

Self-Similarity and Long Range Dependence on the Internet: A Second Look at the Evidence, Origins and Implications

Wei-Bo Gong

Dept. of Electrical Engineering
University of Massachusetts
Amherst, MA 01003
gong@ece.umass.edu

Yong Liu

Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
yongliu@cs.umass.edu

Vishal Misra

Dept. of Computer Science
Columbia University
New York, NY 10027
misra@cs.columbia.edu

Don Towsley

Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
towsley@cs.umass.edu

Abstract

In this paper we critically reexamine some of the long standing beliefs regarding self-similarity and long range dependence (LRD) on the Internet. Power law tails have been conjectured to be a cause of LRD. In this paper, we reexamine the claims regarding heavy tails. We first examine the generative models for the heavy tail phenomena, both in terms of the fragility of some proposed mechanisms to modeling perturbations as well as the weak statistical evidence for the mechanisms. Next, we take a look at some of the implications of LRD in key performance aspects of Internet algorithms. Finally, we present an alternative model explaining the LRD phenomena of Internet traffic. We argue that the multiple time-scale nature of the generation of traffic and transport protocols make the observation of LRD inevitable.

1 Introduction

Long range dependence made a dramatic appearance in network traffic modeling, with the publication of the seminal paper by Leland et al. in 1994 [1]. The paper clearly established that simple Poisson models inherited from the telephony world were not appropriate for packet traffic, and one needed new approaches and paradigms for modeling. Network traffic was shown to be “fractal”, or more accurately long range dependent (LRD), and models and explanations were sought for this surprising behavior. The paper also postulated that heavy tails (specifically power laws) of some distributions associated with the generation of traffic were the likely cause of LRD.

Researchers focused their efforts on detecting power laws, and sure enough, power laws were discovered for web file sizes, web site connectivities, and the router connection degrees, see, for example, [2, 3, 4, 5]. In this paper we focus on the importance of the power laws in the context of LRD.

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 as the origins of power law phenomena, including reflected or drifted multiplicative processes, exponentially stopped geometric Brownian motion, self-organized criticality, preferential growth, and highly optimized tolerance among others [3, 6, 7, 8, 9]. There is also a group of researchers advocating log-normal rather than power law distribution as the correct description of the data (see, for example, [10]). The debate has focused on the “tail” behavior since, for example, the primary difference between a log-normal distribution and a power law distribution lies in the tail. We believe that this debate on the nature of the tail (power law vs. log-normal) 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.

Second, all mechanisms aimed at explaining the power law or log-normal tail are fragile in the sense that minor perturbations to their assumptions 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 next point, which is that engineers should focus on the behavior of a distribution’s “waist”, that refers to the portion for which we have enough data to summarize the distributional information, not its “tail”.

In the next part of the paper, we present a generative model for network traffic that does not rely on heavy tails to produce LRD-like traffic. We argue that the multiple timescale nature of the generation of traffic, coupled with transport protocol issues, make the appearance of LRD-like behavior *inevitable* for Internet traffic.

The rest of the paper is organized as follows. In the next section, we cover some preliminaries regarding the usage of the terms self-similarity and long range dependence in the context of network traffic modeling and analysis. We present our point of view regarding these definitions. Next, we explore some of the weaknesses concerning the power-law hypothesis as an explanation for LRD in network traffic. We then present our generative model for network traffic and the impact of transport protocols

on its correlation structure. Finally, we present our conclusions.

2 Preliminaries and definitions

We begin by first addressing some semantic issues related to “self-similarity” and LRD. The term “self-similar” is one which is encountered in almost all recent literature dealing with traffic modeling. The usage is often ambiguous, confusing and misleading. There are a variety of terms, we present definitions of some commonly used ones.

Continuous Self-similar process A continuous time process $Y(t)$ is said to be *exactly self-similar* with self-similarity parameter H (the Hurst parameter) if it satisfies the following condition

$$Y(t) \stackrel{d}{=} a^{-H}Y(at), \forall t \geq 0, \forall a > 0, 0 < H < 1.$$

where the equality is in the sense of finite dimensional distributions.

Self-similar time series In the network traffic context, we normally deal with a time series rather than a continuous process. The definition of self-similarity in that context goes as follows. Let $X = \{X(i), i \geq 1\}$ be a stationary sequence. Let

$$X^{(m)}(k) = 1/m \sum_{i=(k-1)m+1}^{km} X(i), k = 1, 2, \dots,$$

be the corresponding aggregated sequence with level of aggregation m obtained by averaging over non-overlapping blocks of size m . Then for all integers m , the following holds for a self-similar process

$$X \stackrel{d}{=} m^{1-H} X^{(m)} \tag{1}$$

Second-order self-similar A stationary sequence is said to be second-order self-similar if $m^{1-H} X^{(m)}$ has the same variance and auto-correlation as X for all m .

Asymptotically Second-order self-similar A stationary sequence is said to be asymptotically second-order self-similar if $m^{1-H} X^{(m)}$ has the same variance and auto-correlation as X as $m \rightarrow \infty$. Asymptotically second-order self-similar processes are also called Long Range Dependent (LRD) processes. This is the property that network traffic exhibits, and “self-similar” is often used interchangeably with LRD processes which is misleading. An equivalent way to describe asymptotic second-order self-similarity is in terms of its power spectral density ψ

$$\psi_X(f) \sim k|f|^{1-2H}, f \rightarrow 0 \tag{2}$$

These processes are also equivalently called $1/f$ processes.

Fractional Brownian motion Fractional Brownian motion (fBm) is an example of a self-similar process. A process $X(t)$ is said to be an fBm if the increment process is normally distributed with mean zero and variance t^{2H} . The covariance of the process is described by

$$E[X(s)X(t)] = \frac{1}{2}(s^{2H} + t^{2H} - |s - t|^{2H})$$

where H is the self-similarity parameter.

Fractional Gaussian noise Fractional Gaussian noise (fGn) is the *increment* process of fBm.

Many publications in literature use the term self-similar and long range dependence interchangeably and loosely and this has led to much confusion. Moreover, the Hurst parameter has also been used as a parsimonious measure to describe the correlation structure of analyzed data, which is also misleading since it must assume that the underlying process is self-similar. For Internet traffic in particular, it has been known for a long time [11] that this is not true since the traffic exhibits *strong* daily and weekly patterns, and the behavior has not changed over the years [12]. This can be observed in Figure 1 which we have reproduced with permission from [12]. Each data point consisted of the traffic load at two backbone links collected over 5 minute intervals. Daily and weekly patterns are clearly visible. Thus, Internet traffic is not *truly* self-similar or LRD, as for long enough timescales the LRD-like behavior disappears. This is not surprising, as most of the analyses that have been carried out in the context of LRD have dealt with relatively short duration packet traces (of the order of an hour), before the periodic behavior sets in. The “LRD” phenomena, thus, extends only over a finite range of timescales, and we believe it is not appropriate to discuss limiting values as is typically done.

There have been attempts to generalize the concept of the Hurst parameter to models that are not self-similar but exhibit a sustained correlation structure over a finite range of time-scales. Among these, the work of [13] introduced the idea of pseudo self-similarity. However, these generalizations, and the resultant terminology, also suffer pitfalls and will not be used in this paper. Instead, we will turn to the frequency domain to classify such processes, and use the classical family of $1/f$ processes [14] as the basis for our definition.

We define sustained correlation structure in terms of the power spectral density (PSD) $\psi(f)$ of a process. A $1/f$ process is generally defined as a process whose empirical PSD is of the form $\psi(f) \approx k/f^\nu$ for some $k > 0$ and $0 < \nu < 2$ over several decades of frequency f . For a long range dependent (LRD) processes, this relationship is true for arbitrary small frequencies, i.e., there is no low frequency roll-off. However, there is a class of processes that behave like LRD processes over a certain range of frequencies and the relationship $\psi(f) \approx k/f^\nu$ with ν previously defined, does *not* hold for some frequency $f < F_1$, but it does hold for all frequencies $F_1 < f < F_2$. We say such processes (which are also $1/f$ processes) have a *sustained correlation structure* over a finite range of time-scales, where the range of time-scales is determined by the frequency interval $[F_1, F_2]$ and is given by $[1/F_2, 1/F_1]^1$. A classical example of such a process is the Markov on-off process. The PSD of this process resembles that of a Weiner process, with a ν value of 2, bounded by the on-off time periods on the low frequency

¹Note that a continuous time process under our definition can have $F_2 = \infty$, but the inertia inherent in any physical system makes F_2 a finite value for any observable process. Moreover, [15] suggests that the small timescale behavior is closer to Poisson or independent.




