On the Tails of Web File Size Distributions*

Weibo Gong, Yong Liu
Dept.of Electrical & Computer Engineering
Univ. of Massachusetts
Amherst, MA 01003

Vishal Misra
Dept. of Computer Science
Columbia University
New York, NY 10027

Don Towsley
Dept. of Computer Science
Univ. of Massachusetts
Amherst, MA 01003

Abstract

Power laws have been observed in various contexts in the Internet. There has been considerable interest in identifying the mechanisms behind these power laws. Most of these have focused on the tail behavior of the distributions. We argue that the the tails and their asymptotic behavior is very hard to substantiate in realistic engineering systems. In this paper we describe some of the proposed mechanisms for producing power law tails. We show that these mechanisms are not particularly robust. Furthermore, we argue that the data ususally available for classifying a distribution is insufficient to classify the tail. Fortunately, the tail has little impact on Internet performance. Thus it is sufficient to focus on mechanisms leading to power law like "waists" of the distributions.

1 Introduction

Power laws have recently been discovered for web file sizes, web site connectivities, and the router connection degrees, see, for example, [5, 4, 19, 7].

These discoveries are important in the study of various protocols and algorithms. For example, the web file size distribution is important for web-server scheduling. More importantly, these discoveries motivate us to identify the mechanisms behind the observed power laws and therefore enable us to design mechanisms to improve the current operational structure of the Internet.

In the past few years there has been a significant amount of research on the issue of power laws. Various mechanisms have been proposed including reflected or drifted multiplicative processes, exponentially stopped geometric Brownian motion, self-organized criticality, preferential growth, and highly optimized tolerance among others [4, 13, 1, 17, 3]. There is also a group of researchers advocating lognormal rather than power law distribution as the correct description of the data (see,

*This research is supported in part by NSF grants ITR-0085848, ANI-9980552, and EIA-0121608, and DARPA Contract DOD F30602-00-0554

for example, [6]). The debate has focused on the "tail" behavior since, for example, the primary difference between a lognormal distribution and a power law distribution lies in the tail. We believe that this debate on the nature of the tail (power law vs lognormal) is of little practical interest or consequence. Expanding on this is the subject of this paper.

Our view on the "tail debate" is as follows: First, we believe that there is never sufficient data to support any analytical form summarizing the tail behavior and therefore any summary could be misleading and dangerous. For instance, one could fit a set of webfile size data as power law (Pareto) tailed, and give the distribution to a simulation engineer to estimate queueing delays. The simulation engineer would generate service times according to the given power law tail and as a consequence, the simulation results will depend heavily on how long simulation is run, as the longer the simulation length runs the greater the possibility of generating an unrealistically large service time. Second, all mechanisms aimed at explaining the power law or lognormal tail are fragile in the sense that minor perturbations to their assumptions will lead to different analytical forms for the tail. We will make this concrete in the paper. It is important to understand that such perturbations seem to always exist in engineering settings. For example, in the case of exponentially stopped Geometric Brownian motion one needs a stopping time that has an asymptotic exponential tail, which can never be the case in engineering settings since such a stopping time has to end at some finite value. In general, many of the conclusions in the above mentioned studies are statements about the asymptotic tail behavior, which never occur in practice.

Given the difficulty to identify the tail of a distribution from limited data, the fragility of mechanisms to explain different tails, what is one to do as an engineer. Fortunately, as we will observe in this paper, there is little evidence that the tail impacts the design of various algorithms and infrastructure on the Internet. This leads to our final point, which is that engineers should focus on the behavior of the "waist" of a distribution, not its "tail".

The rest of the paper is organized as follows. In Section 2, we introduce the terminology used in the rest of the paper. We describe mechanisms to generate long-tailed distributions in Section 3. Section 4 identifies the fragilities in these mechanisms that can lead to the wrong analytical tail distribution. Next, in Section 5 we look at the difficulty in estimating the tail from a limited dataset. In Section 6 we explore the impact of power law tails on configurations of various algorithms in the Internet. Section 7 explores the concept "power law waist" - causes and effects on traffic characteristics. Finally, we present our conclusions in Section 8.

