A ROBUST STATISTICAL ESTIMATION OF INTERNET TRAFFIC

YOUSRA CHABCHOUB, CHRISTINE FRICKER, FABRICE GUILLEMIN, AND PHILIPPE ROBERT

Abstract. In this paper a new method of estimating flow characteristics in the Internet is developped. For sampled data, a new set of random variables (referred to as observables) is defined, they can be easily evaluated from sampled data. By adopting a convenient mouse/elephant dichotomy also *dependent on the traffic*, it is shown how these variables give a *robust* statistical information of long flows. A mathematical framework is developed to give an estimation of the errors of the method. As an application, it is shown how one can estimate the number of long TCP flows when only sampled traffic is available. The algorithm proposed is tested against experimental data collected from different types of IP traffic.

Contents

1. Introduction

We investigate in this paper how to empirically characterize voluminous flows in the Internet, in particular the tail behavior of their size. It is commonly observed in the technical literature and in real experiments that the total size (in packets or bytes) of such flows has a heavy tailed distribution. In practice, however, this characterization holds only for very large values of the flow size and in order to accurately estimate the tail of the size probability distribution a large number of voluminous flows is necessary. To increase the sample size when empirically estimating probability distribution tails, one is led to increase the length of the observation period. But the counterpart is that the distribution of the flow size cannot be longer described by means of simple probability distributions of the Pareto type for example. This is due to the fact that traffic is generated by a wide variety of applications which give rise to flows with very different characteristics.

Date: December 22, 2008.
Key words and phrases. Flow statistics, Statistical Models, Pareto distribution, Poisson Approximation.

Actually, numerous flow size distributions have been proposed in the literature in order to model voluminous flows as well as their superposition properties; see Feldman *et al.* [11, 12], Duffield *et al.* [7], Crovella and Bestravos [6], Hohn and Veitch [13] or Krunz and Makowski [14].

Robust Statistics. We develop in this paper an alternative method to get robust statistics described by a unique heavy tailed distribution, namely a Pareto distribution: statistics are done during successive time windows of limited length (instead of one time window for the whole trace). The advantage of this method is that, with a careful procedure, a simple statistical characterization is possible and it turns out to be quite robust as shown by our experiments on various sets of traffic traces. The intuitive reason for considering short time periods is that in short time windows, volumes of flows exhibit only one major statistical mode (typically a Pareto behavior). In larger time windows, different modes due to the wide variety of flows in IP traffic necessarily appear. (See Feldman *et al.* [12].) This approach allows us to establish a robust statistical characterization of flows which can be used to infer information from sampled traffic as it will be seen. The counterpart is that the statistics of the *total* size of a flow (obtained when considering the complete traffic trace) cannot be computed directly in this way since the trace is cut into small pieces.

An algorithm is proposed to get the statistical representation of voluminous flows when all the packets of the trace are available. A special care has been devoted to the choice of constants: Length of the observation window, Definition of TCP flows referred to as elephants, ... This is, in our view, one important aspect which is quite often neglected in the literature: the procedure invoked to estimate flow statistics should not depend on some hidden pre-processing of the trace.

Application to Sampled Traffic. The basic motivation for developing such heuristics is to infer flow characteristics from sampled data. This is notably the case for sampling processes such as 1-out-of-k sampling implemented by CISCO's NetFlow [5], which greatly degrades information on flows. What we advocate in this paper is that it is still possible to infer relevant characteristics on flows from sampled data if some characteristic of the flow size can be described in a robust way by means of a simple Pareto distribution. By using the robust statistical representation described above, we propose a method of inferring the number of long flows from sampled traffic. It relies on a new set of random variables, referred to as observables and computed in successive time intervals with fixed length. Specifically, these random variables count the number of flows sampled once, twice or more in the successive observation windows. The properties of these variables can be obtained through robust characteristics, in particular *mean values* of variables instead of *remote quantiles* of the tail distribution, which are much more difficult to accurately estimate. By developing a convenient mathematical setting (Poisson approximation methods), it is moreover possible to show that quantities related to the observables under consideration are close to Poisson random variables with an explicit bound on the error. This Poisson approximation is the key result to estimate the total number of long flows.

The organization of the paper is as follows. A statistical description of long TCP flows is presented in Section 2, this representation is tested against five exhaustive sets of traffic traces: three from the France Telecom (FT) commercial IP network carrying ADSL traffic and two others from Abilene network. In Section 4, variables called observables are introduced, their mathematical properties are analyzed in light of Poisson approximation methods. The results developed in this section are crucial to infer the statistics of an IP traffic from sampled data. Experiments with the five sets of sampled traces used in this paper are presented and discussed in Section 5. Some concluding remarks are presented in Section 6.