Figure 1: Week long traffic load in bytes at two different links on the Sprint backbone

range (F_1). However, unlike the Wiener process its spectrum flattens out for frequencies below F_1 and it does not suffer from the infinite variance problem associated with LRD processes. The Markov on-off process is a well known short range dependent (SRD) process, however one could easily be misled if only a short finite frequency range is observed, as the process can possess sustained correlation structure over the corresponding time-scales.

3 The fragility and importance of tails

3.1 Terminology

We now turn our attention to the study of power law tails and file size distributions. Before we begin, we introduce terminology that will be used in the remainder of this section. 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 $\alpha > 0$. We refer to the tail of a power law distribution as *power tail* (or, *heavy tail*, *long tail*).

The Pareto distribution is the canonical power law distribution. It 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

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

3.2 Mechanisms for Generating Long-tailed Distributions

In this section, we briefly introduce several mechanisms which have been proposed to generate long-tailed distributions, both lognormal distributions and power law distributions.

3.2.1 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

$$Y_n = Y_{n-1}X_n, \quad n = 1, \dots \quad (3)$$

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

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

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.

3.2.2 Mixture Mechanisms

It is proved in [8] that a Geometric Brownian Motion (GBM) with randomly distributed stopping time has double-Pareto distribution, although the GBM with fixed stopping time has lognormal distribution. Consider a Geometric Brownian Motion (GBM) with initial state X_0 :

$$dX = \mu X dt + \sigma X dw, \quad (4)$$

where w is a standard Brownian Motion process. We are interested in the case where $\mu \geq \sigma^2/2$. At any fixed time t ,

$$Y_t \triangleq \ln(X_t) \sim N(\ln X_0 + (\mu - \sigma^2/2)t, \sigma^2 t)$$

X_t assumes a lognormal distribution with probability density function

$$f_{X_t}(x, t) = \frac{1}{\sqrt{2\pi t}\sigma x} \exp\left\{\frac{-(\ln x - \ln X_0 - (\mu - \sigma^2/2)t)^2}{2\sigma^2 t}\right\} \quad (5)$$

Let T be a nonnegative random variable with density $f_T(t)$, $t \geq 0$. Reed [8] showed that if T is exponentially distributed with parameter λ , $f_T(t) = \lambda e^{-\lambda t}$, the observation of the GBM made at time T , $X_T \triangleq X(T)$, has a double Pareto distribution

$$f_{X_T}(x) = \int_0^\infty \lambda e^{-\lambda t} \times f_{X_t}(x, t) dt = \begin{cases} \frac{\beta\theta}{\beta+\theta} \left(\frac{x}{X_0}\right)^{\theta-1} & x < X_0 \\ \frac{\beta\theta}{\beta+\theta} \left(\frac{x}{X_0}\right)^{-\beta-1} & x \geq X_0 \end{cases} \quad (6)$$

where β and $-\theta$ ($\beta, \theta > 0$) are two roots of the characteristic equation

$$\frac{\sigma^2}{2} z^2 + \left(\mu - \frac{\sigma^2}{2}\right) z - \lambda = 0$$

β determines the power coefficient of $f_{X_T}(x)$ when $x \rightarrow \infty$.

$$\beta = \frac{\sqrt{(\mu - \sigma^2/2)^2 + 2\sigma^2\lambda} - (\mu - \sigma^2/2)}{\sigma^2} \quad (7)$$

For the discrete case, Mitzenmacher [16] investigate the tail behaviors when the stopping time has geometric distribution. The tail of a geometric mixture of lognormal distributions can be approximated by a power tail. This is due to the special form of the stopping time distribution, being either exponential or geometric.

3.2.3 Optimization-based mechanisms.

Constrained optimization has been proposed as a mechanism behind many observed power law like data sets [3]. The basic idea is that optimization of an expected performance index will decrease the probabilities of events contributing significantly to such a mean value, and will 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.

3.3 Fragility of Analytical Tail distributions

Analytical forms of the tail behavior are useful 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.

3.3.1 The multiplication mechanism - fragility of the lognormal distribution.

Recall the multiplication mechanism as defined by (3). For any finite n , the Central Limit theorem states that $\log Y_n$ approximately assumes a distribution with both mean and variance proportional to n . If we choose $E[X_n] = 1$, then $E[\ln X_n] < \ln E[X_n] = 0$ (by Jensen's inequality). It can be shown that $\ln Y_n$ converges to $-\infty$ in probability. Therefore Y_n converges to 0 in probability while $\lim_{n \rightarrow \infty} E[Y_n] = E[Y_0]$.

Suppose that we add a reflection of the form that Y_n is never allowed to fall below a threshold $\xi > 0$. In other words, we modify the construction to be:

$$Y_n = \max\{Y_{n-1}X_n, \xi\}, \quad n = 1, \dots$$

Gabaix, [17], has shown that Y_n converges to an random variable with a Pareto distribution as $n \rightarrow \infty$. The smaller the reflection bound ξ , the slower the convergence. With finite stopping time n , the distribution of Y_n is determined by $\{Y_0, \xi, n\}$. We illustrate the convergence of the distribution of Y_n by the following example: we choose $Y_0 = 1$, $\xi = 1/2^8$, $X_n = 2^{m_n}$, where $\{m_n\}$ are *i.i.d.* random variables with probability mass function:

$$P(m = l) = \begin{cases} \frac{1}{3} & l = -1 \\ \frac{1}{2} & l = 0 \\ \frac{1}{6} & l = 1 \end{cases}$$

Let $Z_n \triangleq \log_2(Y_n/\xi)$, then the evolution of Z_n is

$$Z_{n+1} = \max\{Z_n + m_{n+1}, 0\}$$

Its probability mass function at time n can be solved from the corresponding Markov chain and the initial condition $Z_0 = \log_2(Y_0/\xi)$. As $n \rightarrow \infty$, Z_∞ follows a geometric distribution independent of Z_0 , and thus Y_0 and ξ . Then the distribution of Y_∞ has a power law, with ξ as its smallest value. Figure 2 compares CCDFs of the multiplication process with and without reflection at different stopping time: $n = 64, 128, 192, 256, \infty$. For each stopping time n , the distribution of the multiplication process without reflection ($\xi = 0$) can be calculated as follows: let $W_k \triangleq \log_2(Y_k)$, then

$$W_{k+1} = W_k + m_{k+1} \tag{8}$$

We can solve for the probability mass function of W_n from the Markov chain defined by (8) and $W_0 = \log_2(Y_0)$. From those figures, we can see for finite stopping time and small reflection bound, X_n tends to have a tail similar to lognormal, even though in the limit X_n will eventually converge to a power tail.

Therefore, the multiplication process for generating a lognormal distribution is fragile to a perturbation on the reflection; if $\xi = 0$, it produces a lognormal distribution and otherwise a Pareto distribution.

3.3.2 The mixture mechanism

We revisit the mechanism described in Section 3.2.2. We have the following result:




Figure 2: Reflective Multiplication Process with Finite Stopping Times

Theorem 1 Let T be a nonnegative random variable with density $f_T(t)$ and let X_T denote $X(t)$ when $t = T$. Then

1. $t \prec -\ln f_T(t)$ implies that $f_{X_T}(x) \prec x^{-\alpha}$, $\forall \alpha > 1$
2. $-\ln f_T(t) \prec t$ implies that $x^{-\alpha} \prec f_{X_T}(x)$, $\forall \alpha > 1$
3. $f_T(t) \sim e^{-\lambda t}$ implies that $f_{X_T}(x) \sim x^{-\alpha}$ for some α determined by λ .

See Appendix A for proof.

Thus we observe that the distribution of $X(T)$ is sensitive to the tail behavior of $f_T(t)$; 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.

3.3.3 Fragility of the optimization mechanism

Optimization leads to a power law. While an optimization process does lead to a power law tail with a suitably chosen objective function and under appropriate assumptions, it appears to be a fragile mechanism against the assumptions. For example, one needs to know exactly the a priori distribution model in order to perform the optimization procedure leading to a power law tail. Moreover these a priori distribution models often need to have infinite tails themselves. These assumptions are difficult to satisfy in an engineering system.

3.4 Difficulty of Estimating the Tail

The “waist” of the lognormal distribution and the power law distribution often look very similar. When the variance is large, 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 a constant slope; the CCDF of a lognormal 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 is hard to determine which one is a better fit.

To further illustrate this difficulty, suppose we have a data-set of file sizes: $\{X_i, i = 1, 2, \dots, M\}$. Assume all samples in the data-set 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 estimate of $G(x)$ with variance $Var[\bar{G}(x)] = (1 - G(x))G(x)/M$. (Note: the variance of the empirical CCDF $\bar{G}(x)$ is finite at any point x as long as X assumes a legitimate distribution, even if the variance of X itself is infinite, as for a Pareto distribution.) When M is large, $\bar{G}(x)$ 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

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

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:

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

where M_x is the expectation of the number of files with size greater than x . W keeps increasing when M_x decreases. W will blow up when $M_x \leq a^2$.




