see [1] for a survey). Many proposed explanations of this phenomenon are based on the assumption that the distribution of transfer times in the network is long-tailed [2] [3] [4] [5]. In turn, this assumption is based on the assumption that the distribution of file sizes is long-tailed [6] [7].

We examine these assumptions, looking at data from a variety of systems, including many of the datasets originally presented as evidence of long-tailed distributions.

*

Email address: allen.downey@olin.edu (Allen B. Downey).

Section 2 evaluates existing methods for identifying long-tailed distributions, and proposes a new statistical method for classifying distributions. Section 3 applies this methodology to empirical distributions of file sizes from a variety of systems. We find that the distribution of file sizes tends to be lognormal, in local file systems and in the World Wide Web. This tendency is strongest in large datasets that aggregate many file systems.

Section 4 discusses the implications of this result on current explanations of self-similarity, and presents alternative explanations. The remaining sections evaluate these alternatives by examining the distributions of interarrival times (Section 5), transfer times (Section 6) and burst durations (Section 7).

We find that there is little evidence that the distribution of interarrival times is long-tailed. Similarly, there is only ambiguous support for long-tailed transfer times. On the other hand, there is some evidence that bursts of file transfers in both ftp and HTTP are long-tailed. We investigate this possibility and its causes.

2 Methodology

A fundamental problem in this area of inquiry is the lack of methodology for identifying a long-tailed distribution based on a sample. For explanatory models of self-similarity, the relevant definition of “long-tailed” is a distribution with polynomial tail behavior; that is

$$P\{X > x\} \sim cx^{-\alpha} \text{ as } x \rightarrow \infty \quad (1)$$

where X is a random variable, c is a location parameter, and α is a shape parameter. When α is less than 2, the distribution has infinite variance, which is also required for these models to produce self-similarity.

2.1 cdf test

There are several characteristics we expect to see in a sample from a long-tailed distribution. If we plot the complementary cumulative distribution function (ccdf) on a log-log scale, we expect a straight line, at least in the tail.

Figure 1 shows the ccdf of samples from lognormal and Pareto distributions with similar tail behavior ($n=10\,000$). There is an obvious disparity in the bulk of the distribution (below the 90th percentile) but the tails overlap.




Fig. 1. ccdf of samples from lognormal and Pareto distributions with similar tail behavior.

The definitive characteristic of the long-tailed distribution is that its steepness does not increase in the extreme tail. It continues, with constant slope, to the limit of the sample (where it is increasingly jagged as the values become sparse).

Most prior claims about long-tailed distributions are based on these kinds of observations. We call this visual examination the “ccdf test.” As this example demonstrates, there are distributions like the lognormal that are not long-tailed, but whose ccdfs appear long-tailed, at least to a point. The definitive characteristic of these distributions is that the ccdf eventually drops away with increasing slope.

2.2 Using *aest*

Crovella and Taquq developed a tool called *aest* that estimates the slope parameter of a Pareto distribution based on a sample [8]. They propose a graphical technique that can “show the segment of the tail over which heavy-tailed behavior appears to be present.”

We applied *aest* to the samples in Figure 1. For the Pareto sample, the actual parameter is 1.42 and the estimate from *aest* is 1.33, which is reasonably accurate. For the lognormal sample the estimate is 1.42. The graphical output for the two samples is similar. For both distributions *aest* identifies points that show long-tailed behavior.

Thus, as its authors acknowledge, *aest* cannot distinguish long-tailed and lognormal distributions based on samples. Nevertheless, we find it useful for estimating the parameter of a sample that is hypothetically Pareto and use it to implement the curvature test presented below.




Fig. 2. ccdf of file sizes from a university web server.

2.3 Model fitting

A standard way to choose among alternative models is to estimate parameters to fit the data and choose whichever model yields the better goodness of fit. This approach may not be appropriate for this problem. For both models, conventional estimators (moment-matching or maximum likelihood) do not necessarily yield the model that is the best match for the tail behavior. Also, it is not obvious how to measure goodness of fit.

For example, Figure 2 shows the ccdf of 15 160 files on a web server at the University of Calgary, from traces collected by Arlitt and Williamson [9].

We fitted a Pareto model using `aest` to estimate α and choosing the lower bound by eye. We fitted a lognormal model by calculating moments under a log transformation.

The bulk of the distribution clearly fits the lognormal model better, but the tail is closer to the Pareto model. By conventional goodness-of-fit measures, the Pareto model is a better fit. Nevertheless, the measured distribution clearly displays the characteristic behavior of a non-long-tailed distribution: increasing slope in the extreme tail. So in this case quality of fit may be misleading. Although the fitted models are useful for visual comparison, they do not provide a mechanical, quantitative way to identify long-tailed distributions.

