

A Pragmatic Approach to Dealing with High-Variability in Network Measurements

Walter Willinger
AT&T Labs
Research
walter@research.att.com

David Alderson
California Institute
of Technology
alderd@cds.caltech.edu

Lun Li
California Institute
of Technology
lun@cds.caltech.edu

ABSTRACT

The Internet is teeming with high variability phenomena, from measured IP flow sizes to aspects of inferred router-level connectivity, but there still exists considerable debate about how best to deal with this encountered high variability and model it. While one popular approach favors modeling highly variable event sizes with conventional, finite variance distributions such as lognormal or Weibull distributions, Mandelbrot has argued for the last 40 years that there are compelling mathematical, statistical, and practical reasons for why infinite variance distributions are natural candidates for capturing the essence behind high variability phenomena. In this paper, we elaborate on Mandelbrot's arguments and present a methodology that often allows for a clear distinction between the two approaches. In particular, by requiring the resulting models to be resilient to ambiguities (i.e., robust to real-world deficiencies in the underlying network measurements) and internally self-consistent (i.e., insensitive with respect to the duration, location, or time of the data collection), we provide a rigorous framework for a qualitative assessment of the observed high variability. We apply the proposed framework to assess previously reported findings about measured Internet traffic and inferred router- and AS-level connectivity. In the process, we also discuss what our approach has to say about recent discussions concerning network traffic being Poisson or self-similar and router-level or AS-level connectivity graphs of the Internet being scale-free or not.

Categories and Subject Descriptors

G.3 [Probability and Statistics]: Distribution functions

General Terms

Measurement, Experimentation, Verification

Keywords

High variability, power-laws, heavy-tailed distributions, scaling distributions, Pareto distributions, lognormal distributions, borrowing strength, sequential moment plots, model consistency, model robustness

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'04, October 25–27, 2004, Taormina, Sicily, Italy.

Copyright 2004 ACM 1-58113-821-0/04/0010 ...\$5.00.

1. INTRODUCTION

Informally, high variability is a phenomenon by which a set of observations takes values that vary over orders of magnitude, with most observations taking small values, with a few observations attaining extremely large values with non-negligible probabilities, and with intermediate-sized observations occurring with appreciable frequencies. A trademark of high variability data is that their sample standard deviation is in general enormous, implying a largely uninformative sample mean as the latter fails to adequately describe the location of the bulk of the observed values. For more than a decade, the Internet has been a rich source for such high variability phenomena, ranging from traffic-related quantities such as, for example, IP flow size [8, 30], TCP connection duration [16], or packet interarrival times of TELNET connections [33], to connectivity-related aspects of the physical infrastructure of the Internet (e.g., number of connections per router [15]) or of virtual networks (e.g., the node degree of inferred Autonomous System, or AS-level maps [15]) or overlay networks constructed at the upper layers of the protocol stack (e.g., the in- or out-degree of the nodes in the Web graph [4]). While the ubiquitous nature of high variability in the Internet is well documented and generally accepted, there is an ongoing debate about how best to capture, model, and ultimately explain observed high variability phenomena. Representative for this debate are the re-occurring discussions about the virtues of using Lognormal versus Pareto-type distributions or, more generally, finite variance versus infinite variance distributions for describing highly variable observations. Many of the arguments made in the existing networking literature (e.g., see [11]) parallel earlier ones encountered in areas such as finance, economics, or biology and detailed by Mandelbrot in [28].

The purpose of this paper is to demonstrate that an appropriate treatment of network-related measurements is capable of adding a new component to this debate that has been largely missing in the past, but helps in advancing network modeling. To this end, we propose a concrete framework that attempts to exploit the qualitative and quantitative properties typically encountered when dealing with network-related measurements more fully than commonly-used methods and requires the resulting models to be *internally self-consistent* and *resilient to ambiguities*. Here, by resilient to ambiguities, we mean that the resulting models are required to be as robust as possible to the deficiencies and ideosyncracies associated with actual network measurements. By insisting on internally self-consistent models, we require that the modeling process yield consistent models when modifying the given data set to reflect shorter (or longer) periods over which the measurements were collected. The main reason for aiming for self-consistent models or model classes is their statistically appealing feature of being by and large insensitive to aspects of the data that are concerned with

measuring one and the same quantity over varying time periods, at different points in time, at different points in the network, or under different networking conditions.

We illustrate the proposed framework in the context of IP flow and Internet connectivity measurements and demonstrate how conventional model fitting is used as a stepping stone and not as an end in itself when deciding whether or not a given model satisfies our desired requirements. In particular, we re-assess previously reported findings in, for example, [11, 12] and [15] concerning Lognormal Internet traffic characteristics and Pareto-type node degrees in router-level graphs, respectively. In the process, we rely on known statistical techniques for Lognormal distributions, exploit somewhat less-known techniques for inferring Pareto-type distributions, and elaborate on largely unknown data-analytic methods and principles advocated a long time ago by Mandelbrot [27] and Tukey [41] for dealing with finite/infinite variance distributions and large data sets, respectively. Specifically, we will focus on and illustrate Tukey’s principles of “borrowing strength (from large data sets)” and “broadening the basis” as well as Mandelbrot’s simple but powerful method of “sequential (moment) plots.” We also show how the resulting Internet modeling approach differs from current practice, we discuss why anybody interested in network simulation, modeling, and analysis should care about the implications of this proposed approach, and we elaborate in detail on what our approach has say about re-occurring claims about network traffic being Poisson rather than self-similar (see for example [5, 23]) and about the recently reported findings that inferred router-level and AS-level connectivity graphs of the Internet are scale-free (e.g., see [1, 45]).

The remainder of the paper is structured as follows. In Section 2, we introduce subexponential distributions and appropriate refinements as efficient and parsimonious models for describing high-variability phenomena and elaborate on their mathematical, statistical, and practical properties. The requirement of internally self-consistent models is the focus of Section 3, while Section 4 deals with our second requirement, namely that the resulting models should be resilient to ambiguities in the underlying measurements. After illustrating these requirements in the context of measured IP traffic and inferred Internet connectivity data, we conclude in Section 5 by discussing a number of implications of the proposed framework for Internet modeling.

2. A MATHEMATICAL FRAMEWORK FOR MODELING HIGH VARIABILITY

We introduce in this section the class of subexponential distributions which provides a rigorous and convenient mathematical framework for dealing with high variability phenomena. To further distinguish between finite versus infinite variance distributions, we consider a subclass of the subexponential distributions, called heavy-tailed or scaling distributions, and elaborate on some key mathematical properties of the latter. For a more comprehensive treatment of these topics, we refer to a survey on subexponential distributions by Goldie and Klüppelberg in [20] and to early works by Mandelbrot on scaling distributions reproduced in [28].

2.1 Heavy-tailed or scaling distributions

Focusing throughout this paper on non-negative random variables X , let $F(x) = P[X \leq x]$, $x \geq 0$, denote the *cumulative distribution function (CDF)* of X and $\bar{F}(x) = 1 - F(x)$ the *complementary CDF (CCDF)*. A typical feature of commonly-used distribution functions is that their (right) tails decrease exponentially fast, implying that all moments, including exponential moments,

exist and are finite. In practice, this property ensures that X exhibits low variability and thus concentrates tightly around its mean. To describe in a mathematically convenient way high variability phenomena, we introduce next the class of subexponential distribution functions.¹ Following [20], we call F (or X) *subexponential* if

$$\lim_{x \rightarrow \infty} \frac{P[X + Y > x]}{P[X > x]} = 2,$$

where Y is an independent copy of X . This definition can be shown to be equivalent to

$$\lim_{x \rightarrow \infty} \frac{P[X_1 + \dots + X_n > x]}{P[\max(X_1, \dots, X_n) > x]} = 1 \text{ for some (all) } n \geq 2,$$

where X_1, X_2, \dots are iid non-negative random variables with distribution function F . This shows that in contrast to low variability distributions, the sum of n iid subexponential random variables is likely to be large if and only if their maximum is likely to be large, and accounts for the non-negligible probability that there will be extremely large values in a subexponential sample. This definition also implies that for subexponential distributions, we have

$$\bar{F}(x)/e^{-\epsilon x} \rightarrow \infty \text{ for all } \epsilon > 0;$$

that is, the (right) tail of a subexponential distribution decays more slowly than any exponential, implying that all exponential moments of a subexponential are infinite. Well-known examples of subexponential distributions include the Lognormal, Weibull, Pareto, and certain stable laws, while the Gaussian, exponential, and Gamma are examples that are not in the class of subexponentials.

To distinguish between subexponential distributions whose regular moments can also be infinite, we next consider the subclass of subexponentials consisting of the heavy-tailed or scaling distributions. To this end, a subexponential distribution function $F(x)$ or random variable X is called *heavy-tailed* or *scaling* if for some $0 < \alpha < 2$

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

where $0 < c < \infty$.² The parameter α is called the *tail index*. For $1 \leq \alpha < 2$, F has infinite variance but finite mean; for $0 < \alpha < 1$, F has not only infinite variance, but also infinite mean. In general, all moments of F of order $\beta \geq \alpha$ are infinite. For example, while both the Lognormal and Pareto have infinite exponential moments, all regular moments of the former are finite, while the regular moments of order β of the latter are infinite for all $\beta \geq \alpha$. Heavy-tailed distributions are called *scaling distributions* because the sole response to conditioning is a change in scale; that is, if X is heavy-tailed with index α and $x > w$, the conditional distribution of X given that $X > w$ satisfies

$$P[X > x | X > w] = \frac{P[X > x]}{P[X > w]} \approx c_1 x^{-\alpha},$$

which—at least for large values of x —is identical to the (unconditional) distribution $P[X > x]$, except for a change in scale. Scaling distributions are also called *power law distributions*, and we will use below the notions of scaling, heavy-tailed, and power law distributions interchangeably and only insist that the right tail of the

¹It is sometimes convenient to consider the slightly more general class of *long-tailed distribution functions* [20], but for the purpose of this paper, this generalization is not needed.

