

A comparative study of different heavy tail index estimators of the flow size from sampled data

Patrick Loiseau
Patrick.Loiseau@ens-lyon.fr

Paulo Gonçalves
Paulo.Goncalves@ens-lyon.fr

Pascale Primet
Vicac-Blanc
Pascale.Primet@ens-lyon.fr

INRIA
UNIVERSITÉ DE LYON

ABSTRACT

In this article, we address the problem of estimating the tail parameter of a flow size distribution from sampled packet traffic. Based on synthetic data, we perform a systematic comparison of several estimators proposed in the literature. In the course, we propose a variant to an existing method which takes into account some statistical *a priori* on the expected distribution. This adapted estimator shows a significantly improved performance, as compared to the others.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations—*Network monitoring*; G.3 [Probability and Statistics]

General Terms

Measurement, Theory

Keywords

Heavy Tail Distribution, Packet Sampling, Estimation, Grid

1. INTRODUCTION

1.1 Motivation

Grids are distributed systems, based on shared computational, storage and visualisation resources interconnected by long distance networks. Compared to clusters, they introduce new scales in terms of heterogeneity and number of co-operative equipments, users' community size, number of inter-dependant processes, processing capacities, bandwidth, etc. The large distances between computation entities leading to large delays, together with the increasing possibility of losing packets turn communications performance into a major challenge in grid networks.

For a decade, research on grid performances have essentially been based on internet transport protocols such as TCP or

UDP. However, the particular topology of grid networks and the specificity of grid applications makes the grid context very different from the internet's one. For example, our grid testbed, Grid5000, an experimental grid platform gathering a total of 5000 CPUs is based on a dumbbell core topology interconnecting 9 sites geographically distributed in France with very high speed access links. The access rate of computing nodes in such environment is 1Gb/s. Some resources are interconnected by 10Gb/s links. In grid networks, the aggregation level is often quite low.

This gives rise to a few interrogations : are these internet protocols adapted to grid applications ? do they guarantee optimal Quality of Service (QoS) and security ? what are the influence of the different parameters of these protocols on the performance of specific applications ? To answer these questions, traffic characteristics have to be studied in grid context.

Traffic characteristics have already been studied in the internet for a decade. Long Range Dependence (LRD) and self similarity have been observed in internet traces. Typical flow characteristics such as heavy tail have also been observed. Then, the modelling of the traffic has become a very active field of research. A few theoretical and empirical results arose about traffic characterisation [14, 10]. But these results are mainly based on the traffic observed in the internet. Are these results still valid in the grid context ? To what extent can they be adapted to grids ?

Lots of methods (see [1] and references within, [9, 5, 8, 13] for traffic characterisation are based on the observation of the entire traffic (*i.e.* every packets are picked).

However, such methods are very challenging in very high speed networks because of memory and CPU consumption issues. For example, in the worse case where we have to deal with a 64 Bytes packet stream reaching the maximal bandwidth of 10Gb/s, the time available to process a packet is about 50 ns. Moreover, to stock a 56 Bytes header for each packet would need to stock more than 1GB/s. It is then necessary to sample, *i. e.* to pick only a sub-sample of the packets going through the link. When observing only a sub-sample of the traffic, the estimation of the flow size distribution tail parameter α and LRD parameter H is harder. The major question addressed in this paper is the estimation of the tail parameter α from sampled traffic.

1.2 Contribution and outline

The problem we address in this article has been tackled in many different ways by various authors, and despite its relative simplicity, no close form solution was found. In all cases, inevitable assumptions of different natures yield only approximated solutions with their own advantages and drawbacks. Our first objective is therefore to compare these different approaches, assessing the accuracy of the corresponding α 's estimators on synthesized, thus totally controlled flow streams. Then, in the course of our systematic study, we were led to propose slight adaptations to an existing method, namely the *scaling method*, to exploit some *a priori* information concerning the expected flow size distributions. Based on the same data set, we empirically demonstrate that compared to its companion forms, the resulting estimator shows significantly improved performances, both in terms of bias and of variance.

Layout of our article is as follows: In section 2, we clearly expose the estimation problem and develop its mathematical formulation. Basic definitions are introduced and the usual notations are posed. Section 3 itemizes some of the proposed methods to solve the α estimation problem. We start with statistical approaches aimed either at maximizing the likelihood function with an EM algorithm [6, 7] or at (approximately) inverting the conditional probabilities [7]. Our methodological contribution comes here, where we derive a parameterized inversion which is adapted to the statistical *a priori* we have on the flow size distributions. A fourth approach relies on a non-parametric wavelet-based estimator of the distributions tail exponent [8]. The last estimator we tested is based on a totally different approach, and resorts to stochastic counting techniques [2]. Section 4 reports experimental results obtained from numerical simulations. To properly evaluate the bias and the variance of each estimate, we advisedly used a flow stream simulator which allows for tuning traffic parameters that deem relevant to our study. Section 5 ends with conclusive remarks and our future work plan.

2. FORMULATION AND NOTATIONS

2.1 Notations

In this study, we focus on flow-oriented analyses to characterize traffic distributions.

Our (traffic) time series correspond to a succession of contiguous IP packets with variable sizes, observed over a time window of duration T . This segment is analyzed independently from the rest of the traffic.