2.4 Percentile-percentile plots

A percentile-percentile plot (P-P plot) shows how well the rank statistics of a sample match a model distribution. For each value that appears in the sample, the P-P plot shows the actual rank of the value versus the expected rank of




Fig. 3. P-P plot on complementary log axes.




Fig. 4. Complementary P-P plot with alternate model.

the value in the model. Perfect agreement with the model yields the 45-degree line from the origin.

To examine the tail behavior, we can plot the complementary probabilities on log axes. Figure 3 shows this transformation for the Calgary dataset. There is a clear divergence from the model in the tail of the distribution (the left side of the figure).

Although P-P plots are useful for visualizing discrepancies like this, they are not useful for identifying long-tailed distributions. The reason is that, unlike the cdf test, P-P plots depend on parameter estimation. A P-P plot tests the fit of a specific model, not a family of models.

For Figure 3, we estimated the lognormal model by calculating moments. But if we choose a different model to match the tail behavior, we get Figure 4, which shows good agreement with the tail of the distribution. Even with a sample that actually comes from a Pareto distribution, it is often possible to find a lognormal model that fits the tail.

We conclude that P-P plots do not contribute additional discriminatory power beyond what we get from the cdf test. Furthermore, they suffer a serious disadvantage—the need to estimate parameters.

2.5 Curvature test

The characteristic that distinguishes the tail behavior of Pareto and lognormal distributions is curvature. In this section we propose a way of using measured curvature as a statistic to identify long-tailed distributions.

In general it is difficult to estimate the slope and curvature of a sample, because numerical differentiation tends to amplify noise. Fortunately, for this application, we don't need to know the curvature as a function. Instead, we estimate the average curvature of the tail by computing a numerical first derivative and fitting a line to it. The slope of that line reflects the overall shape of the cdf; we call this statistic the “tail curvature.”

By itself this number is meaningless, but we can compare it to the tail curvature of a Pareto sample, the curvature of which should be near zero. Then we can use a simulation to estimate the probability of seeing the given curvature in a sample from a Pareto distribution.

We formulate the problem as a hypothesis test, where the null hypothesis is that the observed distribution is a sample from a Pareto distribution. Assuming that we have a sample of n points, the procedure is

- (1) Compute the tail curvature of the sample.
- (2) Use `aest` to estimate the slope parameter, α , of the sample. The location parameter does not affect the measured tail curvature so there is no need to estimate it.
- (3) Test the hypothesis is that the sample came from a Pareto distribution with parameter α .
 - (a) Generate 1000 samples with n points from a Pareto distribution with slope parameter α . For each sample, calculate the tail curvature. Calculate μ , the mean curvature of the 1000 samples.
 - (b) Calculate d , the difference between the curvature of the original sample and μ .
 - (c) Count the number of samples, out of 1000, that have a curvature that differs from μ by as much as d . This count is an estimate of the p-value for the null hypothesis.
 - (d) If the p-value falls below a threshold of confidence, we can reject the null hypothesis.

We use a one-sided test, since the alternate hypothesis is that the sample

comes from a distribution with higher tail curvature than the Pareto.

We have tested this method on synthetic samples with a range of sample sizes. For $n = 10\,000$, we can set a threshold on the tail curvature so that 95% of the Pareto samples are classified correctly. Applying that threshold to the lognormal samples, we correctly reject the null hypothesis 93% of the time. For $n = 40\,000$ the test accepts 99% of the Pareto distributions while rejecting 99% of the lognormal distributions. At least for synthetic data, this method discriminates strongly between lognormal and Pareto distributions, even in cases where the cdfs are visually indistinguishable.

One limitation of this method is that it depends on `aest` to estimate the parameter of the Pareto distribution. For a given dataset, it is possible that a different parameter would yield a higher p-value. We have extended this method by searching for the parameter that maximizes the p-value, and found that the results are very insensitive to the estimated parameter.

The estimated tail curvature depends on several details of implementation, but because it is only used in a context of comparison, these details can be chosen almost arbitrarily with little qualitative effect on the results. One important parameter is the lower bound where the tail is considered to begin. In this paper, the lower bound is at $Pr\{X > x\} = 2^{-4}$; approximately the 94th percentile. Our reference implementation is available from <http://allendowney.com/research/longtail/>.

3 File sizes

In this section we survey prior studies that have looked at measured file sizes and presented evidence that the distribution is long-tailed.

