To appear at MASCOTS'01

The structural cause of file size distributions

Allen B. Downey
Wellesley College
adowney@wellesley.edu

Abstract

We propose a user model that explains the shape of the distribution of file sizes in local file systems and in the World Wide Web. We examine evidence from 562 file systems, 38 web clients and 6 web servers, and find that this model is an accurate description of these systems. We compare this model to an alternative that has been proposed, the Pareto model. Our results cast doubt on the widespread view that the distribution of file sizes is long-tailed; we discuss the implications of this conclusion for proposed explanations of self-similarity in the Internet.

Keywords: *File sizes, lognormal distribution, long-tailed distribution, self-similarity.*

1. Introduction

Numerous studies have reported traffic patterns in the Internet that show characteristics of self-similarity (see [17] for a survey). Most proposed explanations are based on the assumption that the distribution of transfer times in the network is long-tailed [19] [18] [22] [12]. In turn, this assumption is based on the assumption that the distribution of file sizes is long-tailed [16] [9].

We contend that the distribution of file sizes in most systems fits the lognormal distribution. We support this claim with empirical evidence from a variety of systems and also with a model of user behavior that explains why file systems tend to have this structure.

We argue that the proposed model is a better fit for the data than the long-tailed model, and furthermore that our user model is more realistic than the explanations for the alternative.

We conclude that there is insufficient evidence for the claim that the distribution of file sizes is long-tailed. This result creates a problem for existing explanations of self-similarity in the Internet. We discuss the implications and review alternatives.

1.1. What does "long-tailed" mean?

In the context of self-similarity, a long-tailed distribution must have a hyperbolic tail; that is

$$Pr[X \geq x] \sim cx^{-\alpha} \quad \text{as} \quad x \rightarrow \infty \tag{1}$$

where X is a random variable, c is a constant, and α is a shape parameter that determines how long-tailed the distribution is.

Other definitions of long-tailed are used in other contexts. This definition is appropriate for us because it describes the asymptotic behavior that is required to produce self-similarity [20]. By this definition, the lognormal distribution is not long-tailed [19].

2. Distribution of file sizes

In this section we propose a model of the operations that create new files and show that a simulation of this model yields a distribution of file sizes that is a good match for the distributions that appear on real file systems.

Then we show that the simulator is equivalent to a numerical method for solving a partial differential equation (PDE). We show that the solution of this PDE is the analytic form of the result of the simulation. Finally, we use the analytic form to estimate the parameters of observed distributions and measure the goodness-of-fit of the model.

We find that the model describes real file systems well. We conclude that the user behavior described by the model explains the observed shape of the distribution of file sizes.

2.1. User model

Thinking about how users behave, we can list the most common operations that create new files:

copying: The vast majority of files in most file systems were created by copying, either by installing software (operating system and applications) or by downloading from the World Wide Web.

translating and filtering: Many new files are created by translating a file from one format to another, compiling, or by filtering an existing file.

editing: Using a text editor or word processor, users add or remove material from existing files, sometimes replacing the original file and sometimes creating a series of versions.

Thus we assert that many file-creating operations can be characterized as linear file transformations: a process reads a file as input and generates a new file as output, where the size of the new file depends on the size of the original.

This assertion suggests a model for the evolution of a file system over time: assume that the system starts with a single file with size s^*, and that users repeat the following steps:

1. Select a file size, s, at random from the current distribution of file sizes.

2. Choose a multiplicative factor, f, from some other distribution.

3. Create a new file with size fs and add it to the system.

It is not obvious what the distribution of f should be, but we can make some assumptions. First, we expect that the most common operation is copying, so the mode of f should be 1. Second, thinking about filtering and translations, we expect that it should be about as common to double the size of a file or halve it; in other words, we expect the distribution of f to be symmetric on a log axis.

In Section 2.4 we show that, due to the Central Limit Theorem, the shape of this distribution has little effect on the shape of the resulting distribution of sizes. For now we will choose a distribution of f that is lognormal with the mode at 1 and an unspecified variance.

There are, then, two parameters in this model, the size of the original file, s^*, and the standard deviation of the distribution of f, which we call γ. The parameter s^* determines the mode of the final distribution; γ controls the dispersion.

Next we see if this model describes real systems. In November 1993 Gordon Irlam conducted a survey of file systems. He posted a message on Usenet asking UNIX system administrators to run a script on their machines and mail in the results. The script uses the `find` utility to traverse the file system and report a histogram of the sizes of the files. The results are available on the web [15].

In June 2000 we ran this script on one of our workstations, a Pentium running Red Hat Linux 6.0. There were 89937 files on the system, including the author's home directory, a set of web pages, the operating system and a few applications.

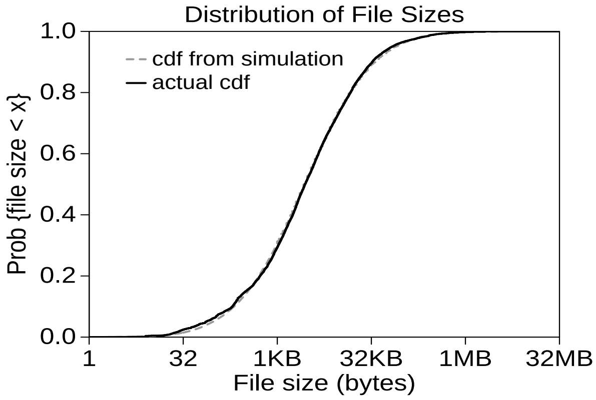

Figure 1. cdf of file sizes on a UNIX workstation.