²A more general definition involving regularly varying tails is possible [20], but such a generalization makes applying and inferring scaling behavior cumbersome.

distribution satisfies property (1). One of the most publicized features of scaling distributions which follows trivially from (1) is that their CCDF, when plotted on a log-log scale, appears as a straight line, at least asymptotically.

2.2 Invariance properties of scaling distributions

Scaling distributions enjoy a number of invariance properties that (sometimes uniquely) characterize them. We follow here the presentations in [28], show that scaling distributions are essentially invariant under transformations such as aggregation, mixture, maximization, and marginalization, and discuss some practical implications of this invariance property.

2.2.1 Aggregation

The classical central limit theorem (CLT) is often cited as the reason for the ubiquity with which Gaussian (normal) distributions occur in nature. In its standard form, the classical CLT states that if X_1, X_2, \dots are iid random variables with distribution function F (where F has finite mean μ and variance 1), and if $S(n) = X_1 + X_2 + \dots + X_n$ denotes the n^{th} partial sum, then in the limit as $n \rightarrow \infty$,

$$\frac{S(n) - n\mu}{\sqrt{n}} \rightarrow N(0, 1),$$

where $N(0, 1)$ is the standard Gaussian (normal) distribution having mean 0 and variance 1. More general versions of this statement are available and can be found, for example, in [17].

A somewhat less well-known version of the CLT goes as follows. Let X_1, X_2, \dots be iid random variables with scaling distribution function F with $1 < \alpha < 2$ (implying finite mean but infinite variance). Again, define $S(n)$ to be the n^{th} partial sum. Then in the limit as $n \rightarrow \infty$,

$$\frac{S(n) - n\mu}{\sqrt[n]{n}} \rightarrow S_\alpha,$$

where S_α is a *stable distribution* with index α . Again, more general versions of this non-classical CLT are available. Here, a random variable U is said to have a stable distribution with index $\alpha \in (0, 2]$ if for any $n \geq 2$, there is a real number d_n such that

$$U_1 + U_2 + \dots + U_n = n^{1/\alpha}U + d_n,$$

where U_1, U_2, \dots, U_n are independent copies of U . For $0 < \alpha < 2$, α -stable laws are scaling distributions with index α in the sense of definition (1); for the extreme case $\alpha = 2$, we recover the Gaussian as a special case of the stable distributions. For a detailed treatment of stable distributions, we refer to [37].

Together, these results show that the Gaussian and scaling distributions are both invariant under aggregation. More precisely, the classical and non-classical CLTs state that the stable distributions are the only fixed points of the renormalization group transformation (i.e., aggregation) and that Gaussian distributions are, in fact, a very special case (i.e., $\alpha = 2$).

2.2.2 Maximizing Choices

Consider the case of n independent random variables denoted X_1, X_2, \dots, X_n and assume that their distribution functions are scaling distributions with the same tail index parameter α , but possibly with different scale coefficients $c_i > 0$; that is,

$$P(X_i > x) \approx c_i x^{-\alpha} \text{ for } (1 \leq i \leq n).$$

For $1 \leq k \leq n$, define the random variables M_k to be the k -th

successive maxima given by

$$M_k = \max(X_1, X_2, \dots, X_k).$$

Using that $P(M_k \leq x) = \prod_{1 \leq i \leq k} P(X_i \leq x)$, it is easy to show that for large x ,

$$P[M_k > x] \approx c_{M_k} x^{-\alpha},$$

where $c_{M_k} = \sum_{1 \leq i \leq k} c_i$. Thus, the k -th successive maxima of scaling distributions are also scaling, with the same tail index α , but different scale coefficients than the individual X_i 's.

As for the converse (i.e., M_k is scaling only if the X_i 's are scaling), for the invariance-up-to-scale to hold formally, the distributions of the X_i 's need not follow the scaling distribution exactly. In fact, a result from extreme value theory identifies the invariant distributions as the Frechet distributions and characterizes the distributions of the X_i 's that are in the domain of attraction of the Frechet distribution. The Frechet distribution is defined by $P[M > x] = 1 - \exp(-x^{-\alpha})$, $x > 0$ and is clearly scaling for large x . As a consequence, the individual X_i 's must be so close to being scaling distributions as to be scaling for all practical purposes. In this sense, scaling distributions are the only distributions that are invariant under the transformation of maximization. In particular, Gaussian distributions lack this invariance property.

2.2.3 Weighted Mixtures

As before, assume that X_1, X_2, \dots, X_n are n independent random variables with scaling distribution functions F_i , all with the same tail index parameter α , but possibly with different scale coefficients $c_i > 0$. Consider the *weighted mixture* W_n of the X_i 's, and denote by p_i the probability that $W_n = X_i$. It is easy to show that

$$P[W_n > x] = \sum p_i P[X_i > x] \approx c_{W_n} x^{-\alpha},$$

where $c_{W_n} = \sum p_i c_i$ is the *weighted average* of the separate scale coefficients c_i . Thus, the distribution of the weighted mixture of scaling distributions is also scaling, with the same tail index α , but a different scale coefficient than the individual X_n 's.

Mathematically, the converse (i.e., W_n is scaling only if the X_i 's are scaling) holds only to a first approximation. In fact, in the limit as $n \rightarrow \infty$, the invariant "distributions" for W are of the form $P[W > x] = cx^{-\alpha}$, $x \geq 0$, which are improper distribution functions because they yield an infinite total probability. However, for all practical purposes, the X_i 's are typically restricted by some relation of the form $0 < a \leq x$ which results in perfectly well-defined (conditional) distribution functions of the scaling type. With these qualifications, scaling distributions are the only distributions that are invariant under the transformation of weighted mixture.

2.2.4 Marginalization

Recall that stable distributions are trivially scaling. For the sake of completeness, we note that the stable distributions, like the Gaussian, have natural extensions to the multivariate case. Indeed, the multivariate stable distributions can be characterized as being those for which every linear combination of the coordinates has a (scalar) stable distribution. We call this transformation marginalization and refer to [37] for an in-depth treatment of stable distributions and their properties.

2.3 Scaling distributions: More normal than Normal

Aggregation, mixture, maximization, and marginalization are transformations that occur frequently in natural and engineered systems and are inherently part of many measured observations that are

collected about these systems. For example, aggregate incomes are more widely collected and reported than each type of income separately; the flow or file/document sizes transmitted across the Internet and observed at a particular link within the network are naturally a mixture of distributions of the different file/document sizes residing on the various file/Web servers; for historical data such as natural or technological disasters (i.e., droughts, floods, earthquakes, hurricanes, blackouts, nuclear accidents), the fully recorded and commonly available observations reflect a maximizing choice and correspond to the exceptional (i.e., largest, or most catastrophic) events; and the marginalization transformation is relevant for dealing with a variety of multidimensional economic quantities. The ubiquity with which scaling distributions are observed in such data suggests that scaling is widespread among the constituents underlying these frequently encountered transformations. In this sense, scaling distributions should be considered the norm rather than the exception and should not require “special” explanations.

However, there is an implicit tradeoff between Gaussians being the norm for low variability data and scaling distributions being the norm for high variability data. In the former case, the (traditional) CLT imposes only minimal conditions on the distribution of the individual constituent (i.e., finite variance), but as a result, invariance properties can only be obtained for aggregation and marginalization. In contrast, for high variability data, the relevant CLT requires the (right) tail of the distribution of the individual constituents to decay at a certain rate, and as a result of this more restrictive assumption, the individual constituents are not only invariant under aggregation and marginalization, but also under maximization and weighted mixtures. The pragmatic approach to dealing with high variability data advocated in this paper then consists of viewing Gaussians as the natural null hypothesis for low variability data, where variance estimates exist, are finite, and converge robustly to their theoretical value as the number of observations increases. Similarly, it views scaling distributions as the natural and parsimonious null hypothesis for high variability data, where variance estimates tend to be ill-behaved and converge either very slowly or fail to converge altogether as the size of the data set increases. In addition, it fully exploits the different invariance properties exhibited by low versus high variability data. After illustrating its two main ingredients in Sections 3 and 4 below, a “recipe” for our proposed pragmatic approach to dealing with high-variability data is given in Section 5.

3. REQUIREMENT 1: MODEL CONSISTENCY

Increasingly, conventional model fitting is experiencing the dilemma that when faced with large data sets or with data having non-traditional characteristics (e.g., high variability), standard goodness-of-fit tests to select among alternate models are in general inadequate and fail to choose the “best” model. In this section, we suggest an alternative approach to distinguishing between competing models that uses conventional model fitting not as an end in itself, but applies it iteratively to increasingly larger subsets of the data set at hand and checks for self-consistency among the resulting models.

