Evidence for long-tailed distributions in the Internet

Allen B. Downey
Computer Science Department
Wellesley College
Wellesley, MA 02481
adowney@wellesley.edu

Abstract—**We review evidence that Internet traffic is characterized by long-tailed distributions of interarrival times, transfer times, burst sizes and burst lengths. We propose a new statistical technique for identifying long-tailed distributions, and apply it to a variety of datasets collected on the Internet. We find that there is little evidence that interarrival times and transfer times are long-tailed, but that there is some evidence for long-tailed burst sizes. We speculate on the causes of long-tailed bursts.**

I. Introduction

Numerous studies 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].

In previous work we presented evidence that the distribution of file sizes is lognormal, and not long-tailed [8]. If this claim is correct, it invites alternative explanations of self-similarity. There are several possibilities:

1. The distribution of interarrival times might be long-tailed. There is some evidence for this possibility, but also evidence to the contrary [2] [9] [10] [11].
2. 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.
3. Even if the length of individual transfers is not long-tailed, the length of bursts might be. From the network's point of view, burst sizes might be more relevant than transfer sizes.
4. Several groups have argued that TCP retransmission and/or congestion control are sufficient to produce self-similarity in network traffic, and that it is not necessary to assume that size or interarrival distributions are long-tailed [12] [13] [14] [15].
5. Another possibility is that network traffic is not truly self-similar. In the M/G/∞ model, if the distribution of service times is lognormal, the resulting count process is not self-similar and not long-range dependent [2], but over a range of time scales it may be statistically indistinguishable from a truly self-similar process.

Since much of this discussion is about long-tailed distributions, the next section discusses existing and new methods for identifying long-tailed distributions based on a sample.

The following sections discuss the first three options above, reviewing prior claims and reexamining proposed evidence.

We find that there is little evidence that the distribution of interarrival times (by any definition) is long-tailed. Similarly, there is only ambiguous support for long-tailed transfer times. We propose a structural model that leads us to expect transfer times, like file sizes, to be lognormal.

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.

II. 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

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

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.

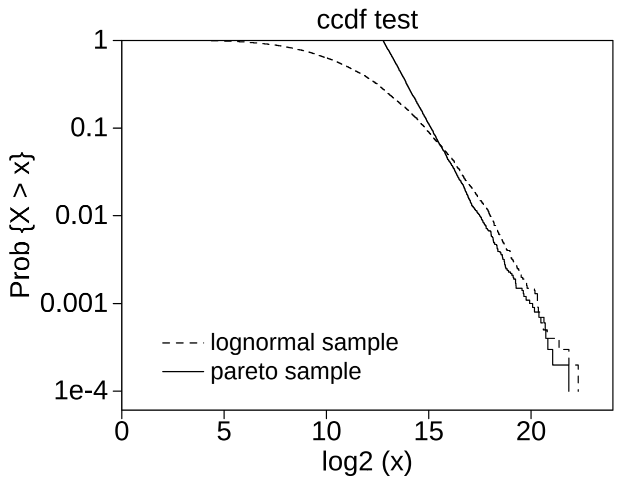

Fig. 1. ccdf of samples from lognormal and Pareto distributions with similar tail behavior.

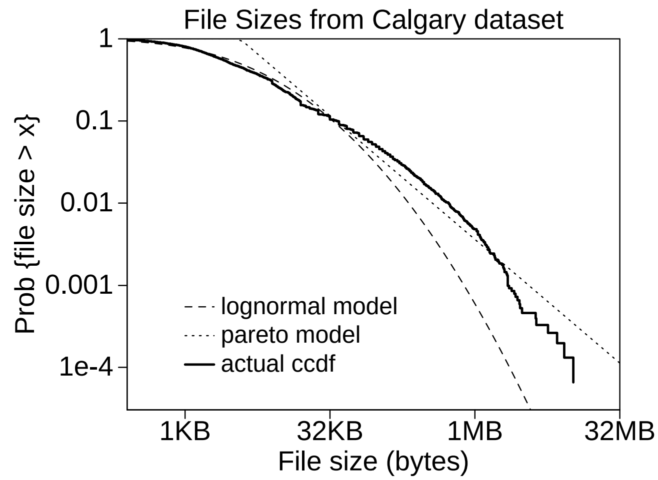

Fig. 2. ccdf of file sizes from a university web server.

The following sections discuss methods for identifying long-tailed distributions.

A. ccdf test

There are several empirical behaviors 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 behavior, and at least out to the boundary of the measurement.

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

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 ccdf can appear long-tailed, at least to a point. The definitive characteristic of these distributions is that the ccdf eventually drops away with increasing slope.

B. Using `aest`

Crovella and Taqqu developed a tool called `aest` that estimates the slope parameter of a Pareto distribution based on a sample [16]. 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. In the Pareto sample it identifies 271 such points; in the lognormal sample it identifies 312 points. By this (admittedly coarse) measure, the lognormal sample appears to be *more* long-tailed than the Pareto sample.

We conclude that `aest` cannot distinguish long-tailed and lognormal distributions based on samples.

C. Model fitting

A standard way to choose among alternate 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] and available from the the Internet Traffic Archive (`http://ita.ee.lbl.gov`).

We fitted a Pareto model using `aest` to estimate α, and choosing the lower bound by eye. We fitted a lognormal model by conventional moment-matching.

The bulk of the distribution clearly fits the lognormal model better, but the tail behavior is harder to characterize. 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 distribu-

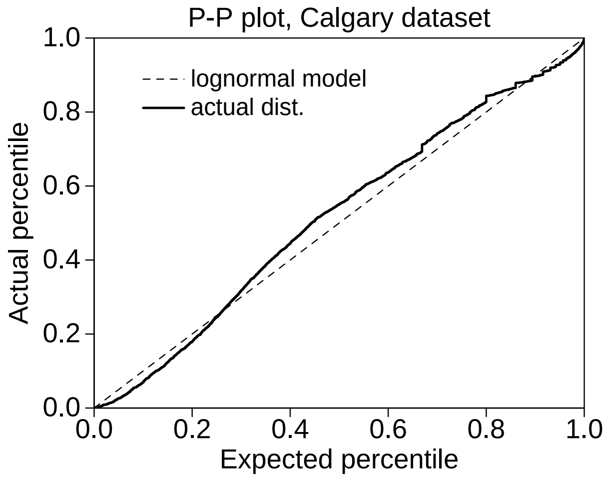

Fig. 3. Percentile-percentile plot of sizes from Calgary dataset.