2 Terminology

We introduce terminology that will be used in the remainder of the paper. First, consider two functions $f, g : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. We say that $f \sim g$ iff $\lim_{x\rightarrow\infty} f(x)/g(x) = c$, $0 < c < \infty$. We say that $f \prec g$ iff $\lim_{x\rightarrow\infty} f(x)/g(x) = 0$.

The complementary cumulative distribution function (CCDF)of a random variable X is defined as $\bar{F}_X(x) = P(X \geq x)$. We say that X is characterized by a *power law distribution* if $\bar{F}_X(x) \sim x^{-\alpha}$, where $c > 0$, $\alpha > 0$. The Pareto distribution is the canonical power law distribution. It has has probability density function

$$f(x) = \frac{\alpha b^\alpha}{x^{\alpha+1}}, \quad x \geq b, \quad \alpha > 0, \quad b > 0$$

If $\alpha \leq 1$, it has infinite mean; if $1 < \alpha \leq 2$, it has finite mean and infinite variance; if $\alpha > 2$, both the mean and variance are finite. The corresponding CCDF. is $\bar{F}_X(x) = (b/x)^\alpha$.

A positive random variable X is said to be lognormally distributed with two parameters μ and σ^2 if $Y = \ln X$ is normally distributed with mean μ and variance σ^2. The probability density function of X is

$$f_X(x) = \begin{cases} \frac{1}{\sqrt{2\pi}\sigma x} \exp\{-\frac{(\ln x - \mu)^2}{2\sigma^2}\} & x > 0 \\ 0 & x \leq 0 \end{cases}$$

3 Mechanisms for Generating Long-tailed Distributions

Multiplication Mechanism. Consider a sequence of independent and identically distributed positive valued random variables (rvs) $\{X_i\}_{i=1}^\infty$ with finite first two moments. Let Y_n be defined as

$$Y_n = Y_{n-1}X_n, \quad n = 1, \ldots \tag{1}$$

with $Y_0 = 1$. Taking the logarithm of both sides yields

$$\ln Y_n = \sum_{i=1}^{n} \ln X_i, \quad n = 1, \ldots$$

which, by the central limit theorem, converges in distribution to a normally distributed random variable. Consequently, Y_n converges in distribution to a lognormally distributed random variable.

Mixture Mechanisms. Let $\{Y_t\}$ be an independently distributed sequence of lognormally distributed random variables where $\ln X_t$ has mean 0 and variance t as might be produced by a geometric brownian motion process. Let T be a nonnegative rv with density $f_T(x)$, $x \geq 0$. Reed [17] showed that if T is exponentially distributed with parameter λ, $f_{X(T)}(x) \sim x^{-1-\lambda}$ (here $X(T) = X_T$).

Optimization-based mechanisms. Constrained optimization has been proposed as a mechanism behind many observed power law like data set [4]. The basic idea is that optimization of an expected performance index would decrease the probabilities of events contributing significantly to such mean value, and increase the probabilities of certain rare events correspondingly. This leads to heavy tailed distributions under appropriate assumptions. For example, it has been shown in a simplified forest fire control model that an optimal resource allocation indeed leads to power law tail under certain assumptions, which includes that the form of the pre-optimization distribution tail is known exactly.

4 Fragility of Analytical Tail distributions

Analytical forms of the tail behavior are nice for many purposes. However they could be fragile to the perturbations in the assumptions leading to the tail formulas. We give several examples below to illustrate this point.

The Monkey Text Problem. One of the oldest power laws characterizes the usage of English words. Briefly, the frequency of word usage plotted against rank exhibits a power law.... Zipf is usually accredited for this discovery. To study Zipf's law researchers has developed the so-called "Monkey text" model in which each of M letter keys is struck with a certain probability. In [10], Li showed that when such striking probabilities are exactly equal then the frequency of the words against the rank of the words exhibits a power law. He pointed out that this is actually quite straightforward. Basically the occurrence probability of a word with length L is $[M(M+1)^L]^{-1}$. The rank of such a word, denoted by $r(L)$, in terms of occurrence frequency is proportional to M^L since we have

$$\sum_{l=0}^{L-1} M^l < r(L) \leq \sum_{l=0}^{L} M^l$$

Combining these two facts together we see that the frequency or probability against the rank in a perfect monkey text is a power law. However this power law is "shallow" in Li's words, since the key trick here is the exponential stretching performed by the rank-length transformation $r(L) \approx M^L$. More importantly, as R. Perline pointed out in [16], if the hitting probabilities for different letters are different, then one would obtain an asymptotic lognormal distribution for the frequency-rank relation instead. The high sensitivity of the distribution tail is due to the multiplicative nature of the model. For example, if one letter has a slightly smaller probability than other letters, then very long words consisting of this particular letter only will generate very small probabilities due to the multiplication, which makes the tail much "thinner" than in the equal hitting probability case. Therefore we see that Zipf's law in the monkey text context is very fragile when subject to perturbations in the letter striking probabilities.