2. Statistical Properties of Flows

This section is devoted to a statistical study of the size of flows in a limited time window of duration Δ. The goal of this section is show that some *robust* statistical behavior of voluminous flows can quite generally be exhibited, i.e., for various sets of traffic traces.

2.1. **Assumptions and Experimental Conditions.**

The sets of traces used for testing theoretical results. For the experiments carried out in the following sections, several sets of traces will be considered: Commercial IP traffic, namely ADSL traces from the France Telecom (FT) collect network, and traffic issued from campus networks (Abilene III traces). Their characteristics are given in Table 1.

Table 1. Characteristics of traffic traces considered in experiments.

Name	Nb. IP packets	Nb. TCP Flows	Duration
ADSL Trace A	271 455 718	20 949 331	2 hours
ADSL Trace B Upstream	54 396 226	2 648 193	2 hours
ADSL Trace B Downstream	53 391 874	2 107 379	2 hours
Abilene III Trace A	62 875 146	1 654 410	8 minutes
Abilene III Trace B	47 706 252	1 826 380	8 minutes

The Abilene traces 20040601-193121-1.gz (trace A) and 20040601-194000-0.gz (trace B) can be found at the url http://pma.nlanr.net/Traces/Traces/long/ipls/3/.

Time Windows. Traffic will be observed in successive time windows with length Δ. In practice, the quantity Δ can vary from a few seconds to several minutes depending upon traffic characteristics on the link considered. On the one hand, Δ has to be chosen sufficiently large so that sufficiently many packets arrive in time intervals of duration Δ to derive robust estimations. On the other hand, it should not be too large so that the statistical properties (a Pareto distribution in our case) can be identified, i.e., so that the statistics are unimodal, see Section 2.2 for a detailed discussion.

The ideal value of Δ actually depends on the targeted application. For the design of network elements considering the flow level (e.g., flow aware routers, measurement devices, etc.), it is necessary to estimate the requirements in terms of memory to store the different flow descriptors. In this context, Δ may be of the order of a few seconds. The same order of magnitude is also adapted to anomaly detection, for instance for detecting a sudden increase in the number of flows. For the computation of traffic matrices, Δ can be several minutes long (typically 15 minutes). In our study, the "adequate" values for Δ are of the order of several seconds.

Mice and Elephants. With regard to the analysis of the composition of traffic, in light of earlier studies on IP traffic (see Estan and Varghese [10], Papagiannaki *et al.* [15] or Ben Azzouna *et al.* [1]), two types of flows are identified: small flows (referred to as mice) and voluminous flows (referred to as elephants). In commercial IP traffic, this simple traffic decomposition is justified by the predominance of peer to peer traffic giving rise to either signaling (mice) or file downloads (elephants).

This dichotomy may be more delicate to verify in a different context than the one considered in Ben Azzouna *et al.* [1]. For LAN traffic, for example, there may be very large amounts of data transferred at very high speed. As it will be seen in the various IP traces used in our analysis, the distinction between mice and elephants has to be handled with care. In particular, to reach our goal, it is dependent on the type of traffic considered. The distinction between the constants depending on the trace and "universal" constants is, in our view, a crucial issue. It amounts to precisely stating which constants *depend* on traffic. This aspect is generally (unduly in our opinion) neglected in traffic measurement studies. In particular, the variable Δ and the dichotomy mice/elephants are dependent on the trace, as explained in the next section.

2.2. **Heavy Tails.** The fact that the distribution of the size S of a voluminous TCP flow is heavy tailed is a folk result.Although the heavy tailed property of the size of voluminous flows is commonly admitted, little attention has been paid to identify properly a class of heavy tailed distributions so that the corresponding parameters can be estimated for an arbitrary traffic trace with a significant duration.

One of the reasons for this situation is that the most common heavy tailed distributions $G(x) = \mathbb{P}(S \geq x)$ (e.g., Pareto, i.e., $G(x) = C/x^\alpha$ for $x \geq b$, or Weibull, i.e., $G(x) = \exp(-\nu x^\alpha)$) have a very small number of parameters and consequently a limited of number of possible degrees of freedom for the distribution of the sizes of flows. For this reason, such a distribution can hardly represent the statistics of the total number of packets transmitted by a flow in a trace of arbitrary duration.

