

# Maximum Likelihood Estimation of the Flow Size Distribution Tail Index from Sampled Packet Data

Patrick Loiseau<sup>2</sup>, Paulo Gonçalves<sup>1,2</sup>, Stéphane Girard<sup>1,3</sup>, Florence Forbes<sup>1,3</sup>,  
Pascale Vicat-Blanc Primet<sup>1,2</sup>

<sup>1</sup> INRIA Rhône-Alpes, France

<sup>2</sup> Université de Lyon, ENS Lyon, France

<sup>3</sup> Laboratoire Jean Kuntzmann, Grenoble Universities, France

{Patrick.Loiseau, Paulo.Goncalves, Pascale.Primet}@ens-lyon.fr

{Stephane.Girard, Florence.Forbes}@inrialpes.fr

## ABSTRACT

In the context of network traffic analysis, we address the problem of estimating the tail index of flow (or more generally of any group) size distribution from the observation of a sampled population of packets (individuals). We give an exhaustive bibliography of the existing methods and show the relations between them. The main contribution of this work is then to propose a new method to estimate the tail index from sampled data, based on the resolution of the maximum likelihood problem. To assess the performance of our method, we present a full performance evaluation based on numerical simulations, and also on a real traffic trace corresponding to internet traffic recently acquired.

## Categories and Subject Descriptors

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

## General Terms

Measurement, Theory

## Keywords

Network Monitoring, Traffic Measurement, Heavy-Tailed Distribution, Packet Sampling, Maximum Likelihood Estimation, Expectation-Maximization Algorithm

## 1. MOTIVATIONS

Comprehension and prediction of network traffic is a central preoccupation for the Internet community, because it is an important step in the improvement of the Quality of Service (QoS). In the last decade, Long Range Dependence (LRD) of aggregated traffic time series has been exhibited as a characteristic capable of influencing the QoS in some

particular situations [27, 23], and it has been shown that a possible origin for this characteristic is the heavy-tail nature of the flow size distributions [26, 31]. This makes of LRD parameter and tail index estimation a primary issue. Moreover, as the flow size distribution has been shown to be heavy-tailed in many situations (see *e.g.* [11]), it is of primary interest to estimate the tail index parameter, in order to have a complete knowledge of the system. Indeed, this knowledge can be very useful in problems such as characterization of the network resource usage, dimensioning and scheduling.

Most proposed methods for estimating the LRD parameter and the tail index (see [3] and references therein, [30, 18, 12, 25, 17, 10]) rely on the full traffic observation, implying every packets capture.

However, with very high speed networks, these methods turn very demanding in terms of memory resources and CPU consumption. It then becomes necessary to sample the packet stream, retaining (deterministically or randomly) only a subpart of the aggregated traffic going through the link.

As LRD parameter estimation essentially involves large-scale aggregation levels, its precision does not significantly deteriorates with sub-sampled packets traffic. The situation is more intricate with the estimation of the heavy-tail index of a flow size distribution, from incomplete data. It is this challenging question that we address in the present paper.

The more general problem of inferring the original flow size distribution from sampled data has been extensively studied in the literature [14, 19, 21, 28, 33]. However, and despite its relatively simple formulation, no conclusive estimation procedure or closed form estimate expression has been derived so far. All approaches undertake simplifying assumptions that lead to approximate solutions only. In this paper, we derive the exact solution of the maximum likelihood problem, and provide an analytic expression for the heavy-tail exponent estimate from a sub-sampled packet series. We also develop a corresponding estimation algorithm whose performance is compared to the theoretical Cramér-Rao bound and confronted to the performance of different approximate solutions.

Without giving an exhaustive bibliography on the subject (the interested reader can see the references within the cited papers), let us mention that the impact of sampling has recently been studied in many other contexts: In [24] and [15], the authors tackle the problem of identifying elephant

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.

SIGMETRICS/Performance'09, June 15–19, 2009, Seattle, WA, USA.

Copyright 2009 ACM 978-1-60558-511-6/09/06 ...\$5.00.

flows from sampled data. In [4], the authors consider the problem of ranking the largest flows on a link under packet sampling. In [6] and [20], the authors study the impact of sampling on anomaly detection methods. Furthermore, in [19], the authors point out that *flow sampling* (*i.e.* the sampling decision is made on the flows, and every packets of a picked flow are collected) is a way to bypass the major difficulties inherent to the packet sampling in the flow statistics estimation. As a way to combine the good statistical properties of *flow sampling* with the low computational cost of *packet sampling*, the authors in [32] propose *Dual Sampling*. However, as it remains the most simple sampling procedure, and as it is implemented in many routers [9], we will only consider packet sampling in this work.

The article is organized as follows: Section 2 exposes the mathematical formulation of the problem and sets the definitions and notation. Section 3 lists some of the previously proposed solutions to the problem of tail index estimation from sampled observations, and classes those into two categories: the two-steps solutions where the whole original flow size distribution is inferred first, whereupon the tail index is deduced; and the one-step solutions where the tail index of the original distribution is directly estimated from the sampled flow size distribution. In Section 4, we state and explicitly solve the maximum likelihood formulation of the problem. The proposed MLE solution is given an enlightening interpretation in terms of a geometric solution that we intuitively proposed in a former work. We complete this section with the Cramér-Rao bound derivation. Section 5 shows the results of numerical simulations to assess and rate MLE performance with respect to that of previous approaches; and then compares the performance of the MLE to the performance of the other estimators on the basis of a real internet traffic trace.

## 2. PROBLEM DEFINITION AND NOTATION

In this paper, we analyze traffic time series consisting of a succession of TCP-IP packets, observed over a time duration  $T$ . The strict notion of a flow still fuels an active debate, which falls out of the scope of our present discussion. Therefore, and without loss of generality, we adopt the consensual definition that flows are reconstituted by regrouping packets sharing the same protocol, the same source and destination IP addresses, and the same source and destination ports. The flow size is then defined as the number of packets therein. It is a discrete random variable that we denote  $X$ . The original flow size distribution reads:

$$P_X(X = i) = \phi_i, i \in \mathbb{N}. \quad (1)$$

It can be empirically estimated as the normalized frequency of flows with size  $i$  that are observed over the period  $T$ .

A *heavy-tailed* random variable  $Z$  with tail index  $\alpha$  is drawn from a distribution of the form

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

where  $L$  is a slowly varying function, *i.e.*  $L(tz)/L(z) \rightarrow 1$  as  $z \rightarrow \infty$  for any  $t > 0$ . For discrete random variables, the Zipf law is the paradigm of heavy-tailed distributions:

$$P(X = i; \alpha) = \frac{i^{-(\alpha+1)}}{\zeta(\alpha+1)}; \quad (3)$$

where

$$\zeta(\alpha+1) = \sum_{k=1}^{\infty} k^{-(\alpha+1)} \quad (4)$$

is the Riemann zeta function. The Zipf distribution is the discrete counterpart of the Pareto distribution. In our study, we implicitly consider original flows whose sizes are i.i.d. random variables drawn from a heavy-tailed distribution of the form (3). Our goal is then to estimate the tail exponent  $\alpha$  of this underlying distribution.

Estimating the  $\alpha$  index of heavy-tailed distributions from a set of independent realizations of the random variable  $X$  is a classical problem, broadly treated in the literature for both continuous and discrete variables. For instance, Hill [18] and Nolan [25], proposed maximum likelihood estimators (MLE) for Pareto and alpha-stable laws respectively, whereas Seal in [30], derived the discrete MLE counterpart for a Zipf distribution, recently revisited in [10].