The multiplication mechanism - fragility of the lognormal distribution. Recall the multiplication mechanism as defined by 1. Suppose that we add a reflection of the form that Y_n is never allowed to fall below a threshold $\epsilon > 0$. In other words, we modify the construction to be:

$$Y_n = \max\{Y_{n-1}X_n, \epsilon\}, \quad n = 1, \ldots .$$

Gabaix, [8] has shown that Y_n converges to an rv with a Pareto distribution as $n \to \infty$. Therefore, the multiplication process for generating a lognormal distribution i very fragile to a perturbation on the reflection; if $\epsilon = 0$, it produces a lognormal distribtuion and otherwise a Pareto distribution.

The mixture mechanism We revisit the mechanism described in Section 3. We have the following result (see [18] for proof).

Theorem 1 *Let T be a nonnegative rv with density $fp_T(x)$ and let $Y(T)$ denote Y_t when $t = T$. Then*

1. $\lambda x \prec -\ln f_T(x)\ \forall \lambda > 0$ *implies that* $f_{X(T)}(x) \prec x^{-\alpha}\ \forall \alpha > 1$,
2. $-\ln f_T(x) \prec \lambda x\ \forall \lambda > 0$ *implies that* $x^{-\alpha} \prec f_{X(T)}(x)\ \forall \alpha > 1$,

Thus we observe that the distribution of $Y(T)$ is sensitive to the tail behavior of p_T; if it deviates from an exponential decay, the power law disappears. This points to the fragility of a power law to the assumptions regarding the specific mixture.

Fragility of the optimization mechanism Optimization leads to a power law. While optimization processes do lead to power law tails with suitably choosen objective functions and under appropriate assumptions, it appears to be a fragile mechanism against the assumptions. For example, one may need to know exactly the apriori distribution model in order to perform the optimization procedure leading to power law tail. Moreover these apriori disibution models often need to have infinite tails themselves. These assumptions are difficult to satisfy in an engineering system.

5 Difficulty of Estimating Tail

The "waist" of the lognormal distribution and power law distribution could look very similar. When the variance is big, the body of a lognormal CCDF is very close to a straight line in the log-log plot. Given a dataset of real web file sizes, we can fit most of the samples into either a lognormal model or power law model by tuning the parameters. The dramatic difference between these two distributions lies in the tail of their CCDFs. In a log-log plot, the CCDF of a power law decays with constant slope; the CCDF of alognormal distribution decays with increasing slope. However in any set of sample data, we only have a small number of samples in the extreme tail, which means it would be hard to determine which one is a better fit.
To further illustrate, suppose we have a dataset of file sizes: $\{X_i, i = 1, 2, \cdots, M\}$. Assume all samples in the dataset are i.i.d samples of some random variable X with CCDF $G(x)$. We can estimate $G(x)$ by the empirical CCDF $\bar{G}(x) \triangleq \frac{1}{M}\sum_{i=1}^{M} I(X_i > x)$, where $I(A)$ is the indicator function of set A. $\bar{G}(x)$ is an unbiased estimation of $G(x)$ with variance $Var[\bar{G}(x)] = (1 - G(x))G(x)/M$. When M is big, $\bar{G}(x)$ should approach a normal distribution $N(G(x), \sqrt{\frac{(1-G(x))G(x)}{M}})$. The α-confidence interval for $\bar{G}(x)$ is

$$\big[\max\{0, G(x) + a\sqrt{\frac{(1-G(x))G(x)}{M}}\}, \quad G(x) + b\sqrt{\frac{(1-G(x))G(x)}{M}}\big]$$

where $\int_a^b e^{-x^2/2}dx = \sqrt{2\pi}\alpha$ and $a < 0, b > 0$. When we plot the confidence interval in the log-log plot and let $G(x) \ll 1$, the width of the confidence interval is:

$$W \approx \log(1 + \frac{b}{\sqrt{M_x}}) - \log(1 + \frac{a}{\sqrt{M_x}})$$

where M_x is the expectation of number of files with size greater than x. W keeps increasing when M_x decreases. W will blow up when $M_X \le a^2$.
We plotted 95% confidence interval in Figure 1 for both the pareto and the lognormal model used in [6] to fit Crovella's web file size dataset [5]. We observe that the confidence intervals of both models diverge when file sizes grow. At the tail, the two confidence intervals have a large overlap, which makes it difficult to distinguish them.

6 Effect of tail behavior on the design and dimensioning of algorithms and protocols on the Internet:

An important factor to consider in support of our arguments is what effect does the extreme tail of a job size distribution have on the configuration or design of existing mechanisms on the Internet. There don't appear to be any design problems or adjustment to existing algorithms or protocols

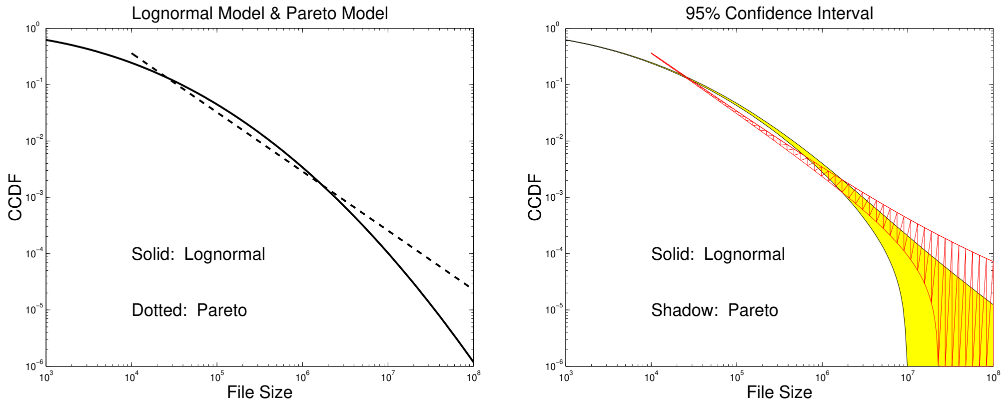

Figure 1: Confidence Interval for Pareto and Lognormal Models

that depend upon the fact whether a (job size) distribution has infinite variance or not, whether the tail is described by a power law or not. We look at two particular cases where, at first glance, the heavy tailed nature of filesize distribution can lead to a significant change to the existing design.

6.1 Buffer dimensioning at routers

There have been numerous studies investigating the effects on the queue length of a multiplexer fed by a heavy-tailed input [14], [9]. All studies show that the tail of the queue length distribution is significantly heavier for a heavy tailed input as opposed to a light tailed or exponential tailed input. For a finite buffer, this implies that to maintain performance guarantees in terms of drop probability buffers of significantly larger sizes would be required with a heavy tailed input as opposed to a light tailed one. However, all the studies performed assumed an *open-loop* transport protocol. Files, on the Internet, are (invariably) transported by TCP which is a *closed-loop* transport protocol. The arguments fall apart when we consider the closed loop case, and loss rates in fact do not depend on the tail of the flow length (file size) distribution. Buffer dimensioning on the Internet is done more in tune with the bandwidth-delay product of the output link of the router (typical values are 2-8 times bw*delay). A major consideration is to limit the amount of queueing delay for TCP flows, since the delay affects the overall throughput of a feedback based flow control protocol like TCP. Hence the buffers are kept relatively small, and these small buffers then filter out the tail effects. This implies that even if we consider open-loop transport protocols, the performance difference between a heavy tailed input and a light tailed one is likely to be negligible because of the small buffers. TCP itself, is configured or tuned according to fine timescale behavior, of the orders of a typical round trip time. Longer timescale effects, where the extreme tail of a distribution is likely to show up, are not considered at all in the design of TCP.