3.1 Conventional model fitting: An end in itself

In simplified terms, conventional model fitting as described in the standard statistics or time series analysis literature proceeds in four steps. It starts by considering a given data set “as is”, that is, all the available observations are taken at once and at face value. This

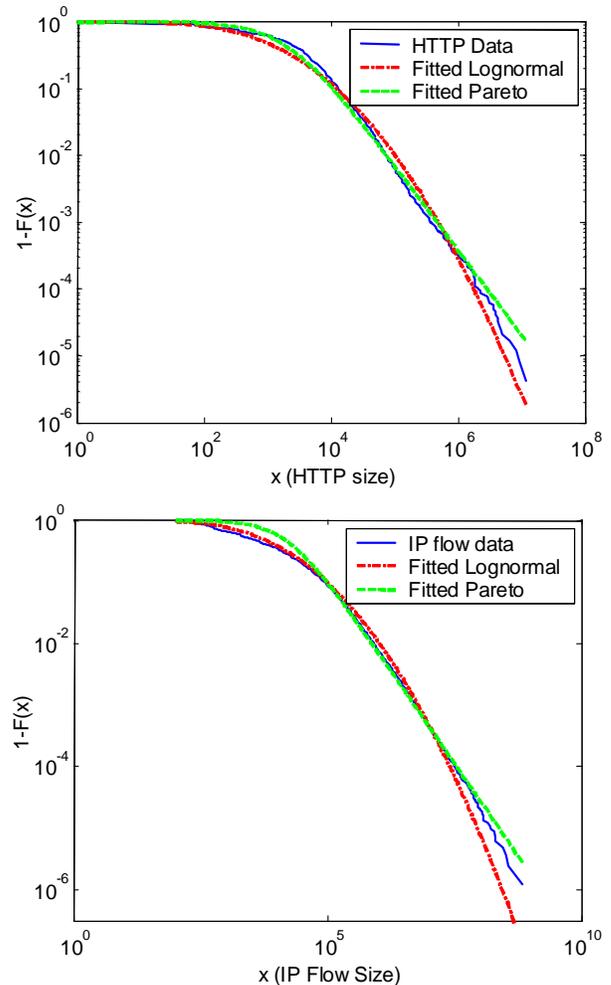


Figure 1: Model fitting by example: CCDF plots of fitted Lognormal, fitted Pareto, and original data; (top) for HTTP data set and (bottom) IP flow data set.

is followed by selecting parameterized models or model classes that are deemed appropriate for the data at hand. In a third step, the full data set is used to estimate the necessary model parameters, and the last step consists of selecting the model that fits the data “best” according to some goodness-of-fit criterion.

Figure 1 illustrates steps 1–3 of this approach with two different data sets and two different models. The first data set is from [44] and consists of some 240,000 HTTP connection sizes (in bytes) collected at LBL’s WAN (in- and outbound) for a 24-hour period in June of 1996. The second data set consists of some 800,000 IP flow sizes (in bytes) derived from the AUCKLAND IV trace [32] and collected over a 4-day period, June 9–12, 2001. For both data sets, we consider the 2-parameter Lognormal(μ, σ) distribution and the 2-parameter Pareto(β, α) distribution of the second kind (e.g., see [22]). Fitting of the Lognormal was done by conventional moment-matching techniques, and fitting of the Pareto involved the “naive” tail index estimate (i.e., slope of fitted straight line through the tail of the CCDF, where the CCDF is plotted on a log-log scale). While more sophisticated parameter estimation techniques could be used, the outcome of this standard model fitting exercise is highly predictable. Reasonable models will provide a reasonable fit, with more highly parameterized models typically yielding a better fit

than more parsimonious ones. Moreover, because of the voluminous data sets and some “unusual” features in the data (e.g., extreme values that are genuine and cannot be dismissed as outliers; possible dependencies), commonly-used goodness-of-fit measures to choose among comparable candidate models generally fail to identify the “best” model. For example, models that are excellent approximations tend to be rejected in large samples, no matter how closely they seem to fit the data, resulting in similar discussions as, for example, in [11] about whether Lognormal or Pareto is a better model for a range of Internet traffic-related quantities.

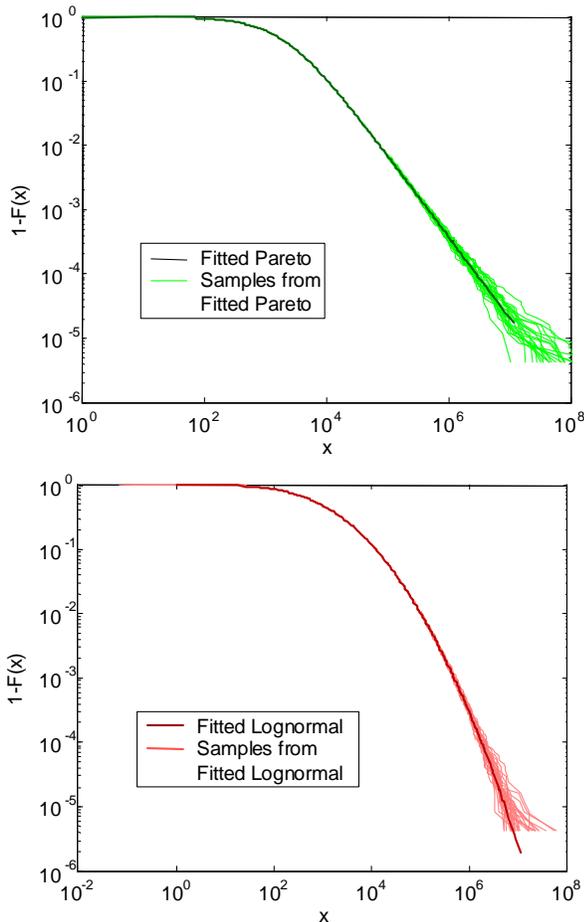


Figure 2: Samples from fitting distributions for the HTTP data set: (top) CCDF plots of 20 samples from fitted Pareto and (bottom) CCDF plots of 20 samples from fitted Lognormal.

A plausible conclusion that could be drawn from Figure 1 is that there seems little difference in using either a Lognormal or Pareto model to describe the two data sets and that given the available data, there is not much else to say. This view seems to be confirmed by Figure 2, where we consider the HTTP data set (the same results hold for the IP flow data set) and the two fitted distributions from Figure 1 (top), and show the CCDFs of 20 independent samples generated from the fitted Pareto (top) and from the fitted Lognormal (bottom). Together, Figures 1 and 2 show that the variability in the tail of the CCDF of the data is consistent with that of samples drawn from either of the two distributions, yet another indication that the two models seem to describe the data set about equally well. In view of G. P. E. Box’s comment that “all model are wrong, but some models are useful”, conventional model fitting applied to network measurements offers increasingly less guidance as to

which models are indeed useful and has left Internet modeling in a rather precarious state, where the same set of measurements are fitted with very different, but apparently equally “good” models, which in turn can give rise to completely opposite explanations and theories for the same observed phenomenon (see Section 5).

3.2 Beyond conventional model fitting: Borrowing strength

To find a way out of the above dilemma, we first note that taking a data set “as is” in step one of the described model fitting process is somewhat arbitrary. For example, in the case of the HTTP data set, we may just as well have ended up with only 1 hour, or half a day, or maybe even with two days worth of measurements, depending on the circumstances under which this measurement effort took place. Thus instead of viewing a given data set as “static”, we propose taking a more “dynamic” view of the data at hand and apply Tukey’s principle of “*borrowing strength from large data sets*” [41] in practice. To this end, let D denote the original data set of size N , start with an initial subset D_0 of D of size N_0 , and consider successively larger nested subsets $D_1 \subset D_2 \subset \dots \subset D_n$ of length $N_0 < N_1 < \dots < N_n$, with $N_n \approx N$. The main motivation for taking this dynamic view of the data set D is that it allows for a careful exploration of the consistency of an assumed model (e.g., a Lognormal or Pareto distribution) as the number of observations increases. In particular, making the commonly-used assumption that one and the same (unknown) underlying process generated the data at hand in the first place, increasing the number of observations as we examine the sets D_0 through D_n should have only the following two main effects. First, the parameter estimates of the fitted model M_i should stabilize, and second, their accuracy expressed in terms of the widths of their corresponding 95% confidence intervals CI_i should improve in such a way that ultimately (i.e., as i tends to n), the confidence intervals CI_i should become roughly nested, with $CI_i \supseteq CI_{i+1}$.

To examine whether the fitted models M_i are indeed self-consistent, we combine Tukey’s borrowing strength principle with Mandelbrot’s “*sequential (moment) estimate plots*” [28]. The latter is simply a method that plots the “running (moment) estimates”; that is, the value of a model parameter estimate or a moment estimate of the data is plotted as a function of the number of observations used in the estimation of the parameter/moment. For example, Figure 3 shows the sequential standard deviation plot for the HTTP and IP flow data sets, respectively. Clearly, while for any fixed n , the sample standard deviation $s(n)$ always exists and is finite, $s(n)$ does not seem to converge as n increases, suggesting that it is conceivable to assume that the second moment does not exist, i.e. the data set is a sample from an underlying infinite variance distribution. To compare and become more familiar with interpreting sequential moment plots, Figures 3(top) and (bottom) also show sequential standard deviation plots for (i) a random permutation of the observations in the two data sets, (ii) same-sized samples generated from the fitted Lognormal in Section 3.1, (iii) same-sized samples generated from the fitted Pareto in Section 3.1, and (iv) same-sized samples generated from fitted exponential distributions. The sequential standard deviation plots for the randomized samples (i) are not inconsistent with the non-existence of a second moment, and there are noteworthy differences between (ii)-(iv). In particular, standard deviation estimates of typical samples generated from a Lognormal with a given standard deviation tend to be far off the true value; while this is fully expected for Pareto samples where second moment estimates are known to diverge, it is troublesome in the Lognormal case because moments estimates are known to converge to their true values.