Figure 3: Confidence Interval for Pareto and Lognormal Models

We plot the 95% confidence interval in Figure 3 for both the Pareto and the lognormal model used in [10] to fit Crovella’s web file size data-set [2]. We observe that the confidence intervals of both models diverge when file sizes grow. At the extreme tail, the two confidence intervals have a large overlap, which makes it difficult to distinguish them.

This suggests that one really should take into consideration the information richness when investigating distributional properties based on empirical data. Any inference should be made carefully in regions where the information is sparse. This point is echoed by Marron *et al.* in their papers addressing the variability of the heavy tailed durations in Internet traffic [18, 19]. They divided the tail of the distribution of Internet flows’ durations into three regions: “extreme tail” – the part of the tail that is beyond the last data point (thus with no information at all in the data); “far tail” – the part of the tail where there is some data present, but not enough to reliably understand distributional properties; and finally, “moderate tail” – the part of tail where the distributional information in the data is “rich”. It provides a framework for duration distribution analysis in Internet traffic. They also observed the wobbles in the tail of the duration distribution: the tail index varies and does not stabilize. Their analysis in [19] generalized the theory for heavy tail durations leading to long range dependence.

3.5 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 the effect that the extreme tail of a job size distribution has on the configuration or design of existing mechanisms on the Internet. There do not appear to be any design problems or adjustment to existing algorithms or protocols 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.

3.5.1 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 [20], [21], [22]. 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. This was highlighted in simulation a study [23]. TCP 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. Study in [24] also suggests that TCP moderates the effect of heavy tailed file sizes.

Buffer dimensioning on the Internet is done typically in tune with the bandwidth-delay product of the output link of the router (typical values are 2-8 times bandwidth*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, that share the link/buffer with TCP, the tail effects of the performance are largely filtered out due to the small buffers. A recent study [25] also shows that at small time scales network traffic can be well represented by the Poisson model.

3.5.2 The danger of extrapolation for simulation studies

Extrapolating a distribution without evidence can have serious consequences on simulation studies. For example, in [26], Gong *et al.* have developed the MacLaurin series of the average waiting time in $GI/GI/1$ queue. Assume the service time B_n can be expressed as θX_n , where $\{X_n\}$ is i.i.d. with $E[X_n] = 1$, and $\{X_n\}$ are independent of θ . Also assume the inter-arrival density $f(t)$ has a strictly proper rational Laplace transform. Then the average waiting time can be expressed as

$$E[W](\theta) = \sum_{i=0}^{\infty} \sum_{k=1}^{i+1} \frac{1}{k} \frac{f_{*(k-1)}^{(i)}(0)}{(i+2)!} E[(X_1 + \dots + X_k)^{i+2}] \theta^{i+2}, \quad (9)$$

where $f_{*k}^{(n)}(0) = \sum_{i=0}^{n-1} f^{(i)}(0) f_{*(k-1)}^{(n-1-i)}(0)$. The distribution knowledge for the inter-arrival and service time enter the average waiting time in different ways. While the inter-arrival density affects the waiting time mostly through its derivatives around the origin, high moments of the service time are very important. Long tail distributions are poor models for average waiting time, queue length and their variance etc. One could fit a set of web-file 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 it runs the greater the possibility of generating an unrealistically large service time. Simulations with heavy-tailed workloads have been

investigated in [27]. It has been shown that simulations with power law distributions are difficult to converge to steady state and even at steady state there is often high variability.

3.6 The Effect of a Power Law Tail on Traffic

In this section we study the impact of file size distributions on the traffic correlation structure. In particular, we evaluate how the tail of the file size distribution affects the LRD of the resulting network traffic. We perform two tests, one “visual” test, similar to the one shown in [1] that strikingly established the failure of Poisson modeling, and another more formal, statistical test. Because of its simplicity, we use the $M/G/\infty$ input process to model web traffic. An $M/G/\infty$ input process $b(t)$ characterizes the number of busy servers at time t 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 [28]. The auto-covariance of $b(t)$ has been established as:

$$\Gamma(h) \triangleq \text{cov}(b(t+h), b(t)) = \lambda \sum_{r=h+1}^{\infty} P[\sigma \geq r] \quad (10)$$

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 random variable 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

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

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

$$P[\sigma_1 \geq r+1] = P[s_1 > r] = \begin{cases} 1 & 0 \leq r \leq \lfloor c \rfloor \\ (\frac{c}{r})^\alpha & r \geq \lfloor c \rfloor + 1 \end{cases} \quad (11)$$

$$P[\sigma_2 \geq r+1] = P[s_2 > r] = \begin{cases} 1 & 0 \leq r \leq \lfloor c \rfloor \\ (\frac{c}{r})^\alpha & \lfloor c \rfloor + 1 \leq r \leq \lfloor x_0 \rfloor \\ ke^{-\beta r} & r \geq \lfloor x_0 \rfloor + 1 \end{cases} \quad (12)$$

The decay in the auto-covariance function for the Pareto-body and exponential tail process is given by

$$\Gamma(h) - \Gamma(h+1) = \begin{cases} 0 & 0 \leq h \leq \lfloor c \rfloor \\ (\frac{c}{h})^\alpha & \lfloor c \rfloor + 1 \leq h \leq \lfloor x_0 \rfloor \\ he^{-\beta h} & h \geq \lfloor x_0 \rfloor + 1 \end{cases} \quad (13)$$

Thus, we see that the auto-covariance of the process decays *polynomially* over a finite timescale, defined exactly by the Pareto body of the file-size distribution. Now, in [29], the authors showed that a polynomial decaying function over a finite time scale can be fitted with arbitrary accuracy by a finite mixture of exponentials (the result was for a continuous function rather than the discrete auto-covariance as in our case, but we can view the continuous function as a limiting case of the $M/G/\infty$ arrival process as the time slot t goes to 0). The wavelet estimator [30] computes the power spectral density of a process over a *finite* frequency range, governed by the length of the data and sampling frequency. Thus, if we applied the wavelet test to a time series (traffic process) generated by a Pareto body-exponential tail distribution, it would be *statistically indistinguishable* from a process generated by a pure Pareto tail distribution over a finite frequency range. This is demonstrated in the following simulation experiment:

The auto-covariance functions of the $M/G/\infty$ processes $\{b_1(t), b_2(t)\}$ can be computed from (10). For this experiment, we set the arrival rate $\lambda = 1$, Pareto parameters $\alpha = 1.5$, $c = 5$, the turning point for the exponential tail $x_0 = G_1^{-1}(0.001)$, and the rate of exponential tail $\beta = 0.1$. Figure 4 shows the two CCDFs in log-log scale.

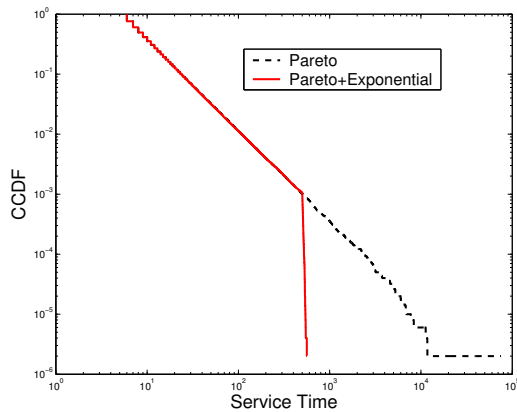


Figure 4: CCDF of the Simulated Service Time

We simulate $\{b_1(t), b_2(t)\}$ by generating 500,000 arrivals for each service time distribution and running them through the server system. Auto-Covariances of $b_1(t)$ and $b_2(t)$ are plotted in Figure 5. In Figure 6 and 7 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

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

We observe that the shapes of both $b_1(t)$ and $b_2(t)$ are similar on timescales from 1 to 1000, although the magnitudes decrease with the increase of timescale, which is consistent with the self-similar properties described in Section 2, such as equation (1). This suggests that a power tail is not necessary for visual self-similarity within a finite range. Next we perform the wavelet test to study the impact of the tail.