6.2 Web server scheduling

Recently there has been a lot of interest [2] in changing the job scheduling of web servers to exploit the underlying power law distribution of job sizes. Specifically, Shortest Remaining Processor Time (SRPT) has been proposed as an alternative to Processor Sharing (PS) or round-robin scheduling. In [2] the authors investigate analytically the benefits of SRPT over PS under a "heavy tailed" job size distribution assumption. A persuasive argument is made for SRPT over PS. However, a closer look at their results reveal that the "heavy tailed distributions" that they considered are not really heavy tailed, but instead have a "heavy tailed property" (HT property) which is defined as follows in the paper:

> One key property of heavy-tailed distributions and (many) Bounded Pareto distributions is that a tiny fraction ($< 1\%$) of the very largest jobs comprise over half of the total load. We will refer to this as the heavy tailed property (HT property)

All their results hold for distribution which have this "HT property". However, this is not a heavy tailed property, but is really a high variance property. This property is exhibited, for instance, by a simple hyperexponential distribution. A large subset of this class of distributions is light-tailed, and in fact the distribution considered in [2] is the lightest tail of all: it is a bounded distribution with compact support, i.e. *no tail* at all! Thus, we again observe that the power law decay of the tail of the distribution does not play a critical part in the design of algorithms, the results are equally applicable for distributions with a "heavy enough" waist or high variance.

7 The Effect of a Power Law "Waist" on Traffic

In this section we study the impact of web file size distributions on the traffic correlation structure. In particular, we evaluate how the tail of the web file size distribution affect the self-similarity of the resulting web traffic. Because of its simplicity, we use the $M/G/\infty$ input process to model web traffic. An $M/G/\infty$ input process $b(t)$ is the busy server process of a discrete time infinite server system fed by Poisson arrivals of rate λ and with generic service time σ distributed according to $F_\sigma(x)$. One can think of this as modeling the arrival of web sessions according to a Poisson process where each session desires to transfer a file of size X and is given a rate of one unit per second. The $M/G/\infty$ input process has been shown to be versatile and tractable [15]. The auto-covariance of $b(t)$ has been established as:

$$\Gamma(h) \triangleq cov(b(t+h), b(t)) = \lambda \sum_{r=h+1}^{\infty} P[\sigma \geq r] \tag{2}$$

To obtain strict sense self-similar traffic, σ must have a power tail, i.e., $\bar{F}_\sigma(x) \propto 1/x^\alpha$. In order to evaluate the impact of the tail of $\bar{F}_\sigma(x)$ on the auto-correlation of $b(t)$, we generate two $M/G/\infty$ input processes $\{b_1(t), b_2(t)\}$ with two service time distributions: $\sigma_1 = \lceil s_1 \rceil$, where r.v. s_1 follows a pareto distribution, with CCDF $\bar{F}_{s_1}(x) = (c/x)^\alpha$, $x > c$, and $\sigma_2 = \lceil s_2 \rceil$, where the distribution of s_2 is the mixture of a pareto body and an exponential tail. The CCDF of s_2 is

$$\bar{F}_{s_2}(x) = \begin{cases} (c/x)^\alpha & \text{for } c < x < x_0 \\ ke^{-\beta x} & \text{for } x \geq x_0 \end{cases}$$