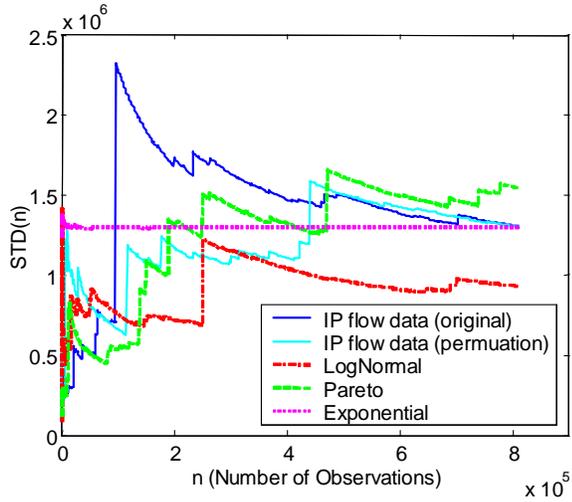
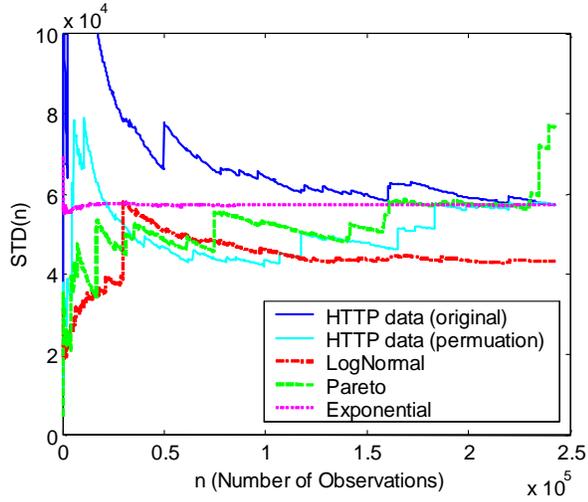


Figure 3: Sequential standard deviation plot for (top) HTTP data set and (bottom) IP flow data set: original data, random permutation of original data, fitted Lognormal, fitted Pareto, and fitted Exponential.

The observed lack of convergence of the sequential standard deviation plot for the data in Figure 3 translates directly into inconsistencies of models that assume finite moments upfront, either implicitly or explicitly. To illustrate, for the HTTP data set, Figure 4(top) shows the sequential estimates $\hat{\sigma}(i)$ of the parameter $\sigma(i)$ of fitted Lognormal models M_i , together with their 95% confidence intervals CI_i (here we used $n_i = 1000 * i$). We observe that the parameter estimates $\hat{\sigma}(i)$ don't seem to converge and that successive 95% confidence intervals are so small so as to have little chance to overlap. In short, while for any fixed i , the fitted Lognormal model M_i appears to provide an adequate fit for the data set D_i , when viewed together, the disadvantage of using Lognormal distributions to fit our data sets becomes evident; the resulting models M_i are clearly inconsistent with one another, and while increasing the number of observations produces more accurate parameter estimates, an apparent lack of convergence of the latter renders the more precise estimates useless. As a reminder that the apparent minor differences in $\hat{\sigma}(i)$ translate into very substantial differences for the standard deviation estimates $\hat{s}(i)$ of the Lognormal variable M_i itself, Figure 4(bottom) shows the sequential estimates $\hat{s}(i)$ and their approximate 95% confidence intervals, where

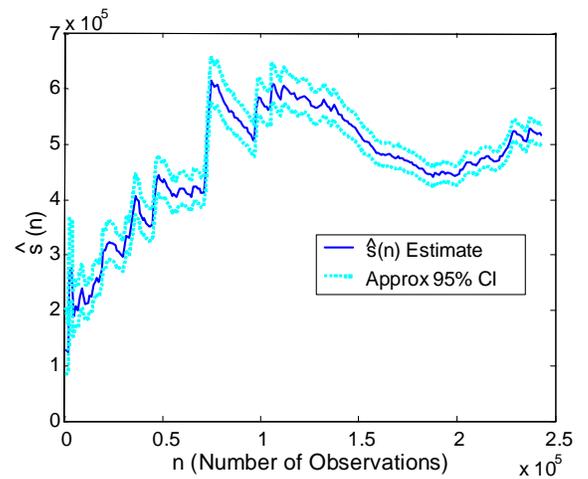
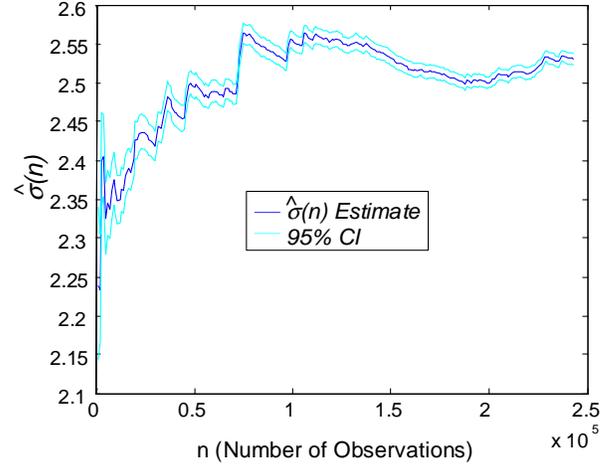


Figure 4: Sequential estimate plots for HTTP data set: (top) Sequential $\hat{\sigma}$ plot with 95% CIs, and (bottom) sequential \hat{s} plot with approximate 95% CIs.

$\hat{s}(i) = \sqrt{e^{2\hat{\mu}(i) + \hat{\sigma}^2(i)}(e^{\hat{\sigma}^2(i)} - 1)}$, with $(\hat{\mu}(i), \hat{\sigma}(i))$ the MLE estimates of the parameters $(\mu(i), \sigma(i))$ of the Lognormal model M_i [22]. Figure 5 shows the same plots for the IP flow data set, with very similar conclusions. To quote Mandelbrot [28, p. 21], “when exactitude is elusive, it is better to be approximately right than certifiably wrong.” For the data sets at hand, using the proposed framework shows that fitting Lognormals is a case of being “certifiably wrong.”

We next apply our approach to show that fitting Pareto models to our data is indeed a case of being “approximately right.” To this end, Figure 6 shows the sequential estimates $\hat{\alpha}(i)$ of the tail index parameter α of fitted Pareto models M_i , together with their 95% confidence intervals CI_i , where we used again $n_i = 1000 * i$. More precisely, we use here the well-known *Hill estimator* to estimate the tail index α of a Pareto distribution and exploit the fact that Hill’s estimator is asymptotically normal to compute approximate 95% confidence intervals for $\hat{\alpha}(i)$. For details about Hill’s estimator, conditions under which it is asymptotically normal, and an expression for the 95% confidence intervals, see for example [34]. The contrast between Figures 4 and 5 and Figure 6 is telling. Not only is there evidence that the tail index estimates $\hat{\alpha}(i)$ converge as i increases to the full size of the data sets, but their corresponding

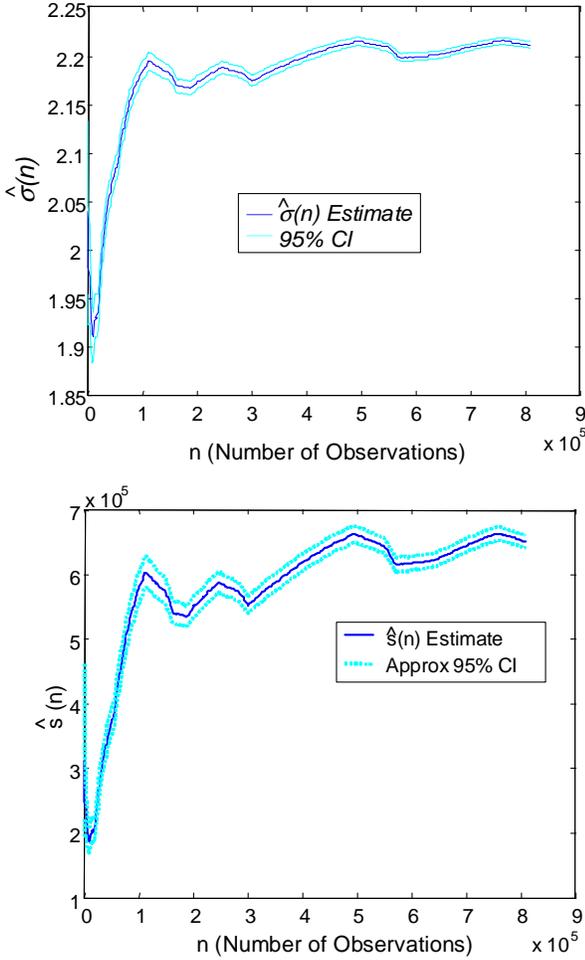


Figure 5: Sequential estimate plots for IP flow data set: (top) Sequential $\hat{\sigma}$ plot with 95% CIs, and (bottom) sequential \hat{s} plot with approximate 95% CIs.

CI_i 's are such that the fitted Pareto models M_i are by and large fully consistent with one another. We take this as strong evidence that in the case of our data, Pareto models are not only “useful” but in fact “better” than Lognormal models. In this sense, model consistency is a powerful requirement and represents an effective criterion for selecting among otherwise comparable alternate models. It also benefits tremendously from the availability of voluminous data sets.