As a matter of fact, if a traffic trace is sufficiently long, some non stationary phenomena may arise and the diversity of file sizes may not be captured by one or two parameters. For example, with a Pareto distribution, the function $x \to G(x)$ in a log-log scale should be a straight line. The statistics of the file sizes in the traces used in our experiments are depicted in Figure 1 and 2 for an ADSL traffic trace from the France Telecom backbone IP collect network and for a traffic trace from Abilene network, respectively.

Figure 1 and 2 clearly show that for the two traffic traces considered, the file size exhibits a multimodal behavior: at least *several* straight lines should be necessary to describe properly these distributions. These figures also exhibit the (intuitive) fact that has been noticed in earlier experiments: the longer the trace is, the more marked is the multimodal phenomenon. (See Ben Azzouna *et al.* [2] for a discussion.)

The *key observation* when characterizing a traffic trace is the fact that if the duration Δ of the successive time intervals used for computing traffic parameters is appropriately chosen, then the distribution of the size of the main contributing flows in the time interval can be represented by a Pareto distribution. More precisely, there exist Δ, B_{min}, B_{max} and $a > 0$ such that if S is the number of packets

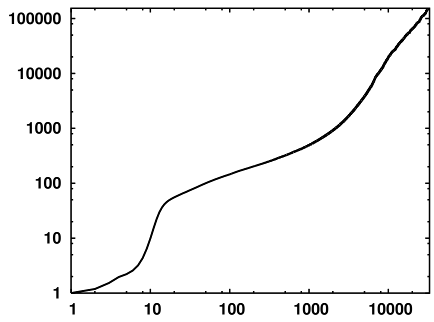

FIGURE 1. Statistics of the number of packets S of a flow for ADSL A (2 hours): the quantity $-\log(\mathbb{P}(S > x))$ as a function of $\log(x)$.

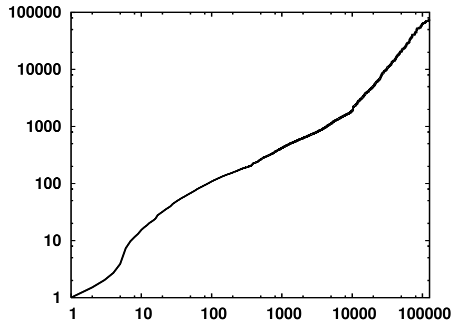

FIGURE 2. Statistics of the number of packets S of a flow for ABILENE A trace (8 minutes): the quantity $-\log(\mathbb{P}(S > x))$ as a function of $\log(x)$.

transmitted by a flow during Δ, then

$$\mathbb{P}(S \geq x \mid S \geq B_{min}) = \left(\frac{B_{min}}{x}\right)^a, \text{ for } B_{min} \leq x \leq B_{max}, \tag{1}$$

and furthermore the proportion of long flows of size greater than B_{max} is less than 5%. The parameter B_{min} is usually referred to as the location parameter and a as the shape parameter.

In other words, if the time interval is sufficiently small then the distribution of the number of packets transmitted by a long flow has one dominant Pareto mode and therefore can be characterized in a robust way. The algorithm used to validate this result is described in Table 2. It is run from the beginning of the trace; in practice a couple of minutes is sufficient to obtain results for the constants Δ, B_{min}, B_{max}. The algorithm is of course valid when the total trace is available for at least an interval of several minutes. In the case of sampled traffic for which this algorithm cannot be used, another method will be proposed in Section 2.

The quantity B_{min} defines the boundary between mice and *elephants* in the trace. A *mouse* is a flow with a number of packets less than B_{min}. An elephant is

TABLE 2. Algorithm for Identifying Δ and the Pareto Distribution.

— Δ is fixed so that at least 1000 flows have more than 20 packets.
— B_{max} is defined as the smallest integer such that less than 5% of the flows have a size greater than B_{max}.
— A Least Square Method is performed to get a linear interpolation in a log-log scale of the distribution of sizes between B_{min} and B_{max}. The constant B_{min} is chosen as the smallest integer such that the distance with the approximating straight line is less than 2.10^{-3}. The slope of the line gives the value of the parameter a.

a flow such that its number of packets during a time interval of length Δ is greater than or equal to B_{min}. By definition of B_{max}, flows whose size is greater than B_{max} represent a small fraction of the elephants.

Experimental results for the ADSL A and Abilene A traffic traces are displayed in Figures 3 and 6, respectively. The same algorithm has been run for the ADSL trace B Upstream and Downstream as well as for the Abilene III B trace. The benefit of the algorithm is that the number of packets in elephants can always be approximated by a unimodal Pareto distribution if the duration of Δ is adequately chosen by using the algorithm given in Table 2. Results are summarized in Table 3.

TABLE 3. Statistics of the elephants for the different traffic traces.