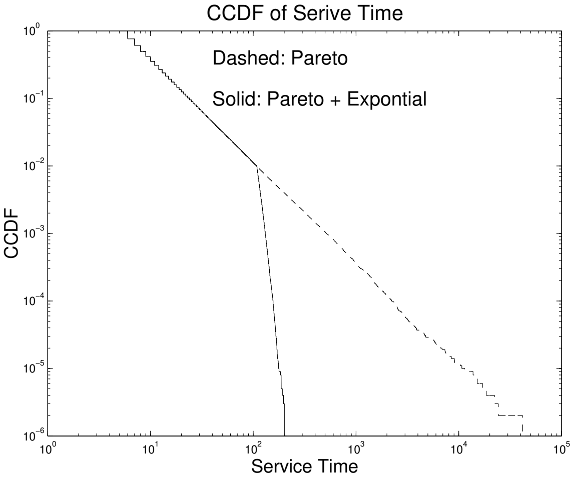

Figure 2: CCDF of Service Time

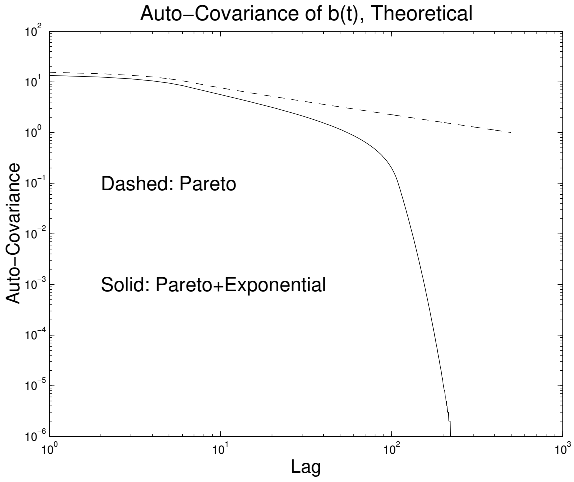

Figure 3: Auto-Cov of $b_1(t)$, $b_2(t)$

where $k = c^\alpha e^{\beta x_0}/x_0^\alpha$.
Then CCDFs of $\{\sigma_1, \sigma_2\}$ are

$$P[\sigma_1 \geq r+1] = P[s_1 > r] = \begin{cases} 1 & 0 \leq r \leq \lfloor c \rfloor \\ (\frac{c}{r})^\alpha & r \geq \lfloor c \rfloor + 1 \end{cases} \tag{3}$$

$$P[\sigma_2 \geq r+1] = P[s_2 > r] = \begin{cases} 1 & 0 \leq r \leq \lfloor c \rfloor \\ (\frac{c}{r})^\alpha & \lfloor c \rfloor + 1 \leq r \leq \lfloor x_0 \rfloor \\ ke^{-\beta r} & r \geq \lfloor x_0 \rfloor + 1 \end{cases} \tag{4}$$

The auto-covariance functions of the$M/G/\infty$ processes $\{b_1(t), b_2(t)\}$ can be computed from (2) For this experiment, we set the arrival rate $\lambda = 1$, pareto parameters $\alpha = 1.5$, $c = 5$, the turning point for the exponential tail $x_0 = G_1^{-1}(0.01) = 107.7$, and the rate of exponential tail $\beta = 0.1$. Figure 2 shows the two CCDFs in log-log scale. $\Gamma_1(h)$ and $\Gamma_2(h)$ are plotted in Figure 3.

We simulate $\{b_1(t), b_2(t)\}$ by generating $1000,000$ arrivals for each service time distribution and running them through the server system.

In Figure 4 and 5 we show the evolution of $b_i(t)$ over different time scales. On time scale k, $b_i(t)$ is averaged in windows of length k

$$\bar{b}_i^k(n*k) = \frac{1}{k}\sum_{j=1}^{k} b_i((n-1)*k+j) \quad i = 1, 2; \quad n = 1, 2, \cdots$$

We observe that the shapes of both $b_1(t)$ and $b_2(t)$ are similar on timescales from 1 to 1000. This suggests that a power tail is not necessary for visual self-similarity within finite range.