Figure 1 shows the cumulative distribution function (cdf) of the sizes of these files, plotted on a log x-axis. 877 files with length 0 are omitted.

The figure also shows the cdf of file sizes generated by the simulation, using parameters that were tuned by hand to yield the best visual fit. Clearly the model is a good match for the data, suggesting that this model is descriptive of real systems.

2.2. The analytic model

There are two problems with this model so far: first, it provides no insight into the functional form of these distributions, if there is one; second, it does not provide a way to estimate the model parameters.

A solution to both problems comes from the observation that the simulator is effectively computing a numerical solution to a partial differential equation (PDE). By solving the PDE analytically, we can find the functional form of the distribution.

The PDE is the diffusion, or heat equation: $u_t = k^2 u_{xx}$, where $u(x,t)$ is the probability density function of file sizes as a function of x, which is the logarithm of the file size, and t, which represents time since the file system was created. k is a constant that controls the rate of diffusion.

The range is from 0 to ∞. The initial condition in time is the delta function with a peak at the initial file size, $\zeta = log(s^*)$. It is not obvious what the boundary condition at $x = 0$ should be, since the model does not include a meaningful description of the behavior of small files. Fortunately, the boundary behavior is usually irrelevant, so we can choose either $u(0,t) = 0$ or $u_x(0,t) = 0$.

In either case the solution is approximately

$$u(x,t) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{x-\zeta}{\sigma}\right)^2\right] \qquad (2)$$

where $\sigma = \sqrt{2kt}$.

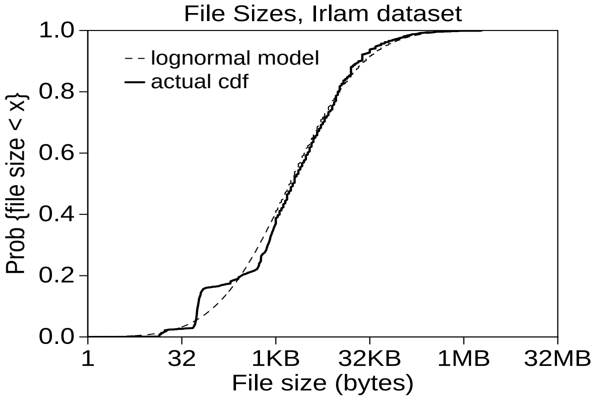

Figure 2. cdf of file sizes on a machine that participated in the Irlam survey.

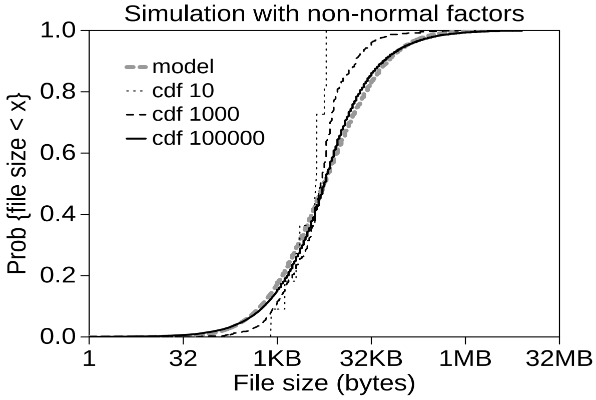

Figure 3. A file system simulation using an empirical distribution of multiplicative factors.

In other words, the distribution under a log transform is Gaussian with mean ζ and standard deviation that increases with k, the rate of diffusion, and t, time. The model of user behavior provides no way to estimate k, or even to map t onto real time, so we treat σ as a free parameter.

This observation leads to an easy way to estimate ζ and σ. Given a list of file sizes s_i, we can use the mean of $x_i = log(s_i)$ as an estimate of ζ, and the standard deviation of x_i as an estimate of σ. In the next section we use this technique to fit analytic models to a collection of datasets.

2.3. Is this model accurate?

Irlam's survey provides data from 656 machines in a variety of locations and environments. Of these, we discarded 43 because they contained no files with non-zero length, and an additional 52 because they contained fewer than 100 files.

For the remaining 561 file systems, we estimated parameters and compared the analytic distributions with the empirical distributions. As a goodness-of-fit metric we used the Kolmogorov-Smirnov statistic (KS), which is the largest vertical distance between the fitted and actual cdfs, measured in percentiles. The KS statistic is not affected by the log transform of the x-axis.

In the best case KS is 1.4 percentiles. For comparison, the fitted model in Figure 1 has a KS of 2.7. The median value of KS is 8.7, indicating that the typical system fits the model well. In the worst case KS is 40, which indicates that the model is not a good description of all systems.

To give the reader a sense for this goodness-of-fit, we present the dataset with the median value of KS in Figure 2. The maximum deviation from the model occurs between 32 and 64 bytes, where there appears to be a second mode. This kind of deviation is common, but in general the fitted models describe the tail of the cdf well.