3.3 Beyond borrowing strength

While we have illustrated our approach with two examples where a Pareto model is picked over a Lognormal model, the same method succeeds just as well in selecting a Lognormal over a Pareto when the underlying data is not consistent with a scaling distribution. For example, Figure 7 shows the results of trying to fit a Pareto model to a large sample generated from a Lognormal distribution. We note that in contrast to Figure 6, the tail index estimates $\hat{\alpha}(i)$ steadily increase and don't seem to converge, with obvious implications for the corresponding 95% confidence intervals. Also note that our framework does not only apply to choosing among otherwise comparable distributions, but works just as well for selecting among alternate stochastic process models that are deemed reasonable for having generated the time series at hand in the first place. In fact,

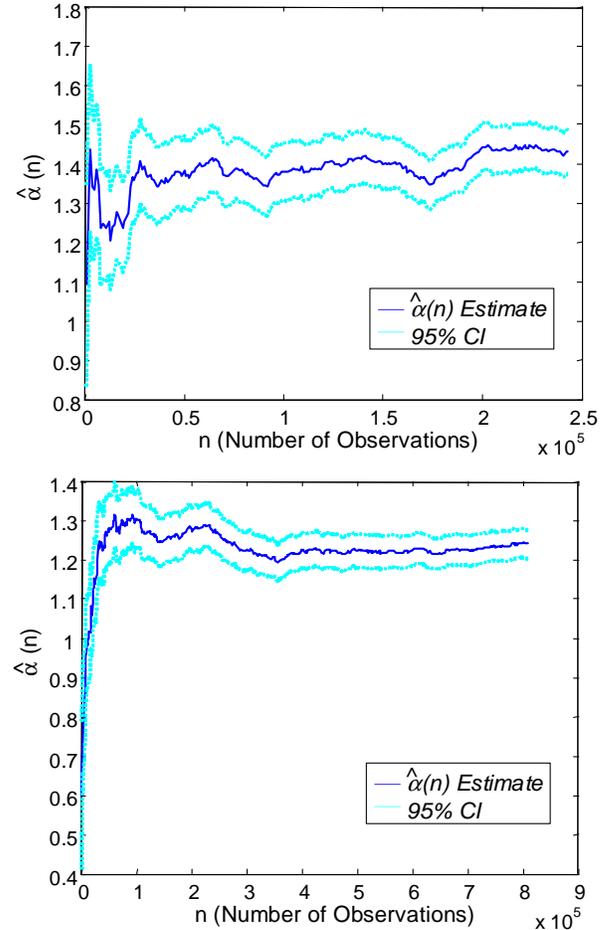


Figure 6: Sequential $\hat{\alpha}$ estimate plot, with corresponding 95% confidence intervals: (top) HTTP data set and (bottom) IP flow data set.

insisting on model consistency turns out to provide an especially powerful and elegant tool for determining whether the time series in question is consistent with long-range dependence or whether a short-range dependent process is a “better” model for the data (for a concrete example, see [18]). Of course, there will always be situations where our approach will fail to identify the “best” model among competing candidates, but in this case, it almost certainly will be able to reveal whether the fitted candidate models are all uniformly “good” or “bad” with respect to the model consistency requirement.

While we advocate here that future Internet modeling efforts should adhere more faithfully to Tukey’s “borrowing strength” principle and thus to making model consistency a general requirement, the networking community already practices another data analysis principle that is also attributed to Tukey and is called “*broadening the basis*”. While related to “borrowing strength”, “broadening the basis” refers more explicitly to attempts on generalizing a finding by drawing on a wider variety of data [13], collected under similar or even dissimilar conditions, at different points in space and time. Thus, in the networking context, broadening the basis is an approach that attempts to find law-like relationships that describe not a single set of measurements, but apply to many data sets collected from the same (or a similar) network or perhaps from very different networks, at different places within the network, over dif-

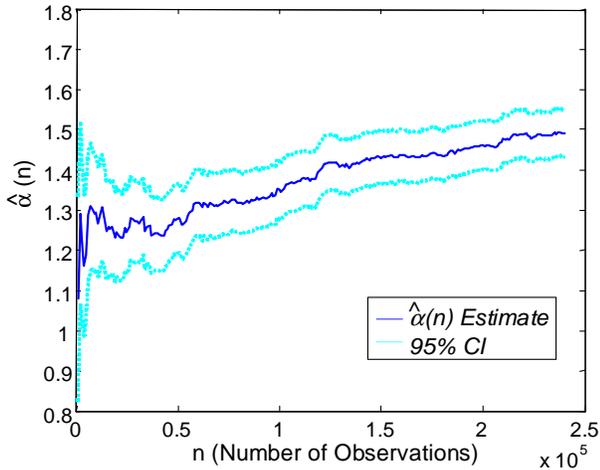


Figure 7: Sequential $\hat{\alpha}$ estimate plot with 95% CIs for a large sample of Lognormal observations.

ferent period of time, and under varying networking conditions. IP flows are a perfect example where applying this principle has produced overwhelming evidence in favor of the scaling property of the size distribution of these basic constituents of aggregate network traffic.

4. REQUIREMENT 2: RESILIENCE TO AMBIGUITY IN THE DATA

While measured Internet traffic traces are generally free of major ambiguities³, there are other network-related quantities of interest which cannot be measured directly and where errors or other ambiguities in the available measurements can lead to premature or erroneous conclusions. In this section, we show how the invariance properties of scaling distributions discussed in Section 2 can be exploited to ensure that fitted models are to some degree resilient to ambiguities in the data. We illustrate our technique with examples of router-level and AS-level Internet connectivity measurements.

4.1 Scaling properties to the rescue

The invariance properties of the scaling distributions discussed in Section 2 have a number of very practical implications for scientific modeling in general. For example, they make scaling distributions insensitive to a wide range of ambiguities that occur when measuring various quantities (see for example [3]). Ambiguities commonly exist in levels of aggregation (e.g., grouping into classes, choice of time segment), changing environments (e.g., entries or exits from a population, varying growth rates, different time segments), differences in accounting (e.g., treatment of multiple authorship) or measuring (e.g., off sets in clock times), etc. As a result of such robustness properties, the power of empirical studies can be vastly expanded by demonstrating, for example, that a fitted model is not only self-consistent in the sense discussed in section 3 but is in addition insensitive to the ambiguities that are often inherent in the process of obtaining the raw measurements in the first place. Moreover, the inherent robustness properties of the scaling distributions greatly facilitate scientific discovery, because they essentially ensure that detecting and identifying scaling laws

³For example, current sampling techniques applied to flow-level measurements are one source of inaccuracies when analyzing flow data, but it is known [14] that the scaling property of flow size distributions is insensitive to this type of ambiguity in the data.

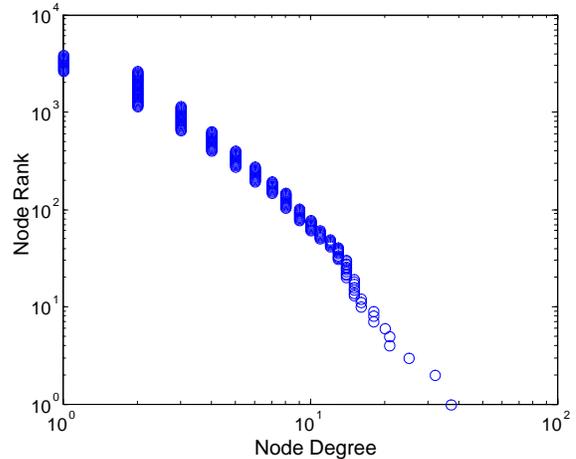


Figure 8: Degree distribution for Rout-95 data [31]. This data set was comprised of 266703 traceroute-based hop records, which translated to 3888 nodes and 5012 edges, spanning dozens of ISPs.

in real data is not only feasible, but can be fully expected to succeed despite imperfect measurements as well as a wide range of ambiguities associated with the actual processes of measuring, accounting, and reporting the data. Properties of a system that require perfect measurements and tolerate no such ambiguities are highly unlikely to be useful, let alone be discovered.

4.2 Ambiguities in measurements of Internet connectivity

For a number of technological, organizational, and economic reasons, the physical connectivity of the Internet as well as of some of its virtual connectivity does not lend itself to direct measurements. Thus it needs to be inferred from other types of measurements, and therein lies the potential source for more or less ambiguity associated with network topology data. Although the elaborate nature of the network protocol suite means that there are a multitude of possible measurements that can be made at different layers, each type of measurement has its own strengths, weaknesses, and idiosyncrasies, and results in possibly distinct, yet fundamentally incomplete, views of the network topology. Even for data within a particular layer, there is no single place from which one can obtain a complete picture of network topology, owing in large part to the thousands of smaller networks that comprise the Internet and are under their own administrative control. This challenge is further exacerbated by a reluctance on the part of commercial network owners and operators to share topology information, particularly when it reveals information about the relationship with their customers or competitors. Nonetheless, the two network topologies that have received significant attention from experimentalists are the *router-level topology* (representing a type of machine-level connectivity) and the *AS topology* (representing organizational interconnectivity between subnetworks) of the Internet.