The rigorous definition of a flow is a delicate task which depends on the required degree of refinement. For instance, a timeout is commonly considered which fixes the maximum acceptable time interval between two consecutive packets of a same flow. As for us, we will disregard this timeout, and adopt an oversimplified definition of a flow, as the set of packets occurring in period T with the same source and destination IPs, the same source and destination ports, and the same protocol (TCP or UDP). As it introduces no difference between bulk data transfer, control traffic, workloads, etc., the flow notion such defined is not application oriented.

The flow size (or flow length) is a random variable noted X ,

and corresponds to the number of packets of a given flow. The original flow size distribution reads:

$$P_X(X = i) = \phi_i \propto \mathbb{E}\{f_i\}, \quad (1)$$

where \mathbb{E} stands for the ensemble mean or expectation operator, and f_i is the frequency of flows of size i observed during period T .

A random variable Z is called *heavy tailed* with tail parameter α if

$$P_Z(|Z| \geq z) = z^{-\alpha} L(z), \quad (2)$$

where L is a slowly varying function, *i.e.* $L(tz)/L(z) \rightarrow 1$ as $z \rightarrow \infty$ for any $t > 0$. In the sequel, we systematically assume original flow size distributions with heavy tail, and endeavor to estimate the tail exponent α .

There exist many efficient ways, parametric or not, for estimating the asymptotic distribution decay from independent realizations of a random variable X . For instance, the Hill [9] and the Nolan [13] estimators are maximum likelihood estimators in the case of Pareto and Alpha stable distributions respectively. More simply also, the tail exponent α can be associated to the asymptotic slope obtained when performing a linear regression of the distribution in a log-log plot.

However, memory consumption and CPU limitations, may prohibit to collect and to analyze every packets going through a high speed link. In those situations, we are compelled to apply a sampling procedure in order to reduce the amount of treated data. Naturally then, neither the random variable X , nor its distribution P_X , is directly observable anymore.

The most straightforward way for down-sampling the packets acquisition is to sequentially retain one packet every $N \in \mathbb{N}$, where $1/N$ is the sampling rate. This method, called deterministic sampling, is implemented in Cisco routers, [4]. However, a random sampling procedure, which would consist in randomly selecting a packet with a probability $p = 1/N$, drastically simplifies theoretical developments, even if the two methods are equivalent when a sufficiently large number of intertwined flows is assumed [3]. Thereafter, we assume probabilistic sampling for theoretical issues, although it is a deterministic sampling that is experimentally used.

After sampling, the observed variable becomes the sampled flow size, that is the number of packets belonging to same initial flow that were picked. We designate by Y this random variable and its distribution is noted:

$$P_Y(Y = j) \propto \mathbb{E}\{g_j\}, \quad (3)$$

where g_j is the frequency of the sampled flows with size j observed in a time window of length T . Let us stress that the sampling procedure has a double impact on the distribution P_Y : not only the actual flow size i scale down to a sampled flow size $j \leq i$, but also the number of occurrences f_i of original flows (of a given size i) reduces to a smaller number of effective observations.

The goal of this paper is then to compare some possible approaches to estimate the tail parameter α of an original distribution P_X from the distribution P_Y of observed sampled flows.

2.2 Mathematical formulation

Sampled flow distribution. Assuming random sampling with probability $p = 1/N$, the probability for a sampled flow to be of size j given that it comes from an original flow of size i is governed by a binomial law:

$$P_{Y|X}(Y = j|X = i) = B_p(i, j) = \binom{i}{j} p^j (1-p)^{i-j}. \quad (4)$$

Then, the complete probability formula permits to express the sampled flow sizes distribution in terms of the actual distribution:

$$\begin{aligned} P_Y(Y = j) &= \sum_{l=j}^{\infty} B_p(l, j) P_X(X = l) \\ &= \sum_{l=j}^{\infty} B_p(l, j) \phi_l, \end{aligned} \quad (5)$$

a key relation which must be inverted to get the original flow size distribution.

Maximum likelihood solution. Whenever it is analytically feasible, a classical solution to this estimation problem is to find an estimate $\hat{\phi} = \{\hat{\phi}_i, i = 1, \dots, \infty\}$ that maximizes the conditional probability $P(\mathbf{g}|\phi)$ to observe the particular realization $\mathbf{g} = \{g_j, j = 1, \dots, \infty\}$. For practical reasons, this conditional probability is often replaced by the so-called *log-likelihood function* $\mathcal{L}(\phi) = \log P(\mathbf{g}|\phi)$, and the resulting maximum likelihood estimate denoted $\hat{\phi}^{\text{ML}}$, is solution of the following constrained optimization problem:

$$\hat{\phi}^{\text{ML}} = \operatorname{argmax}_{\phi \in \Delta} \mathcal{L}(\phi), \quad (6)$$

where the constraint $\phi \in \Delta = \{\phi : \phi_i \geq 0 \forall i, \sum_i \phi_i = 1\}$ compels $\hat{\phi}^{\text{ML}}$ to be an admissible probability density function.

Implicitly, the likelihood function $\mathcal{L}(\phi)$ assumes that no original flow is missed by the sampling procedure, meaning that we observe at least one sampled packet per flow. If it is so, the log-likelihood can be written as:

$$\mathcal{L}(\phi) = \sum_{j \geq 1} g_j \log \sum_{i \geq j} \phi_i c_{ij}, \quad (7)$$

where $c_{ij} = B_p(i, j)$.