Irlam's survey also includes data from one DOS machine. For this dataset the KS statistic is 5, which indicates that this model fits at least one non-UNIX file system.

We conclude that this model is a good description of many real systems, with the qualification that for some purposes it might be more accurate to extend the model by including a mixture of lognormal modes.

2.4. Is this model realistic?

The assumptions the model makes about user behavior are

1. The file system starts with a single file.
2. New files are always created by processing an existing file in some way, for example by copying, translating or filtering.
3. The size of the derivative file depends on the size of the original file.

The first assumption is seldom true. Usually a new file system is populated with a copy of an existing system or part of one. But in that case the new file system fits the model as well as the old, and evolves in time the same way.

The second assumption is not literally true; there are many other ways files might be created. For example, a new file might be the concatenation of two or more existing files, or the size of a new file might not depend on an existing file at all.

The third assumption is based on the intuition that many file operations are linear; that is, they traverse the input file once and generate output that is proportional to the size of the input. Again, this is not always true.

In addition, the simulation assumes that the distribution of multiplicative factors is lognormal with mean 1, but we can relax this assumption. As long as the logarithms of the multiplicative factors have two finite moments, the distribution of file sizes converges to a lognormal distribution, due to the the Central Limit Theorem.

To derive this result, recall that as the simulation proceeds, the size of the nth file added to the system depends on the size of one of its predecessors, and the size of the predecessor depends on the size of one of *its* predecessors, and so on. Thus, the size of the nth file is

$$s_n = s^* \cdot f_1 \cdot f_2 \cdot ... \cdot f_m \tag{3}$$

where m is the number of predecessors for the nth file and the f_i are the random multiplicative factors. Taking the logarithm of this equation yields

$$log(s_n) = log(s^*) + log(f_1) + log(f_2) + ... + log(f_m) \tag{4}$$

In log space, the distance of the nth file from the mean is the sum of m random variables, each of which is roughly normally-distributed. The Central Limit Theorem says that as m goes to infinity, the distribution of this sum converges to normal, provided that the logarithms of the f_i are independent and have two finite moments. Therefore the distribution of the sum (Equation 4) is normal, and the distribution of the product (Equation 3) is lognormal.

To demonstrate this effect with a realistic workload, we collected a sample of multiplicative factors from a single user (the author) and a single application (emacs). When emacs updates a file, it creates a backup file with the same name as the original, postpended with a tilde. In the author's file system there are 989 pairs of modified and original files. For each pair we computed the ratio of the current size to the backup size.

The distribution of ratios is roughly symmetric in log space, although there is a small skew toward larger values (it is more common for files to grow than shrink). Also, the distribution is significantly more leptokurtotic than a Gaussian (more values near the mode).

These deviations have little effect on the ultimate shape of the size distribution. Figure 3 shows a simulation starting with $\zeta = 11$ and using the observed ratios as multiplicative factors. The black curves show the cdf of file sizes after 10, 1000, and 100000 files were created. The dashed gray line shows a lognormal model fitted to the final curve. The simulated distribution converges to a lognormal.

We conclude that, even if the user model is not entirely realistic, it is robust to violations of the assumptions.

3. The Pareto model of file sizes

Several prior studies have looked at distribution of file sizes, in both local file systems and the World Wide Web. The consensus of these reports is that the tail of the distribution is well-described by the Pareto distribution.

To explain this observation, 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 such that the distribution of file sizes obeys a power law [6] [23].

In this section we review prior studies and compare the Pareto model to the lognormal model. We find that the lognormal model is a better description of these datasets than the Pareto model, and conclude that there is little evidence that the distribution of file sizes is long-tailed.

Furthermore, we believe that the diffusion model is more realistic than the HOT model. The HOT model is based on the assumption that the material available on a web page is "a single contiguous object" that the website designers are free to divide into files, and that they do so such that the the files with the highest hit rates are the smallest. We believe that constraints imposed by the content determine, to a large extent, how material on a web page is divided into files. Also, while web designers give some consideration to minimizing file sizes and transfer times, there are other objectives that have a stronger effect on the structure of web pages.

Finally, a major limitation of the HOT model is that it does not explain why local file systems exhibit the same size distributions as web pages, when local file systems are presumably not subject to the kind of optimization Carlson and Doyle hypothesize.

3.1. Evidence for long tails

The cdf of the Pareto distribution is

$$Pr[X < x] = 1 - \left(\frac{x}{k}\right)^{-\alpha} \qquad k > 0, \alpha > 0, x \geq k \tag{5}$$

The parameter α determines the shape of the distribution and the thickness of the tail; the parameter k determines the lower bound and the location of the distribution.

This distribution satisfies Equation 1, so the Pareto distribution is considered long-tailed, as is any distribution that is asymptotic to a Pareto distribution.

Unfortunately, there is no rigorous way of identifying a long-tailed distribution based on a sample. The most common method is to plot the complementary cumulative distribution function (ccdf) on a log-log scale. If the distribution is long-tailed, we expect to see a straight line or a curve that is asymptotic to a straight line.