When the variable  $X$  is not directly observable, but only a thinned version of it, things turn more complicated and the proposed solutions not so conclusive. That is precisely the case when memory and CPU consumption issues compel to sample the packets streams. Thus, the observed flow size, defined as the number of sampled packets within a flow, is a random variable on its own, different from the original flow size  $X$ , and following the new distribution:

$$P_Y(Y = j) = \eta_j. \quad (5)$$

The simplest way to practically perform packet sampling consists in sequentially picking one packet every  $K \in \mathbb{N}$ , where  $p = \frac{1}{K}$  is the sampling rate. However, for theoretical development purposes, it is much simpler to consider random sampling, which consists in randomly picking every packet with a probability  $p$ . The two methods have been proved equivalent when a sufficiently large number of intertwined flows is assumed [7]. Thereafter, a probabilistic sampling with sampling rate  $p$  is always assumed. Then, the conditional probability that a sampled flow of size  $j$  comes from an original flow of size  $i \geq j$  is simply governed by a binomial law:

$$P_{Y|X}(Y = j|X = i) = B_p(i, j) = \binom{i}{j} p^j (1-p)^{i-j}, \quad (6)$$

From this, we get the sampled flow size distribution expressed in terms of the original flow size distribution:

$$P_Y(Y = j) = \eta_j = \sum_{i=j}^{\infty} B_p(i, j) \phi_i. \quad (7)$$

Let us notice that equation (7) is properly normalized provided that all sampled flow sizes, including  $j = 0$ , can be observed. As this is generally unrealistic in practice<sup>1</sup>, we are led to introduce the minimal observable size of a sampled flow:  $j_{\min}$  (so that  $P(Y = j) = 0$  if  $j < j_{\min}$ ), and the correct re-normalization of the sampled flows size distribu-

<sup>1</sup>With TCP protocol, the number of missed flows can be estimated as the difference between the total number of expected flows, given by the counting of observed SYN packets divided by  $p$ , and the number of actually observed flows.

tion is:

$$\begin{aligned} P_Y(Y = j) = \eta_j &= \frac{\sum_{i=j}^{\infty} B_p(i, j) P_X(X = i)}{\sum_{j'=j_{\min}}^{\infty} \sum_{l=j'}^{\infty} B_p(l, j') P_X(X = l)} \\ &= \frac{\sum_{i=j}^{\infty} B_p(i, j) \phi_i}{\sum_{j'=j_{\min}}^{\infty} \sum_{l=j'}^{\infty} B_p(l, j') \phi_l}, \end{aligned} \quad (8)$$

for all  $j \geq j_{\min}$ . The normalization factor naturally reduces to 1 when  $j_{\min} = 0$ .

The problem we address in this paper is the following: how can we estimate the tail index  $\alpha$  of the original flow size distribution  $P_X$  from a finite set of sampled flow size observations  $\{y_k\}_{k=1, \dots, m}$ ?

### 3. RELATED WORK

In this Section, we provide a comprehensive survey of the major methods existing to estimate the flow size distribution tail index from sampled data, clarifying their different nature and their common features.

One can tackle this estimation issue from two distinct perspectives: (a) by applying a *standard* tail index estimator to the flow size distribution which has been priorly estimated from the sampled observations (2-steps methods); or (b) by directly inferring the tail index from the sampled observations (without estimating the underlying distribution). These approaches are respectively described in Sections 3.1 and 3.2.

#### 3.1 Two-steps methods

In the presence of fully observed data, tail index estimation has been largely investigated and the 2-steps procedures share the same estimation choice given below (Section 3.1.1). Regarding the inference of the unknown flow size distribution, several possibilities are listed in Section 3.1.2.

##### 3.1.1 Tail exponent estimation

Let us assume that we were able to properly estimate the original flow size distribution beyond a minimal size  $i_{\min}$  from the observed sampled flows. If this latter obeys a power law model of the form (2), estimating the tail index  $\alpha$  is not plagued by the sampling obstacle anymore, and *classical* estimators can directly apply. In this study, we choose to use a Hill-type method [18], whose good performance is notoriously reckoned by the statistical community. The Hill estimator from the inferred original flow size distribution  $\{\hat{\phi}_i\}_{i=i_{\min}, \dots, \infty}$  reads:

$$\hat{\alpha}_{\text{Hill}} = \left( \sum_{i=i_{\min}}^{\infty} \hat{\phi}_i \ln \frac{i}{i_{\min}} \right)^{-1}. \quad (9)$$

A slightly modified Hill estimation procedure has been proposed in [10] as a more appropriate procedure than the *classical* Hill estimator for discrete random variables:

$$\hat{\alpha}_{\text{Discrete}} = \left( \sum_{i=i_{\min}}^{\infty} \hat{\phi}_i \ln \frac{i}{i_{\min} - \frac{1}{2}} \right)^{-1}. \quad (10)$$

Note that these two procedures leads to the same estimation if  $i_{\min}$  is sufficiently large.

Next Section details some proposed methods to estimate the original flow size distribution from a sampled packet sequence.

#### 3.1.2 Flow size distribution inference

*Inverse approximation using an a priori.*

The basic idea behind inverse approximation is to recover the original distribution via the sum and product rules:

$$P_X(X = i) = \sum_{j=0}^{\infty} P_{X|Y}(X = i|Y = j) P_Y(Y = j), \quad (11)$$

where the conditional probability  $P_{X|Y}$  needs first to be estimated. To this end, we use the Bayes formula where we initialize the unknown distribution  $P_X$  to an *a priori* distribution  $P_X^{ap}$ , and get:

$$\begin{aligned} P_{X|Y}(X = i|Y = j) &= \frac{P_{Y|X}(Y = j|X = i) P_X^{ap}(X = i)}{P_Y(Y = j)} \\ &= \frac{B_p(i, j) P_X^{ap}(X = i)}{\sum_{l=j}^{\infty} B_p(l, j) P_X^{ap}(X = l)}. \end{aligned} \quad (12)$$

In this approach, the estimation accuracy of  $P_X$  essentially depends on a relevant choice of  $P_X^{ap}$ .

*Uniform a priori with rectangular approximation of the conditional probability – scaling method.* The simplest *a priori* distribution that we can plug in equation (12) corresponds to a uniform *a priori*:  $P_X^{ap}(X = i) = C$ ,  $\forall i$ , which yields the conditional probability:

$$P_{X|Y}(X = i|Y = j) = \frac{B_p(i, j)}{\sum_{l=j}^{\infty} B_p(l, j)} = p \cdot B_p(i, j). \quad (13)$$

We can further simplify this expression, approximating the binomial function by a simplistic rectangular window:

$$\begin{aligned} P_{X|Y}(X = i|Y = j) &= p, \text{ for } i = \frac{j}{p} + 1, \dots, \frac{j+1}{p} \\ &= 0, \text{ otherwise,} \end{aligned} \quad (14)$$

which finally leads to the original distribution estimate:

$$\hat{\phi}_i = p \eta_{\lfloor ip \rfloor}, \quad \forall i. \quad (15)$$

This expression was originally proposed in [14] and called *scaling estimator*, as it simply corresponds to the sampled flow size distribution, re-scaled by a factor  $p$ .

*Zipf a priori with geometric mean approximation.* As we are interested in heavy-tailed distributions, a more natural choice for the *a priori* distribution  $P_X^{ap}$  is the Zipf law of equation (3) with pre-fixed tail index  $\alpha^{ap}$ . Under this assumption, the conditional probability (12) becomes:

$$P_{X|Y}(X = i|Y = j) = \frac{B_p(i, j) i^{-(\alpha^{ap}+1)}}{\sum_{l=j}^{\infty} B_p(l, j) l^{-(\alpha^{ap}+1)}}. \quad (16)$$

Then, as suggested in [22], instead of distributing the binomial mass  $B_p(i, j)$  uniformly on the interval  $[jp^{-1} + 1, (j+1)p^{-1}]$ , we concentrate all the conditional probability mass corresponding to a given  $j$ , on a unique point denoted  $\langle i \rangle_{(\alpha^{ap})}(j)$ :

$$\begin{aligned} P_{X|Y}(X = i|Y = j) &= 1, \text{ for } i = \langle i \rangle_{(\alpha^{ap})}(j) \\ &= 0, \text{ otherwise.} \end{aligned} \quad (17)$$

It was heuristically proposed in [22] to choose  $\langle i \rangle_{(\alpha^{ap})}(j)$  as the geometric mean of the sequence  $i = j, \dots, \infty$ , weighted by the conditional probability (16):

$$\langle i \rangle_{(\alpha^{ap})}(j) = \exp \left( \frac{\sum_{i=j}^{\infty} \ln(i) B_p(i, j) i^{-(\alpha^{ap}+1)}}{\sum_{l=j}^{\infty} B_p(l, j) l^{-(\alpha^{ap}+1)}} \right). \quad (18)$$

A possible justification for this approximation lies in the fact that geometric means are *naturally adapted* to hyperbolic functions such as the power law decay of heavy-tailed distributions. In Section 4, we will demonstrate that the expression  $\langle i \rangle$  consistently arises when deriving the exact solution of the maximum likelihood estimator.