In practice however, the assumption that no flow was withdrawn hardly holds. In that case, it is necessary to first calculate the conditional probability ϕ'_i that a flow is of size i , knowing that at least one of its packets was observed. Only then, can we solve the maximum likelihood (6) where $\mathcal{L}(\phi')$ is adapted accordingly.

Analytical solutions for the maximum likelihood principle are generally difficult to find. But, as we will see in the next section, there exist an iterative algorithm, called Expectation-Maximization (EM), which permits to numerically solve the constrained optimization problem of equation (6) and to approximate the maximum likelihood estimator $\hat{\phi}^{\text{ML}}$.

Inverse approximation. As an alternative to ML estimation, we can strive to directly invert the relation (5), and express

the conditional probability $P_{X|Y}$ using the Bayes Formula:

$$\begin{aligned} P_{X|Y}(X = i|Y = j) &= \frac{P_{Y|X}(Y = j|X = i)P_X(X = i)}{P_Y(Y = j)} \\ &= \frac{B_p(i, j)P_X(X = i)}{\sum_{l=j}^{\infty} B_p(l, j)P_X(X = l)}. \end{aligned} \quad (8)$$

In this expression, it is now necessary to set P_X to an *a priori* distribution, in order to infer the searched original distribution via the marginal relation:

$$P_X(X = i) = \sum_{j=0}^{\infty} P_{X|Y}(X = i|Y = j)P_Y(Y = j). \quad (9)$$

Needless to say that the estimation accuracy of P_X severely depends on a consistent choice (guess) for the *a priori* distribution.

Tail exponent estimation. Let us assume for now, that we were able to get a satisfactory estimation of P_X . The corresponding complementary cumulative density function reads:

$$F(i) = \Pr(X \geq i). \quad (10)$$

Under the heavy tail hypothesis of equation (2), we have

$$F(i) \underset{i \rightarrow \infty}{\sim} \left(\frac{b}{i}\right)^\alpha, \quad (11)$$

with b some real constant. Taking the logarithm, this asymptotic relation becomes linear:

$$\lim_{i \rightarrow \infty} \log F(i) = -\alpha \cdot \log(i) + \log(b), \quad (12)$$

and α is simply estimated by a linear regression over an appropriately chosen scale range.

3. DIFFERENT ESTIMATORS

3.1 Statistical inference

This section presents a first class of possible tail exponent estimators, which consist in estimating the original flow size distribution first, and then apply the linear regression of equation (12) to get an estimate of the parameter α .

MLE via EM algorithm. As we mentioned it, the EM (for Expectation-Maximization) algorithm permits to numerically solve the maximum likelihood problem of equation (6). This approach was first used in [7] to estimate original flow size distributions from sampled traffic, and later on in [11] with some sensible adaptations.

The EM algorithm was introduced in [6], and a detailed description can be found in [12]. This versatile algorithm is an iterative procedure that converges to the maximum likelihood estimate (6) when missing data prohibit its straightforward calculation. Observed data are said incomplete and viewed as an observable function of the so-called complete-data. Let f_{ij} denote the frequency of original flows of size i which give rise to sampled flows of size j . The set of the f_{ij} for every i and every j is the complete data. The observation g can simply be recovered from the complete data : $g_j = \sum_{i \geq 1} f_{ij}$. Assuming the complete data to be known, we form the complete data log-likelihood :

$$\mathcal{L}_c(\phi) = \sum_{i \geq j \geq 1} f_{ij} \log \phi_i c_{ij}. \quad (13)$$

After an initialisation (choose some initial flow length distribution $\phi^{(0)}$, let us say for example a uniform distribution), the algorithm iterates the following two steps:

E step. Form the expectation $\mathcal{Q}(\phi, \phi^{(k)})$ of the complete data log-likelihood, conditionally to the observation, considering the distribution at the k -th iteration $\phi^{(k)}$:

$$\mathcal{Q}(\phi, \phi^{(k)}) = \mathbb{E}_{\phi^{(k)}} \{ \mathcal{L}_C(\phi) | \mathbf{g} \} \quad (14)$$

M step. Define and determine:

$$\phi^{(k+1)} = \operatorname{argmax}_{\phi \in \Delta} \mathcal{Q}(\phi, \phi^{(k)}). \quad (15)$$

In our case, a direct calculation of the expectation of step E, and utilisation of the Lagrange multipliers to maximize Q under the same constraint as in (6), lead to the expression :

$$\phi_i^{(k+1)} = \frac{\phi_i^{(k)}}{\sum_{j \geq 1} g_j} \sum_{j=1}^i \frac{g_j c_{ij}}{\sum_{l=j}^I \phi_l^{(k)} c_{lj}}. \quad (16)$$

As it is often the case with EM algorithms, this estimation is computationally expensive. Moreover, as we will report in the section 4, the choice of a timely stopping criterion, has a crucial impact on our particular estimate $\hat{\phi}^{\text{ML}}$.