Figure 5: Auto-Covariance of $M/G/\infty$ Process

Energy-scale plot is an efficient tool to identify self-similarity and estimate Hurst parameter for time series [31]. It plots a time series' energy $\{y_j\}$, the logarithm of the variance of the wavelet coefficients, at different time scales. It has been shown in [32] that y_j is a good spectral estimator of the time series at the frequency determined by time scale j . The confidence interval of the estimates increases as we increase the scale, as noted in [31]. This is due to the fact that it is essentially a power spectral density (PSD) estimator, and at higher scales it is estimating low frequency components. The finiteness of the data-length leads to increasingly inaccurate estimates of low frequency components. The linear region in the energy-scale plot suggests self-similarity and the slope of the linear region is an estimate of the Hurst parameter. Figure 8 compares the energy-scale curves of $b_1(t)$ and $b_2(t)$. As we can observe, the two plots are indistinguishable from each other up to scale (octave) 11. From Figure 8(a), between octave 6 and 12, the Hurst parameter of $b_1(t)$ can be estimated as $H = \frac{1+0.616}{2} = 0.808$. According to [33], the H parameter of a $M/G/\infty$ process with Pareto service time can be calculated as $H = \frac{3-\alpha}{2}$, where α is the Pareto parameter. We used $\alpha = 1.5$ to generate $b_1(t)$. Therefore the theoretical value of H is 0.75. Energy-scale plot gives us a good estimate for the H parameter of $b_1(t)$. At the same time, if we look at the Energy-scale plot of $b_2(t)$ and estimate its self-similarity index between octave 6 and 12, we will get a similar estimate of $H = \frac{1+0.597}{2} = 0.798$. If we had enough data points to only reliably estimate up to scale 12, the presence or absence of a heavy tail makes no impact on our estimate of the self-similarity index. It is the “power waist” that is the crucial factor here. To distinguish the two processes, one would need enough data-points to obtain reliable estimates of the PSD at low frequencies where the two processes deviate.


In this section, we have argued against focusing on the tails of file size distributions in order to study, model and analyze network traffic. We have demonstrated that “waist” behavior is equally, if not more, important. In the next section, we take it a step further, and develop a Markovian model for traffic exhibits LRD-like behavior.




(a) Timescale= 1



(b) Timescale= 10




(c) Timescale= 100




(d) Timescale= 1000


Figure 6: Sample Path of $b_1(t)$




(a) Timescale= 1



(b) Timescale= 10



(c) Timescale= 100



(d) Timescale= 1000

Figure 7: Sample Path of $b_2(t)$

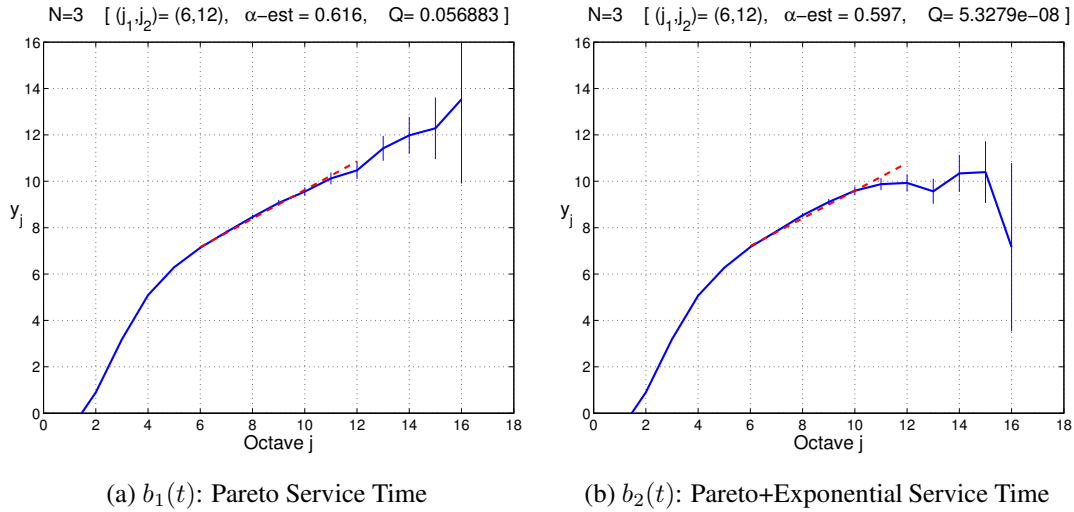


Figure 8: Energy-Scale plot of $M/G/\infty$ Process

4 A Markovian model for LRD traffic

We now present a model, that is Markovian in its components, and yet exhibits LRD-like correlation and spectral behavior. The model was originally proposed by us in [34]. The crucial observation in developing the model is that layered nature of generation of traffic has an important role in determining the correlation structure for individual traffic sources.

4.1 The life of a packet

Let's consider the transmission of a packet between a sender-receiver pair on a network. The packet is transmitted when a number of things simultaneously happen. Firstly, a session has to be in progress between a sender-receiver pair. Next, within the session the particular application (e.g., web browsing) intermittently requests/supplies data. The sender then starts sending the data, usually according to a reliable transport protocol like TCP. The release of packets to the network is decided by the flow and congestion control mechanism. Finally, if the link layer is shared, then packets are allowed on shared access medium only one at a time. For instance, if the access medium is wireless, then packet transmission is governed by the 802.11 mechanism, with intermittent transmission (if the channel is free) and back off stages (if the channel is busy). As is evident, a number of different conditions have to be *true*, largely independent of each other, for a packet to appear on the network. Not only that, the timescales at which the events are occurring are disparate. The transport layer timescale corresponds to typical round trip times, that are sub-second. Application layer on-off events (e.g., browsing, downloading, thinking) are of the order of seconds to minutes, whereas session layer events last from minutes to hours (e.g., user login sessions). Clearly, one should expect multiple timescale behavior from network traffic if it was generated using this mechanism. In the following section, we demonstrate more precisely the multiple timescale nature of the traffic.




Figure 9: A 3-level Hierarchical On-Off Process

4.2 Hierarchical On-Off Process

Definition 1.1 An n -level Hierarchical on-off process (HOP) $Y(t)$ is defined by

$$Y(t) = \prod_{i=1}^n X_i(t) \quad (14)$$

where each $X_i(t)$ is an independent on-off process.

Graphically, the process is shown in Figure 9. The system has an output of 1 only when it is in the leftmost state, i.e. all the on-off processes are on.

We have made no assumption on the nature of the on-off processes besides from independence till now. Let's assume the component on-off processes to be Markovian and examine its spectral properties.

Definition 1.2 A Markovian Hierarchical On-Off Process (MHOP) is a HOP where the component processes are Markovian.

We model our traffic as an aggregation of MHOPs. We also assume in our model that the timescales of different On-Off processes in an MHOP are disparate. We note that an MHOP can be described by a Markov process with augmented state. For example, let $n = 3$. Then pictorially, the Markov chain is illustrated in Figure 10. The output of the process is 1 only when it's in the state 111. It could then also be equivalently thought of as a Hidden Markov Model, where the observation process has 7 Markov states mapped into one state ("off") and we directly observe the remaining eighth state ("on").

We now explain how or why such a process exhibits behavior similar to a $1/f$ process.




Figure 10: Markov chain representation of 3-level MHOP

4.3 Spectral Properties of the MHOP

The autocorrelation function of a Markovian on-off process is given by

$$R_x(\tau) = p_{on}(1 - p_{on})e^{-(\lambda+\mu)|\tau|} + p_{on}^2 \quad (15)$$

Where λ is the transition rate from off to on and μ the rate in the reverse direction. They are related to p_{on} via the relation $p_{on} = \lambda/(\mu + \lambda)$. For the i th process, denote $\lambda_i + \mu_i$ by ν_i , p_{on}^2 by k_{i1} and $p_{on}(1 - p_{on})$ by k_{i2} . The correlation of a product of two independent processes is then be just the product of the individual correlations, i.e.,

$$\begin{aligned} R_{x_1 x_2} &= R_{x_1} \cdot R_{x_2} \\ &= (k_{11} + k_{12}e^{-\nu_1|\tau|})(k_{21} + k_{22}e^{-\nu_2|\tau|}) \\ &= k_{11}k_{21} + k_{12}k_{21}e^{-\nu_1|\tau|} + k_{11}k_{22}e^{-\nu_2|\tau|} + k_{12}k_{22}e^{-(\nu_1+\nu_2)|\tau|} \end{aligned}$$

The Fourier transform of the autocorrelation function gives the power spectral density, and thus we have

$$S_{x_1x_2}(f) = k_{11}k_{21}\delta(f) + \frac{2k_{12}k_{21}\nu_1}{(2\pi f)^2 + \nu_1^2} + \frac{2k_{11}k_{22}\nu_2}{(2\pi f)^2 + \nu_2^2} + \frac{2k_{12}k_{22}(\nu_1 + \nu_2)}{(2\pi f)^2 + (\nu_1 + \nu_2)^2}$$