But empirical ccdfs are often misleading. First, there are other distributions, including the lognormal, that appear straight up to a point and then deviate, dropping off with increasing steepness. Thus, a ccdf that appears straight does

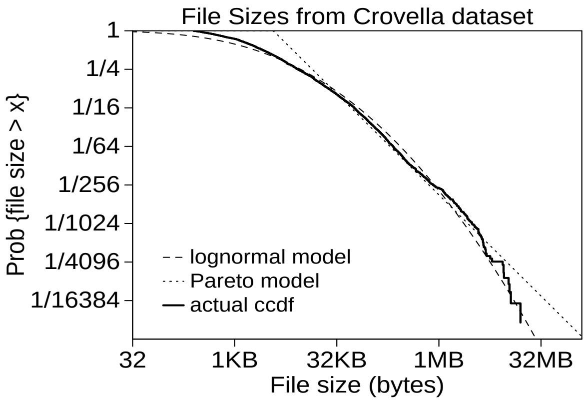

Figure 4. ccdf of file sizes from Crovella dataset.

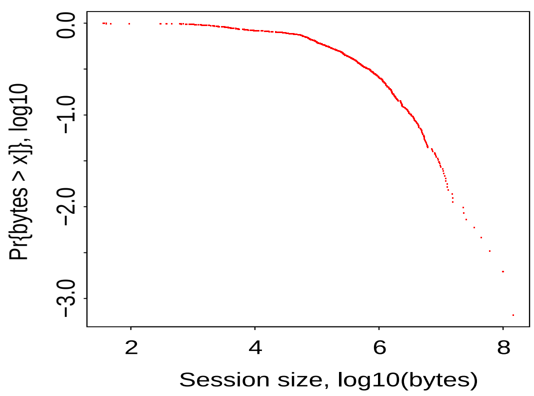

Figure 5. ccdf of session sizes from Feldmann dataset (reproduced from [13]).

not necessarily indicate a long tail. Also, log-log axes amplify the extreme tail so that only a few files tend to dominate the figure. What appears to be a significant feature in the extreme tail might be the result of just a few files.

One alternative is to use Q-Q plots or P-P plots to compare the measured distributions to potential models. For the datasets in this section, Q-Q plots are not useful because they are dominated by the few largest values. P-P plots are useful, but we did not find that P-P plots produced additional useful information for these data.

Crovella and Taqqu have proposed an additional test based on the aggregate behavior of samples [8]. These techniques are useful for estimating the parameter of a synthetic Pareto sample, but they are not able to distinguish a Pareto sample from a lognormal sample with similar mean and variance (see their Figures 5 and 8).

In the following sections we use ccdfs to evaluate the evidence that size distributions are long-tailed.

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 [10] [7] [9]. From these traces they extracted the unique file names and plotted the ccdf of their sizes.

We obtained these traces from their web pages and performed the same analysis, yielding 36208 unique files. Figure 4 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 lower bound, k, is 3800, which we chose to be the best match for the ccdf, and visually similar to Figure 8 in [7].

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 ccdf and its divergence from the model for files larger than 4MB.

The lognormal model in the figure has parameters $\zeta = 11$ and $\sigma = 3.3$. This 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 rather than continuing in a straight line. We conclude that this dataset provides greater support for the lognormal model than for the Pareto model.

Feldmann et al. argue that the distribution of Web session sizes is long-tailed, based on data they collected from an ISP [13]. 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 5. 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." Since the ccdf is curved throughout, it is not clear what they are referring to. In our opinion, this distribution exhibits the characteristic tail behavior of a lognormal distribution. We conclude that this dataset provides no support for the Pareto model.

3.3. File sizes on the web, server's view

Between October 1994 and August 1995, Arlitt and Williamson [3] collected traces from web servers at the University of Waterloo, the University of Calgary, the University of Saskatchewan, NASA's Kennedy Space Center, ClarkNet (an ISP) and the National Center for Supercomputing Applications (NCSA).

For each server they identified the set of unique file names and examined the distribution of their sizes. They