In contrast with the scaling method of [14], the proposed Zipf *a priori* depends on the tail exponent  $\alpha^{ap}$  that is precisely to be estimated. In [22], the authors suggest an iterative procedure which uses the tail index estimated at step  $k - 1$  to set the *a priori*  $\alpha^{ap}$  of step  $k$ .

### Maximum likelihood estimation of the original distribution.

In [14] also, Duffield et al. tackle the direct estimation of the  $\phi_i$ 's in equation (1), solving the maximum likelihood formulation of the problem. However, as this approach is highly sensible to the variance of the observations (i.e. to the sampled flow size frequencies), they show that maximizing the likelihood function yields negative frequencies. Therefore, imposing a positive constraint on the  $\phi_i$ 's, they resorted to an iterative Expectation-Maximization (EM) algorithm [13] whose output converges towards the ML solution. Notwithstanding its persuasive interest, the difficulty to define a relevant criterion to stop the EM iterations, induces oscillating phenomena that drastically distort the tail decay of the inferred distribution, and lead to highly biased tail index estimates. For this reason, we discard this method from our numerical evaluation in Section 5.

This approach from Duffield et al. has recently been improved in two different ways: In [28], the authors propose a method taking advantage of the protocol informations contained in a packet (mainly the TCP sequence numbers) to improve the accuracy of the estimation. In [33], the authors use a mixture model for the original flow size distribution, separating small and large flows, and jointly estimate the mixture coefficient and the distribution. However, as we consider here a rigorous Zipf law for the flow size distribution, and we assume no knowledge about the protocol informations, these methods, despite their natural interest fall out of the scope of this paper.

### Expansion of the probability generating function.

In [19], the authors propose an original spectral approach to recover the probability densities of the original flows size from the thinned packets distributions. The central idea relies on the one-to-one correspondence between the density function and the probability generating function of a random variable. Then, due to the analyticity properties of this latter, theoretical results from complex analysis can directly apply to infer the original flows size distribution from the probability generating function of the sampled packets time series. Two distinct methods are devised: The first one constructs a power series expansion of the thinned probability generating function about the origin, and finds its non-trivial *analytic continuation* to the entire analytic domain of the original flows size random variable. The second one is based on the *Cauchy integral formula*, whose evaluation on any closed contour including the origin can lead to the desired  $\phi_i$ 's.

According to the author's own evaluation, both methods perform fairly well only for values of the sampling rate  $p > 0.5$ . As we are interested in a sparser packet thinning

(typically  $p < 0.1$ ), we will not proceed with this approach. This limitation notwithstanding, its elegant theoretical formulation is deemed to be stressed here.

## 3.2 Direct tail index estimation

Although more straightforward, methods which do not imply prior estimation of the underlying heavy-tailed flow size distribution (or characteristic function), but directly deduce the tail index  $\alpha$  from the sampled packet series have received much less attention. To our knowledge, the only existing approach of this kind was proposed in [8].

The method relies on stochastic counting. Under the same sampling conditions, let  $W_k$  be a random variable defined as the number of sampled flows observed  $k$  times during a given observation period  $\Delta$ . Using a Poisson approximation, and assuming (i) that the total number of packets is much larger than  $p^{-1}$ , and (ii) that the number of flows is large enough, the authors in [8] analytically prove that a relation between  $E\{W_k\}$  and  $E\{W_{k+1}\}$  holds, leading to the following estimate of  $\alpha$ :

$$\hat{\alpha}_k = (k + 1) \left( 1 - \frac{E\{W_{k+1}\}}{E\{W_k\}} \right) - 1, \text{ for } k \geq k_0. \quad (19)$$

In this expression,  $k_0$  is a threshold defined to ensure that the counting process only comprises flow sizes lying in the power law decay of the distribution, and beyond which  $\hat{\alpha}_k$  converges to the expected value of the tail index.

In practice, the duration  $\Delta$  is divided into non-overlapping shorter time intervals of size  $\Delta/M$ . For each segment  $m = 1, \dots, M$ , the counts  $W_k^{(m)}$  simulates independent realizations of the random variable  $W_k$ . The empirical mean  $M^{-1} \sum_{m=1}^M W_k^{(m)}$  then substitutes the ensemble average  $E\{W_k\}$  in (19).

## 4. MAXIMUM LIKELIHOOD ESTIMATION OF THE TAIL INDEX

We now elaborate on the direct estimation of the tail index  $\alpha$ , from a statistical angle. This Section constitutes the main contribution of our work.

### 4.1 Formulation

In the sequel, we assume that the original flow sizes distribution follows a heavy-tailed Zipf law of the form (3). Under this condition, the sampled flow size distribution (7) becomes:

$$P_Y(Y = j|\alpha) = \frac{1}{\zeta(\alpha + 1)} \sum_{i=j}^{\infty} B_p(i, j) i^{-(\alpha+1)}, \quad (20)$$

where all sampled flow sizes, including  $j = 0$ , are implicitly observed ( $j_{\min} = 0$ ). Then, bypassing the unstable estimation of the underlying original flows size distribution  $P_X$ , we can directly express the log-likelihood function as:

$$\mathcal{L}(\alpha) = -n \ln \zeta(\alpha + 1) + n \sum_{j=0}^{\infty} \eta_j \ln \left( \sum_{i=j}^{\infty} B_p(i, j) i^{-(\alpha+1)} \right), \quad (21)$$

where  $n$  is the number of observed sampled flow sizes. In practice, when not all the sampled flow sizes are observed (i.e.  $j_{\min} > 0$ ), it is the properly normalized form (8) that needs to be adopted and accordingly, equations (20) and (21) take on the form given in Appendix A.

Formally, the maximum likelihood estimate of the tail index  $\alpha$  is solution of the following maximization problem:

$$\widehat{\alpha}_{\text{ML}} = \arg\max_{\alpha} \mathcal{L}(\alpha), \quad (22)$$

and is asymptotically unbiased.

## 4.2 Resolution and interpretation

Differentiating the log-likelihood function (21) and equaling the result to zero readily brings out the quantity  $\langle i \rangle_{(\alpha)}(j)$  that was intuitively introduced in [22] (see Section 3.1.2):

$$\frac{\zeta'(\alpha+1)}{\zeta(\alpha+1)} = - \sum_{j=0}^{\infty} \eta_j \ln \langle i \rangle_{(\alpha)}(j). \quad (23)$$

A closed-form solution to this equation is not available and we adopted a fixed-point resolution technique. We numerically checked that such a fixed point iteration converges toward the maximum likelihood estimate  $\widehat{\alpha}_{\text{ML}}$  within a reasonable number of iterations (see Section 5.1.2). In addition, following the lines of [14], we found that solving this ML problem via an Expectation-Maximization algorithm leads to the exact same iterated procedure, whose solution then coincides with the MLE.

This striking accordance draws an indirect justification for the geometric transformation of relation (18). Indeed, let us suppose that, regardless of the choice  $\alpha^{ap}$ , the transformed variable  $\langle i \rangle_{(\alpha^{ap})}(Y)$  does follow a heavy-tailed Pareto distribution with tail exponent  $\alpha$ . This assertion can be justified as follows. Firstly, given that the random variable  $X$  follows a Pareto distribution with tail exponent  $\alpha$ , the random variable  $Y$  is asymptotically heavy-tailed with the same exponent  $\alpha$ . This is proved in two different ways in [19] and [7]. In [19], the proof is based on the study of the generating functions of  $X$  and  $Y$  with a Tauberian Theorem (see [5, p. 333]). In [7], the proof is based on a Berry-Esséen Theorem (see [16, p. 542]). Then, given that the random variable  $Y$  is asymptotically heavy-tailed with tail exponent  $\alpha$ , equation (44) of appendix B clearly shows that the transformed variable  $\langle i \rangle_{(\alpha^{ap})}(Y)$  is also asymptotically heavy-tailed with the same exponent  $\alpha$ . Experimentally, we checked that this heavy-tailness of  $\langle i \rangle_{(\alpha^{ap})}(Y)$  indeed holds when considering only the observations  $y_k$  larger than some threshold  $j_{\min}$  rarely exceeding 3 or so. Under this hypothesis, maximization of the corresponding maximum likelihood principle leads to the equation

$$\frac{\chi'(\alpha+1, j_{\min})}{\chi(\alpha+1, j_{\min})} = - \sum_{j=j_{\min}}^{\infty} \eta_j \ln \langle i \rangle_{(\alpha^{ap})}(j), \quad (24)$$

where

$$\chi(\alpha+1, j_{\min}) = \sum_{j=j_{\min}}^{\infty} (\langle i \rangle_{(\alpha^{ap})}(j))^{-(\alpha+1)}. \quad (25)$$