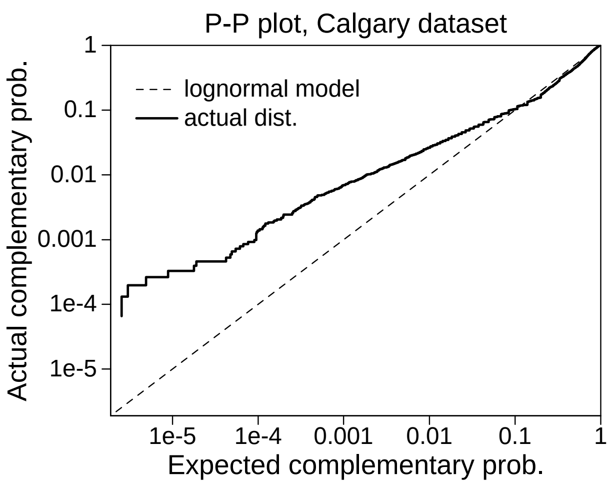

Fig. 4. P-P plot on complementary log axes.

tion: increasing steepness in the extreme tail. So in this case quality of fit may be misleading.

Although the fitted models are useful as a reference point, they do not provide a mechanical, quantitative way to identify long-tailed distributions.

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

Figure 3 shows a P-P plot for the Calgary dataset. It shows that there is some deviation from the model throughout the distribution, but that overall the agreement is good.

To examine the tail behavior, we can plot the complementary probabilities on log axes. Figure 4 shows this transformation for the same dataset. Here the divergence

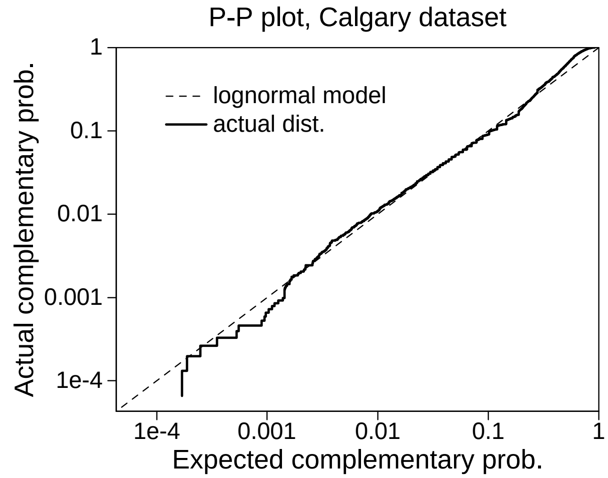

Fig. 5. Complementary P-P plot with alternate model.

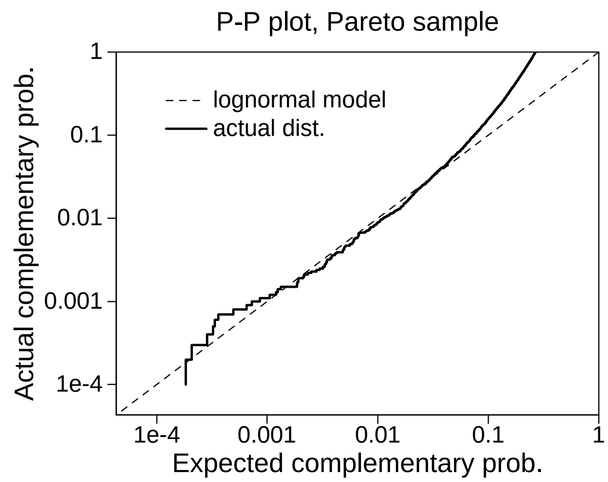

Fig. 6. Complementary P-P plot of a Pareto sample.

of the tail from the model is more apparent.

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

The interpretation of a P-P plot depends on the choice of the parameters. For Figure 4, we estimated the lognormal model by moment-matching. But if we choose an alternate model to match the tail behavior, we get Figure 5, which shows very good agreement with the tail of the distribution.

In fact, it is often possible to find a lognormal model that fits the tail, even if the sample actually comes from a Pareto distribution. Figure 6 shows the complementary P-P plot of the Pareto sample from Figure 1, using the parameters of a lognormal model with a similar tail. The discrepancy in the bulk of the distribution is clear, but the tail behavior is difficult to distinguish from the model.

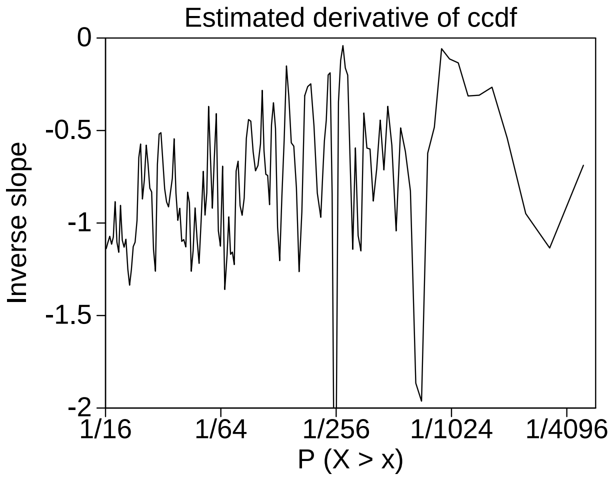

Fig. 7. Numerical derivative of the ccdf of a lognormal sample.

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

E. Curvature

Looking at Figure 1, the characteristic difference in the tail behaviors 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. For example, we applied a three-point estimate of the first derivative to the lognormal sample from Figure 1. The result is Figure 7. The horizontal axis is $P(X > x)$ on a log scale. The vertical axis is the estimated value of $dx/dP(X > x)$, which is the inverse slope of the ccdf.

Clearly the result is noisy, so differentiating this function again in order to estimate curvature is out of the question. Fortunately, we don't need to know the curvature as a function; it is sufficient to estimate the overall curvature in the ccdf, which is just the change in slope.

Figure 7 shows a clear upward trend, meaning that the slope of the ccdf is becoming more negative. We can confirm the trend statistically simply by fitting a line to the estimated derivative. In this case the measured trend is 0.0620. We call this statistic the "tail curvature."

By itself the number is meaningless, but we can compare it to the tail curvature of the Pareto sample, which is 0.00836, more than 7 times smaller. We expect the curvature of a sample from a Pareto distribution to be near zero.