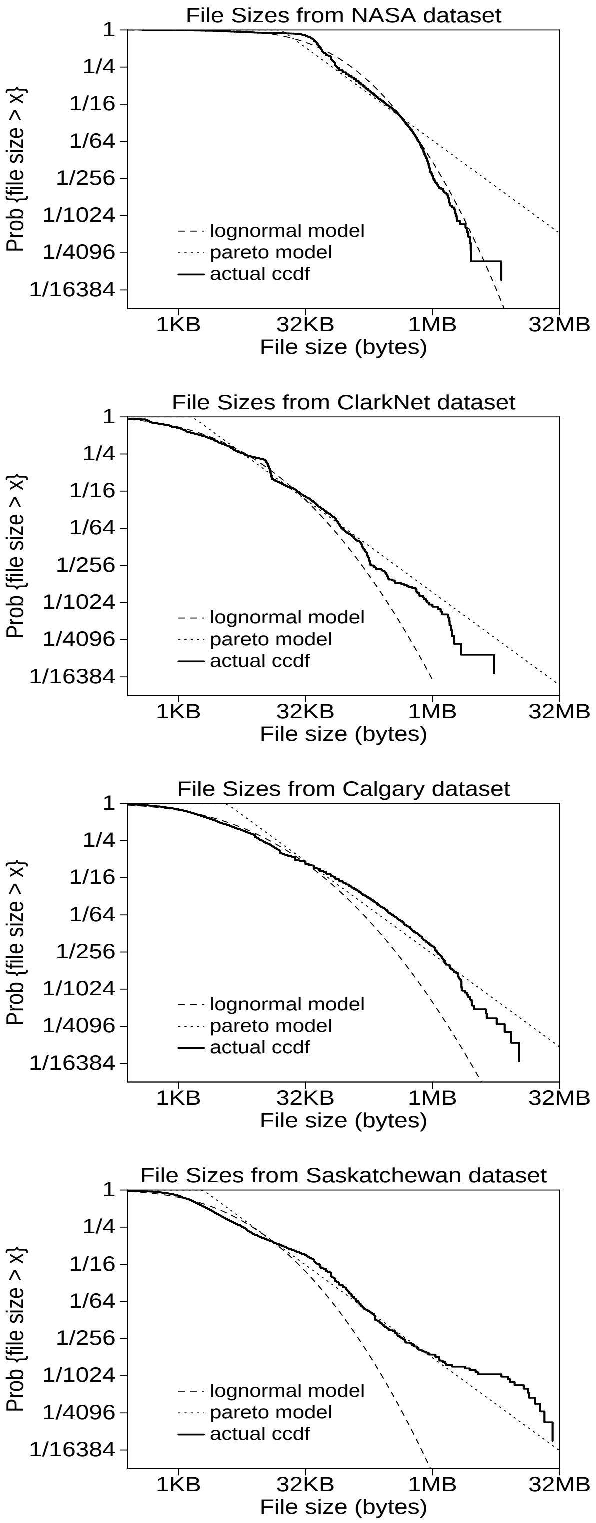

Figure 6. ccdfs for the datasets collected by Arlitt and Williamson.

report that these data sets match the Pareto model, and they give Pareto parameters for each dataset, but they do not present evidence that these models fit the data.

Four of these traces are in the Internet Traffic Archive (`http://ita.ee.lbl.gov`). We processed the traces by extracting each successful file transfer and recording the name and size of the file.

To derive a set of distinct files, we treated as distinct any log 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 is significantly different from the numbers in [3], so our treatment of this dataset may not be identical to theirs. Nevertheless, our ccdfs are visually similar to theirs.

We estimated the Pareto parameter for each dataset using `aest` [8]. The resulting range of values is from 0.97 to 1.02. We estimated the lower bounds by hand to yield the best visual fit for the ccdf. We estimated lognormal parameters for each dataset using the method in Section 2.2. Figure 6 shows these models along with the actual ccdfs.

The results are difficult to characterize. For the NASA dataset the lognormal model is clearly better. For the Saskatchewan dataset the Pareto model is clearly better. For the other two the ccdf lies closer to the Pareto model, but both curves show the characteristic behavior of the lognormal distribution, increasing steepness.

We believe that this curvature is indicative of non-long-tailed distributions. The claim that a distribution is long-tailed is a statement about how we expect it to behave as file sizes go to infinity. In these datasets, the increasing steepness of the tails does not lead us to expect the tail to continue along the line of the Pareto model.

Although the Saskatchewan dataset provides some support for the Pareto model, overall these datasets provide little evidence that the distribution of file sizes is long-tailed.

Arlitt and Jin collected access logs from the 1998 World Cup Web site [2] and reported the distribution of file sizes for the 20728 "unique files that were requested and successfully transmitted at least once in the access log." The raw logs are available in the Internet Traffic Archive, but we (gratefully) obtained the list of file sizes directly from the authors.

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 $\alpha = 1.37$. Again, we chose a lower bound by hand to match the ccdf and estimated lognormal parameters analytically.

Figure 7 shows the ccdf along with the two models. For files smaller than 128KB, the lognormal model is a slightly better fit. For larger files, neither model describes the data well.

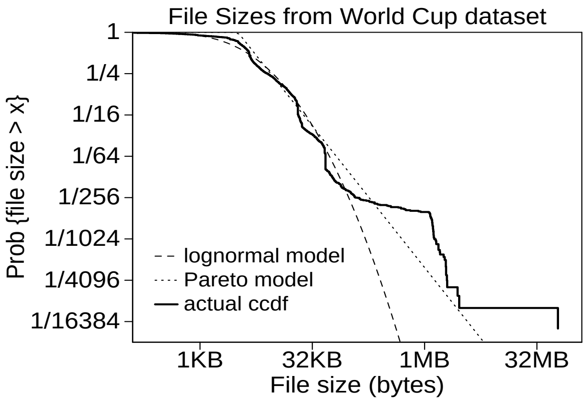

Figure 7. ccdf of file sizes from World Cup dataset.

Again, this dataset gives us little reason to expect the distribution to continue along the line of the Pareto model. Except for a single 64 MB file, the extreme tail is dropping off very steeply, which is consistent with a non-long-tailed distribution. We conclude that this dataset does not support the claim that the distribution of file sizes is long-tailed.

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 [1]. They plot the cdf of file sizes and show that a lognormal model fits it very 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."