In appendix B, we show that for large values of  $j_{\min}$ :

$$\chi(\alpha+1, j_{\min}) \simeq p \cdot \zeta(\alpha+1, \langle i \rangle_{(\alpha)}(j_{\min})), \quad (26)$$

where

$$\zeta(\alpha+1, i_{\min}) = \sum_{i=0}^{\infty} (i + i_{\min})^{-(\alpha+1)}, \quad (27)$$

which clearly implies:

$$\frac{\chi'(\alpha+1, j_{\min})}{\chi(\alpha+1, j_{\min})} = \frac{\zeta'(\alpha+1, \langle i \rangle_{(\alpha)}(j_{\min}))}{\zeta(\alpha+1, \langle i \rangle_{(\alpha)}(j_{\min}))}. \quad (28)$$

This shows that applying directly a maximum likelihood principle to the transformed random variable  $\langle i \rangle_{(\alpha^{ap})}(Y)$  for the tail index estimation with a heavy-tailed Pareto hypothesis (*i.e.* using equation (24), and iterating), leads to the exact same procedure as applying a fixed point method to solve the initial maximum likelihood problem (equation (23)).

In appendix B, we also show that for large values of  $j_{\min}$ :

$$\chi(\alpha+1, j_{\min}) \simeq p \cdot \frac{(\langle i \rangle_{(\alpha)}(j_{\min}))^{-\alpha}}{\alpha}. \quad (29)$$

Plugging equation (29) into equation (24) yields after differentiation to

$$\widehat{\alpha} = \left( \sum_{j=j_{\min}}^{\infty} \eta_j \ln \frac{\langle i \rangle_{(\alpha^{ap})}(j)}{\langle i \rangle_{(\alpha^{ap})}(j_{\min})} \right)^{-1}, \quad (30)$$

which is a classical Hill estimation applied to the random variable  $\langle i \rangle_{(\alpha^{ap})}(Y)$  (see equation (9)). Note that as previously mentioned, a classical Hill estimation (equation (9), [18]) and a modified estimation (equation (10), [10]) leads to the same result because in our case the practical values of  $\langle i \rangle_{(\alpha^{ap})}(j_{\min})$  are sufficiently large.

## 4.3 Properties of the MLE

One additional feature of our approach is that being statistically well-based, theoretical properties of the proposed estimator are accessible and allow us to evaluate its performance in terms of bias and variance.

As previously mentioned, the ML estimator is asymptotically unbiased. In addition, we can derive the theoretical Cramér-Rao bound fixing, for a given sample size, the minimum variance that an estimator can achieve, that is:

$$\text{Var}(\widehat{\alpha}) \geq \frac{1}{\mathcal{I}(\widehat{\alpha})}, \quad (31)$$

with Fisher information:

$$\mathcal{I}(\widehat{\alpha}) = -E \left\{ \frac{\partial^2}{\partial \alpha^2} \mathcal{L}(\alpha) \right\} \Big|_{\alpha=\widehat{\alpha}}. \quad (32)$$

We recall that an unbiased estimator is said efficient if its variance attains the Cramér-Rao bound, and if such estimator exists, then the MLE is necessarily efficient.

In our case, straightforward differentiation of equation (21) gives

$$\begin{aligned} \mathcal{I}(\alpha) = & n \left( \frac{\zeta''(\alpha+1)\zeta(\alpha+1) - \zeta'^2(\alpha+1)}{\zeta^2(\alpha+1)} \right. \\ & + E_Y \left\{ \left( \frac{\sum_{i=1}^{\infty} \ln i B_p(i, Y) i^{-(\alpha+1)}}{\sum_{l=1}^{\infty} B_p(l, Y) l^{-(\alpha+1)}} \right)^2 \right\} \\ & \left. - E_Y \left\{ \frac{\sum_{i=1}^{\infty} (\ln i)^2 B_p(i, Y) i^{-(\alpha+1)}}{\sum_{l=1}^{\infty} B_p(l, Y) l^{-(\alpha+1)}} \right\} \right), \quad (33) \end{aligned}$$

where, again  $j_{\min} = 0$  is implicit. A different value of  $j_{\min}$  essentially modifies the first term of this sum, where the Riemann zeta functions undergo the same change as in equation (8). The equation corresponding to  $j_{\min} > 0$  is given in Appendix A (eq. (36)). In equation (36), both  $n$  and its multiplicative factor depend on  $j_{\min}$ . However, as illustrated by the plots of Figure 1, the decrease of  $n$  with  $j_{\min}$  always dominates the non monotonous variations of the other term. Therefore, the Cramér-Rao bound, as the Fisher information, is essentially controlled by  $n$ , the number of observed sampled flows.

Note that the minimal index in the summations have been set to 1 in these equations instead of the realization  $y_k$  of the random variable  $Y$ . Since  $B_p(i, j)$  is equal to zero for  $i < j$ , the sum remains unchanged.

Notably, we find that the second term of the rhs of the sum simply reduces to  $E_Y\{\ln^2\langle i \rangle_{(\alpha)}(Y)\}$ , where  $\langle i \rangle_{(\alpha)}$  was heuristically defined in (18).

## 5. RESULTS

### 5.1 Performance evaluation using numerical simulations

#### 5.1.1 Simulation scheme

We numerically evaluate the performances of the maximum likelihood estimator  $\widehat{\alpha}_{\text{ML}}$  derived in the previous Section, and draw up a comparative study with the other estimators itemized in Section 3. Our study relies on synthetic traffic generated under Matlab<sup>®</sup>. This allows us to flexibly adjust the different influencing parameters such as the tail exponent of the prescribed flow size distributions.

The traffic simulator we designed reproduces the aggregated traffic generated by  $N_{\text{sources}} = 100$  homogeneous sources on a core link. Each source emits packets' bursts according to a ON/OFF model: two consecutive flows (ON period) are separated by an idle time (OFF period).

The flow size distribution is prescribed to a Zipf law (3), with tail index  $\alpha$ . In addition, we fixed the following traffic characteristics:

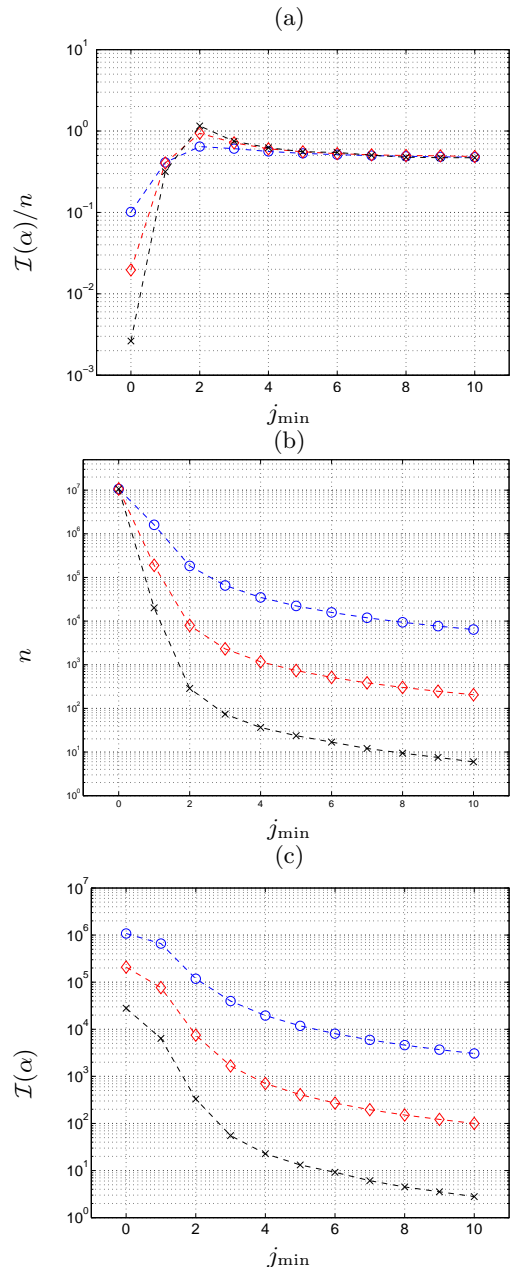
- OFF periods are exponentially distributed. OFF and ON durations have the same mean;
- Each source rate is set to 10 Mb/s, resulting in a mean aggregated traffic of 500 Mb/s;
- The packet size is constant and fixed to 1500 Bytes;
- Experimental data correspond to a stationary packets series generated over a  $T = 300$  s period.