Observe the last two terms in the above expression. They can be rewritten as

$$(k_{22})(k_{11}\nu_2 + k_{12}(\nu_1 + \nu_2)) \cdot \frac{(2\pi f)^2 + (\nu_1 + \nu_2)\nu_2\rho}{((2\pi f)^2 + \nu_2^2)((2\pi f)^2 + (\nu_2 + \nu_1)^2)}$$

where

$$\rho = \frac{k_{11}(\nu_2 + \nu_1) + k_{12}\nu_2}{k_{12}(\nu_2 + \nu_1) + k_{11}\nu_2}$$

The above corresponds to the power spectral density of the response of a linear time invariant (LTI) system to white noise excitation. The LTI system has two poles at ν_2 and $\nu_1 + \nu_2$ and a zero at $\sqrt{\nu_2(\nu_1 + \nu_2)}\rho$ that lies between the two poles. In our model, we have assumed that the timescales of operations of the component processes are well separated. In physical terms, it is reflecting the fact that different layers like application, transport and link have very different timescales of operations, ranging from seconds to microseconds. Thus $\nu_2 \gg \nu_1$, which implies that both the poles and the zero are bunched very close together. This causes what is known as a *pole-zero cancellation* and we can approximate² the system with one having a single pole and no zero. Making that approximation and placing the resultant pole at ν_2 , we get the power spectral density of the process as

$$S_{x_1x_2} \approx k_{31}\delta(f) + \frac{2k_{32}\nu_1}{(2\pi f)^2 + \nu_1^2} + \frac{2k_{33}\nu_2}{(2\pi f)^2 + \nu_2^2} \quad (16)$$

where we have absorbed all the multiplicative constants into new constants k_{31} , k_{32} and k_{33} . The last two terms can be rewritten as

$$2(k_{32}\nu_1 + k_{33}\nu_2) \frac{(2\pi f)^2 + \nu_1\nu_2\rho}{((2\pi f)^2 + \nu_1^2)((2\pi f)^2 + \nu_1^2)} \quad (17)$$

where

$$\rho = \frac{k_{32}\nu_2 + k_{33}\nu_1}{k_{33}\nu_2 + k_{32}\nu_1}$$

Again viewing it as a linear system, this corresponds to the response of a system that has two poles at ν_1 and ν_2 and a zero at $\sqrt{\nu_1\nu_2}\rho$ that lies between the two poles. Unlike the previous time, this time the poles and zero are well separated and there is no cancellation. However, the presence of the zero between the two poles leads to an interesting phenomena, which is the reason why MHOPs exhibit $1/f$ -like behavior over a certain frequency range. The Bode (double-log) plot for a pole in a linear system corresponds to a flat line until the corner frequency (the pole) and then a straight line with a slope of -2 beyond that. The Bode plot of a zero is a flat line till the corner frequency and then

²Note that we are making the approximation to only explain the phenomena, the plots shown subsequently plot the exact spectrum.




Figure 11: Power Spectral Density for 2-level MHOP and Self-Similar Process

a straight line with slope +2 beyond that. The effect of having a zero between two poles is that the rate of decay of the power spectral density slows down in the frequency region bounded on either side by the poles. If the zero is in exactly the middle of the two poles on the Bode plot (i.e. is the geometric mean of the two) then the decay of the power spectral density is like that of a $1/f$ process, which is LRD. If the zero shifts to the left or right, then the decay is of the form $1/f^\gamma$, $\gamma < 1$ and $\gamma > 1$ respectively. This is shown in Figure 11. The two poles and the zero are marked out on the plot. We chose $\lambda = \mu$ for both levels and $\nu_2 = 10\nu_1$. Between the frequencies bounded by the two poles, the decay of the power spectral density slows down, and goes down with a slope less than 1. A reference power spectral density of a self-similar process corresponding to a Hurst parameter of 0.88 is plotted alongside. Clearly, the 2-level MHOP gives a close approximation to the self-similar process in the frequency range bounded by the poles. Thus, if we have an n -level MHOP, it would correspond to a cascade of such pole-zero systems and would thus tend to approximate the self-similar spectrum over the range of ν 's of all the component processes. To illustrate that, we plot a similar figure (Figure 12) for a 3-level process, again marking out the poles and zeros and plotting a reference self-similar process. Again, the approximation to the self-similar process over the range of ν 's (poles) is striking.

The range of self-similar processes that they approximate (via the Hurst parameter), corresponds to the range observed in network traffic [1], [31]. By changing the value of the *on* probability p in our component processes (thereby changing the λ/μ ratio) we can move the zero around between the poles and have considerable freedom in the range of self-similar slopes that we can approximate. Now that we have seen the capability of MHOPs to approximate the *spectrum* of self-similar processes over a range, the next natural question is that do these processes exhibit the same scaling behavior? We ran




Figure 12: Power Spectral Density for 3-level MHOP and Self-Similar Process

simulations and present the results in a following section.

4.4 Simulation and Analysis of MHOP

To compare the behavior of MHOP with the Bellcore data set, we generated samples which were equal in length and with the same mean value as the various traces analyzed in [1], [31]. While it is well known that self-similar processes can be approximated by Markovian processes, it is usually associated with a state explosion problem. In contrast, here we use only *three* layers for the MHOP model. The three layers (more appropriately, the timescales) could correspond to the transport, application and session layers. For the plots shown, we used $\nu_2 = 10\nu_1$ and $\nu_3 = 30\nu_2$. For all the component processes, we had $\lambda/\mu = 0.8$. The timescale for ν_1 was chosen to be 50 milliseconds. We used an aggregation of 16 independent such processes³. We'll look at three plots, proposed previously, which indicate "self-similarity" at certain timescales. The first one is the so-called "visual-proof".

This is similar to the plot shown in [1]. We selected random sections of the simulated traffic and the August Bellcore dataset and plotted them side by side at the same resolution level. Similar to the graphs plotted in [1], the plots at the highest resolution level have the same random noise term added to both to avoid the visually jarring quantization effect. This is depicted in Figure 13. Visually at least, the simulated MHOP process exhibits a similar bursty behavior as the Bellcore dataset over all the

³Note that while different application and session layers might be independent of each other, it is reasonable to expect some correlation between the transport layer behavior of individual traffic sources. We will address the issue of transport layer behavior in a later section