Inverse approximation – Uniform a priori (scaling method). The simplest *a priori* we can plug in relation (8) to directly evaluate the expression (9), is to start with X uniformly distributed, i.e. $P_X(X = i) = C, \forall i$. Then, the conditional probability (8) becomes:

$$P_{X|Y}(X = i|Y = j) = \frac{B_p(i, j)}{\sum_{l=j}^{\infty} B_p(l, j)} = \frac{B_p(i, j)}{N}, \quad (17)$$

which can even be further simplified, if $B_q(i, j)$ is coarsely approximated by a rectangle window with support $jN \leq i < (j+1)N - 1$. The resulting form:

$$P_{X|Y}(X = i|Y = j) = \begin{cases} 1/N, & \text{for } i = jN, \dots, (j+1)N - 1 \\ 0, & \text{otherwise} \end{cases} \quad (18)$$

leads to a very simple expression of the inferred original density :

$$\hat{f}_i = \frac{1}{N} g_{\lceil \frac{i}{N} \rceil}, \quad \forall i. \quad (19)$$

This estimate conveys the naive idea that all the g_j sampled flows of size j exclusively come from original flow sizes i comprised between jN and $(j+1)N$, and hence, that it suffices to re-scale \mathbf{g} by the sampling rate $1/N$ to get \mathbf{f} .

Along the same lines, the estimation accuracy of the f_i 's can substantially improve if we refine the choice of our *a priori*. In particular, as we are expecting heavy tail distributions, it seems reasonable to integrate this knowledge from the beginning, and to consider an *a priori* P_X that is itself an heavy tail distribution. It is this approach that we develop in our next section.

Inverse approximation – Pareto a priori. Let us then consider an *a priori* distribution P_X that follows a Pareto law

of the form of (2), with the tail exponent α^{ap} arbitrarily set. We substitute this *a priori* in (8) which becomes:

$$P_{X|Y}(X = i|Y = j) = \left(\frac{B_p(i, j)}{i^{1+\alpha^{ap}}} \right) / \left(\sum_{l=j}^{\infty} \frac{B_p(l, j)}{l^{1+\alpha^{ap}}} \right), \quad (20)$$

and we denote a_{ij} this conditional probability.

In the *scaling method*, this conditional probability was then approached by a rectangle function centered on the interval $[jN, (j+1)N - 1]$. By construction, this approximation necessarily yields piecewise constant estimates $\hat{\mathbf{f}}$, with constant step widths that singularly disrupt the linear regression of of equation (12). Rather than duplicating the same \hat{f}_i at N different locations of i , we propose to choose for each j a unique value of i , that we call $\langle i \rangle$ thereafter, on which all the conditional probability mass will symbolically concentrate:

$$P_{X|Y}(X = i|Y = j) = \begin{cases} 1, & \text{for } i = \langle i \rangle \\ 0, & \text{otherwise.} \end{cases} \quad (21)$$

It is clear that the specific choice of $\langle i \rangle$ is determinant in our method, and the perhaps most intuitive choice would be to consider the mode of the conditional distribution : $\langle i \rangle = \operatorname{argmax}_i P_{X|Y}(X = i|Y = j)$. However, this latter is not sensitive to the $1/x$ decay of the distribution and does not reveal a pertinent choice for estimating α . Instead, we recommend to set $\langle i \rangle$ to the geometric mean of the sequence $i = j, \dots, \infty$, weighted by the conditional probability a_{ij} :

$$\langle i \rangle = \exp \left(\frac{\sum_{i=j}^{\infty} a_{ij} \ln(i)}{\sum_{i=j}^{\infty} a_{ij}} \right). \quad (22)$$

Although we have no theoretical argument to support this particular choice, an unweighted geometrical mean shows natural as it preserves structures of the type $y = 1/x^\alpha$, since then $\langle y \rangle = 1/\langle x \rangle^\alpha$.

3.2 Wavelet based estimation

In [8] was introduced a non parametric estimator of the bound $Q \in \mathbb{R}$ such that all the moments of order $q < Q$ of a random variable Z are finite. It was also recalled that this bound Q controls the asymptotic decay of $P_Z(|Z| \geq z) \sim z^{-Q}$ as z goes to infinity. As for heavy tail distributions, we have $Q = \alpha$, the estimation of Q provides a consistent estimator of the tail exponent α .

Without detailing the minors, it is a wavelet based estimator that measures the regularity at the origin of the characteristic function $\chi(u) = \mathbb{E}e^{iuZ}$ of the random variable Z . Applied to a series $\{Z_n, n = 0, \dots, N-1\}$ of i.i.d. realizations of a Pareto variable Z (Eq. (11)), the resulting kernel estimator reads:

$$W(s) = \frac{1}{N} \sum_{n=0}^{N-1} \Psi(s \cdot Z_n) \sim s^\alpha, \text{ as } s \rightarrow 0^+, \quad (23)$$

where Ψ stands for the Fourier transform of some properly chosen wavelet. Compared to more standard ones, this estimator presents the advantage to perform reasonably well with small data sets.

Returning to our specific problem, it is proved in [3] that asymptotically, P_X and P_Y share the same asymptotic be-

havior, that is:

$$\lim_{j \rightarrow \infty} \frac{P(Y \geq j)}{P(X \geq Nj)} = 1. \quad (24)$$

Thus, we propose to directly apply the estimator defined in (23) to the random variable Y corresponding to the sampled flow size, and to restrict the analysis to a scale range $s_{\min} \leq s \leq s_{\max}$ that matches the asymptotic behavior of P_Y (or equivalently of P_X) where the flows of largest size prevail.

3.3 Stochastic counting

The method briefly sketched in this section was introduced in [2], and relies on a probabilistic counting of original flow sizes from the effective count of the sampled flows.