The tail index  $\alpha$  takes on five possible values: 1.1, 1.3, 1.5, 1.7 and 1.9. For each of those, we generated 50 independent time series, that we randomly thinned afterwards, imposing three different values for the sampling rate  $p$ : 1/10, 1/100 and 1/1000. For each combination of  $\alpha$  and  $p$ , bias and variance of all studied estimators are empirically evaluated from the 50 independent realizations.


#### 5.1.2 Comparison between the different estimators

Under the experimental conditions described above, Figure 2 displays the statistical performances of the different tail index estimators: the scaling method, the inverse approximation with a Zipf *a priori*, the stochastic counting and the MLE. Every experimental result was obtained after we manually *optimized* each methods' parameters, so that it reaches the minimum mean square error (MSE).

As a general remark, let us notice that all estimators' performance systematically degrades with  $\alpha$ . We put forward two causes to explain this. In the one hand, the estimation difficulty is inherently increasing with growing scarcity of large flows (as  $\alpha$  gets close to 2). On the other hand, as we consider a fixed length observation window, the number of original flows during this period grows as the tail distribution gets lighter (i.e.  $\alpha$  goes towards 2). These two



**Figure 1: Variations of the Fisher information (eq.(36)) with  $j_{\min}$  for three different values of  $p$ : ( $\circ$ , blue)  $p = \frac{1}{10}$  – ( $\diamond$ , red)  $p = \frac{1}{100}$  – ( $\times$ , black)  $p = \frac{1}{1000}$ . The three graphs shows the variations of: (a) the pre-factor of the Fisher information – (b)  $n$  (the number of sampled flows) – (c) the total Fisher information. The number of original flows is fixed to  $N = 10^7$  and  $\alpha$  is fixed to 1.5. The graphs are obtained with the numerical simulation described in Section 5.1.**



**Figure 2: Comparison of the different tail index estimation methods for five values of  $\alpha$ : 1.1, 1.3, 1.5, 1.7, 1.9 (a small horizontal shift is introduced for clarity purposes): (blue) MLE – (red) Zipf *a priori* method – (green) scaling method – (magenta) stochastic counting method. The error bars are centered on the bias and their length correspond to one standard deviation from the mean.**

competing phenomena should somehow balance. However, as we mentioned it,  $j_{\min} > 0$  is an important tuning parameter in all the methods, that forces to discard the small sampled flows. Given that its effect is more penalizing for large values of  $\alpha$  than for small ones, the number of observed flows effectively drops when  $\alpha$  tends to 2, hence an increasing estimation variance.

**Scaling method.** This estimator shows a systematic bias which increases both with the tail index  $\alpha$  and as the sampling rate  $p$  gets smaller. As reported in [22], this poor performances certainly come from the crude uniform *a priori* density choice and from the finite support boxcar approximation. In practice then, the scaling method cannot be reliably used when  $p \leq 1/10$ .

**Inverse approximation with Zipf *a priori*.** We had to select a sensible criterion to stop this iterative procedure: Convergence is supposedly attained when the difference between two consecutive estimates of  $\alpha$  becomes smaller than 0.005. In practice this leads, in most cases, to a number of iterations between 10 to 100 iterations.

As expected, compared to the scaling method, a more appropriate choice of the *a priori* law, along with a more adapted approximation of the binomial mass, sensibly reduces the estimation bias and variance. For  $p \geq 1/100$ , the bias even stabilizes with  $\alpha$ . Still, for  $p \leq 1/1000$ , the influence of  $j_{\min} > 1$  becomes too penalizing and the estimates turn rapidly unreliable for large values of  $\alpha$ .

**Stochastic counting method.** Practical relevance of equation (19) depends on a correct choice of the threshold  $k_0$ . As for the other methods, this parameter was systematically tuned so as to minimize the MSE for each combination of the pair  $\alpha$  and  $p$ . The observation period  $\Delta$  is set to  $\Delta = 5$  s.

Compared to the previous approach, both bias and variance go up for almost all configurations. Yet, they remarkably remain steady as  $p$  goes from 1/10 to 1/100. Even more, the variance remains roughly constant with  $\alpha$  and almost unchanged when the sampling rate falls to 1/1000. Its

relatively poor performances notwithstanding, this striking stability is a valuable asset that prompts the use of stochastic counting with sparse thinning. Moreover, the extreme simplicity of the method allows for a responsive implementation at a very low computational cost.

**MLE.** For the sake of fairness, we discarded from this analysis the particular choice  $j_{\min} = 0$ , as it would involve non observable data, only retrievable from a deeper TCP packet inspection. Then, as the thorough MLE study of the next Section will show, MSE systematically increases with  $j_{\min}$ , and so we keep this index constant and equal to 1 in the following experiments.

It is clear from Figure 2, that MLE outperforms all the other methods. The variance of estimation is not only an order of magnitude below that of the inverse approximation with a Zipf *a priori*, but also remains perfectly acceptable at very loose sampling rates. More precisely, we attain a precision up to the second decimal for  $p = 1/100$ , and up to the first decimal in the worse case corresponding to  $p = 1/1000$  and  $\alpha = 1.9$ . Regarding the bias, the estimates of Figure 2 show no visible deviation from the theoretical values of  $\alpha$ . The next Section reports on a more systematic evaluation of the MLE solution.

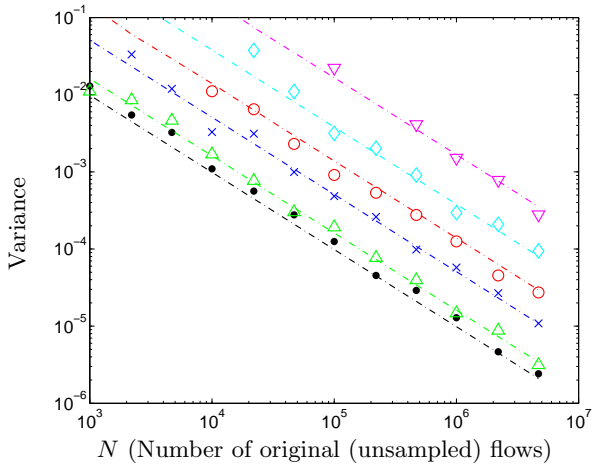
### 5.1.3 MLE performances

Maximum likelihood estimators are, by nature, asymptotically unbiased, as the numerical simulations of Figure 2 seem to confirm. We undertook a complementary series of experiments to precisely evaluate the evolution of the vanishing bias when the number  $N$  of original flows grows to infinity, and for different sampling rates  $p$ . Fixing  $\alpha = 1.5$ , we then observed that beyond a number of original flows  $N \geq 10^6$  the bias stays below 0.003 for all thinning cases ( $p \geq 1/1000$ ) and for  $j_{\min} = 0$ . Under the same conditions though, it raises to 0.005 if, shifting  $j_{\min}$  to 1, we discard the smallest sampled flows and reduce the number of effective observations. Numerical approximations needed to implement the MLE can also partially explain the residual bias.

We derived in Section 4.3 the Cramér-Rao bound associated to the estimation of the tail index  $\alpha$ , from a sequence of  $n$  sampled flows. As  $n$  varies with  $j_{\min}$ , ( $n = N$  if  $j_{\min} = 0$ ,  $n \leq N$  otherwise), it is empirically estimated for  $N$ ,  $p$  and  $j_{\min}$  fixed, and then used to evaluate the theoretical bound of inequality (31). For different values of  $p$  and  $j_{\min}$ , Figure 3 plots as a function of  $N$ , the empirical variances obtained from numerical simulations. Experimental points overlay almost perfectly with the theoretical limits, and prompt to the conclusion that the proposed MLE is efficient, even though we have no rigorous proof of this claim.