3.1 *File sizes on the web, server's view*

Between October 1994 and August 1995, Arlitt and Williamson [9] collected traces from web servers at the the University of Calgary, the University of Saskatchewan, NASA's Kennedy Space Center, and ClarkNet (an ISP). For each server they identified the set of unique file names and examined the distribution of their sizes. They report that these distributions match the Pareto model, and they give parameters for each dataset, but they do not present evidence that these models fit the data.

We obtained their traces and extracted the name and file size of each successful HTTP transfer. To derive a set of distinct files, we treated as distinct any log




Fig. 5. ccdfs for the datasets collected by Arlitt and Williamson.

entries that had the same name but different sizes, on the assumption that they represent successive versions. Whether we use this definition of “distinct” or the alternative, we found a number of distinct files that differs from the original report, so our treatment of this dataset may not be identical to theirs. Nevertheless, our cdfs are visually similar to theirs.

We estimated the Pareto parameter for each dataset using `aest`; the range of the parameters is from 0.97 to 1.02. We estimated the location parameter by hand to yield the best visual fit for the cdf (this parameter has no effect on the curvature test). We estimated lognormal parameters for each dataset by calculating moments. Figure 5 shows these models along with the actual cdfs.

The results are difficult to characterize. For the NASA dataset the lognormal model is clearly better. For the Saskatchewan dataset the Pareto model is better. The other two cdfs lie closer to the Pareto model, but both show the characteristic curvature of the lognormal distribution.

For these cases, the curvature test provides some insight. The tail curvature of the Calgary dataset is 0.141, which has a negligible p-value under the Pareto model. This means that we can reject the hypothesis that the data are a sample from a Pareto distribution. The tail curvature of the ClarkNet dataset is 0.023, which has $p > 0.5$, which means we cannot reject the hypothesis that the sample comes from a long-tailed distribution.

Although the ClarkNet and Saskatchewan datasets provide some support for the Pareto model, the evidence is weak and mixed.

Arlitt and Jin collected access logs from the 1998 World Cup Web site [10] and reported the distribution of file sizes for the 20 728 “unique files that were requested and successfully transmitted at least once in the access log.” They report that the bulk of the distribution is roughly lognormal, but that the tail of the distribution “does exhibit some linear behavior” on log-log axes. They estimate a Pareto model for the tail, with $\alpha = 1.37$.

We obtained the file sizes from this dataset and fitted lognormal and Pareto models, shown in Figure 6. For files smaller than 128KB, the lognormal model is a slightly better fit. For larger files, neither model describes the data well.

The curvature test does not support either model; the tail curvature is -0.075, which has a low p-value for the Pareto and lognormal models. These results are typical for small datasets from a single file system; tail behavior is often idiosyncratic and hard to characterize.

Arlitt, Friedrich and Jin did a similar analysis of more than 16 million unique HTML files transferred by the Web proxy server of an ISP [11]. They plot the




Fig. 6. ccdf of file sizes from World Cup dataset.




Fig. 7. Distribution of file sizes from a Web proxy server (reproduced from [11]).

ccdf of file sizes and show that a lognormal model fits it well. They also show the ccdf on log-log axes and claim that “since this distribution does exhibit linear behavior in the upper region we conclude that it is indeed heavy-tailed.”

That figure is reproduced here as Figure 7. We do not see linear behavior in the ccdf; it shows the characteristic curvature of a non-long-tailed distribution.

3.2 File sizes on the web, client’s view

Crovella et al. presented one of the first measurements of file sizes that appeared to be long-tailed. In 1995 they instrumented web browsers in computer labs at Boston University to collect traces of the files accessed [12] [13] [14].

We obtained the dataset presented in [13] and identified as W95 in [14] and extracted the unique file names, yielding 36 208 unique files. Figure 8 shows the resulting ccdf along with a Pareto model and a lognormal model. The slope of the Pareto model is 1.05, the value reported by Crovella et al. The location parameter is 3800, which we chose to be the best match for the ccdf, and visually similar to Figure 8 in [13].




Fig. 8. cdf of file sizes from Crovella dataset.




Fig. 9. cdf of file sizes from Web browser logs and a two-mode lognormal model.

The Pareto model is a good fit for file sizes between 4KB and 4MB, which includes about 25% of the files. Based on this fit, Crovella et al. argue that this distribution is long-tailed. At the same time, they acknowledge two disturbing features: the apparent curvature of the cdf and its divergence from the model for files larger than 4MB.