4.2.1 Router-level connectivity

Because most ISPs consider their router-level connectivity to be proprietary, quantities such as a router's node degree cannot be measured directly and need to be inferred by other means. To this end, the measurements-of-choice are traceroute samples, and coaxing from them the quantities of interest typically requires significant effort and involves more or less sophisticated heuristics for

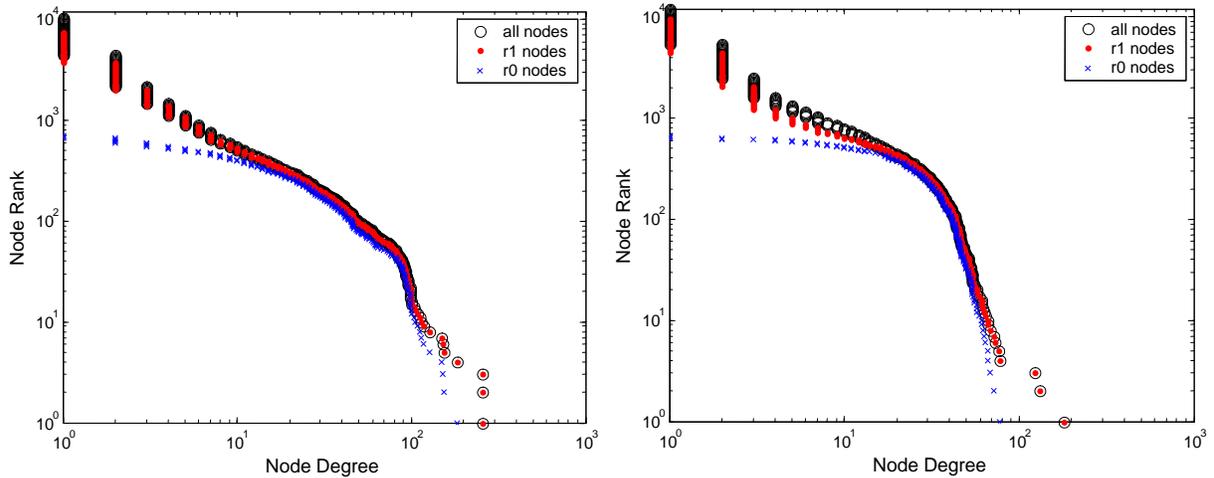


Figure 9: (left) Node degree distribution for AS 1239 from Rocketfuel data, comprised of 700 “radius 0” nodes, 7337 “radius 1” nodes, and 10333 nodes at all levels. (right) Node degree distribution for AS 7018 from Rocketfuel data, 656 “radius 0” nodes, 9429 “radius 1” nodes, and 11799 nodes at all levels.

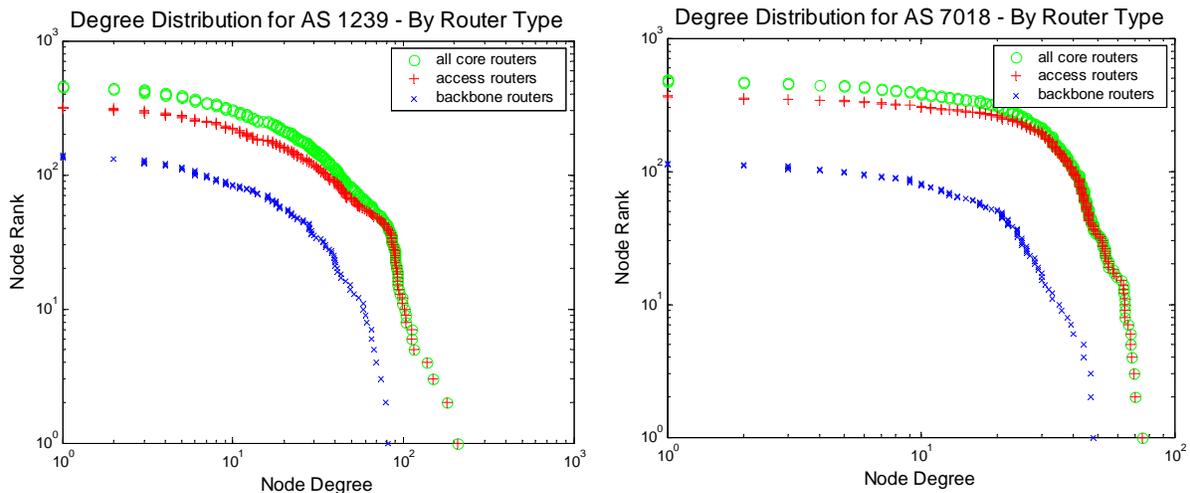


Figure 10: Router-level degree distribution by core router type: (left) Rocketfuel data for AS 1239. (right) Rocketfuel data for AS 7018.

interpreting the data and transforming them into approximate node degree values. It is well-known among networking researchers that the inferred node degrees are extremely ambiguous, and fitted models (e.g., Pareto node degree distribution) that are not resilient to some of the most serious ambiguities are likely to be of little scientific value.

The router-level graph reflects one-hop connectivity between routers running the IP. Information about the connectivity between routers can be inferred from *traceroute* experiments which record successive IP-hops along paths between selected network host computers (see for example, the Mercator [21], Skitter [9], and Rocketfuel [36] projects). Ongoing research continues to reveal more and more idiosyncrasies of traceroute data and shows that their interpretation requires great care and diligent mining of other available data sources. On the one hand, a recent study [25] pointed out a problem inherent to traceroute experiments that can give rise to a potential bias when inferring node degrees. On the other hand, the dependence between the layer-2 and layer-3 issues associated with

router-level connectivity can be another potential source of confusion and/or ambiguity.

When modeling and explaining the router-level connectivity of the Internet, much of the recent emphasis has been on establishing that the router level node degree distribution is a scaling distribution; see for example Faloutsos et al. [15], who based their empirical findings on traceroute data that had been collected earlier by Pansiot and Grad [31] (denoted “Rout-95” in [15] and shown in Figure 8). In contrast, in a more recent study reported in [36], traceroute experiments are used to obtain detailed and sophisticated mappings of a number of router-level ISP topologies, and the resulting data represents the current state-of-the-art in router-level connectivity. In particular, by leveraging additional information available from BGP they are able to conduct more focused traceroute experiments that result in better coverage of the ISP topology under study, and by using knowledge about the way in which IP addresses are commonly assigned within an ISP point of presence (PoP), they are able to perform more accurate alias resolution be-

tween IP addresses and routers. The results based on these more carefully collected Rocketfuel data suggest that while router-level ISP topologies have node degrees that exhibit high variability, they are not consistent with scaling distributions as previously claimed (see Figures 9(left) and (right)).

Before exploiting some of the mathematical properties of scaling distributions to resolve the confusion caused by these observations, we illustrate that the ambiguity of the Rocketfuel data has become manageable in the sense that while still imperfect, fitted models based on these data can be expected to be by and large resilient to the remaining imperfections. The fact that the Rocketfuel data are not perfect either should come as no surprise. For example, since traceroute measurements capture connectivity only at the IP layer, the presence of different technologies at layer-2 (e.g. ATM or Ethernet) can greatly obscure the real connectivity of the physical routing devices (see [40] for additional examples). However, the question of interest is whether or not the remaining imperfections or ambiguities can have a major impact on the model fitting process. Here, we take the Rocketfuel data as is⁴, and we ask what a more networking-based understanding has to say about the reasons behind the plots in Figures 9(left) and (right). We observe first that ISP routers are categorized in Rocketfuel by their “radius” to the network core (i.e. core nodes have radius 0, nodes one hop out from the core have radius 1, and so on). Since some of the traces collected may penetrate deeply into a customer’s network, there is a practical question of how much of the data should be included in the picture of the ISP. The analysis of node degree distribution in [36] was based on data in the “r1” data set (containing nodes of radius 0 and 1), which at the time was believed to provide the most representative sample. However, plots 9(left) and (right) show the vast difference in overall degree distribution between the core nodes and the broader network. The core nodes exhibit much lower variability in their node degree distribution while including nodes further toward the edge has the effect of making the distribution more variable. Thus, knowledge about the underlying techniques used to collect the data, when coupled with engineering knowledge of the network, provide a viewpoint that is significantly more robust than a view that takes the basic data at face value.

Our understanding of the observed structure of the router-level connectivity of these ISPs is further strengthened if we consider the *role* of the router within a given ISP. For example, some routers serve as *backbone* routers (providing connectivity within or between POPs) while others serve as *access routers* (providing connectivity from the ISP to its customers). Inferring this role is possible from the DNS information associated with a particular router, since routers are typically named according to their function⁵. A simple separation of core nodes for each network into these two categories provides insight into the design of the topology, namely that access routers are greater in number and tend to have more connections than backbone routers (see Figure 10(left) and (right)). Again, since this data reflect connectivity at the IP-layer where there may be many virtual connections sharing the same link, we might expect the variability in physical connectivity to be much less. On the one hand, this suggests that the actual “by-router-type” degree distributions exhibit even less variability than shown in Figure 10.

⁴Another level of self-consistency would be to bring in additional resources to scrutinize and validate the Rocketfuel data itself. For example, is the connectivity data consistent with the technological and economic factors that have been recently suggested to shape the way that an ISP thinks about building its topology [26]?

⁵For example, Sprintlink (AS 1239) publishes its router naming convention. See <http://www.sprintlink.net/faq/namingconvention.html> for additional details.

On the other hand, this observation also implies that if there is in fact high variability in observed physical connectivity, it is caused by high variability in connectivity at the edge of an ISP’s network and not within its core.