Now, analyzing the plots for  $N$  fixed, the variance of estimation naturally increases as the sampling process gets looser (i.e.  $p$  gets smaller). It also increases with  $j_{\min}$ . This is clear from Figure 1 (c), which shows that the Fisher information is a decreasing function of  $j_{\min}$  and thus the Cramér-Rao bound is an increasing function of  $j_{\min}$ . Let us recall here that the dominant effect in the deterioration of the variance when  $j_{\min}$  increases is the consequent decrease of the number of observed flows (see Section 4.3 and Figure 1).

Compared to the variance amplitude, the bias of estimation is clearly negligible. Then, regarding the mean square error (defined as the sum of squared bias and variance), it is primarily governed by the variance and behaves like



**Figure 3: Variance of the MLE estimator against number of original flows considered ( $\alpha = 1.5$ ) for different values of  $N$  and  $j_{\min}$ : ( $\bullet$ , black)  $p = \frac{1}{10}$ ,  $j_{\min} = 0$  - ( $\Delta$ , green)  $p = \frac{1}{10}$ ,  $j_{\min} = 1$  - ( $\times$ , blue)  $p = \frac{1}{100}$ ,  $j_{\min} = 0$  - ( $\circ$ , red)  $p = \frac{1}{100}$ ,  $j_{\min} = 1$  - ( $\diamond$ , cyan)  $p = \frac{1}{1000}$ ,  $j_{\min} = 0$  - ( $\nabla$ , magenta)  $p = \frac{1}{1000}$ ,  $j_{\min} = 1$ . The dash-dotted plots represents the theoretical lower bound of the variance (Cramér-Rao bound).**

$\mathcal{O}(N^{-1})$ . In particular, experimental results show that for a number of original flows larger than  $10^6$ , the MSE does not exceed  $10^{-4}$  when  $p \geq \frac{1}{100}$ , leading to a two decimal accuracy on the tail index estimate. Obviously, this outstanding precision, albeit very loose sampling, stems from a perfect match between data and the Zipf model. In the more general case though, where distributions are only asymptotically heavy-tailed, we are led to choose a larger value for  $j_{\min}$ . As a consequence, the effective number of sampled observed flows reduces accordingly and the performances of the MLE estimator can notably degrade.

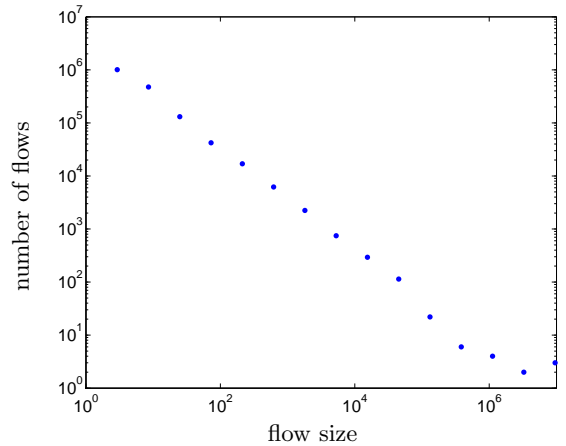
Finally, let us stress that the computational cost of MLE, is an important drawback that can seriously hamper its use with real time constraints.

## 5.2 Confrontation to real traces

To evaluate the robustness of our maximum likelihood estimation of  $\alpha$  in the context of real traffic traces, we estimate  $\alpha$  from an artificially sampled trace of internet traffic, captured on a real network link.

Our trace comes from a university environment: the 1Gb/s output link of the campus site of *École Normale Supérieure de Lyon* is monitored. The bidirectional traffic (input and output traffic) going through this link is optically splitted and captured with an appropriate DAG card [1] (see [29] for a more complete description of the system). The trace we consider here was acquired on March 4, 2007, 16:30 pm to 17:30 pm. It mainly consists in HTTP traffic, achieving a mean throughput of 87.7 Mb/s.

The trace is first processed with *ipsumdump*, a program developed at UCLA [2], able to read the dag format and to summarize TCP/IP dump files into a self-describing binary format. The output file is then processed with IPTools, an original set of tools that is able to reconstruct the flows from the packet trace, based on their definition as groups



**Figure 4: Flow size distribution for the internet traffic in a university context.**

of packets sharing the same source and destination IPs, as well as the same source and destination ports. In addition, a *timeout* can be set such that two consecutive packets of a flow cannot be separated by more than *timeout*. A *timeout* has been chosen and set to 1 s.

The corresponding flow size distribution displayed in Figure 4 in log-log coordinates, clearly shows a heavy tail behavior, and a Zipf model reasonably fits the data, provided that we discard flows of size smaller than some threshold  $i_{\min}$ . In our experiment, setting  $i_{\min}$  to 35, we obtained a maximum likelihood estimate of the tail index equal to  $\hat{\alpha} = 0.9047$ . We will use this value as our benchmark to compare the different sampling approaches.

To evaluate the accuracy of the different methods of estimating the flow size distribution tail index from sampled data, we artificially perform sampling on the previously described traffic trace with three different sampling rates:  $p = 1/10$ ,  $p = 1/100$  and  $p = 1/1000$ . Again, flows of size smaller than some threshold  $j_{\min}$  are discarded for the estimation of  $\alpha$ . This threshold is set as follows:  $j_{\min}$  is the smallest value of  $j$  verifying  $B_p(i_{\min}, j) < \varepsilon$  where  $j \geq pi_{\min}$  and  $\varepsilon$  is a threshold set to 0.05. Thus, every considered flow of size  $j \geq j_{\min}$  is guaranteed to be the sampled version of an original flow of size  $i \geq i_{\min}$  with a probability greater than  $(1 - \varepsilon)$ . Practically, in our case, the values of  $j_{\min}$  are  $j_{\min} = 7$  for  $p = 1/10$  and  $j_{\min} = 2$  for  $p = 1/100$  and  $p = 1/1000$ . For the estimations based on the stochastic counting method, the parameters  $k_0$  and  $\Delta$  are set to  $k_0 = j_{\min}$  and  $\Delta = 30$  s respectively.

Table 1 reports the estimated tail indices obtained with the different methods previously considered. Although all the methods give estimates of  $\alpha$  roughly coherent with the expected value, two methods clearly stand out.

Firstly, the MLE clearly outperforms all the other methods in any situation. This is a direct consequence of the maximum likelihood principle which yields an adapted estimator, and it shows the relevance of the Zipf model utilization. Then, the inverse approximation method using a Zipf *a priori* with a geometric mean approximation yields the closest estimates. This concordance is naturally interpreted by the fact that, as shown in Section 4.2, the Zipf *a pri-*



$p$	MLE	Zipf $a$ $p$ . with geom. mean approx.	Scaling	Stochastic counting
1	0.9047	-	-	-
1/10	0.9196	0.9281	0.9861	0.8413
1/100	0.9216	0.9935	1.2741	0.7050
1/1000	0.9572	1.0042	1.3160	1.0407

**Table 1: Tail index parameter estimation for the internet trace.**

*ori* with geometric mean approximation method shares the same basic estimation quantities ( $\langle i \rangle_{(\alpha^{ap})}(j)$ ) as the MLE, and gives a good approximation of the exact maximum likelihood solution. Table 1 for this method reveals a small bias increasing as the sampling gets looser ( $p$  gets smaller), which is fully consistent with the results of Section 5.1.2 (Figure 2), for small values of  $\alpha$ .

The last two methods give less accurate estimates. In adequacy to the results of Section 5.1.2 (Figure 2) for small values of  $\alpha$ , the scaling method shows a positive bias, increasing as the sampling gets looser. This bias practically makes the use of the scaling method impossible for  $p \leq 1/10$ . The stochastic counting method seems to give reasonable estimates for any value of  $p$ . However, let us stress here that this method relies on an appropriate choice of the observation period  $\Delta$ , which has to be made with respect to a bias-variance tradeoff (a too large value of  $\Delta$  deteriorates the estimation of  $E\{W_k\}$ , thus increasing the variance, whereas a too small value of  $\Delta$  might introduce an important bias). An inappropriate choice might make the estimation unstable and hamper practical use of this method.

Finally, let us stress again that the outstanding performance of the MLE comes at a high computational cost, whereas the stochastic counting method has a very low cost.