Let us consider a traffic segment of length T , as described in section 2.1. This segment is itself chopped into M' smaller packets sequences of length $T' < T$. P designates the total number of packets, and K , the total number of flows, lying within the window T' . Let us redefine the variable g_j as the number of flows from which exactly j packets are sampled over the period T' . Then :

$$\mathbb{E}\{g_j\} = K \mathbb{E} \left\{ \binom{\frac{P}{N}}{j} \left(\frac{X}{P} \right)^j \left(1 - \frac{X}{P} \right)^{\frac{P}{N} - j} \mathbb{1}_{\{X \geq j\}} \right\}. \quad (25)$$

It is demonstrated in [2] that, when P/N is large and X/P is small, the following approximation holds :

$$\mathbb{E}\{g_j\} \sim K \mathbb{E} \left\{ e^{-X/N} \frac{X^j}{N^j j!} \right\}, \quad (26)$$

and corresponds to the Poisson approximation of the binomial in (25). Moreover, if X is distributed according to a Pareto law (11), and for K sufficiently large, then:

$$\mathbb{E}\{g_j\} \sim K \frac{P^j}{j!} \alpha b^\alpha \cdot \int_b^\infty x^{j-\alpha-1} e^{-px} dx. \quad (27)$$

Finally, assuming there exists some threshold j_0 delimiting an asymptotic region where the influence of small flows can be neglected, we end up with a very simple estimator of α :

$$\hat{\alpha} = (j+1) \left(1 - \frac{\mathbb{E}\{g_{j+1}\}}{\mathbb{E}\{g_j\}} \right) - 1, \text{ for } j \geq j_0. \quad (28)$$

In practice, $\mathbb{E}g_j$ is replaced by its empirical estimate $\sum_{n=1}^{M'} g_j^{(n)} / M'$, where $g_j^{(n)}$ comes from the n -th short time window of length T' (and similarly for $\mathbb{E}g_{j+1}$).

4. NUMERICAL ANALYSIS

4.1 Experimental setup

We want to compare the accuracy of the different estimators, for different values of α . Our systematic study relies on synthetic traffic generated with Matlab[®], so that we can flexibly control the tail exponent of the prescribed flow size distributions.

Traffic simulator. The simulator we designed reproduces a packet traffic on an aggregated link. A packets stream is randomly generated with a time-stamp, a flow identifier and a packet size associated to each packet.

To prescribe the flow size distribution, we chose a standard compound model which combines :

- a uniform distribution, for $i \leq 10$ (small flows),

- a Pareto distribution (11) with tail parameter α , for $i > 10$ (large flows).

With this model, in the “most heavy-tail” case (i.e. $\alpha = 1.1$), 99% of the flows have size less than 23 packets.

In addition to the flow size distribution, we fix the following traffic characteristics:

- total trace duration is set to 2 hours,
- maximal bandwidth of the link is set to 1 Gbps with a mean load equal to 0.43 (i.e. the mean throughput is 430 Mbps),
- packet size can be either constant and equal to 64 or 1514 Bytes, or uniformly distributed between these two values,
- flows inter-arrival time follows a gamma distribution,
- packets arrival time follows a Poisson point process.

While they should not influence the flow size distribution directly, these characteristics have been carefully chosen so as to reproduce commonly measured traffic features. More importantly though, we checked that they guarantee a good mixing of the flows, which is a necessary condition to hypothesize probabilistic sampling.

Methodology. We generate traffic traces for 5 different values of α : 1.1, 1.3, 1.5, 1.7 and 1.9. In accordance with section 2.1, each trace is divided into non-overlapping segments of length T , corresponding to supposedly independent traffic realizations. For the window size, we test two different widths, $T = 20$ s and $T = 100$ s, leading to $M = 360$ and to $M = 72$ “independent” time series, respectively. The devised choice of T and of the Pareto law parameters warrants that only a few flows last longer than the observation period T , and spread over several adjacent windows. Still, as the largest sampled flows measured within each segment are liable to come from actual flows whose size exceeds the window length, they are systematically considered as truncated and thus discarded from our analysis.

Finally, we sample the traffic time series at two different sampling rates, corresponding to $N = 10$ and to $N = 100$.

Each tail exponent estimator of section 3 yields a sequence of M estimates $\hat{\alpha}$'s corresponding to the M traffic time-series. Statistical performance of the estimators (bias and precision of estimation) are then associated to the empirical means and the empirical variances of the corresponding sequences.

4.2 Results

MLE via EM algorithm. Figure 1(a) represents the maximum likelihood estimates of P_X , solved with the EM algorithm. Despite an apparently good match between estimates and prescribed distribution, this method calls for some important remarks. The distribution body is poorly estimated as small flows are very likely to be missed by the sampling procedure. Regarding tail estimation, oscillations appear and intensify with N . These are due to an excessive number of iterations, that we were not able to automatically adapt to all our experimental parameters. Indeed, finding a satisfying stopping criterion, operational for all sampling rates N , all durations T and all tail exponents α showed a very