This type of insight into the engineering details of the network shows that the qualitative features of Rocketfuel-based node degree distributions of router-level ISP topologies are consistent with networking reality and can be explained in terms of underlying networking design principles, technology, and economic considerations. In turn, they can be expected to be resilient to any potentially remaining ambiguities in the data. Thus since the router-level topology of the Internet’s core can be roughly viewed as the aggregation and/or some weighed mixture of individual ISP router-level topologies, and since the latter are not scaling, the former cannot be scaling either. In short, the Rocketfuel study implicitly suggests that the original claims of scaling distributions for router-level node degrees *in the core of the network* are either due to ambiguities in the original data or due to the influence of additional data not captured by Rocketfuel (i.e. corporate and residential edge networks). In either case, claims of scaling in the router-level core of the Internet collapse under scrutiny with more diligently collected and interpreted data.

4.2.2 AS-level connectivity

In models of the AS topology, each node represents an autonomous system (AS) while a link between two nodes indicates the presence of a “peering relationship” between the two ASs, reflecting a mutual willingness to carry or exchange traffic. In this representation, a single “node” (e.g., AT&T) represents potentially hundreds or thousands of routers as well as their interconnections. Although most large ASes have several peering points with other networks, the use of BGP data to infer connectivity means that one is collapsing possibly hundreds of router-level connections all into a single link between two ASes. In this sense, the AS graph is a “virtual” graph representing peering (i.e., business) relationships among ASs (i.e., businesses) and is expressively not a representation of any physical aspect of the Internet structure. Directly measuring AS connectivity is infeasible. The measurements that form the basis for inferring this connectivity consist of BGP routing table snapshots collected, for example, by the University of Oregon Route Views Project [35].

Significant attention has been directed toward discovering the structural features of the AS “graph” and speculating on what these features imply about the large-scale properties of the Internet. Scaling distributions entered by way of empirical studies, first reported by Faloutsos et al. [15]. The findings presented in [15] consist of observed power law relationships between the node outdegree (i.e., number of outgoing connections) and node rank; between number of nodes and outdegree; between the number of node pairs within a neighborhood (in hops) and neighborhood size; and between the eigenvalues of the adjacency matrix of the inferred AS graph and the rank of the AS.

However, AS connectivity data as inferred from BGP is subject to considerable ambiguity. For example, due to the way BGP routing works, snapshots of BGP routing tables taken at a few vantage points on the Internet over time are unlikely to uncover and capture all existing connections between ASs. Other ambiguities that are of concern in this context have to do with the dynamic nature of AS level connectivity, whereby new ASs can enter and existing ASs can leave, merge, or split at any time. To illustrate the degree of this ambiguity, we compare the inferred AS degree distribution for the week of 5 May 2001 as reported by Route Views with the perspective afforded by two additional data sets reported in [6] that

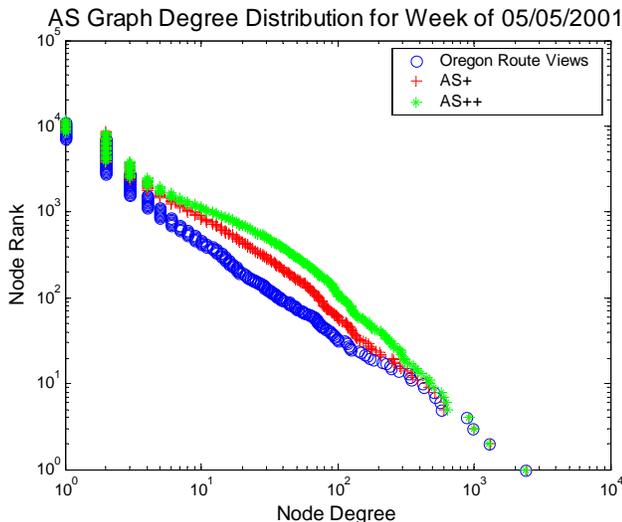


Figure 11: Degree distribution for the Internet at the AS-level: original Oregon Route Views data and enhanced data incorporating additional information sources.

augment Route Views information with independent information⁶ (see Figure 11). Both data sets start with the Route Views data and then add (1) full BGP table dumps from a dozen additional BGP sources, (2) BGP summary information from a large number of Internet Looking Glass sites, and (3) routing policy information from Internet Routing Registry databases. While bringing in these additional sources of data creates in general additional ambiguities, diligent mining of this new information can minimize their impact (see [6]). The first data set denoted here by “AS+” augments the Route Views data with (1) and (2), while the second data set combines “AS+” and (3) and is denoted by “AS++”.

The results from [6] suggest that AS graphs inferred from the Route Views data typically miss between 20-50% or even more of the existing AS connections. Furthermore, in view of our present understanding of scaling distributions, the conjectured/observed power law relationships reported in [15] are intriguing, but should not come as a surprise; after all, the heterogeneity of the Internet is also reflected at the AS level and corresponds to a high variability/infinite variance property of companies (i.e., ASs) that is well-documented in the literature. That the qualitative nature of the original findings of scaling distributions for the node degrees in inferred AS graphs can be expected to hold despite the known deficiencies and ambiguities in the original data is testimony of the degree of robustness to messy and incomplete data exhibited by scaling distributions. However, strict power law relationships for inferred node degrees and related quantities are clearly too brittle and cannot be expected to be robust to missing data on the order of more than 50%. The empirical findings of [6] show that the qualitative nature of the original claims about the node degree distribution of inferred AS graphs is correct, while the quantitative description requires a change, with scaling distributions replacing the originally proposed strict power law distributions.

⁶The findings in [6] are based on data for nine consecutive weeks starting in March 2001, while the data set for 5 May 2001 is representative of all nine weeks.

5. DISCUSSION

We have proposed a framework for Internet modeling that is tailor-made for Internet-related measurements, where high variability, voluminous data sets and varying degrees of ambiguity in the data are the rule rather than the exception. To account for these features, we advocate an approach that results in models that show strong self-consistency properties and are resilient to a range of known or unknown ambiguities in the data. We achieved this by relying on known invariance properties that distinguish scaling distributions from non-scaling distributions in theory and practice and by combining Tukey’s “borrowing strength” principle with Mandelbrot’s “sequential (moment) estimate plot” technique.

In terms of a concrete “recipe”, our proposed approach consists of applying at least one of the following two main procedures. **Procedure 1** (“Exploit the size of the data set”) aims at checking the stability/accuracy of model parameter estimates as a function of the number of observations, consists of computing sequential parameter estimate and confidence interval plots, and provides insight into the (in)consistency properties of a proposed model or model class. **Procedure 2** (“Exploit the semantic content of the measurements”) is intended to check whether observed high variability in a given data set is (in)consistent with an underlying scaling distribution, consists of slicing and dicing the given data into different components, and sheds light on the nature of the observed high variability by appealing to invariance properties of scaling distributions under a number of commonly-used transformation involving the obtained components. Note that while Procedure 1 is in general easy to execute, the applicability of Procedure 2 depends on whether or not the available measurements contain sufficient semantics to perform a sensible “decomposition.”

When applied to Web and IP flow measurements, our results provide a number of technical and practical reasons for preferring scaling distributions over their Lognormal counterparts. In view of these reasons and the fact that aggregate traffic on a link is made up of a superposition of IP flows, mathematics tells us that measured link traffic has to exhibit long-range dependence and is perforce (asymptotically) self-similar (e.g., see [24, 39]), in agreement with a wide range of traffic studies reported in the literature over the last 10 years. In particular, the ubiquitous scaling property of IP flow size distributions reveals a fundamental inconsistency in reported findings of aggregate traffic being Poisson [5, 23], because mathematical arguments that respect the networking-specific relationship between packets and flows preclude a Poisson packet-level dynamics in the presence of scaling distributions for IP flows sizes. In fact, self-similar traffic models—while clearly imperfect, especially at fine time scales—are presently the only traffic models that, over moderate–large time scales, are self-consistent in the statistical sense discussed above, and also self-consistent in the sense of the packet-flow relationship dictated by the network architecture.

Also note that measured Internet traffic can be sliced and diced in many different ways (e.g., by source- and/or destination IP address or prefix, source- and/or destination port numbers, protocols, application-specific attributes), resulting in a number of different traffic constituents. When viewing the traffic on a link as an aggregate, or a weighted mixture, or a combination of aggregate and weighted mixture of many different or a few major such constituents, the scaling property of IP flow size distributions together with the invariance properties of scaling distributions (see Section 2) predict that scaling distributions are widespread among the individual constituents, including the size of the files or documents that reside at the different servers. A popular interpretation of this finding is the “mice/elephant” view of network traffic—while most files (“mice”) have few packets, most packets are in large files (“ele-

phants”). From a general networking perspective, but especially for designing future network congestion protocols, an important question is whether this mice/elephant view is likely to be an invariant of much of the future network traffic, and the work in [46] suggests a positive answer. It does so by considering a number of abstract models of Web layout design with varying levels of analytic tractability and realism with respect to modeling of structure, hyperlinks, and user behavior. Scaling distributions for Web file sizes arise then consistently as solutions to a formulation of Web layout as an optimal design problem, consistent with measured data and irrespective of the different modeling assumptions. In this sense, [46] not only completes the original explanation of self-similar Internet traffic through the connection with Web files and Web layout, but in doing so, it also supports a view of Internet traffic where the observed features at the IP, TCP, and application layers are fully consistent with one another. The work also suggests that the separation of Internet traffic into mice and elephants is likely to be an intrinsic aspect (an “invariant” in the sense of [19]) of future traffic scenarios. To compare, there have been parallel efforts to [46] that have emphasized the use of Lognormal distributions for modeling IP flow sizes [12, 29]. While these efforts have been able to explain Lognormals through generative models based on multiplicative processes, the interpretation of the latter remains unclear and the efforts have resulted in inconsistent views of Internet traffic at the application and IP or TCP layers.