## 6. CONCLUSION

In the context of network traffic analysis, a common assumption is that relevant characteristics of traffic time series originate from the heavy-tailed nature of flow size distributions. In this paper we considered Zipf models as a paradigm of heavy-tailed distributions in order to address a number of central questions in a theoretically well-based framework. Zipf distributions depend on a so-called tail index parameter whose estimation from realistic data is not straightforward. More particularly, we addressed the challenging question of estimating this index when the traffic data is not fully observed but only made of a sampled population of individuals (packets). Most existing methods, in this case, are based on approximations induced by simplifying assumptions to make the algorithms tractable.

A first contribution of this paper was then to provide better insights into the nature of these approximations and to highlight the relationships between the different methods. In particular, we provided a theoretical justification for the (binomial) weighted geometric mean, a heavy-tailed data mapping initially proposed in [22] in a ad-hoc manner. A second contribution was then to go beyond approximated solutions by deriving the exact maximum likelihood estimation of the tail index. Our analytic solution was clearly shown to outperform other estimation variants and we reported very good results on simulated data as an illustration.

For higher practical capabilities, further investigation is required to make our maximum likelihood approach more efficient. Although theoretically sub-optimal some approximated methods are much faster in term of computational time and the possibility to find a good compromise between optimality and computational efficiency should be considered. Indeed, a computationally efficient estimation procedure using sampled data could prove very useful in the context of real time adaptive protocols and network mechanisms, where the time constraint makes long computation based on the entire traffic observation impossible. Furthermore, the design of more robust estimators is also an interesting direction of research as maximum likelihood approaches are likely to be very sensitive to data model mismatches.

Finally, even if our ML estimator was developed within the specific context of network monitoring, it can readily apply to other situations of the same kind. For instance, this is the case with social networks where individuals are clustered into groups of heavy-tailed distributed sizes, and while only a cross-section of the population is observed.

## 7. ACKNOWLEDGMENTS

Experiments presented in this paper were carried out using the Grid'5000/Aladdin experimental testbed, an initiative from the French Ministry of Research, INRIA, CNRS and RENATER and other contributing partners (see <https://www.grid5000.fr>). Traces have been captured by *MetroFlux* system based on GtreNET equipment provided by the AIST (Japan). The *MetroFlux* system was partially developed within EC-GIN European project (FP6-2006-IST-045256 grant).

The authors also want to thank the physics laboratory of *École Normale Supérieure de Lyon* for providing us with internet traces as well as with tools for trace analysis (IP-Tools).

Finally, the authors want to thank anonymous reviewers for their valuable comments that helped us improving this paper.

## 8. REFERENCES

- [1] <http://www.endace.com/>.
- [2] Ipsumdump.  
<http://www.cs.ucla.edu/~kohler/ipsumdump/>.
- [3] P. Abry and D. Veitch. Wavelet analysis of long-range dependent traffic. *IEEE Trans. on Info. Theory*, 44(1):2–15, January 1998.
- [4] Chadi Barakat, Gianluca Iannaccone, and Christophe Diot. Ranking flows from sampled traffic. In *CoNEXT '05: Proceedings of the 2005 ACM conference on Emerging network experiment and technology*, pages 188–199, New York, NY, USA, 2005. ACM.
- [5] N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular variations*. Cambridge University Press, Cambridge, UK, 1987.
- [6] Daniela Brauckhoff, Bernhard Tellenbach, Arno Wagner, Martin May, and Anukool Lakhina. Impact of packet sampling on anomaly detection metrics. In *IMC '06: Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, pages 159–164, New York, NY, USA, 2006. ACM.

- [7] Yousra Chabchoub, Christine Fricker, Fabrice Guillemin, and Philippe Robert. Deterministic versus probabilistic packet sampling in the internet. In *ITC'20*, June 2007.
- [8] Yousra Chabchoub, Christine Fricker, Fabrice Guillemin, and Philippe Robert. A robust statistical estimation of internet traffic. preprint, June 2008.
- [9] Cisco. Netflow. see [http://www.cisco.com/en/us/products/ps6601/~products\\_ios\\_protocol\\_group\\_home.html](http://www.cisco.com/en/us/products/ps6601/~products_ios_protocol_group_home.html).
- [10] Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-law distributions in empirical data, June 2007. arXiv:0706.1062v1.
- [11] Mark E. Crovella and Azer Bestavros. Self-similarity in world wide web traffic: evidence and possible causes. *IEEE/ACM Trans. Netw.*, 5(6):835–846, 1997.
- [12] Mark E. Crovella and Murad S. Taqqu. Estimating the heavy tail index from scaling properties. *Methodology and Computing in Applied Probability*, 1(1):55–79, July 1999.
- [13] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the *em* algorithm. *J. Roy. Stat. Soc., Series B (Method.)*, 1977.
- [14] Nick Duffield, Carsten Lund, and Mikkel Thorup. Estimating flow distributions from sampled flow statistics. In *SIGCOMM*, 2003.
- [15] Cristian Estan and George Varghese. New directions in traffic measurement and accounting. In *SIGCOMM*, 2002.
- [16] William Feller. *An introduction to probability theory and its applications*, volume II. John Wiley & Sons, third edition, 1971.
- [17] Paulo Gonçalves and Rudolf Riedi. Diverging moments and parameter estimation. *J. American Stat. Assoc.*, 100(472):1382–1393, December 2005.
- [18] Bruce M. Hill. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics*, 3(5):1163–1174, September 1975.
- [19] Nicolas Hohn and Darryl Veitch. Inverting sampled traffic. *IEEE/ACM Trans. Netw.*, 14(1):68–80, 2006.
- [20] Ryoichi Kawahara, Tatsuya Mori, Noriaki Kamiyama, Shigeaki Harada, and Shoichiro Asano. A study on detecting network anomalies using sampled flow statistics. In *SAINT-W '07: Proceedings of the 2007 International Symposium on Applications and the Internet Workshops*, page 81, Washington, DC, USA, 2007. IEEE Computer Society.
- [21] Weijiang Liu, Jian Gong, Wei Ding, and Guang Cheng. A method for estimation of flow length distributions from sampled flow statistics. In *ICOIN*, 2006.
- [22] Patrick Loiseau, Paulo Gonçalves, and Pascale Primet Vicat-Blanc. A comparative study of different heavy tail index estimators of the flow size from sampled data. In *MetroGrid Workshop, GridNets*, New York, USA, October 2007. ACM Press.
- [23] Michel Mandjes and Nam Kyoo Boots. The shape of the loss curve and the impact of long-range dependence on network performance. Tinbergen Institute Discussion Papers 01-051/4, Tinbergen Institute, May 2001.
- [24] Tatsuya Mori, Masato Uchida, Ryoichi Kawahara, Jianping Pan, and Shigeki Goto. Identifying elephant flows through periodically sampled packets. In *IMC '04: Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, pages 115–120, New York, NY, USA, 2004. ACM.
- [25] J. P. Nolan. Maximum likelihood estimation and diagnostics for stable distributions. In O. Barndorff-Nielsen, T. Mikosh, and S. Resnick, editors, *Lévy Processes : Theory and application*. Birkhäuser, Boston, 2001.
- [26] K. Park, G. Kim, and M. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. In *Int. Conf. on Network Protocols*, page 171, Washington, DC, USA, 1996. IEEE Computer Society.
- [27] Kihong Park, Gitae Kim, and Mark Crovella. On the effect of traffic self-similarity on network performances. In *SPIE International Conference on Performance and Control of Network Systems*, November 1997.
- [28] Bruno Ribeiro, Don Towsley, Tao Ye, and Jean C. Bolot. Fisher information of sampled packets: an application to flow size estimation. In *IMC '06: Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, pages 15–26, New York, NY, USA, 2006. ACM.
- [29] Antoine Scherrer, Nicolas Larrieu, Philippe Owezarski, Pierre Borgnat, and Patrice Abry. Non-gaussian and long memory statistical characterizations for internet traffic with anomalies. *IEEE Trans. Dependable Secur. Comput.*, 4(1):56–70, 2007.
- [30] H. L. Seal. The maximum likelihood fitting of the discrete Pareto law. *Journal of the Institute of Actuaries*, 78:115–121, 1952.
- [31] Murad S. Taqqu, Walter Willinger, and Robert Sherman. Proof of a fundamental result in self-similar traffic modeling. *SIGCOMM Comput. Commun. Rev.*, 27(2):5–23, 1997.
- [32] Paul Tune and Darryl Veitch. Towards optimal sampling for flow size estimation. In *IMC '08: Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, pages 243–256, New York, NY, USA, 2008. ACM.
- [33] Lili Yang and George Michailidis. Sampled based estimation of network traffic flow characteristics. In *INFOCOM*, May 2007.