Those figures are reproduced here in Figure 8. We do not see any sign of linear behavior in the ccdf. In fact, it clearly exhibits increasing steepness throughout, which is the characteristic behavior of a non-long-tailed distribution.

We conclude that this dataset provides strong support for the lognormal model and no support for the Pareto model.

3.4. Hybrid models of file sizes

Both Barford et al. and Arlitt et al. have proposed hybrid models that combine a lognormal distribution with a Pareto tail [5] [4] [1] [2].

Figure 9 is a reproduction from [4], showing the size distribution of 66998 unique files downloaded by a set of Web browsers at Boston University in 1998 (the W98 dataset), along with a hybrid model.

The hybrid model fits both the bulk of the distribution and the tail behavior, but it is not clear how much of the improvement is due to the addition of two free parameters. Furthermore, the extreme tail still appears to be diverging with increasing steepness from the model.

If we are willing to use a model with more parameters, it is natural to extend the lognormal model to include more

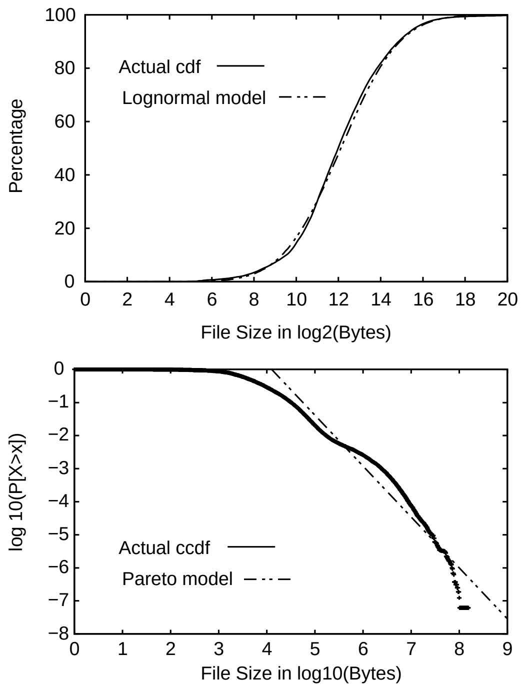

Figure 8. Distribution of file sizes from a Web proxy server (reproduced from [1]).

than one mode. Figure 10 shows a two-mode lognormal model chosen to fit the W98 dataset. This model is a better fit for the data than the hybrid model.

For this example we performed an automated search for the set of parameters—the mean and variance of each mode and the percentage of files from the first mode—that minimized the KS statistic. There are more rigorous techniques for estimating multimodal normal distributions, but they are not necessary for our purpose here, which is to find a lognormal model that fits the data well.

For the other datasets in this paper there are two-mode lognormal models that describe the tail behavior better than either the Pareto or the hybrid model. For example, Figure 11 shows a lognormal model fitted to the problematic Saskatchewan dataset. It captures even the extreme tail behavior accurately.

We conclude that long-tailed models are not necessary to describe the observed distributions, and therefore that these datasets do not provide evidence that the distribution of file sizes is long-tailed.

3.5. Aggregation

Looking at file sizes on the Internet, we are seeing the mixture of file sizes from a large number of file systems. If the distribution of file sizes on local systems is really log-

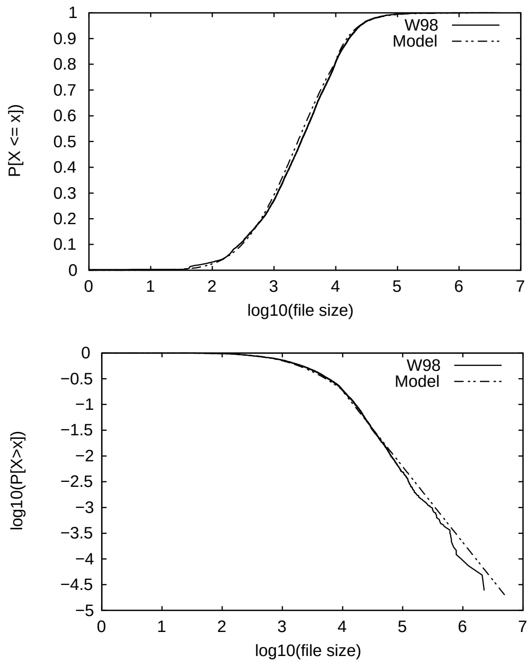

Figure 9. Distribution of unique file sizes from Web browser logs, and a hybrid lognormal-Pareto model. (Reproduced from [4])

normal, then it is natural to ask what happens when we aggregate a number of systems. To address this question, we went back to the Irlam survey and assembled all the data into an aggregate.

In total there are 6,156,581 files with 161,583 different sizes. The size of this sample allows us to examine the extreme tail of the distribution.

Figure 12 shows the ccdf of these file sizes along with lognormal and Pareto models chosen by hand to be the best fit. The lognormal model is a better fit. Throughout the range, the curve displays the characteristic curvature of the lognormal distribution. This dataset clearly does not demonstrate the definitive behavior of a long-tailed distribution.

A bigger data set allows us to see even more of the tail. In 1998 Douceur and Bolosky collected the sizes of more than 140 million files from 10568 file systems on Windows machines at Microsoft Corporation [11]. 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 13 shows the ccdf 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.