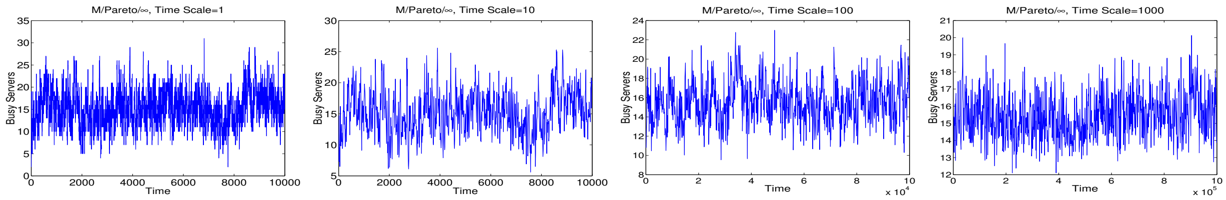

Figure 4: Sample Path of $b_1(t)$

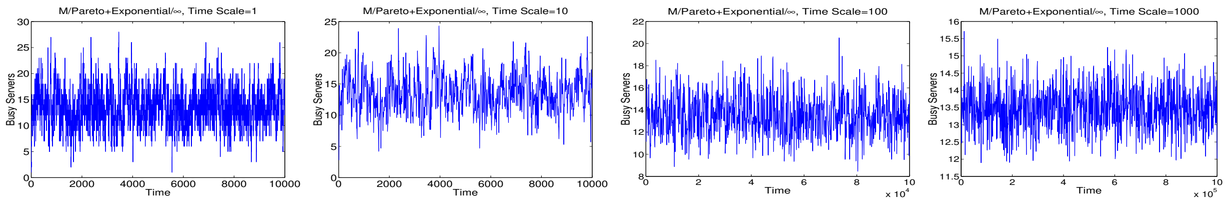

Figure 5: Sample Path of $b_2(t)$

8 Conclusions

In this paper we have argued that the focus on power law tails in the Internet is misguided. First, many mechanisms have been proposed to explain where they come from. However, they are all very fragile, sensitive to the underlying assumptions. Second, it is extremely difficult if not impossible to statistically characterize a distribution tail based on a finite amount of data. Third, in many applications, the tail plays little role in determining design and performance. Instead, it is the power "waist" that does. The latter was illustrated through a simple example.

We are currently investigating mechanisms for predicting power waists. Our preliminary work indicates that a clusteredmultiplicative model may be the source for the "power law" waist that has been observed. This model suggests that hierarchies and proportional growth could be the mechanisms behind the multiplications in the model, and in-homogeneity of the hierarchies and the growth could be handled by appropriate clustering. One important feature of the multiplicative model is that it is similar to the central limit theorem based models where many "small" random effects add to the very robust Gaussian "body". In our case we would say that the observed power law or lognormal "waist" is due to many "small" random effects multiplied together. To this end we would also like to point out that the Gaussian distribution in engineering practice refers to its bell, and not the tail. The latter is also for the mathematical convenience and is as unrealistic as any other tail.

References

[1] P. Bak. How Nature Works: the science of self-organized criticality. OUP. 1997.

[2] N. Bansal, M. Harchol-Balter. "Analysis of SRPT Scheduling: Investigating Unfairness", *Proc. ACM Sigmetrics*, June 2001.

[3] A.L. Barabasi, R. Albert, Emergence of Scaling in Random Networks. Science 286(1999), 509-512. 2000.

[4] J. M. Carlson, J. Doyle, "Highly Optimized Tolerance: A Mechanism for Power Laws in Designed Systems", *Physical Review E, 60, 1412-1427, 1999.*

[5] M. E. Crovella, A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", *IEEE/ACM Transactions on Networking, 5(6):835–846, December 1997.*

[6] A. B. Downey, "The structural cause of file size distributions", *Proc. MASCOTS'01*, Aug. 2001.

[7] M. Faloutsos, P. Faloutsos, C. Faloutsos. "On power-law relationships of the internet topology", *Proc. ACM SIGCOMM'99*, 251-262, Sept. 1999.

[8] X. Gabaix. "Zipf's law for cities: an explanation", *Quarterly Journal of Economics*, Aug. 1999.

[9] P. Jelenkovic, P. Momcilovic. "Capacity Regions for Network Multiplexers with Heavy-tailed Fluid On-off Sources", *Proc. INFOCOM01*, April 2001.

[10] W. Li, "Random texts exhibit Zipf's-law-like word frequency distribution", *IEEE Transactions on Information Theory, 38(6) 1842-1845, 1992*.

[11] A. Medina, "On the Origin of Power Laws in Internet Topologies", *ACM Computer Communications Review, April 2000*.

[12] M. Mitzenmacher, "Dynamic Models for File Sizes and Double Pareto Distributions", Submitted to INFOCOM 2002.

[13] M. Mitzenmacher, "A Brief History of Generative Models for Power Law and Lognormal Distributions", *Proc. ALLERTON'01*, Oct. 2001.

[14] I. Norros. "A storage model with self-similar input", *Queueing Systems*, 16 (1994), 387–396.

[15] M. Parulekar, A. Makowski. "M/G/1 input processes: A versatile class of models for network traffic," *Proc.INFOCOM'97*, Apr. 1997.

[16] R. Perline, "Zipf's Law, the central limit theorem, and the random division of the unit interval", *Physical Review, Volumn 54, Number 1*'

[17] Reed, W. J. 2001. "The double Pareto-lognormal distribution - A new parametric model for size distribution", http://www.math.uvic.ca/faculty/reed/index.html

[18] D. Towsley, etal. "Fragility of Tails in Power Law Distributions", under preparation.

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