The lognormal model is a better fit for the data over most of the range of values, including the extreme tail. Also, it accurately captures the apparent tail behavior, which drops off with increasing steepness. The curvature test confirms the visual evaluation; the tail curvature is 0.050, which has a probability $p < 0.03$ under the Pareto model. We conclude that this dataset provides greater support for the lognormal model than for the Pareto model.

3.3 Multimodal models

In this section we present two datasets where the lognormal model does not match the observed tail behavior, but a two-mode lognormal model does.




Fig. 10. ccdf of file sizes from the Saskatchewan dataset and a two-mode lognormal model.

In 1998, Barford et al. instrumented Web browsers at Boston University and collected the sizes of 66 998 downloaded files [14]. They fit a hybrid model to these data, with a lognormal body and a Pareto tail, and suggest that the distribution is long-tailed.

Figure 9 shows the ccdf of this dataset along with a two-mode lognormal model. For this example we performed an automated search for the set of parameters that minimizes the Kolmogorov-Smirnov statistic. There are more rigorous techniques for estimating multimodal normal distributions, but they are not necessary for our purpose here, which is to demonstrate a lognormal model that fits the data well.

Although this dataset also fits the Pareto model, the curvature test suggests that the lognormal model is a better characterization. The tail curvature is 0.044, which is similar to the curvature of the lognormal model (0.045), but unusual for the Pareto model ($p < 0.002$).

The two-mode lognormal model also works well for the Saskatchewan dataset, shown in Figure 10. Of course, the two-mode model has more parameters, so a better fit is not surprising. Nevertheless, it is still reasonably parsimonious (5 parameters), and the quality of the fit is quite good. These examples show that, even in cases where the Pareto model fits well, there are non-long-tailed models that are just as realistic, and that avoid the analytic difficulties of long-tailed distributions (infinite moments and long transients [15]).

3.4 Aggregation

Looking at file sizes on the Internet, we are seeing the mixture of file sizes from a large number of file systems. To characterize this aggregate distribution, we consider two large datasets: a survey of 655 UNIX systems collected by




Fig. 11. ccdf of file sizes in the Irlam survey.




Fig. 12. ccdf of file sizes in the Microsoft survey.

Gordon Irlam in 1993 [16] and a survey of 10 568 Windows machines collected by Douceur and Bolosky in 1998 [17].

The Irlam survey contains 6 156 581 files with 161 583 different sizes. The size of this sample allows us to examine the extreme tail of the distribution. There are 169 files bigger than 32MB and 16 files bigger than 128MB. The largest file is 377MB.

Figure 11 shows the ccdf of these file sizes along with lognormal and Pareto models. The lognormal model is a better fit. Throughout the range, the ccdf displays the characteristic curvature of the lognormal distribution. The tail curvature is 0.024, which has probability $p < 0.001$ under the Pareto model.

Douceur and Bolosky collected the sizes of more than 140 million files from 10 568 machines at Microsoft Corporation [17]. They report that the bulk of the distribution fits a lognormal distribution, and they propose a two-mode lognormal model for the tail, but they also suggest that the tail fits a Pareto distribution.

Figure 12 shows the cdf of file sizes from this dataset along with three models we chose to fit the tail: a lognormal model, a Pareto model and a two-mode lognormal model.

Again, the tail of the distribution displays the characteristic curvature of a non-long-tailed distribution. The simple lognormal model captures this behavior well, although it is offset from the actual distribution. The two-mode lognormal model fits the entire distribution well. The tail curvature is 0.029, which has probability $p < 0.001$ under the Pareto model.

We conclude that the lognormal model is sufficient to describe the aggregate distributions that result from combining large numbers of file systems, and that these datasets provide little evidence that the overall distribution of file sizes is long-tailed.

To summarize this section, we find that many of the datasets that have been presented as evidence of long-tailed file sizes actually support the lognormal model. Although some small datasets appear to fit the Pareto model, larger datasets that aggregate files from many systems show the characteristics of a non-long-tailed distribution.

4 Self-similar network traffic

Many current explanations of self-similarity in the Internet are based on the assumption that some network metric—either transfer times, interarrival times, or burst sizes—is long-tailed.

One of these explanatory models is an $M/G/\infty$ queue in which network transfers are customers and the network is an infinite-server system [2] [3]. If the distribution of service times is long-tailed then the number of customers in the system is an asymptotically self-similar process.

Willinger et al. propose an alternative that models users as ON/OFF sources in which ON periods correspond to network transmissions and OFF periods correspond to inactivity [4]. If the distribution of the lengths of these periods is long-tailed, then as the number of sources goes to infinity, the superposition of sources yields an aggregate traffic process that is fractional Gaussian noise, which is self-similar at all time scales.