## APPENDIX

### A. FORMULAS OF THE SAMPLED FLOW SIZE DISTRIBUTION, LOG-LIKELIHOOD AND FISHER INFORMATION FOR AN ARBITRARY VALUE OF $j_{\min}$

With an arbitrary value of the minimal sampled size observed  $j_{\min}$ , the sampled flow size distribution (eq. (20)) becomes when using the proper normalization (eq. (8)):

$$P_Y(Y = j|\alpha) = \frac{\sum_{i=j}^{\infty} B_p(i, j) i^{-(\alpha+1)}}{\sum_{j=j_{\min}}^{\infty} \sum_{l=j}^{\infty} B_p(l, j) l^{-(\alpha+1)}}, \quad (34)$$

The log-likelihood function can then be written:

$$\begin{aligned} \mathcal{L}(\alpha) = & -n \ln \left( \sum_{j=j_{\min}}^{\infty} \sum_{l=j}^{\infty} B_p(l, j) l^{-(\alpha+1)} \right) \\ & + n \sum_{j=j_{\min}}^{\infty} \eta_j \ln \left( \sum_{i=j}^{\infty} B_p(i, j) i^{-(\alpha+1)} \right). \end{aligned} \quad (35)$$

Finally, by differentiation of eq. (35), the Fisher information for an arbitrary  $j_{\min}$  is obtained:

$$\begin{aligned} \mathcal{I}(\alpha) = & n \left( \frac{\left( \sum_{j=j_{\min}}^{\infty} \sum_{i=j}^{\infty} \frac{(\ln i)^2 B_p(i, j)}{i^{(\alpha+1)}} \right) \left( \sum_{j=j_{\min}}^{\infty} \sum_{i=j}^{\infty} \frac{B_p(i, j)}{i^{(\alpha+1)}} \right)}{\left( \sum_{j=j_{\min}}^{\infty} \sum_{l=j}^{\infty} B_p(l, j) l^{-(\alpha+1)} \right)^2} \right. \\ & - \frac{\left( \sum_{j=j_{\min}}^{\infty} \sum_{i=j}^{\infty} \ln i B_p(i, j) i^{-(\alpha+1)} \right)^2}{\left( \sum_{j=j_{\min}}^{\infty} \sum_{l=j}^{\infty} B_p(l, j) l^{-(\alpha+1)} \right)^2} \\ & + E_Y \left\{ \left( \frac{\sum_{i=1}^{\infty} \ln i B_p(i, Y) i^{-(\alpha+1)}}{\sum_{l=1}^{\infty} B_p(l, Y) l^{-(\alpha+1)}} \right)^2 \right\} \\ & \left. - E_Y \left\{ \frac{\sum_{i=1}^{\infty} (\ln i)^2 B_p(i, Y) i^{-(\alpha+1)}}{\sum_{l=1}^{\infty} B_p(l, Y) l^{-(\alpha+1)}} \right\} \right). \end{aligned} \quad (36)$$

### B. ASYMPTOTIC RELATIONS BETWEEN SOME NORMALIZATION SUMS

In this appendix, we show a few relation between the normalization sums of the following distributions: Zipf, Pareto, distribution of  $\langle i \rangle_{(\alpha^{ap})}$  assumed to be algebraically decreasing. The tail index of these distributions is denoted  $\alpha$ .

Let us first stand and recall a few notations:

$$\zeta(\alpha + 1, i_{\min}) = \sum_{i=0}^{\infty} (i + i_{\min})^{-(\alpha+1)} = \sum_{i=i_{\min}}^{\infty} i^{-(\alpha+1)}, \quad (37)$$

$$\psi(\alpha + 1, x_{\min}) = \int_{x_{\min}}^{\infty} x^{-(\alpha+1)} dx = \frac{x_{\min}^{-\alpha}}{\alpha}, \quad (38)$$

$$\chi(\alpha + 1, j_{\min}) = \sum_{j=j_{\min}}^{\infty} (\langle i \rangle_{(\alpha^{ap})}(j))^{-(\alpha+1)}, \quad (39)$$

where

$$\langle i \rangle_{(\alpha^{ap})}(j) = \exp \left( \frac{\sum_{i=j}^{\infty} \ln(i) B_p(i, j) i^{-(\alpha^{ap}+1)}}{\sum_{l=j}^{\infty} B_p(l, j) l^{-(\alpha^{ap}+1)}} \right). \quad (40)$$

Considering that  $\zeta(\alpha + 1, i_{\min})$  is the Riemann sum approximating the integral  $\psi(\alpha + 1, i_{\min})$ , we can easily see that:

$$\zeta(\alpha + 1, i_{\min} - 1) \leq \psi(\alpha + 1, i_{\min}) \leq \zeta(\alpha + 1, i_{\min}). \quad (41)$$

As we also have:

$$\zeta(\alpha + 1, i_{\min}) - \zeta(\alpha + 1, i_{\min} - 1) = \frac{1}{i_{\min} - 1} \xrightarrow{i_{\min} \rightarrow \infty} 0, \quad (42)$$

we conclude that

$$\zeta(\alpha + 1, i_{\min}) \underset{i_{\min} \rightarrow \infty}{\simeq} \psi(\alpha + 1, i_{\min}). \quad (43)$$

We give now an approximation of  $\langle i \rangle_{(\alpha^{ap})}(j)$  when  $j$  is large: First note that the function  $B_p(i, j) i^{-(\alpha^{ap}+1)}$  (considered as a function of  $i$ ) mainly takes non null values in an interval centered on the value  $i = \frac{j}{p}$  (where the function is maximal) and of width a few times  $\sqrt{j p q}$ . If  $j$  is large, the function  $\ln(i)$  is approximately constant in this interval, equal to  $\ln(\frac{j}{p})$ . Then we can rewrite  $\langle i \rangle_{(\alpha^{ap})}(j)$  by putting the  $\ln$  function out of the summation:

$$\begin{aligned} \langle i \rangle_{(\alpha^{ap})}(j) & \simeq \exp \left( \ln \left( \frac{j}{p} \right) \frac{\sum_{i=j}^{\infty} B_p(i, j) i^{-(\alpha^{ap}+1)}}{\sum_{l=j}^{\infty} B_p(l, j) l^{-(\alpha^{ap}+1)}} \right) \\ & \simeq \frac{j}{p}. \end{aligned} \quad (44)$$

Note that this approximation can be used to reduce drastically the computational cost for the computation of  $\langle i \rangle_{(\alpha^{ap})}(j)$  for large values of  $j$ .

Direct plug of equation (44) into equation (39) leads to the approximation

$$\begin{aligned} \chi(\alpha + 1, j_{\min}) & \simeq \left( \frac{1}{p} \right)^{-(\alpha+1)} \sum_{j=j_{\min}}^{\infty} (j)^{-(\alpha+1)} \\ & \simeq \left( \frac{1}{p} \right)^{-(\alpha+1)} \zeta(\alpha + 1, j_{\min}). \end{aligned} \quad (45)$$

Using the approximation of equation (43) then leads to

$$\begin{aligned} \chi(\alpha + 1, j_{\min}) & \simeq \left( \frac{1}{p} \right)^{-(\alpha+1)} \frac{j_{\min}^{-\alpha}}{\alpha} \\ & \simeq \frac{1}{p} \frac{(j_{\min}/p)^{-\alpha}}{\alpha} \\ & \simeq \frac{1}{p} \psi(\alpha + 1, \frac{j_{\min}}{p}). \end{aligned} \quad (46)$$

As  $\langle i \rangle_{(\alpha^{ap})}(j_{\min}) \simeq \frac{j_{\min}}{p}$ , the overall conclusion then reads:

$$\begin{aligned} \chi(\alpha + 1, j_{\min}) & \underset{j_{\min} \rightarrow \infty}{\simeq} \frac{1}{p} \psi(\alpha + 1, \langle i \rangle_{(\alpha^{ap})}(j_{\min})) \quad (47) \\ & \underset{j_{\min} \rightarrow \infty}{\simeq} \frac{1}{p} \zeta(\alpha + 1, \langle i \rangle_{(\alpha^{ap})}(j_{\min})). \quad (48) \end{aligned}$$