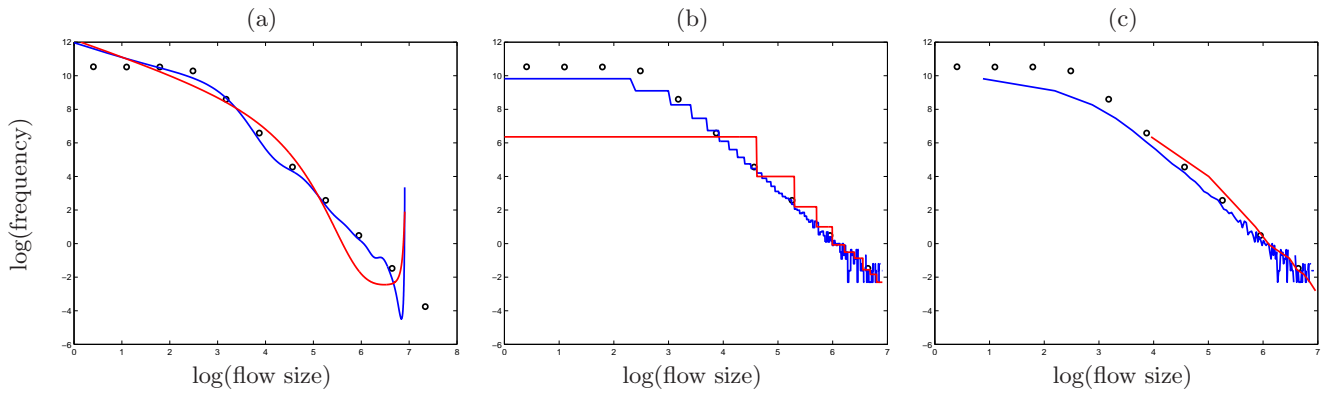


Figure 1: Inferred original distribution for $T = 100$ s and $\alpha = 1.9$. Black dots represent the prescribed distribution. Estimated distributions are obtained: (a) after 50 iterations of EM algorithm, (b) with the scaling method and (c) with the inverse approximation with Pareto a priori method with $\alpha^{ap} = 1$, for $N = 10$ (blue) and $N = 100$ (red).

difficult task. As a result, inferring α from a linear regression of these estimates turns hazardous for it is quite impossible to define a significant regression range where $\log(\hat{f}_i)$ linearly stabilizes. Consequently, the statistical performance, which we deemed not worth reporting here, show aberrant results (diverging mean and variance)...

Moreover, complexity and prohibitive computational cost of EM algorithms lead us to believe that, as it stands, this method is not adapted to our problem of estimating the tail exponent of flow size distributions. However, an EM based approach to directly get the maximum likelihood estimate of α , without estimating the whole original distribution, should perform better. This is under investigation.

Inverse approximation – Uniform a priori (scaling method). Figure 1(b) shows the estimated distribution applying the scaling method of equation (19). As commented in section 3.1, the resulting distributions are piecewise constant on length N intervals (in Fig. 1(b), the steps width decreases with i because of the logarithmic dynamic of axes). As expected, the resolution (i.e. the step width) decreases with N , but still, for large flows, accurate asymptotic decay remains, allowing for estimating the tail parameter α from a linear regression of $\log(\hat{f}_i)$ over some interval $[Nj_{\min}, Nj_{\max}]$ to be determined. The values of j_{\min} and j_{\max} have to be chosen such that the sampled flows we consider are much likely to come from an original flow whose size lies in the tail of the distribution. Let us denote B the supposed value of i that delimits the body from the tail of the original distribution. Then we define j_{\min} and j_{\max} as follows :

- j_{\min} is the smallest value of j verifying $B_p(B, j) \leq \epsilon$ where $j \geq B/N$, and ϵ is a threshold set to 0.01.
- $j_{\max} = j^* - 1$ where j^* is the first j for which $g_j = 0$.

This particular choice of j_{\min} ensures that the probability for a sampled flow of size $j \geq j_{\min}$ to come from an original flow of size $i < B$ is less than ϵ . Although we imposed in our simulations the flow size distribution to be Pareto beyond $i \geq 10$, selecting B obeys a bias-variance trade-off. Indeed, a small value for B entails a small value of j_{\min} , and

may introduce estimation bias due to original flows whose size still lies in the body of the distribution. Conversely, choosing large B leads to large j_{\min} , and reduces the number of points available for regression, thus it leads to an increased variance. We chose $B = 50$. This leads to $j_{\min} = 5$ for $N = 10$ and $j_{\min} = 3$ for $N = 100$.

The choice of j_{\max} leads to discard observed flows of size $j > j^*$. Thus the very end of the tail is not taken into consideration. A possible solution to avoid this problem could be to consider the mean frequencies of observed flows in intervals of geometrically increasing size.

The results obtained with this estimator are summarized in table 1(a) and in figure 2. Whatever the value of α , as the number of observed flows diminishes with T and with the sampling rate, the standard deviation on the estimate $\hat{\alpha}$ increases accordingly. This phenomenon is naturally common to all estimators.

Independently of T and N also, bias and standard deviation increase with α . This can easily be explained by the fact that the larger the α , the rarest the flow sizes with extreme values. This tendency should persist in all methods.

The systematic bias one can observe in all configurations, is probably due to the finite support approximation (18) of the conditional probability. In this framework, for large sampled flows, the corresponding g_j 's are discarded in the estimation of \hat{f}_i . This drawback should be circumvented with the Pareto a priori method.