In their original model the lengths of the ON and OFF periods were identically-distributed, long-tailed random variables. A companion paper proves that the model yields self-similarity if either the distribution of ON periods *or* the distribution of OFF periods is long-tailed [18].

Two subsequent papers have extended this model to include a realistic network topology, bounded network capacity and feedback due to the congestion control mechanisms of TCP [6] [5]. The more realistic models yielded qualitatively similar results.

In the next few sections, we examine the metrics that have been reported to be long-tailed and review the evidence for these reports.

5 Interarrival times

Paxson and Floyd [2] measure the distribution of interarrival times for packets within Telnet connections, and report that “the main body of the observed distribution fits very well to a Pareto distribution ... with shape parameter 0.9, and the upper 3% tail to a Pareto distribution with [shape parameter] 0.95.” They do not show the cdf or explain how they chose these parameters.

This claim is based on traces collected at Lawrence Berkeley Labs during 1-hour intervals in December 1993 and January 1994. Unfortunately, the trace they describe, LBL PKT-1, is not available from the Internet Traffic Archive (ITA). Three other traces from the same dataset are, but they include all TCP packets, not just Telnet packets. The traces have been sanitized, so the contents of the packets, including protocol information, have been removed. As a result, we cannot repeat the original analysis, but we can examine the interarrival times for all TCP packets.

We obtained four packet traces: LBL PKT-3, LBL PKT-4, LBL PKT-5, and DEC PKT-1. The first three are from Lawrence Berkeley Labs; the last from Digital Equipment Corporation (DEC), collected in March 1995. For each trace, we identified connections by the source and destination addresses and the source and destination ports. Traffic from the originator to the responder is considered a different connection from the return traffic. Within each connection, we calculated the time between packet arrivals. Then we computed the cdf of the interarrival times.

The four datasets yield similar distributions, so we aggregated them into a single dataset. Figure 13 shows the resulting cdf, which includes 4 410 851 interarrival times. There is a small mode (0.001% of the data) at 75 seconds, which is the default interval for the TCP keep-alive mechanism.

For intervals smaller than 75 seconds, we agree with Paxson and Floyd that the Pareto model fits this data well. Above 75 seconds, the cdf starts to fall away from the Pareto model with increasing slope. It is possible, though, that this discrepancy is a measurement artifact. Since the traces are only an




Fig. 13. ccdf of interarrival times for TCP packets.




Fig. 14. ccdf of time between web requests from the W95 dataset.

hour long, they cannot contain an interval longer than an hour, and tend to under-represent intervals in proportion to their length.

We conclude that there is some support for the Pareto model of interarrival times, but it may be useful to look at evidence from longer traces.

5.1 HTTP Request Interarrivals

Crovella and Bestavros [19] examine the distribution of times between web requests (OFF times) and report that, although it is long-tailed, it is less long-tailed than the ON time distribution.

Their traces (dataset W95) contain logs from web browsers on 37 workstations. For each browser we compute the time between successive requests and form the ccdf of the interarrival times. Since the distributions were similar for each machine, we aggregated the data into a single ccdf, shown in Figure 14.

In this case, the parameter estimated by `aest` is not a good match for the data, so we show an alternative model chosen by eye.

The Pareto and lognormal models match parts of the distribution, but neither is a good match for the extreme tail. As in the TCP packet datasets, the duration of the trace (42 days) explains some of the behavior of the extreme tail. Some of the longest intervals are on the order of 10 days, which means that they might be underrepresented by a factor of 4. So, if we ignore the extreme tail, the Pareto model may describe this dataset well.

Arlitt and Williamson examined the pattern of accesses to individual files on a web server [9]. For each file that was accessed more than once in their traces, they computed the time between references. They plot the distribution of these interarrival times for each of the servers they studied. They report that these distributions are approximately exponential and independent, but they omit the statistical analysis.

Deng collected traces of WWW requests from users at GTE Laboratories to remote servers, and examined the distribution of times between document requests [20]. He divides the traffic into ON and OFF periods, where an ON period contains a series of requests with interarrival times less than 60 seconds. Interarrival times longer than 60 seconds are considered to be OFF periods. He reports that the distribution of OFF times fits a Pareto distribution, but he does not show the cdf, or compare the Pareto model to the alternatives.

In summary, there are reports that times between WWW requests are long-tailed, but the evidence is inconsistent and inconclusive.

6 Transfer times

Even if file sizes are not long tailed, transfer times might be. The performance of wide-area networks is highly variable in time; it is possible that this variability causes long-tailed transfer times. In this section we investigate the relationship between file sizes and transfer times for HTTP and ftp transfers.