Of course, the calculation depends on our definition of where the "tail" begins. For this dataset, we chose $Pr(X > x) < 1/16$, which we think is a reasonable definition of a tail, and also the point in Figure 1 where the ccdfs intersect.

Using tail curvature as a summary statistic provides a procedure for testing the hypothesis that a sample comes from a distribution with a Pareto tail. Assume we have a sample of n points.

1. Measure 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. The null hypothesis is that the sample came from a Pareto distribution with parameter α.
3. 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.
4. Calculate d, the difference between the curvature of the original sample and μ.
5. 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.
6. If the p-value falls below a threshold of confidence, we can reject the null hypothesis.

Applying this test to the samples in Figure 1:

1. The curvature of the lognormal sample is 0.0620.
2. The estimate of α is 1.42.
3. For 1000 samples from a Pareto distribution with n=10,000 and $\alpha = 1.42$, $\mu = 0.0037$. [1]
4. The sample differs from μ by 0.0583.
5. Out of 1000 samples, 8 differed from μ by as much, so the p-value is 0.008.
6. Thus, with a high degree of confidence we can reject the hypothesis that this sample comes from a distribution with a Pareto tail.

Repeating this procedure with the Pareto sample, we find that the measured tail curvature (0.008363) is not unusual; 85% of samples differ from the mean by as much.

So this procedure classifies the two distributions correctly, even though the tails of their ccdfs are visually similar.

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

[1] In every case we have examined, μ is close to, but slightly above, zero, which is the value we expect theoretically. This makes us suspect that tail curvature is a biased statistic, but this bias does not affect the test procedure.

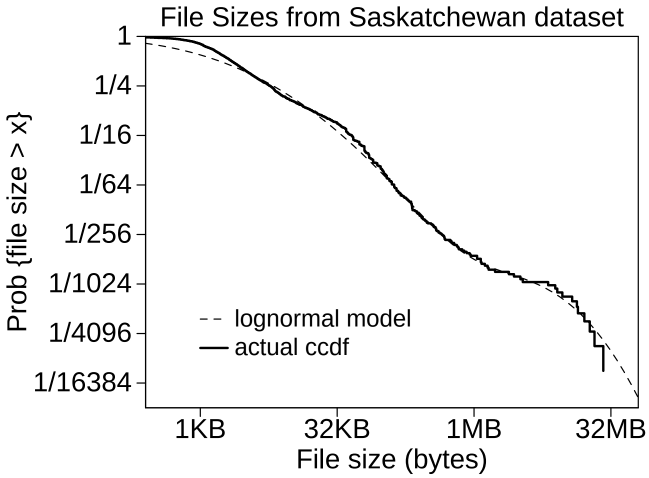

Fig. 8. ccdf of file sizes from a university web server.

99% of the Pareto distributions while rejecting 99% of the lognormal distributions.

For the Calgary dataset, the tail curvature is 0.141, which is much higher than we would expect from a Pareto distribution. In fact, out 1000 Pareto samples, the highest tail curvature is 0.084.

Thus the tail curvature supports our claim, based on the ccdf test, that the tail behavior of this dataset is not characteristic of a long-tailed distribution.

Of course, this test is not infallible. In particular, many multimodal lognormal distributions have a ccdf that is approximately straight. For example, Figure 8 shows the ccdf of file sizes on a web server at the University of Saskatchewan (also collected by Arlitt and Williamson).

The fitted model is a two-mode lognormal chosen by eye to match the tail behavior. Because there is no consistent trend in the steepness of the ccdf, the estimated tail curvature is small and negative. Thus, according to the curvature test, this ccdf is reasonably likely to come from a Pareto distribution (p-value = .08).

If the two-mode lognormal model is right, and there really is an underlying distribution with lognormal tail behavior, then the curvature test fails to identify it.

F. *Explanatory models*

As a practical matter, it may not be possible to identify long-tailed distributions based on samples. For any sample, it is possible to construct a model that is not long-tailed and that matches the behavior of the sample out to the boundary of the observation.

This possibility can be convenient for modelers. For example, Feldmann and Whitt [17] present an algorithm for approximating long-tailed distributions using a mixture of exponentials (a hyper-exponential). They do not claim that this model is explanatory; their intent is to produce a model that is convenient for analysis.

From an instrumentalist point of view, that's all there is to it. Modelers can choose whichever model (long-tailed or not) is sufficiently accurate and tractable for their purposes.

But this view does not provide a satisfying explanation of why file sizes are (or are not) long-tailed. In turn, that leaves us without an explanation of why Internet traffic is self-similar.

The issue may be resolved by explanatory models. For the lognormal and Pareto models of file sizes, there are corresponding structural models that have been proposed.

For the Pareto model, Carlson and Doyle propose a physical model based on Highly Optimized Tolerance (HOT), in which web designers, trying to minimize download times, divide the available information into files so that the most frequent downloads are the smallest. The result is a distribution of file sizes with polynomial tail behavior [18] [19].

For the lognormal model, we have proposed a model of user behavior in which most new files are created by copying, modifying or translating existing files. The result is that file sizes diffuse over time, producing lognormal distributions and mixtures of lognormals [8].

Both of these models are based on unrealistic simplifications of human behavior. Up to a point, this simplicity is a virtue; if a more complicated model is necessary to produce the phenomenon we are interested in, we have less of a sense that we understand the phenomenon.

The question that remains is whether these models are robust to deviations from their assumptions. For example, in the HOT model, it is not clear how many web designers actually tune the contents of their sites for optimal downloading. But if only a few do, or some do only approximately, is that enough to produce long tails?

For its part, the diffusion model omits some common file operations, like concatenation, and ignores any bias users might have against large files. The question is how these behaviors affect the shape of the resulting distribution.

If these questions are answered, and one of these models provides a satisfactory explanation of file size distributions and the other doesn't, then it might be the explanatory model, rather than the data, that allows us to say whether file sizes are long-tailed.

In the rest of this paper we examine evidence for long-tailed distributions using all the tools described in this section.

III. INTERARRIVAL TIMES

Even if the distribution of transfer times is not long-tailed, if the distribution of interarrival times is, then the ON/OFF model yields asymptotic self-similar behavior [20].

Several authors have presented evidence that the distribution of interarrival times is long-tailed. In this section we review these claims.

A. TCP Packet Interarrivals

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 ccdf 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 from the ITA: LBL PKT-3, LBL PKT-4, LBL PKT-5, and DEC PKT-1. The first three are from Lawrence Berkeley Labs; the last is a one-hour trace of TCP traffic between Digital Equipment Corporation (DEC) and the rest of the world. It was 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 ccdf of the interarrival times.