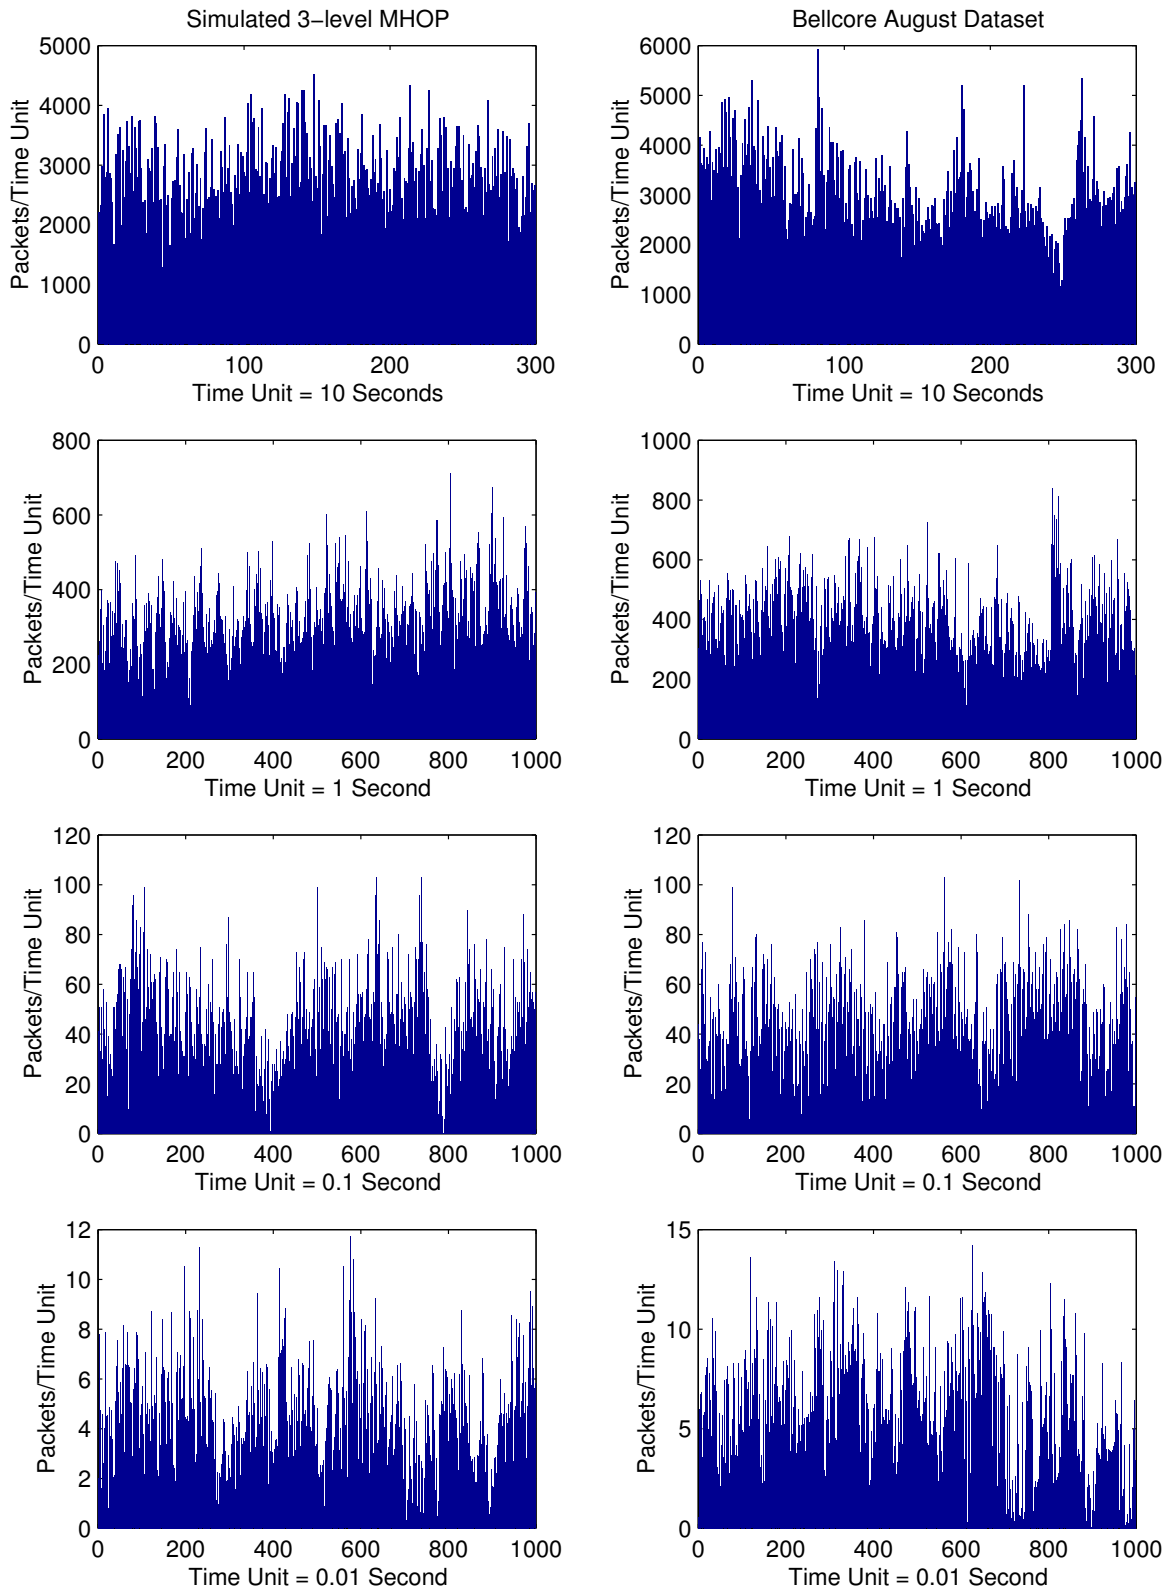


Figure 13: “Visual” proof of Self-similarity




Figure 14: The variance vs. aggregation level plot for the two processes. One of the plots has been shifted down for clarity

levels plotted. Note that the length of the publicly available Bellcore August dataset limits the number of points we can plot at the coarsest resolution level.

The next plot (Figure 14) is the $\log(\text{Variance})$ vs. $\log(\text{aggregation level})$ plots, again similar to those in [1]. Again, the behavior is strikingly similar, with the variances for both processes decaying at the same rate, slower than 1, with increasing aggregation level m .

The final comparison is done via wavelet analysis described earlier. The variance vs. scale plot for the two processes is nearly identical. The two time series are of identical length and hence the confidence intervals for the variances as well as the estimate of the Hurst parameter turn out to be the same. This is depicted in Figure 15.

4.5 The effect of Transport Protocols

In the preceding section, we developed a generative model for traffic that exhibited LRD-like behavior. One of the issues with the model is that when we reach lower layers of the hierarchy, e.g., the transport layer, a Markovian on-off process might be a poor approximation to the actual process, and even independence between sources is not to be expected. Greater than 90% of all Internet traffic is carried by TCP, and a simple on-off model to describe the complex behavior of TCP is inadequate. We address the issue next and demonstrate that the PSD of TCP in fact resembles MHOP.

In [35] we developed a detailed model of the various mechanisms of TCP, and constructed a discrete time Markov process to describe the behavior of TCP traffic. That model can be used as an abstraction to replace a fast timescale component of MHOP.




Figure 15: The variance scale plot for Bellcore Dataset and Simulated Process

The power spectral density function of a discrete-time Markov process can be expressed in term of the eigenvalues of the transition probability matrix. This result was derived in [36] and is given by:

$$\psi(\omega) = \sum_l \frac{\beta_l(1 - \lambda_l^2)}{1 - 2\lambda_l \cos(\omega T) + \lambda_l^2}$$

where ω is the frequency in radians, λ_l is the l -th eigenvalue of the transition probability matrix of the Markov chain, and β_l is the average power contributed by λ_l as defined in [36]. The parameter T represents the time unit of the process, which in the model R , one round trip time.

Note that each eigenvalue λ_l contributes a component to $\psi(\omega)$ over all frequencies. The contribution of λ_l is given by:

$$\psi_l(\omega) = \frac{\beta_l(1 - \lambda_l^2)}{(1 - \lambda_l)^2 + 4\lambda_l \sin^2(\omega T/2)}$$

where the identity $\cos(\omega T) = 1 - 2\sin^2(\omega T/2)$ was applied.

Assuming $\omega T \ll 1$, we have $\sin(\frac{\omega T}{2}) \approx \frac{\omega T}{2}$. As detailed in [35], this assumption is valid for frequencies of interest. We then have

$$\psi_l(\omega) = \frac{\beta_l(1 - \lambda_l^2)/\lambda_0 T^2}{\left(\frac{1 - \lambda_l}{\sqrt{\lambda_l} T}\right)^2 + \omega^2}$$

Observe that the form of the PSD matches what we had for MHOP, i.e. a summation of first order terms, each term defined by an eigenvalue of the transition probability matrix. Thus, if the eigenvalues

are “well separated”, as in MHOP, the process exhibits LRD-like behavior following the analysis of MHOP. Details of the model and its analysis are available in [35], but essentially the value of the eigenvalues depends on the loss probability and round trip time experienced by TCP flows. We reproduce a figure from the paper, that depicts the PSD of the model along with wavelet analysis of simulated traces of the actual TCP protocol using the ns-2 simulator. Note that the traces are for a single long-lived flow, so there is *no power law distribution* involved in the generation of this traffic. We cannot over emphasize this point, that the LRD-like PSD is completely protocol determined and file sizes are of little consequence as long as the transfer is long enough to generate data in the right timescales.

The packet rate time series was generated using the simulation packet trace with a bin size of the average RTT ($R = 100$ ms). Figure 16 shows the results of the wavelet analysis of the normalized time series for different loss probabilities. The solid line indicates the simulation results with a 95% confidence interval denoted by a vertical line over integer values of the time-scale. The dashed line is the PSD of the TCP model, and the *predicted time-scale* marked on the plot depicts the time-scale corresponding to the last eigenvalue of the TCP model. The time-scale is the model prediction where the process starts to lose its correlation structure and the PSD/wavelet energy plot is expected to flatten out.

Our first observation from Figure 16 is that a single TCP flow exhibits sustained correlation over a finite range of time-scales. We note that this correlation structure is present across all loss probabilities. Note that as the loss probability increases, the time-scales over which sustained correlation structure is present increases. For a loss probability of 0.01, the time-scales range from $2R$ to $64R$, while for a loss probability of 0.3 it ranges from $2R$ to $1024R$, which corresponds to a range from 0.2 to 102.4 seconds (almost two minutes). The wavelet plot then flattens out at higher timescales. Coupled with higher layers of the protocol stack and utilizing the MHOP model, it is not surprising that Internet traffic exhibits LRD-like behavior regardless of the tail behavior of the distribution of the data being transferred.