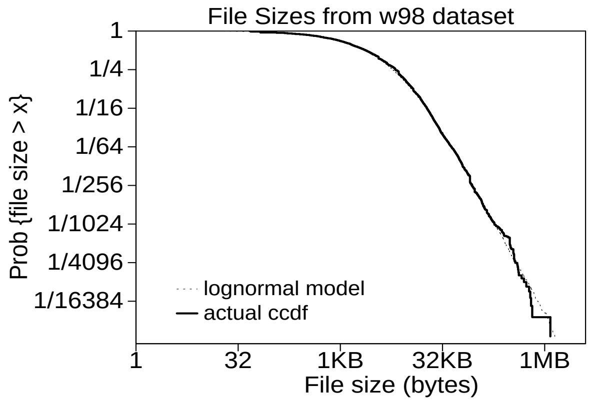

Figure 10. ccdf of file sizes from Web browser logs and a two-mode lognormal model.

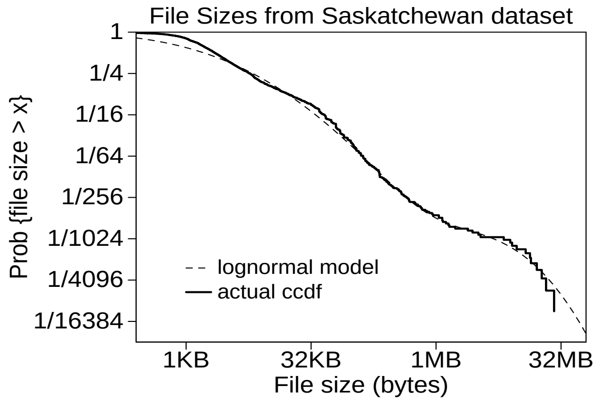

Figure 11. ccdf of file sizes from the Saskatchewan dataset and a two-mode lognormal model.

Again, the tail of the distribution displays the characteristic behavior of a non-long-tailed distribution. The simple lognormal model captures this behavior well, although it is offset from the data. The two-mode lognormal model fits the entire distribution well.

Based on these datasets, we conclude that the lognormal model is sufficient to describe the aggregate distributions that result from combining large numbers of file systems.

4. Self-similar network traffic

Most current explanations of self-similarity in the Internet are based on the assumption that the distribution of file sizes is long-tailed. We have argued that the evidence for this assumption is weak, and that there is considerable evidence that the distribution is actually lognormal and therefore not long-tailed.

In this section we review existing models of self-similar

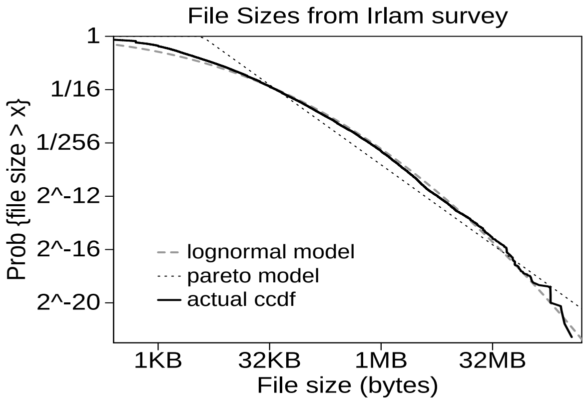

Figure 12. ccdf of file sizes in the Irlam survey.

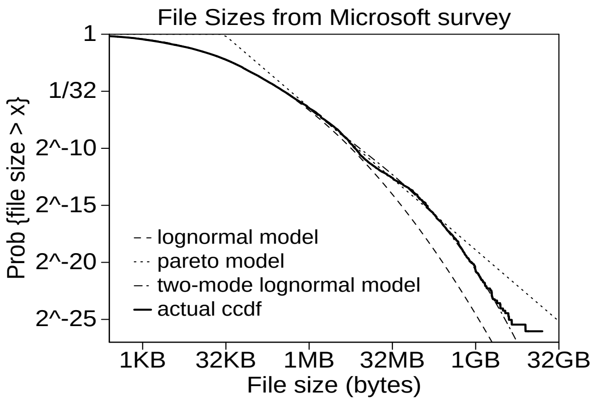

Figure 13. ccdf of file sizes in the Microsoft survey.

traffic and discuss the implications of our findings.

One explanatory model is an M/G/∞ queue in which network transfers are customers with Poisson arrivals and the network is an infinite-server system [19] [18]. In this model, 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 [22]. 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.

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 [16] [12]. In both cases, the more realistic models yielded qualitatively similar results.

All of these models are based on the assumption that the distribution of file transfer times is long-tailed. There is broad consensus that this assumption is true, but there is little direct evidence for it.

Crovella et al. have made the indirect argument that the distribution of transfer times depends on the distribution of available file sizes, and that the distribution of file sizes is long-tailed [9]. However, we have argued that the evidence for long-tailed file sizes is weak.

If the distribution of file sizes is not long-tailed, then these explanations need to be revised. There are several possibilities:

- 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.

- Even if the length of individual transfers is not long-tailed, the lengths of bursts might be. From the network's point of view there is little difference between a TCP timeout during a transfer and the beginning of a new transfer [19].