The four datasets yield similar distributions, so we aggregated them into a single dataset. Figure 9 shows the resulting ccdf, 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 ccdf starts to fall away from the Pareto model with increasing slope. The lognormal model does not fit this extreme tail particularly well either, but qualitatively it may be a better description of the tail behavior.

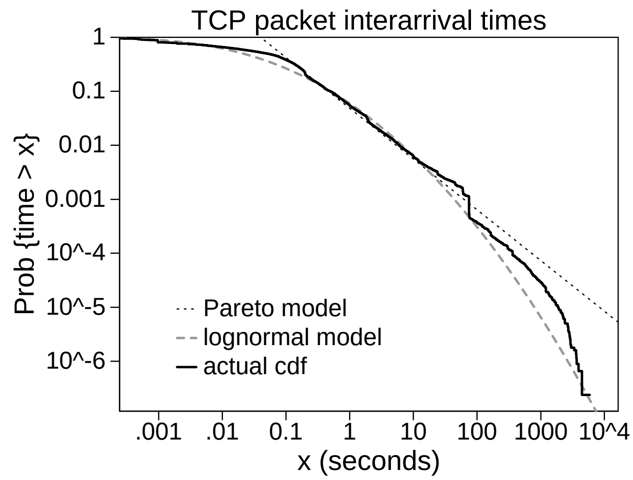

Fig. 9. ccdf of interarrival times for TCP packets.

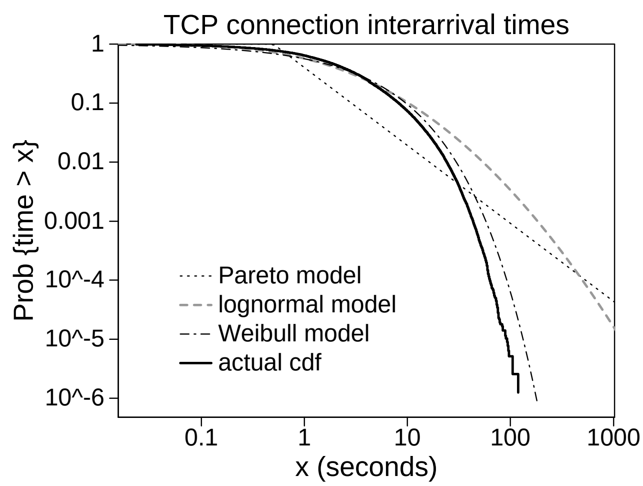

Fig. 10. ccdf of time between TCP connections from LBL traces.

The estimated tail curvature for this ccdf is 0.040, which has a negligible chance of occurring if the underlying model has a Pareto tail ($p < 0.001$).

We conclude that there is some support for the Pareto model of interarrival times, but the extreme tail behavior is characteristic of a non-long-tailed distribution.

B. TCP Connection Interarrivals

Feldmann [21] examines times between TCP connections in traces from several networks and estimates four models to fit the empirical cdfs: Weibull, Pareto, lognormal and exponential. For all datasets, the Weibull distribution is the best fit for the bulk of the distribution. She does not examine the tail behavior or make any claim about whether the distribution is long-tailed in the sense we are

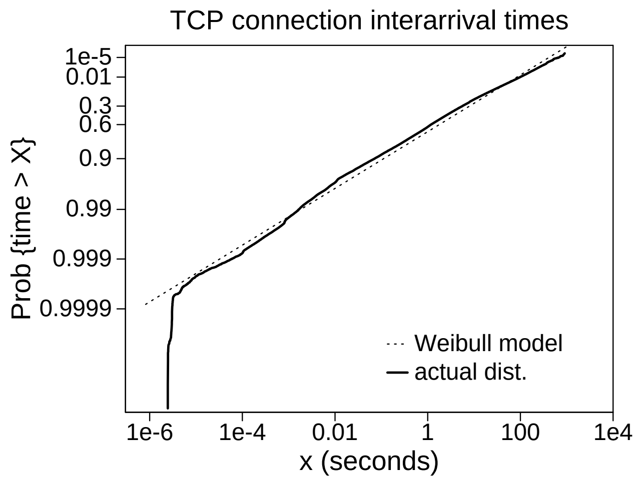

Fig. 11. Weibull test for TCP connection interarrivals.

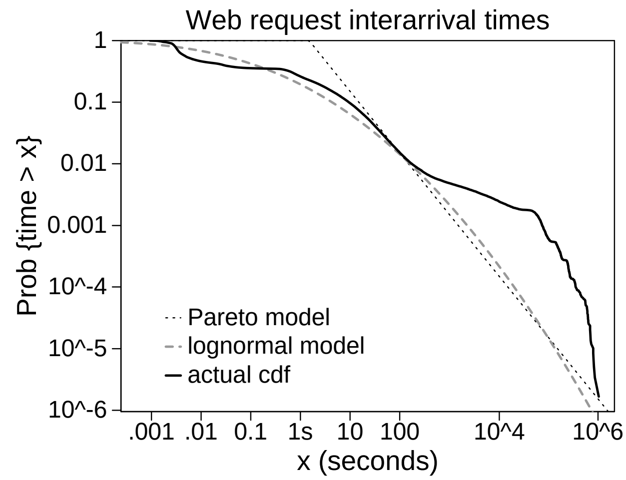

Fig. 12. ccdf of time between web requests from BU traces.

using here.

To investigate the times between TCP connections, we obtained the LBL CONN-7 trace from the ITA and computed the ccdf of the 782,280 interarrival times, shown in Figure 10.

The figure also shows three models fitted to the data: lognormal, Pareto and Weibull. For the Weibull model we used the estimation technique from [21]. The lognormal and Pareto models are not a good fit for the data. The Weibull model is a reasonable description of the tail behavior.

There is another method to test whether a sample fits a Weibull distribution. If we plot $\log \log\{1/ccdf(x)\}$ versus $\log x$, a Weibull distribution yields a straight line [22]. Figure 11 shows the Weibull test for this distribution. Except for the extreme tail, the Weibull model is a good match for the data.

Since the Weibull distribution is not long-tailed, we conclude that this dataset does not support the hypothesis that the distribution of interarrival times for TCP connections is long-tailed.

C. HTTP Request Interarrivals

Crovella and Bestavros [23] 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.