6.1 HTTP transfer times

Crovella and Bestavros examine the distribution of transfer times in the W95 dataset, and report that it is long-tailed [13].

We repeated their analysis of the 135 357 transfers they observed; Figure 15 shows the resulting cdf along with a Pareto model and a lognormal model.




Fig. 15. ccdf of transfer times for web requests from BU traces.




Fig. 16. ccdf of session sizes from Feldmann dataset (reproduced from [21]).

In this case we chose the parameters of the lognormal model by hand.

The lognormal model is a better fit for the ccdf, which has the characteristic curvature of a non-long-tailed distribution. The tail curvature is 0.042, which has probability $p < 0.001$ under the Pareto model. We conclude that this dataset does not support the hypothesis that transfer times are long-tailed.

Feldmann et al. argue that the distribution of Web session sizes is long-tailed, based on data they collected from an ISP [21]. They use the number of bytes transferred during each modem connection as a proxy for bytes transferred during a Web session. The evidence they present is the ccdf in their Figure 3, reproduced here as Figure 16. They do not report what criteria they use to identify the distribution as long-tailed, other than “a crude estimate of the slope of the corresponding linear regions.” In our opinion, this distribution shows the characteristic curvature of a non-long-tailed distribution and does not support the Pareto model.




Fig. 17. cdf of ftp transfer sizes from LBL traces.




Fig. 18. cdf of ftp transfer times from LBL traces.

6.2 ftp transfer times

Paxson [22] examines the sizes of ftp transfers and reports that the distribution fits a lognormal distribution well. He does not discuss the tail behavior.

Using the LBL CONN-7 dataset again, we extracted the 105 542 transfers that used the ftp-data protocol, were successful, and transported a non-zero number of bytes. Figure 17 shows the cdf of these transfer sizes; Figure 18 shows the cdf of transfer times.

We agree with Paxson that the distribution of sizes fits the lognormal distribution, both in the bulk, as shown in his paper, and in the tail, as shown here. The Pareto model also fits the tail well, but the curvature of the tail suggests a non-long-tailed model. The tail curvature is 0.113, which has a negligible p-value under the Pareto model.

On the other hand, the lognormal model does not fit the cdf of transfer times at all. It is not clear whether the Pareto model is better. The slope parameter

estimated by `aest` matches a part of the tail, but for any curve there is likely to be a line that fits as well. Nevertheless, the tail of this cdf is straight enough to suggest a long-tailed distribution. The tail curvature is 0.0011, which is not at all unusual for a Pareto tail ($p > 0.95$).

In prior work [23] we propose a model for the relationship between transfer sizes and transfer times, based on the observed distribution of throughput. We show evidence that the distribution of throughput is lognormal, in which case we can model transfer times as the ratio of two lognormal variables, size and throughput. The result is a lognormal distribution of transfer times.

We also report a strong correlation between transfer size and throughput, $\rho \approx 0.70$, which indicates that large transfers tend to achieve better throughput. The effect of this correlation is to reduce the variance in transfer times. In some cases this effect makes the cdf appear straighter, but we are reluctant to take this shape as evidence of a long-tailed distribution, since the model suggests that transfer times are actually *less* long-tailed than transfer sizes.

7 Transfer Bursts

The motivation for investigating the sizes of transfer bursts is that ON periods in the ON/OFF model might correspond not to individual file transfers, but to periods of network activity interrupted only by network delays and short intervals between files. From the network's point of view there is no difference between a delay caused by a TCP timeout and a delay with the same duration caused by user activity or processing delays.

7.1 *ftp bursts*

Paxson [22] discusses ftp data bursts, which he defines as a sequence of ftp transfers in the same session that are spaced less than 4 seconds apart. He does not show the cdf, but he reports that the largest 5% of the bursts are well-modeled using a Pareto distribution. This claim appears again in subsequent work [2], along with an estimated parameter for the Pareto model.

We applied Paxson and Floyd's analysis to the LBL CONN-7 dataset, grouping consecutive ftp data transfers into bursts if the time between the end of one and the beginning of the next is less than 4 seconds. We found 56 155 such bursts, which is somewhat more than the number in the original paper (48 568). We don't know the reason for the discrepancy. The longest burst comprises 979 connections (which is the same as in the original paper).




Fig. 19. Distribution of burst sizes (total bytes) from LBL traces.




Fig. 20. Distribution of burst lengths (number of transfers) from the LBL traces.