5 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. Multiplicative 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 log-normal “waist” is due to many “small” random effects multiplied together. To this

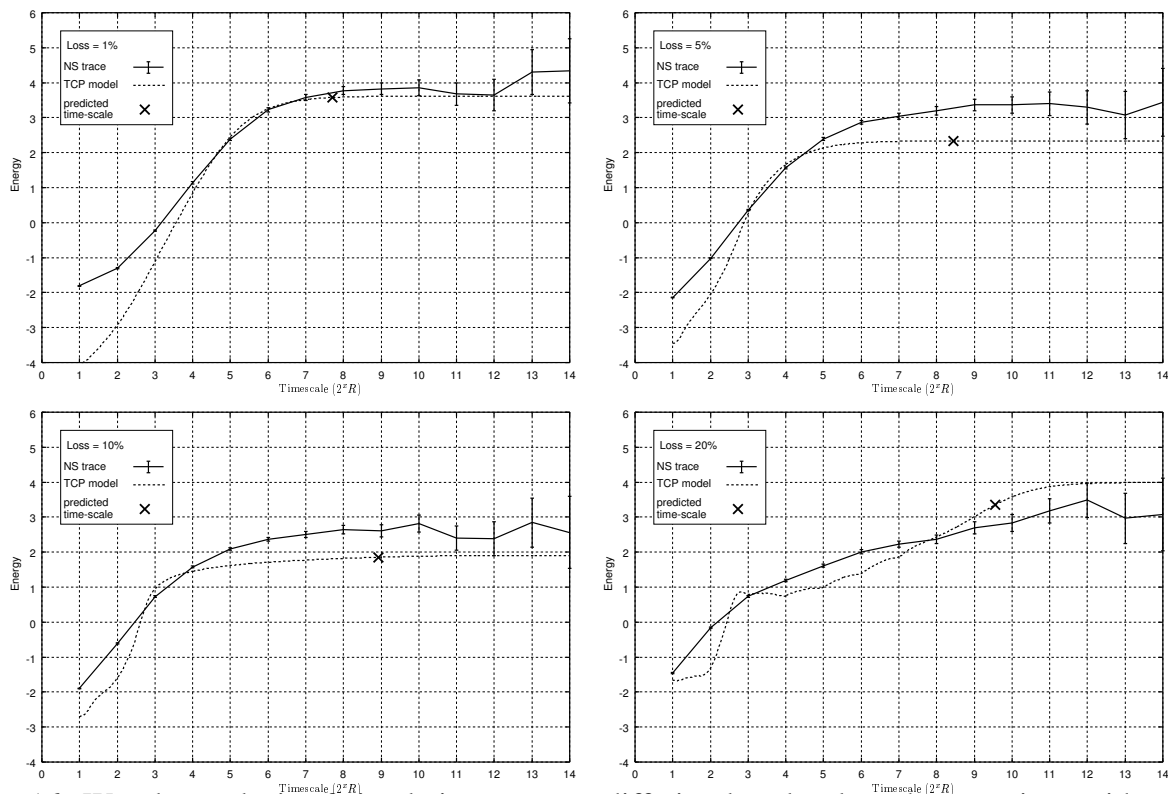


Figure 16: Wavelet analysis of simulation traces at differing loss levels and comparison with a TCP model

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.

In the second half of the paper, we have presented alternate models that explain the LRD-like behavior of Internet traffic. We argue that the generation mechanism of traffic is inherently multiple timescale, and hence sustained correlations are inevitable. Coupled with that is the fact that TCP, the protocol that carries upward of 90% of Internet traffic, generates sustained correlations on its own.

Our point in this paper is that while limiting distributions of filesizes (e.g., power-law or lognormal tails) offer a mathematically convenient and elegant explanation for LRD, they are fraught with problems. The effort, instead, should be on studying the complex nature of generation of traffic and understanding its implications.

References

- [1] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson, “On the self-similar nature of ethernet traffic (extended version),” *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1–15, February 1994.

- [2] M. E. Crovella and A. Bestavros, “Self-similarity in world wide web traffic: Evidence and possible causes,” *IEEE/ACM Transactions on Networking*, vol. 5(6), pp. 835–846, December 1997.
- [3] J. M. Carlson and J. Doyle, “Highly optimized tolerance: A mechanism for power laws in designed systems,” *Physical Review E*, vol. 60, pp. 1412–1427, 1999.
- [4] X. Zhu, J. Yu, and J. Doyle, “Heavy tails, generalized coding, and optimal web layout,” in *Proceedings of IEEE/INFOCOM 01*, April 2001.
- [5] M. Faloutsos, P. Faloutsos, and C. Faloutsos, “On power-law relationships of the internet topology,” in *Proceedings of ACM/SigComm '99*, 1999, pp. 251–262.
- [6] M. Mitzenmacher, “A brief history of generative models for power law and lognormal distributions,” *Internet Mathematics*, 2003.
- [7] P. Bak, Ed., *How Nature Works: the science of self-organized criticality*, Springer-Verlag, 1991.
- [8] W. J. Reed, “The double pareto-lognormal distribution - a new parametric model for size distribution,” <http://www.math.uvic.ca/faculty/reed/index.html>.
- [9] A.L. Barabasi and R. Albert, “Emergence of scaling in random networks,” *Science*, vol. 286, pp. 509–512, 2000.
- [10] A. B. Downey, “The structural cause of file size distributions,” in *Proceedings of IEEE/MASCOTS'01*, 2001.
- [11] Kevin Thompson, Gregory Miller, and Rick Wilder, “Wide-area Internet traffic patterns and characteristics,” *IEEE Network*, November/December 1997.
- [12] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot, “Packet-Level Traffic Measurements from the Sprint IP Backbone,” *IEEE Network*, 2003.
- [13] S. Robert and J. Y. Le Boudec, “On a markov modulated chain with pseudo-long range dependences,” *Performance Evaluation*, vol. 27 – 28, pp. 159 – 173, Oct. 1996.
- [14] G.W. Wornell and A.V. Oppenheim, “Estimation of fractal signals from noisy measurements using wavelets,” *IEEE Transactions on Signal Processing*, vol. 40, no. 3, pp. 611–623, Mar. 1992.
- [15] Jin Cao, William S. Cleveland, Dong Lin, and Don X. Sun, “Internet Traffic Tends Toward Poisson and Independent as the Load Increases,” in *Nonlinear Estimation and Classification*, C. Holmes, D. Denison, M. Hansen, B. Yu and B. Mallick, Ed. Springer, New York, 2002.
- [16] M. Mitzenmacher, “Dynamic models for file sizes and double pareto distributions,” in *preprint*, 2002.

- [17] X. Gabaix, “Zipf’s law for cities: an explanation,” *Quarterly Journal of Economics*, August 1999.
- [18] F. Hernandez-Campos, J. S. Marron, G. Samorodnitsky, and F. D. Smith, “Variable heavy tail duration in internet traffic, part i: Understanding heavy tails,” in *Proceedings of IEEE/MASCOTS’02*, October 2002.
- [19] F. Hernandez-Campos, J. S. Marron, G. Samorodnitsky, and F. D. Smith, “Variable heavy tail duration in internet traffic, part ii: Theoretical implications,” in *Proceedings of 40th Allerton Conference on Communication, Control and Computing*, October 2002.
- [20] I. Norros, “A storage model with self-similar input,” *Queueing Systems*, vol. 16, pp. 387–396, 1994.
- [21] Ashok Erramilli, Onuttom Narayan, and Walter Willinger, “Experimental queueing analysis with long-range dependent packet traffic,” *IEEE/ACM Transactions on Networking*, vol. 4, no. 2, pp. 209–223, 1996.
- [22] P. Jelenkovic and P. Momcilovic, “Capacity regions for network multiplexers with heavy-tailed fluid on-off sources,” in *Proceedings of IEEE/INFOCOM 01*, April 2001.
- [23] Y. Joo, V. Ribeiro, A. Feldmann, A. Gilbert, and W. Willinger, “On the impact of variability on the buffer dynamics in IP networks,” in *Allerton Conference on Communication, Control and Computing*, September 1999.
- [24] Kihong Park, Gitae Kim, and Mark E. Crovella, “On the effect of traffic self-similarity on network performance,” in *Proceedings of the SPIE International Conference on Performance and Control of Network Systems*, 1997.
- [25] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido, “A nonstationary poisson view of internet traffic,” in *Proceedings of IEEE/INFOCOM 2004*, 2004.
- [26] W.B. Gong, S. Nananukul, and A. Yan, “Pade Approximations for Stochastic Discrete Event Systems,” *IEEE Transactions on Automatic Control*, vol. 44, no. 12, pp. 1394–1404, August 1995.
- [27] Mark E. Crovella and Lester Lipsky, “Long-lasting transient conditions in simulations with heavy-tailed workloads,” in *Proceedings of the 1997 Winter Simulation Conference*, 1997.
- [28] M. Parulekar and A. Makowski, “M/G/infinity input processes: A versatile class of models for network traffic,” in *Proceedings of IEEE/INFOCOM’ 97*, April 1997.
- [29] Anja Feldmann and Ward Whitt, “Fitting mixtures of exponentials to long-tail distributions to analyze network performance models,” in *Proceedings of IEEE/INFOCOM (3)*, 1997, pp. 1096–1104.