We obtained their traces, which we call the BU dataset, from their web site. The traces contain logs from web browsers on 37 workstations in public labs at Boston University. 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 12.

As usual, we estimated a lognormal model using moment estimators and a Pareto model using `aest`. In the bulk of the distribution (below the 99th percentile), the lognormal model is a better fit for the data; in the tail, neither model fits the data well. The extreme tail exhibits the characteristic behavior of a non-long-tailed distribution.

The tail curvature is 0.144625, which has a negligible chance of occurring if the underlying model has Pareto tail behavior ($p < 0.001$).

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 claim 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 [24]. 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 ccdf, or compare the Pareto model to the alternatives.

Overall, there is little evidence that the times between WWW requests form a long-tailed distribution.

IV. 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.

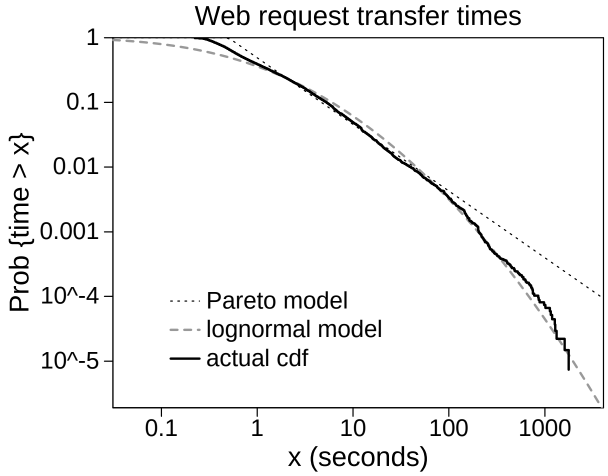

Fig. 13. ccdf of transfer times for web requests from BU traces.

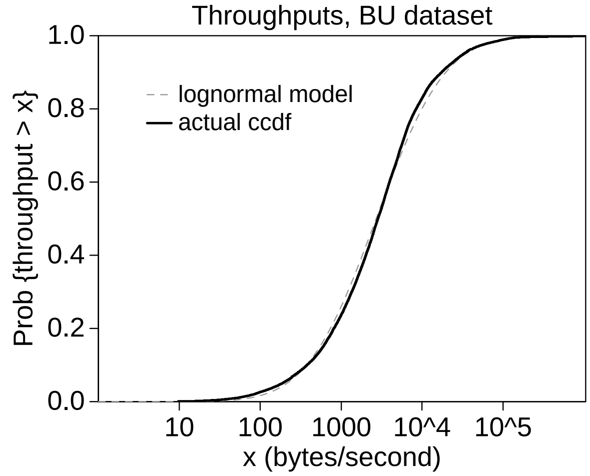

Fig. 14. Distribution of throughputs from BU traces.

In this section we investigate the relationship between file sizes and transfer times for HTTP and ftp transfers.

A. HTTP transfer times

Crovella and Bestavros examine the distribution of transfer times in the BU dataset, and report that it is long-tailed [25].

We performed the same analysis for the 135,357 transfers they observed; Figure 13 shows the resulting ccdf along with a Pareto model and a lognormal model. 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 0.041889, which has a negligible chance of appearing in a sample this size from a Pareto distribution (p-value < 0.001).

We conclude that this dataset does not support the hypothesis that transfer times are long-tailed.

B. HTTP throughput

In this section we investigate the variability of throughput across network paths and time. We find that the distribution of throughput is approximately lognormal. Based on this observation we propose a model to explain the relationship of file sizes and transfer times.

For each web request in the BU dataset we have the size of the file and the transfer time. Dividing size by transfer time yields the throughput of the transfer. As expected, this value varies greatly: different transfers use different network paths, and even on the same path throughput varies over time.

Figure 14 shows the cdf of throughput for the 135,357 transfers in the BU dataset. The distribution fits the lognormal model well.

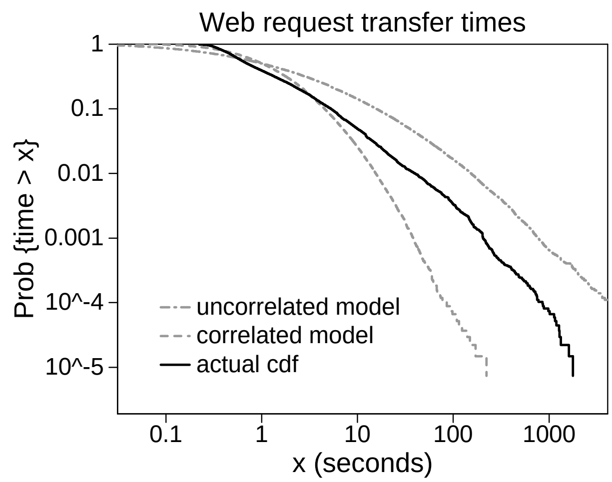

Fig. 15. Simulations of web transfer times.

A lognormal distribution of throughputs suggests a model for the relationship between transfers sizes and times. If files are chosen at random from the set of available files, and the throughput is a random value that is independent of file size, then we can express the transfer time as $t = s/b$, where t is transfer time, s is file size, and b is bandwidth. Thus $\log t = \log s - \log b$.

If s is lognormally distributed, as proposed in [8], and b is lognormally distributed, as seen here, then $\log t$ is the difference of two normal random variables, which is also normal. Thus, t is lognormal.

There are, however, two problems with this model. First, effective bandwidth and size are not independent. In this dataset, the correlation of $\log s$ and $\log b$ is 0.70, which means that larger transfers achieve larger throughput.

To examine the effect of this correlation, we ran a simple simulation of network transfers. The simulation uses the lognormal approximations of the distributions of size and throughput. In the uncorrelated version, it chooses random sizes and throughputs independently and generates a sam-

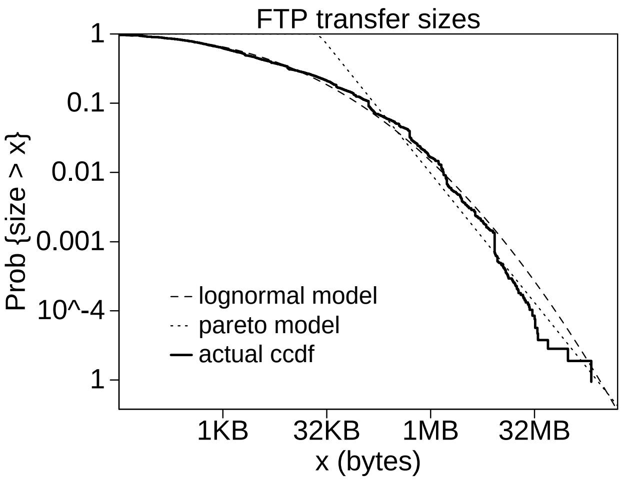

Fig. 16. ccdf of ftp transfer sizes from LBL traces.