Inverse approximation – Pareto a priori. We now apply the inverse approximation with Pareto a priori and approximation (21) of section 3.1. Characteristic estimates of P_X are displayed in figure 1(c), with parameter $\alpha^{ap} = 1$. This value of α^{ap} was tuned heuristically and chosen to yield the best overall estimates for all values of α . In the sequel and in all our experiments, this value is kept fixed. A closer look at figure 1(c), shows that for both samplings $N = 10$ and $N = 100$, the tail decay of the flow size distribution is rea-

sonably rendered, even for large flow sizes. In particular, removal of the stepwise structure, turns easier the identification of a regression range $[(i)_{j_{\min}}, (i)_{j_{\max}}]$ to estimate α . The bounds j_{\min} and j_{\max} are determined as for the uniform approximation. Statistical performance of the resulting estimator are summarized in 1(b), and collected with the others in figure 2. While standard deviation remains close to that of the scaling method, the bias has been considerably reduced and now stabilizes with α (cf. table 1(b)).

Wavelet based estimation As explained in section 3.2, we directly apply the estimator of equation (23) to the random variable Y . We then need to define a scale range, image of the tail asymptotic, over which to perform the linear regression of $\log W(s)$ versus $\log(s)$. According to [8], the lower bound is set to $s_{\min} = \log_2(\frac{1}{Y^{**}})$, where Y^{**} is the second largest value of Y (as explained in section 4.1, the largest value Y^* is removed from analysis). The upper bound should correspond to the frontier between the body and the tail of the distribution. We chose $s_{\max} = \log_2(\frac{1}{Y_{\%}})$ where $Y_{\%}$ is a percentile of the distribution of Y . We considered a 85% percentile for $N = 10$ and a 50% percentile for $N = 100$.

The results of this estimator are summarized in table 1(c) and represented in figure 2.

Despite its relatively good behavior, it is worth noticing that the slope obtained from the linear regression of $\log W(s)$ is highly sensitive to the choice of the regression scale bounds. This instability transposes to the bias and the variance of $\hat{\alpha}$ which can significantly increase when the regression range slightly changes.

Nonetheless, the stringent choice imposed to s_{\max} considerably reduces the number of samples of Y that actually participate to the estimation of α , and yet, the wavelet based estimator is capable of a remarkably accuracy.

Stochastic counting Before applying estimation of expression (28), we need to define j_0 such as for $j \geq j_0$, small flows – which do not follow the Pareto law – have no impact on g_j . In [2], $j_0 = 3$ is suggested for $N = 100$. This value is equal to j_{\min} selected in the previous methods, and therefore we chose $j_0 = j_{\min}$ for any N . Notice that larger values for $j > j_0$ are theoretically admissible, however they experimentally lead to a larger variance of estimation, as then, the number of observed flows decreases.

The performance of the stochastic counting estimator are summarized in table 1(d) and illustrated in figure 2. We observe a systematic bias which decreases with α (no matter what the value of N).

Increasing N though, does not drastically degrade the variance of estimation, which certainly is a clear asset of the method. In addition, its extreme simplicity and its remarkably low computational cost prompt its usage for real time measurements and monitoring.

5. CONCLUSION AND FURTHER WORK

Contrary to a natural first intuition, the utilization of the EM algorithm to infer the MLE of the original distribution does not lead to an accurate estimator of the tail param-

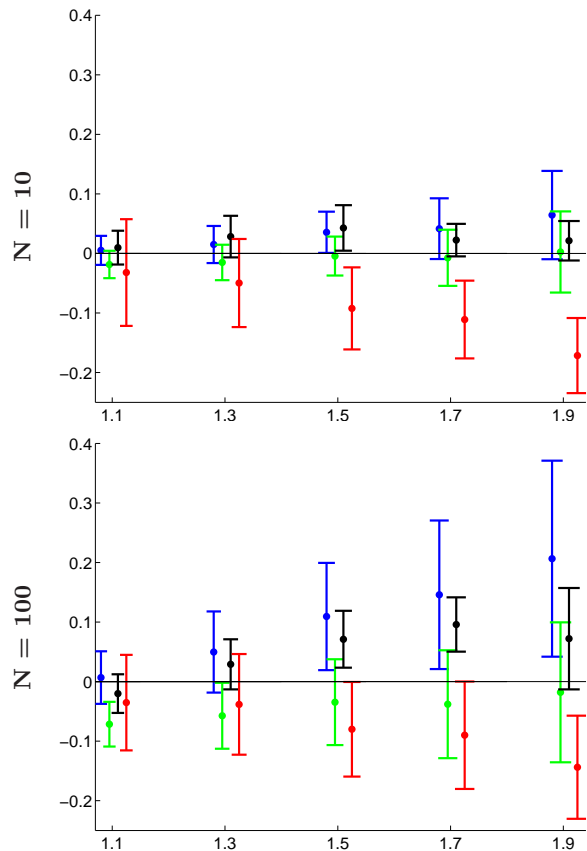


Figure 2: Estimation of α : comparison of the different methods for five values of α : 1.1, 1.3, 1.5, 1.7, 1.9 (a little shift is introduced for clarity purposes). The methods are : scaling method (blue), Pareto *a priori* method (green), wavelet-based method (black), stochastic counting method (red).

ter. However, we are still working on the utilization of the EM algorithm to directly access a ML estimate of the tail parameter.

The method we propose, based on a Pareto *a priori*, proves very promising, as it shows significant improvements compared to the other methods. Despite a lack for theoretical justification for the conditional probability approximation based on a weighted geometric mean, experimental results of the estimator tend to confirm its pertinence.

Although the *stochastic counting method* seems not to be the most accurate one, we need to recall that it definitely relies on the simplest approach. Then, it remains a very interesting method, all the more so as it has little sensitivity to a decreasing of the sampling rate.