When applied to Internet connectivity data, our proposed approach provides evidence that, in agreement with previous claims, AS-level node degrees appear to be consistent with scaling distributions. However, the evidence also suggests that router-level node degrees in the core of the Internet are not consistent with previous claims of scaling distributions. This finding directly contradicts the recently popular “scale-free” network models that suggest the presence of high degree hub-like connections in the central network core. Nonetheless, claims of scaling in the router-level Internet as a whole may be possible if one includes edge networks. For the Rocketfuel data considered, the highest degree routers were at the edge of the ISP, a finding that is consistent with other recent work which suggests that practical constraints and tradeoffs in constructing router-level networks creates the need for traffic aggregation at all levels of hierarchy and that “good” designs are ones where high variability in connectivity is restricted to the network edge [26]. That the technologies in use at the network edge in turn support high variability in connectivity of end users whose variability in population density is approximately scaling [2] also contributes to this thesis. Thus, the availability of more diligently collected and interpreted data is shifting the current research efforts from trying to explain scaling router-level degree distributions to instead understanding the high variability in measurements of IP-level connectivity, with the ultimate objective of developing a modeling approach to router-level Internet connectivity that is truly *explanatory* in the sense of [43]. To this end, important questions to resolve are (1) *To what extent does the use of different technologies below and above IP contribute to the appearance of high connectivity in the router-level graph?* and (2) *To what extent is the high connectivity fundamental to the Internet’s router-level topology?*

Acknowledgments

This paper was inspired from ongoing conversations and collaboration with John Doyle. We would like to thank Vern Paxson for the “old” HTTP data set, Joel Sommers for making the IP flow data set derived from the AUCKLAND VI trace [32] available, Hyunseok Chang for providing the augmented AS data sets “AS+” and “AS++”, as well as Neil Spring and the Rocketfuel folks for help-

ful discussions about their Rocketfuel data. We are also grateful to the anonymous reviewers for valuable suggestions and helpful comments.

6. REFERENCES

- [1] R. Albert, H. Jeong, and A.-L. Barabási. Attack and error tolerance of complex networks, *Nature* 406, pp. 378–382, 2000.
- [2] D. Alderson. Technological and Economic Drivers and Constraints in the Internet’s “Last Mile”, Technical Report CIT-CDS-04-004, California Institute of Technology, 2004.
- [3] A. Bookstein Informetric distributions, Part II: Resilience to ambiguity *Journal of the Amer. Soc. for Information Science*, Vol. 41, No. 5, pp. 376–386, 1990.
- [4] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. Graph structure in the Web, in: *Proc. of 9th Intl. WWW Conference*, pp. 309–320, 2000.
- [5] J. Cao, W. S. Cleveland, D. Lin, and D. 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 (Eds.), Springer-Verlag, New York, 2002.
- [6] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger Towards Capturing Representative AS-Level Internet Topologies *Proc. of ACM SIGMETRICS ’02*, 2002.
- [7] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The Origin of power laws in Internet topologies revisited. *Proc. IEEE INFOCOM ’02*, New York, NY, 2002.
- [8] K. Claffy, H.-W. Braun, and G. Polyzos. A parameterizable methodology for Internet traffic flow profiling, *IEEE Journal on Selected Areas in Communications* 13, pp. 1481–1494, 1995.
- [9] Cooperative Association for Internet Data Analysis (CAIDA), Skitter. Available from www.caida.org/tools/measurement/skitter/.
- [10] A. B. Downey. The structural cause for file size distribution. www.allendowney.com/research/filesizes/cds-tr25-2000.ps.gz, 2000.
- [11] A. B. Downey. Evidence for long-tailed distributions in the Internet. *Proc. 1st ACM SIGCOMM Internet Measurement Workshop ’01*, San Francisco, CA, pp. 229–241, 2001.
- [12] A. B. Downey. Lognormal and Pareto distributions in the Internet. www.allendowney.com/research/longtail/downey03lognormal.pdf, 2003.
- [13] D. Draper, J.S. Hodges, C.L. Mallows and D. Pregibon. Exchangeability and data analysis, *J. R. Statist. Soc. A* 156, pp. 9–37, 1993.
- [14] N. Duffield, C. Lund, and M. Thorup. Estimating flow distributions from sampled flow statistics, *Proc. ACM SIGCOMM ’03*, Karlsruhe, Germany, 2003.
- [15] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. *Proc. ACM SIGCOMM ’99*, Cambridge, MA, pp. 251–262, 1999.
- [16] A. Feldmann, Characteristics of TCP connection arrivals, in: *Self-Similar Network Traffic and Performance Evaluation*, K. Park and W. Willinger (Eds.), J. Wiley&Sons, New York, pp. 367–400, 2000.
- [17] W. Feller. *An Introduction to Probability Theory and its Applications*, Volume II, Second Edition. Wiley&Sons, New York, 1966.

- [18] D.R. Figueiredo, A. Feldmann, B. Liu, V. Misra, D. Towsley, and W. Willinger. On TCP and self-similar traffic. *Performance Evaluation*, 2004 (to appear).
- [19] S. Floyd and V. Paxson. Difficulties in simulating the Internet. *IEEE/ACM Trans. Networking*, Vol. 9, pp. 392–403, 2001.
- [20] C. M. Goldie and C. Klüppelberg. Subexponential distributions. in: *A Practical Guide to Heavy Tails*, R. J. Adler, R. E. Feldman, and M. S. Taqqu (Eds.), Birkhäuser, Boston, 1998.
- [21] R. Govindan and H. Tangmunarunkit. Heuristics for Internet Map Discovery, *Proceeding of IEEE INFOCOM '00*, 2000.
- [22] N.L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*, Volume 1, Second Edition J. Wiley&Sons, New York, 1994.
- [23] T. Karagiannis, M. Faloutsos, M. Molle, and A. Broido. A nonstationary Poisson view of Internet traffic. *Proc. IEEE INFOCOM '04*, 2004.
- [24] T. G. Kurtz. Limit Theorems for Workload Input Models, in: *Stochastic Networks: Theory and Applications*, F.P. Kelly, S. Zachary, and I. Ziedins (Eds.), pp. 119–140, Oxford University Press, 1996.
- [25] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. *Proc. IEEE INFOCOM '03*, 2003.
- [26] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the Internet's router-level topology. *Proc. ACM SIGCOMM '04*, Portland, Oregon, August 30 - September 2, 2004.
- [27] B. B. Mandelbrot. New methods in statistical economics, *Journal of Political Economics* 71, pp. 421–440, 1963.
- [28] B. B. Mandelbrot. *Fractals and Scaling in Finance*. Springer-Verlag, New York, 1997.
- [29] M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, Vol. 1, No. 2, pp. 226–251, 2004.
- [30] P. Newman, T. Lyons, and G. Minshall. Flow labelled IP: A connectionless approach to ATM, *Proc. IEEE INFOCOM '96*, 1996.
- [31] J.-J. Pansiot and D. Grad. On routers and multicast trees in the Internet. *ACM Computer Communication Review*, 28(1):41-50, January 1998.
- [32] Passive Measurement and Analysis (PMA) Project, <http://pma.nlanr.net/>
- [33] V. Paxson and S. Floyd. Wide-area traffic: The failure of Poisson modeling, *IEEE/ACM Trans. on Networking* 3, pp. 226–244, 1995.
- [34] S. I. Resnick. Heavy tail modeling and teletraffic data. *Annals of Statistics* 25, pp. 1805–1869, 1997.
- [35] Route Views, *University of Oregon Route Views Project*, Available at <http://www.antc.uoregon.edu/route-views/>.
- [36] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. *Proc. ACM SIGCOMM '02*, Pittsburgh, PA, pp. 133–145, 2002.
- [37] G. Samorodnitsky and M. S. Taqqu. *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. Chapman and Hall, New York, 1994.
- [38] J. W. Stewart III. *BGP4: Inter-Domain Routing in the Internet*. Addison-Wesley, Boston, 1999.
- [39] M.S. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *Computer Communication Review* 27, pp. 5–23, 1997.
- [40] R. Teixeira, K. Marzullo, S. Savage, and G.M. Voelker. In Search of Path Diversity in ISP Networks, In *Proc. IMC '03*, Miami Beach, Florida. October 27-29, 2003.
- [41] J. W. Tukey. Data analysis and behavioral science or learning to bear the quantitative's man burden by shunning badmandments, in: *The Collected Works of John W. Tukey*, L. W. Jones (Ed.), Vol. III, Wadsworth, Monterey, CA, 1986.
- [42] A. Veres and M. Boda. The chaotic nature of TCP congestion control. *Proc. IEEE INFOCOM '00*, 2000.
- [43] W. Willinger, R. Govindan, S. Jamin, V. Paxson and S. Shenker. Scaling phenomena in the Internet: Critically examining criticality. *Proc. Nat. Acad. Sci.* 99, suppl. 1, pp. 2573–2580, 2002.
- [44] W. Willinger and V. Paxson, Where Mathematics meets the Internet, *Notices of the AMS* 45, pp. 961–970, 1998.
- [45] S.-H. Yook, H. Jeong, and A.-L. Barabási. Modeling the Internet's large-scale topology. *Proc. Nat. Acad. Sci.* 99, pp. 13382-13386, 2002.
- [46] X. Zhu, J. Yu, and J. C. Doyle. Heavy tails, generalized coding, and optimal Web layout. *Proc. IEEE INFOCOM '01*, 2001.