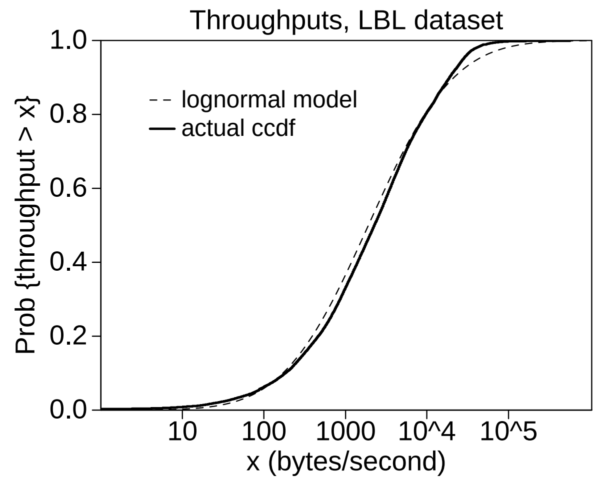

Fig. 18. Distribution of throughputs from LBL traces.

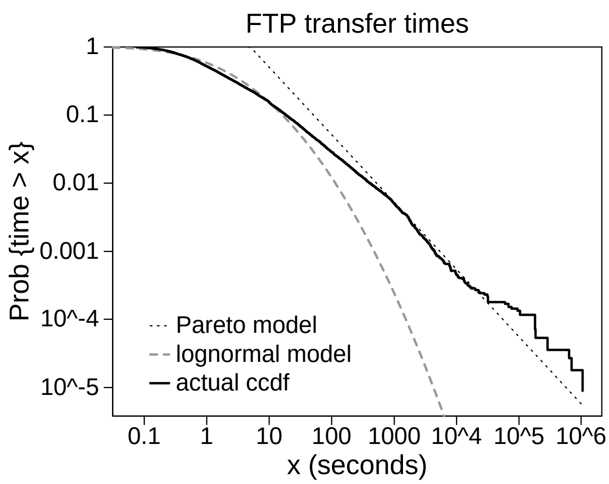

Fig. 17. ccdf of ftp transfer times from LBL traces.

ple of transfer times. In the correlated version, it chooses sizes and throughputs with a correlation of 0.70.

Figure 15 shows the actual distribution of times along with the two simulated models. Neither model is a good match for the data, which suggests that we have not captured the details of the relationship between size and throughput.

Nevertheless, this figure demonstrates the effect that correlation has on the tail of the distribution of transfer times. Since larger transfers tend to achieve higher throughput, the variation in transfer times is compressed. Thus, we expect the distribution of transfer times to be less long-tailed than the distribution of sizes.

In the next section we apply the same analysis to ftp transfers.

C. *ftp transfer times*

Paxson [26] examines the sizes of ftp transfers and reports that the distribution fits a lognormal distribution well. He does not discuss the relationship between transfer sizes and transfer times.

To examine ftp transfer sizes and times, we used the LBL CONN-7 dataset again, and extracted the 105,542 transfers that used the ftp-data protocol, were successful, and transported a non-zero number of bytes. Figure 16 shows the ccdf of transfer sizes; Figure 17 shows the ccdf of transfer times.

We agree with Paxson that the distribution of sizes fits the lognormal distribution, both in the bulk, as shown in [26] 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 0.112604, which has a negligible chance of coming from a Pareto distribution.

On the other hand, the lognormal model does not fit the ccdf 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 ccdf 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-value > 0.95).

D. *ftp throughput*

Again, we explored the relationship between these transfer times and the distribution of sizes. For each connection, we computed the throughput by dividing the transfer size by the transfer time. Figure 18 shows the resulting ccdf.

Again, the distribution is roughly lognormal, although the high end appears to be compressed, possibly by the hardware limitations of the local network. Again, there is a strong correlation between size and throughput: $\rho = 0.73$.

Using the lognormal approximations of the distributions of sizes and throughputs, we simulated 105,542 transfers

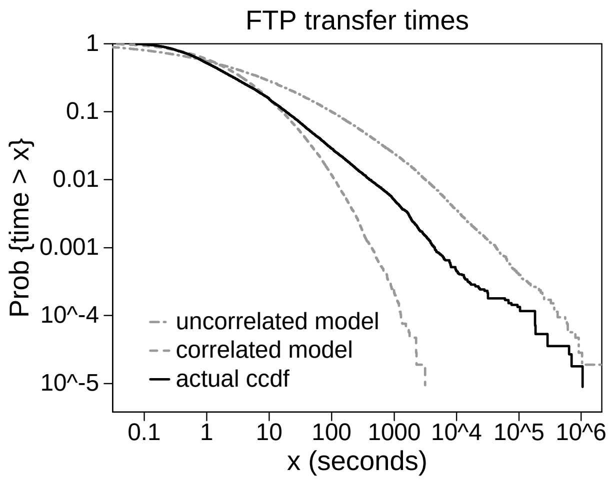

Fig. 19. Simulations of ftp transfer times.

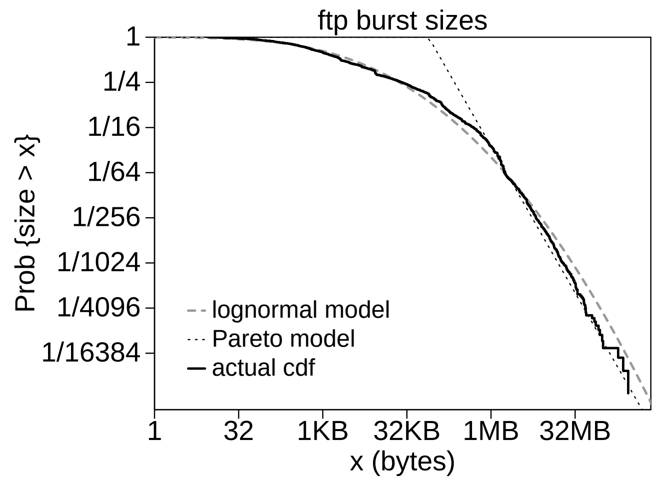

Fig. 20. Distribution of burst sizes from LBL traces.

with and without correlation between size and throughput. Figure 19 shows the results.