Table 1: Results of the estimators for different values of N , T and α . In each case, the estimated value of alpha ($\hat{\alpha}$) and the standard deviation (std) are presented.

(a) : *Scaling method*

T	N	est - α	1.1	1.3	1.5	1.7	1.9
20 s	10	$\hat{\alpha}$	1.08	1.32	1.54	1.76	2.00
		std	0.05	0.05	0.07	0.09	0.12
	100	$\hat{\alpha}$	1.04	1.35	1.64	1.93	2.18
		std	0.08	0.12	0.17	0.25	0.35
100 s	10	$\hat{\alpha}$	1.11	1.32	1.54	1.74	1.96
		std	0.02	0.03	0.03	0.05	0.07
	100	$\hat{\alpha}$	1.11	1.35	1.61	1.85	2.11
		std	0.04	0.07	0.09	0.12	0.16

(b) : *Inverse approximation with Pareto a priori, with $\alpha^{ap} = 1$*

T	N	est - α	1.1	1.3	1.5	1.7	1.9
20 s	10	$\hat{\alpha}$	1.05	1.28	1.49	1.70	1.91
		std	0.04	0.05	0.06	0.08	0.10
	100	$\hat{\alpha}$	0.95	1.20	1.44	1.69	1.91
		std	0.06	0.09	0.12	0.18	0.24
100 s	10	$\hat{\alpha}$	1.08	1.28	1.50	1.69	1.90
		std	0.02	0.03	0.03	0.05	0.07
	100	$\hat{\alpha}$	1.03	1.24	1.47	1.66	1.88
		std	0.04	0.06	0.07	0.09	0.12

(c) : *Wavelet-based method*

T	N	est - α	1.1	1.3	1.5	1.7	1.9
20 s	10	$\hat{\alpha}$	1.11	1.32	1.55	1.71	1.90
		std	0.03	0.04	0.05	0.04	0.04
	100	$\hat{\alpha}$	1.05	1.30	1.54	1.66	1.68
		std	0.04	0.05	0.08	0.11	0.20
100 s	10	$\hat{\alpha}$	1.11	1.33	1.54	1.72	1.92
		std	0.03	0.03	0.04	0.03	0.03
	100	$\hat{\alpha}$	1.08	1.33	1.57	1.80	1.97
		std	0.03	0.04	0.05	0.05	0.09

(d) : *Stochastic counting method*

T	N	est - α	1.1	1.3	1.5	1.7	1.9
20 s	10	$\hat{\alpha}$	1.06	1.25	1.41	1.59	1.73
		std	0.18	0.16	0.14	0.15	0.14
	100	$\hat{\alpha}$	1.06	1.26	1.42	1.61	1.75
		std	0.22	0.19	0.18	0.18	0.19
100 s	10	$\hat{\alpha}$	1.07	1.25	1.41	1.59	1.73
		std	0.09	0.07	0.07	0.07	0.06
	100	$\hat{\alpha}$	1.06	1.26	1.42	1.61	1.76
		std	0.08	0.08	0.08	0.09	0.09

6. REFERENCES

- [1] P. Abry and D. Veitch. Wavelet analysis of long-range dependent traffic. *IEEE Trans. on Info. Theory*, 44(1):2–15, January 1998.
- [2] N. Antunes, Y. Chabchoub, C. Fricker, F. Guillemin, and P. Robert. A new method of inferring flow statistics from sampled data in the internet. preprint, August 2006.
- [3] Y. Chabchoub, C. Fricker, F. Guillemin, and P. Robert. Deterministic versus probabilistic packet sampling in the internet. In *ITC'20*, June 2007.
- [4] Cisco. Netflow. see http://www.cisco.com/en/us/products/ps6601/products_ios_protocol_group_home.html.
- [5] M. E. Crovella and M. S. Taqqu. Estimating the heavy tail index from scaling properties. *Methodology and Computing in Applied Probability*, 1(1):55–79, July 1999.
- [6] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the *em* algorithm. *J. Roy. Stat. Soc., Series B (Method.)*, 1977.
- [7] N. Duffield, C. Lund, and M. Thorup. Estimating flow distributions from sampled flow statistics. In *SIGCOMM*, 2003.
- [8] P. Gonçalves and R. Riedi. Diverging moments and parameter estimation. *J. American Stat. Assoc.*, 100(472):1382–1393, December 2005.
- [9] B. M. Hill. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics*, 3(5):1163–1174, September 1975.
- [10] N. Hohn, D. Veitch, and P. Abry. Cluster processes : a natural language for network traffic. *IEEE Trans. Sig. Proc.*, 51(8):2229–2244, 2003.
- [11] W. Liu, J. Gong, W. Ding, and G. Cheng. A method for estimation of flow length distributions from sampled flow statistics. In *ICOIN*, 2006.
- [12] G. J. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. Wiley's Series in Probability and Statistics. Wiley, New York, 1997.
- [13] J. P. Nolan. Maximum likelihood estimation and diagnostics for stable distributions. In O. Barndorff-Nielsen, T. Mikosch, and S. Resnick, editors, *Lévy Processes : Theory and application*. Birkhäuser, Boston, 2001.
- [14] M. S. Taqqu, W. Willinger, and R. Sherman. Proof of a fundamental result in self-similar traffic modeling. *SIGCOMM Comput. Commun. Rev.*, 27(2):5–23, 1997.