We computed the total number of bytes in each burst; Figure 19 shows the distribution of these burst sizes, along with lognormal and Pareto models. Both models fit the distribution well, but the Pareto model is better. The tail of the cdf is close to straight, right out to the boundary of the measurement.

Since there is some evidence that the distribution of burst sizes is long-tailed, it would be useful to understand how it comes about. A first step is to look at the distribution of burst lengths (the number of transfers in a burst).

Figure 20 shows the distribution of burst lengths and two models. The lognormal model is a good fit for the cdf, which clearly has an increasing slope. The estimated tail curvature is 0.105, which has a negligible p-value under the Pareto model. We conclude that the tail behavior of burst lengths for this dataset is roughly lognormal.




Fig. 21. Distribution of burst lengths (number of transfers) for HTTP connections.

7.2 HTTP bursts

Charzinski [24] collected traces of HTTP activity on a LAN at a German university (Trace A) and a small ISP (Trace B). For both traces he measures the duration of TCP connections that contain one or more HTTP transfers. He measures the durations of these connections and reports that they show polynomial tail behavior. The ccdf of durations is approximately straight, with only a slight deviation in the extreme tail.

He also considers the number of requests per connection, which is analogous to burst length. In both traces the majority of connections (69% and 73%) contain a single request, but some contain hundreds. Figure 21 shows the ccdf of burst lengths for each trace. There are 455 992 connections in Trace A and 739 005 in Trace B.

For Trace A, the lognormal model fits the bulk reasonably well, but not the tail. The Pareto model fits the distribution well between 10 and 100 transfers per connection. The extreme tail deviates from this line, and may be dropping off with increasing slope, but there are very few data at this extreme. We

conclude that this dataset supports the Pareto model, but in this case the parameter is 2.3, so the distribution is not long-tailed in the sense required to explain self-similarity.

For Trace B, the Pareto model is a reasonable fit, but with the estimated parameter, 2.5, the distribution is not long-tailed. The lognormal model fits very well.

Comparing Figures 20 and 21, it is surprising how similar the distributions are, considering that they come from different places, times, and applications. In particular, it is surprising that ftp and HTTP have similar burst behavior, since for ftp transfers, users generally determine what files constitute a burst; for HTTP transfers, web designers do.

Nevertheless, ftp seven years ago and HTTP now have about the same percentage of single-transfer connections, and the distribution of burst lengths has the same shape and range. This observation suggests that there are common characteristics in the way information is organized into files, and the way files are organized into clusters (directories, pages) that are likely to be transported as a burst.

8 Conclusions

We have reviewed techniques for identifying long-tailed distributions, and applied them to datasets that have been reported as long-tailed. Unfortunately, no single test is sufficient to provide convincing evidence of a long-tailed distribution. Looking at previous claims for long-tailed distributions, we find that some are not well supported by the evidence. In other cases, the evidence is ambiguous.

- In our review of published observations we did not find compelling evidence that the distribution of file sizes is long-tailed.
- There is some evidence that interarrival times for TCP packets can be long-tailed, but longer traces are needed to make a stronger claim. We looked at interarrival times for several other kinds of events, and found only inconsistent evidence of long-tailed distributions.
- There is some evidence that the distribution of transfer times can be long tailed.
- The distribution of burst sizes for ftp and HTTP transfers appears to be long tailed. The distribution of burst lengths (number of transfers) is similar for ftp and HTTP traffic and does not appear to be long-tailed.
- In many cases the lognormal model describes the tail behavior of observed distributions as well as or better than the Pareto model. Thus, in cases

where lognormal models are more convenient, they can be used without sacrificing realism.

We have examined some of the causes that have been proposed for self-similar network traffic, and found that file sizes and transfer times, and interarrival times may not be sufficiently long-tailed to cause self-similarity, but burst sizes may be.

As ongoing work, we are examining the relationship between the number of transfers in a burst and the size of the burst, and looking for a fundamental explanation of the distribution of burst lengths.

Acknowledgments

Thanks to Mark Crovella (Boston University), Vern Paxson (ICIR) and Carey Williamson (University of Saskatchewan) for making their datasets available on the Web; Martin Arlitt (Hewlett-Packard) for providing processed data from the datasets he collected; Gordon Irlam for his survey of file sizes, and John Douceur (Microsoft Research) for sending me the Microsoft dataset. Also, many thanks to Joachim Charzinski (Siemens AG) for providing the burst lengths from his traces.

Thanks to Kim Claffy and Andre Broido (CAIDA), Mark Crovella, Rich Wol-ski (University of Tennessee) and Lewis Rothberg (University of Rochester) for reading drafts of this paper and making valuable comments. Finally, thanks to the reviewers who have read and helped improve earlier presentations of this work.