As in the previous example, the models fail to reproduce the actual ccdf, but we believe this simulation suggests an alternate way to interpret the distribution of transfer times. This distribution is approximately lognormal because the distributions of size and throughput are lognormal, but it is compressed somewhat by the upper bound the network imposes on throughput, and by the correlation between size and throughput. This compression makes the ccdf straighter but does not change its extreme tail behavior.

In summary, there is some evidence that the distribution of ftp transfer times is long-tailed, but we believe that there is an alternate explanation for these data that does not require long tails.

V. Transfer Bursts

The motivation for investigating the sizes of transfer bursts is the idea 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.

A. ftp bursts

Paxson [26] 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 ccdf, but he reports that the largest 5% of the bursts are well-modeled using a Pareto distribution. This claim appears again in [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 number reported by Paxson and Floyd).

We computed the total number of bytes in each burst; Figure 20 shows the distribution of these burst sizes, along with lognormal and Pareto models. The parameter estimated by `aest` is 1.0, but the Pareto model in the figure has $\alpha = 1.3$, which is a better fit for the tail.

Both models fit the distribution well, but the Pareto model is better. The tail of the ccdf is close to straight, right out to the boundary of the measurement. Nevertheless, there is some curvature; the measured tail curvature is 0.073, which has a negligible p-value.

We conclude that there is some evidence that the distribution of burst sizes is long-tailed. This evidence is not clear-cut, though, so it would be useful to understand how it comes about.

In the next section we examine the distribution of burst lengths (number of transfers) and its relationship to burst sizes (number of bytes).

B. Burst lengths

A first step is to look at the distribution of burst lengths. If we know the distributions of file sizes and burst lengths, we can model a burst by the following process:

- Choose a value of n from the distribution of burst lengths.
- Choose n sizes from the distribution of sizes and sum them to get the total burst size.

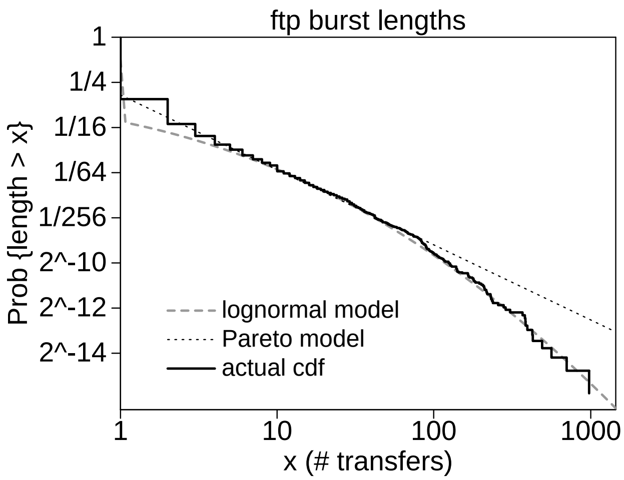

Fig. 21. Distribution of burst lengths (number of transfers) from the LBL traces.

Thus the distribution of burst sizes is a mixture of distributions with parameter n, where n is the number of times the distribution of transfer sizes is convolved with itself.

Figure 21 shows the distribution of burst lengths for the LBL dataset. The lognormal model has two modes: the first mode contains the 85% of bursts that contain a single transfer. The parameters of the second mode were chosen by eye. The lognormal model is a good fit for the ccdf, which clearly has an increasing slope. The estimated tail curvature is 0.104583, which has a negligible p-value.

We conclude that the tail behavior of burst lengths for this dataset is roughly lognormal.

C. HTTP bursts

Charzinski [27] 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 22 shows the ccdf of burst lengths for each trace. There are 455,992 connections in Trace A and 739,005 in Trace B.

Trace A is hard to characterize. The Pareto model estimated by `aest` does not fit well, but there is a long linear section 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.

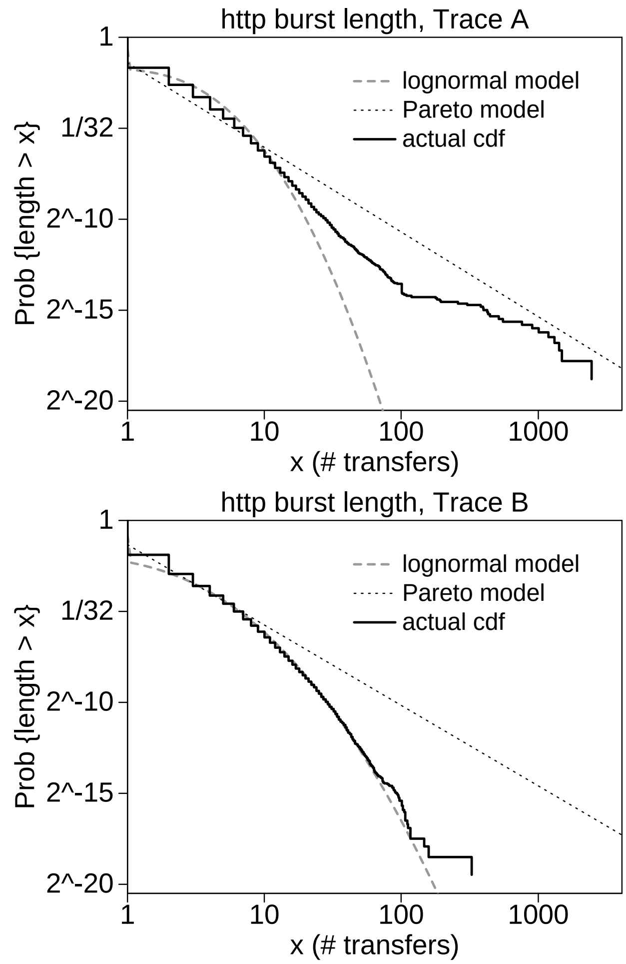

Fig. 22. Distribution of burst lengths (number of transfers) for HTTP connections.

The tail curvature is -0.097, which is unlikely for a Pareto distribution, but even less likely for a lognormal distribution.

The lognormal model, as in the previous section, has two modes. The first mode, at 1, deals with single-transfer connections. The second mode has parameters chosen by eye to fit the bulk of the distribution. It fits the bulk reasonably well, but not the tail.

Trace B is easier to characterize. It looks like a lognormal distribution.