- [30] D. Veitch and P. Abry, “Matlab code for the estimation of scaling exponents,” http://www.cubinlab.ee.mu.oz.au/~darryl/secondorder_code.html.
- [31] D. Veitch and P. Arby, “A wavelet based joint estimator of the parameters of long-range dependence,” *IEEE/ACM Transaction on Information Theory*, April 1999.
- [32] P. Abry, P. Goncalves, and P. Flandrin, “Wavelets, spectrum estimation, 1/f processes,” *Wavelets and Statistics, Lecture Notes in Statistics*, vol. 105, pp. 15–30, 1995.
- [33] M. Parulekar, “Buffer engineering for M/G/Infinity input process,” *Ph.D Thesis, University of Maryland, College Park*, February 2000.
- [34] Vishal Misra and Wei Bo Gong, “A hierarchical model for teletraffic,” in *Proceedings of the 37th Annual IEEE CDC*, 1998, pp. 1674–1679.
- [35] Daniel Ratton. Figueiredo, Benyuan Liu, Vishal Misra, and Don Towsley, “On the Autocorrelation Structure of TCP Traffic,” *Computer Networks:Special Issue on ”Advances in Modeling and Engineering of Long-Range Dependent Traffic*, 2001.
- [36] San qi Li and Chia-Lin Hwang, “Queue response to input correlation functions: discrete spectral analysis,” *IEEE/ACM Transactions on Networking*, vol. 1, no. 5, pp. 522–533, 1993.

A Proof of Theorem 1

From equation (5), we have

$$f_{X_t}(x, t) \prec f_{X_t}(x, s) \prec x^{-\beta}; \quad \forall \beta \geq 0, \quad 0 \leq t < s \quad (18)$$

and

$$\frac{\partial}{\partial t} f_{X_t}(x, t) = \frac{f_{X_t}(x, t)}{2\sigma^2} \left(\frac{(\ln x - \ln X_0)^2}{t^2} - \frac{\sigma^2}{t} - (\mu - \sigma^2/2)^2 \right)$$

$\forall T_0 > 0$, if $(\ln x - \ln X_0) > \max\{\sqrt{2\sigma^2 T_0}, \sqrt{2}T_0(\mu - \sigma^2/2)\}$, then $\frac{\partial}{\partial t} f_{X_t}(x, t) > 0$ for $t \in [0, T_0]$. Thus for any distribution $f_T(t)$,

$$\int_0^{T_0} f_T(t) \times f_{X_t}(x, t) dt < \int_0^{T_0} f_T(t) \times f_{X_{T_0}}(x, T_0) dt < f_{X_{T_0}}(x, T_0) \prec x^{-\beta} \quad \forall \beta \geq 0 \quad (19)$$

(1) $\forall \alpha > 1$, in equation (7), we can choose $\lambda = \lambda_1$, s.t. $\beta_1 > \alpha - 1$. Then

$$\int_0^\infty \lambda_1 e^{-\lambda_1 t} \times f_{X_t}(x, t) dt = \frac{\beta_1 \theta}{\beta_1 + \theta} \left(\frac{x}{X_0} \right)^{-\beta_1 - 1} \prec x^{-\alpha}$$

$t \prec -\ln f_T(t) \Rightarrow \exists T_0$, s.t. $f_T(t) < \lambda_1 e^{-\lambda_1 t}$, for $t > T_0$.

$$\int_{T_0}^\infty f_T(t) \times f_{X_t}(x, t) dt < \int_0^\infty \lambda_1 e^{-\lambda_1 t} \times f_{X_t}(x, t) dt \prec x^{-\alpha} \quad (20)$$

From (19) and (20), we have

$$f_{X_T}(x) = \int_0^{T_0} f_T(t) \times f_{X_t}(x, t) dt + \int_{T_0}^{\infty} f_T(t) \times f_{X_t}(x, t) dt \prec x^{-\alpha}$$

(2) $\forall \alpha > 1$, in equation (7), we can choose $\lambda = \lambda_2$, s.t. $\beta_2 < \alpha - 1$. Thus

$$x^{-\alpha} \prec \int_0^{\infty} \lambda_2 e^{-\lambda_2 t} \times f_{X_t}(x, t) dt = \frac{\beta_2 \theta}{\beta_2 + \theta} \left(\frac{x}{X_0}\right)^{-\beta_2 - 1} \quad (21)$$

$-\ln f_T(t) \prec t \Rightarrow \exists T_0$, s.t. $f_T(t) > \lambda_2 e^{-\lambda_2 t}$, for $t > T_0$.

From equation (19),

$$\int_0^{T_0} \lambda_2 e^{-\lambda_2 t} \times f_{X_t}(x, t) dt \prec x^{-\alpha}$$

together with (21)

$$x^{-\alpha} \prec \int_{T_0}^{\infty} \lambda_2 e^{-\lambda_2 t} \times f_{X_t}(x, t) dt$$

Finally,

$$f_{X_T}(x) > \int_{T_0}^{\infty} f_T(t) \times f_{X_t}(x, t) dt > \int_{T_0}^{\infty} \lambda_2 e^{-\lambda_2 t} \times f_{X_t}(x, t) dt$$

Therefore, $x^{-\alpha} \prec f_{X_T}(x)$

(3) From equation (6),

$$\int_0^{\infty} c_1 e^{-\lambda t} \times f_{X_t}(x) dt = c_2 x^{-\alpha} \quad (22)$$

Given $\lim_{t \rightarrow \infty} \frac{f_T(t)}{c_1 e^{-\lambda t}} = 1$, $\forall \epsilon > 0$, $\exists T(\epsilon)$ s.t. if $t \geq T(\epsilon)$,

$$(1 - \epsilon)c_1 e^{-\lambda t} \leq f_T(t) \leq (1 + \epsilon)c_1 e^{-\lambda t} \quad (23)$$

From (19), $\int_0^{T(\epsilon)} f_T(t) \times f_{X_t}(x) dt \prec x^{-\alpha}$, therefore $\exists X_1(\epsilon)$, s.t. for $x > X_1(\epsilon)$,

$$\left| \int_0^{T(\epsilon)} f_T(t) \times f_{X_t}(x) dt \right| \leq \epsilon x^{-\alpha} \quad (24)$$

Likewise, $\int_0^{T(\epsilon)} c_1 e^{-\lambda t} \times f_{X_t}(x) dt \prec x^{-\alpha}$, $\exists X_2(\epsilon)$, s.t. for $x > X_2(\epsilon)$,

$$\left| \int_0^{T(\epsilon)} c_1 e^{-\lambda t} \times f_{X_t}(x) dt \right| \leq \epsilon x^{-\alpha} \quad (25)$$

Let $X(\epsilon) = \max\{X_1(\epsilon), X_2(\epsilon)\}$, for $x > X(\epsilon)$,

$$\begin{aligned}
& \left| \frac{\int_0^\infty f_T(t) \times f_{X_t}(x) dt}{x^{-\alpha}} - c_2 \right| \\
& \leq \left| \frac{\int_{T(\epsilon)}^\infty f_T(t) \times f_{X_t}(x) dt}{x^{-\alpha}} - c_2 \right| + \left| \frac{\int_0^{T(\epsilon)} f_T(t) \times f_{X_t}(x) dt}{x^{-\alpha}} \right| \\
& \leq \left| \frac{\int_{T(\epsilon)}^\infty c_1 e^{-\lambda t} \times f_{X_t}(x) dt}{x^{-\alpha}} - c_2 \right| + \epsilon \times \left| \frac{\int_{T(\epsilon)}^\infty c_1 e^{-\lambda t} \times f_{X_t}(x) dt}{x^{-\alpha}} \right| + \epsilon \quad \text{by (23) and (24)} \\
& \leq \left| \frac{\int_0^{T(\epsilon)} c_1 e^{-\lambda t} \times f_{X_t}(x) dt}{x^{-\alpha}} \right| + \epsilon \times c_2 + \epsilon \quad \text{by (22)} \\
& \leq \epsilon \times (c_2 + 2) \quad \text{by (25)}
\end{aligned}$$

In other words, $\forall \epsilon > 0, \exists X(\epsilon)$, s.t. for $x > X(\epsilon)$,

$$\left| \frac{\int_0^\infty f_T(t) \times f_{X_t}(x) dt}{x^{-\alpha}} - c_2 \right| \leq \epsilon \times (c_2 + 2),$$

so we have

$$\lim_{x \rightarrow \infty} \frac{\int_0^\infty f_T(t) \times f_{X_t}(x) dt}{x^{-\alpha}} = c_2$$