References

- [1] K. Park, W. Willinger, Self-similar network traffic: An overview, in: Self-Similar Network Traffic and Performance Evaluation, Wiley-Interscience, 2000, Ch. 1, preprint.
- [2] V. Paxson, S. Floyd, Wide-area traffic: the failure of Poisson modeling, IEEE/ACM Transactions on Networking 3 (1995) 226–244.
- [3] M. Parulekar, A. M. Makowski, M/G/ ∞ input process: a versatile class of models for network traffic, Tech. Rep. T.R. 96-59, Institute for Systems Research (1996).
- [4] W. Willinger, M. S. Taqqu, R. Sherman, D. V. Wilson, Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level, in: ACM SIGCOMM'95, 1995, pp. 100–113.

- [5] A. Feldmann, A. C. Gilbert, P. Huang, W. Willinger, Dynamics of IP traffic: a study of the role of variability and the impact of control, in: ACM SIGCOMM'99, 1999, pp. 301–313.
- [6] K. Park, G. Kim, M. E. Crovella, On the relationship between file sizes, transport protocols, and self-similar network traffic, Tech. Rep. BU-CS-96-016, Boston University, short version appeared in Proc. Fourth Int'l Conference on Network Protocols (August 1996).
- [7] M. E. Crovella, M. S. Taqqu, A. Bestavros, Heavy-tailed probability distributions in the World Wide Web, in: A Practical Guide To Heavy Tails, Chapman & Hall, 1998, pp. 3–26.
- [8] M. E. Crovella, M. S. Taqqu, Estimating the heavy tail index from scaling properties, *Methodology and computing in applied probability* 1 (1) (1999) 55–79.
- [9] M. F. Arlitt, C. L. Williamson, Web server workload characterization: the search for invariants, in: ACM SIGMETRICS'96, 1996, pp. 126–137, extended version.
- [10] M. Arlitt, T. Jin, Workload characterization of the 1998 World Cup Web site, Tech. Rep. HPL-1999-35R1, Hewlett-Packard (September 1999).
- [11] M. Arlitt, R. Friedrich, T. Jin, Workload characterization of a Web proxy in a cable modem environment, *ACM Sigmetrics Performance Evaluation Review* .
- [12] C. R. Cunha, A. Bestavros, M. E. Crovella, Characteristics of WWW client-based traces, Tech. Rep. BU-CS-95-010, Boston University (1995).
- [13] M. E. Crovella, A. Bestavros, Self-similarity in World-Wide Web traffic: evidence and possible causes, in: ACM SIGMETRICS'96, 1996, pp. 160–169.
- [14] P. Barford, A. Bestavros, A. Bradley, M. Crovella, Changes in Web client access patterns: Characteristics and caching implications, *World Wide Web 2* (1999) 15–28, special Issue on Characterization and Performance Evaluation.
- [15] G. S. Fishman, I. Adan, How heavy-tailed distributions affect simulation-generated time averages, Tech. Rep. UNC/OR TR03-2, University of North Carolina, Chapel Hill (April 2003).
- [16] G. Irlam, Unix file size survey – 1993, <http://www.base.com/gordon/ufs93.html>, accessed 31 May 2000.
- [17] J. R. Douceur, W. J. Bolosky, A large-scale study of file-system contents, in: ACM SIGMETRICS'99, 1999, pp. 59–70.
- [18] M. S. Taqqu, W. Willinger, R. Sherman, Proof of a fundamental result in self-similar traffic modeling, *Computer Communication Review* 27 (1997) 5–23.
- [19] M. E. Crovella, A. Bestavros, Explaining World Wide Web traffic self-similarity, Tech. Rep. BU-CS-95-015, Boston University (1995).
- [20] S. Deng, Empirical model of WWW document arrivals at access links, in: IEEE International Conference on Communication, 1996.

- [21] A. Feldmann, A. C. Gilbert, W. Willinger, T. Kurtz, The changing nature of network traffic: Scaling phenomena, *ACM Computer Communication Review* 28 (1998) 5–29.
- [22] V. Paxson, Empirically-derived analytic models of wide-area TCP connections, *IEEE/ACM Transactions on Networking* 2 (4) (1994) 316–336.
- [23] A. B. Downey, Evidence for long-tailed distributions in the Internet, in: *ACM SIGCOMM Internet Measurement Workshop*, 2001.
- [24] J. Charzinski, HTTP/TCP connection and flow characteristics, *Performance Evaluation* 42 (2000) 149–162.