Comparing Figures 21 and 22, 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.

Unfortunately, these datasets do not provide a clear picture of the distribution of burst sizes. There are straight segments in these ccdfs that suggest Pareto tails. In that case we have a hint about the origin of long-tailed burst sizes, but we still have to ask why burst lengths are Pareto.

If burst lengths are lognormal, as the preponderance of evidence suggests, then there are two questions to answer: why are burst lengths lognormal, and what effect does this have on the distribution of burst sizes?

If the distributions of burst lengths and transfer sizes are lognormal, it is possible that the result is long-tailed burst sizes. It is possible for a mixture of short-tailed distributions to yield a long-tailed distribution ([22], page 574). Specifically, it is possible that a lognormal mixture of sums of lognormal variates yields a long-tailed distribution. As current work we are exploring this possibility.

VI. Conclusions

- There is little evidence that the distribution of interarrival times, by whatever definition, is long-tailed.
- There is some evidence that the distribution of transfer times can be long-tailed, but we have proposed an alternate explanation for this evidence that does not involve long tails.
- The distribution of burst sizes, for both ftp and HTTP transfers, appears to be long tailed.
- If the distribution of burst lengths is long-tailed, then that explains long-tailed burst sizes. It is an open question why burst lengths should be long-tailed.
- If the distribution of burst lengths is lognormal, which seems more likely, then it is an open question whether lognormal burst lengths can yield long-tailed burst sizes.

Acknowledgments

I would like to thank Vern Paxson (Lawrence Berkeley National Lab), Mark Crovella (Boston University), Carey Williamson (University of Saskatchewan) and Martin Arlitt (Hewlett-Packard) for making their datasets available on the web. Also, many thanks to Joachim Charzinski (Siemens AG) for providing the burst lengths from his traces.

References

[1] Kihong Park and Walter Willinger, "Self-similar network traffic: An overview," in *Self-Similar Network Traffic and Performance Evaluation*, chapter 1. Wiley-Interscience, 2000, Preprint.

[2] Vern Paxson and Sally Floyd, "Wide-area traffic: the failure of Poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, pp. 226–244, 1995.

[3] M. Parulekar and 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] Walter Willinger, Murad S. Taqqu, Robert Sherman, and Daniel 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] Anja Feldmann, Anna C. Gilbert, Polly Huang, and Walter 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] Kihong Park, Gitae Kim, and Mark E. Crovella, "On the relationship between file sizes, transport protocols, and self-similar network traffic," Tech. Rep. BU-CS-96-016, Boston University, August 1996, Short version appeared in Proc. Fourth Int'l Conference on Network Protocols.

[7] Mark E. Crovella, Murad S. Taqqu, and Azer Bestavros, "Heavy-tailed probability distributions in the World Wide Web," in *A Practical Guide To Heavy Tails*, pp. 3–26. Chapman & Hall, 1998.

[8] Allen B. Downey, "The structural cause of file size distributions," in *IEEE MASCOTS'01*, 2001.

[9] Martin F. Arlitt and Carey L. Williamson, "Web server workload characterization: the search for invariants," in *ACM SIGMETRICS'96*, May 1996, pp. 126–137, Extended version.

[10] Paul Barford and Mark E. Crovella, "Generating representative Web workloads for network and server performance evaluation," in *Performance '98/SIGMETRICS '98*, July 1998, pp. 151–160.

[11] Anja Feldmann, Anna C. Gilbert, Walter Willinger, and T.G. Kurtz, "The changing nature of network traffic: Scaling phenomena," *ACM Computer Communication Review*, vol. 28, pp. 5–29, April 1998.

[12] Jon M. Peha, "Retransmission mechanisms and self-similar traffic models," in *IEEE/ACM/SCS Communication Networks and Distributed Systems Modeling and Simulation Conference*, January 1997, pp. 47–52.

[13] András Veres and Miklós Boda, "The chaotic nature of TCP congestion control," in *IEEE INFOCOM'00*, March 2000.

[14] W. Feng and P. Tinnakornsrisuphap, "The adverse impact of the TCP congestion-control mechanism in heterogeneous computing systems," in *International Conference on Parallel Processing (ICPP'00)*, August 2000.

[15] Liang Guo, Mark Crovella, and Ibrahim Matta, "TCP congestion control and heavy tails," in *IEEE INFOCOMM'01*, 2001.

[16] Mark E. Crovella and Murad S. Taqqu, "Estimating the heavy tail index from scaling properties," *Methodology and computing in applied probability*, vol. 1, no. 1, pp. 55–79, 1999.

[17] Anja Feldmann and Ward Whitt, "Fitting mixtures of exponentials to long-tail distributions to analyze network performance models," in *IEEE INFOCOM'97*, 1997.

[18] Jean M. Carlson and John Doyle, "Highly optimized tolerance: a mechanism for power laws in designed systems," *Physics Review E*, vol. 60, no. 2, August 1999.

[19] Xiaoyun Zhu, Jie Yu, and John Doyle, "Heavy tails, generalized coding, and optimal web layout," in *IEEE INFOCOM'01*, April 2001.

[20] Murad S. Taqqu, Walter Willinger, and Robert Sherman, "Proof of a fundamental result in self-similar traffic modeling," *Computer Communication Review*, vol. 27, pp. 5–23, 1997.

[21] Anja Feldmann, "Characteristics of TCP connection arrivals," Tech. Rep., AT&T Labs–Research, 1998.

[22] Norman Johnson, Samuel Kotz, and N. Balakrishnan, *Continuous*

univariate distributions, Vol. 1 & 2, John Wiley and Sons, New York, 1994.

[23] Mark E. Crovella and Azer Bestavros, "Explaining World Wide Web traffic self-similarity," Tech. Rep. BU-CS-95-015, Boston University, 1995.

[24] Shuang Deng, "Empirical model of WWW document arivals at access links," in *IEEE International Conference on Communication*, June 1996.

[25] Mark E. Crovella and Azer Bestavros, "Self-similarity in World-Wide Web traffic: evidence and possible causes," in *ACM SIGMETRICS'96*, May 1996, pp. 160–169.

[26] Vern Paxson, "Empirically-derived analytic models of wide-area TCP connections," *IEEE/ACM Transactions on Networking*, vol. 2, no. 4, pp. 316–336, August 1994.

[27] Joachim Charzinski, "HTTP/TCP connection and flow characteristics," *Performance Evaluation*, vol. 42, pp. 149–162, 2000.