- The distribution of interarrival times might be long-tailed. There is some evidence for this possibility, but also evidence to the contrary [19] [3] [5] [13].

- Two recent papers have argued that the dynamics of TCP 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 [21] [14].

- A final 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 [19], but over a wide range of time scales it may be statistically indistinguishable from a truly self-similar process.

As ongoing work we are investigating these possibilities, trying to explain why Internet traffic appears to be self-similar.

5. Conclusions

- The distribution of file sizes in local file systems and on the World Wide Web is approximately lognormal or a mixture of lognormals. We have proposed a user model that explains why these distributions have this shape.

- The lognormal model describes the tail behavior of observed distributions as well as or better than the Pareto model, which implies that a long-tailed model of file sizes is unnecessary.

- In our review of published observations we did not find compelling evidence that the distribution of file sizes is long-tailed.

- Since many current explanations of self-similarity in the Internet are based on the assumption of long-tailed file sizes, these explanations may need to be revised.

Acknowledgments

Many thanks to Mark Crovella (Boston University) 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. Thanks also to Kim Claffy and Andre Broido (CAIDA), Mark Crovella, Rich Wolski (University of Tennessee) and Lewis Rothberg (University of Rochester) for reading drafts of this paper and making valuable comments. Finally, thanks to the reviewers from SIGMETRICS and MASCOTS for their comments.

References

[1] M. Arlitt, R. Friedrich, and T. Jin. Workload characterization of a Web proxy in a cable modem environment. *ACM Sigmetrics Performance Evaluation Review*, August 1998.

[2] M. Arlitt and T. Jin. Workload characterization of the 1998 World Cup Web site. Technical Report HPL-1999-35R1, Hewlett-Packard, September 1999.

[3] M. F. Arlitt and C. L. Williamson. Web server workload characterization: the search for invariants. In *ACM SIGMETRICS'96*, pages 126–137, May 1996. Extended version.

[4] P. Barford, A. Bestavros, A. Bradley, and M. Crovella. Changes in Web client access patterns: Characteristics and caching implications. *World Wide Web*, 2:15–28, 1999. Special Issue on Characterization and Performance Evaluation.

[5] P. Barford and M. E. Crovella. Generating representative Web workloads for network and server performance evaluation. In *Proceedings of Performance '98/SIGMETRICS '98*, pages 151–160, July 1998.

[6] J. M. Carlson and J. Doyle. Highly optimized tolerance: a mechanism for power laws in designed systems. *Physics Review E*, 60(2), August 1999.

[7] M. E. Crovella and A. Bestavros. Self-similarity in World-Wide Web traffic: evidence and possible causes. In *ACM SIGMETRICS'96*, pages 160–169, May 1996.

[8] M. E. Crovella and M. S. Taqqu. Estimating the heavy tail index from scaling properties. *Methodology and computing in applied probability*, 1(1):55–79, 1999.

[9] M. E. Crovella, M. S. Taqqu, and A. Bestavros. Heavy-tailed probability distributions in the World Wide Web. In *A Practical Guide To Heavy Tails*, pages 3–26. Chapman & Hall, 1998.

[10] C. R. Cunha, A. Bestavros, and M. E. Crovella. Characteristics of WWW client-based traces. Technical Report BU-CS-95-010, Boston University, 1995.

[11] J. R. Douceur and W. J. Bolosky. A large-scale study of file-system contents. In *ACM SIGMETRICS'99*, pages 59–70, May 1999.

[12] A. Feldmann, A. C. Gilbert, P. Huang, and W. Willinger. Dynamics of IP traffic: a study of the role of variability and the impact of control. In *ACM SIGCOMM'99*, pages 301–313, 1999.

[13] A. Feldmann, A. C. Gilbert, W. Willinger, and T. Kurtz. The changing nature of network traffic: Scaling phenomena. *ACM Computer Communication Review*, 28:5–29, April 1998.

[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] G. Irlam. Unix file size survey – 1993, `http://www.base.com/gordoni/ufs93.html`. Accessed 31 May 2000.

[16] K. Park, G. Kim, and M. E. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. Technical Report BU-CS-96-016, Boston University, August 1996. Short version appeared in Proc. Fourth Int'l Conference on Network Protocols.

[17] K. Park and W. Willinger. Self-similar network traffic: An overview. In *Self-Similar Network Traffic and Performance Evaluation*, chapter 1. Wiley-Interscience, 2000. Preprint.

[18] M. Parulekar and A. M. Makowski. M/G/∞ input process: a versatile class of models for network traffic. Technical Report T.R. 96-59, Institute for Systems Research, 1996.

[19] V. Paxson and S. Floyd. Wide-area traffic: the failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, 3:226–244, 1995.

[20] M. S. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *Computer Communication Review*, 27:5–23, 1997.

[21] A. Veres and M. Boda. The chaotic nature of TCP congestion control. In *IEEE INFOCOM'00*, March 2000.

[22] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. In *ACM SIGCOMM'95*, pages 100–113, 1995.

[23] X. Zhu, J. Yu, and J. Doyle. Heavy tails, generalized coding, and optimal web layout. In *IEEE INFOCOM'01*, April 2001.