	ADSL A	ADSL B Up	ADSL B Down	Abilene A	Abilene B
Δ (sec)	5	15	15	2	2
B_{min}	20	29	39	89	79
B_{max}	94	154	128	324	312
a	1.85	1.97	1.50	1.30	1.28
%Elephants in $[B_{min}, B_{max}]$	95	95	85	82	83

It turns out that for commercial (ADSL) traffic, the value of B_{min} is close to 20. This value has been used in earlier studies [1, 2] for classifying ADSL traffic. Note that this is not the case for the Abilene traces, which contain significantly bigger elephants. The two types of traffic are intrinsically different: ADSL traffic is mainly composed of peer to peer traffic (with a huge number of small flows and a few file transfers of limited size because of the segmentation of large files into chunks), while Abilene traffic comprises large file transfers issued from campus networks.

It should be noted that the parameters computed in a time window of length Δ do not give a complete description of the distribution of the *total* number of packets of a flow in a flow, since statistics are done over a limited horizon. To obtain information on the total number of packets, it is necessary to "glue" the statistics from successive time windows of length Δ. This turns out to be a very difficult if not impossible task. In some sense, this is the price to pay to have a robust estimation of the statistics of flows. Nevertheless, as it will be seen in the following, in the case of sampled traffic, these parameters cleverly used give a good estimation on the number of active long flows at a given time.

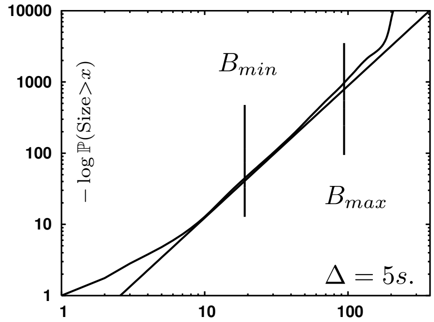

FIGURE 3. Statistics of the flow size (number of packets) in a time interval of length $\Delta = 15$ seconds for the ADSL A trace.

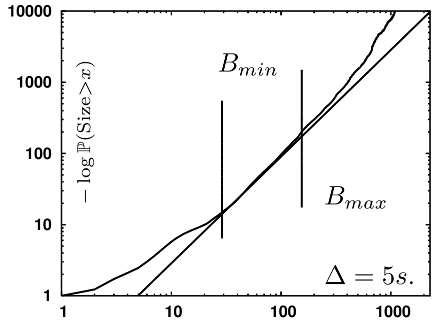

FIGURE 4. Statistics of the flow size (number of packets) in a time interval of length $\Delta = 15$ seconds for the ADSL B Down trace.

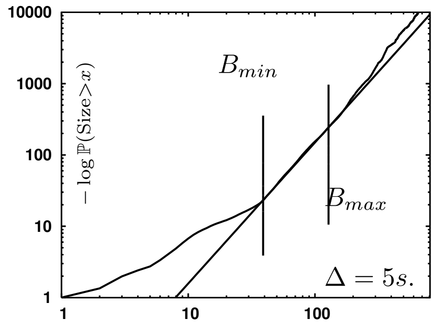

FIGURE 5. Statistics of the flow size (number of packets) in a time interval of length $\Delta = 5$ seconds for the ADSL B Up trace.

3. Sampled Traffic: Assumptions and Definition of Observables

In the previous section, we have designed an algorithm in order to describe voluminous flows by means of a unimodal distribution. Now, we show how to

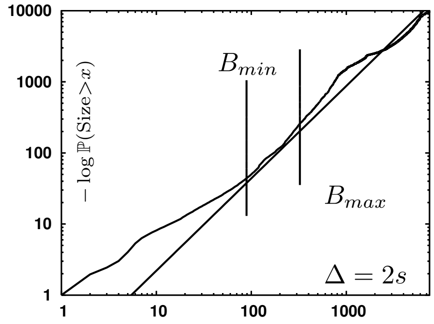

FIGURE 6. Statistics of the flow size (number of packets) in a time interval of length $\Delta = 2$ seconds for the Abilene A trace.

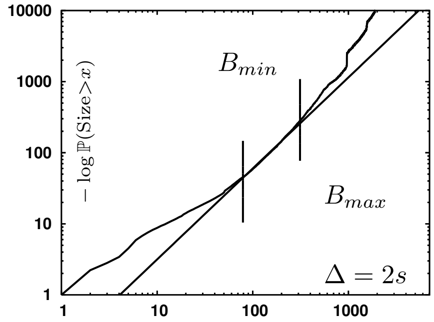

FIGURE 7. Statistics of the flow size (number of packets) in a time interval of length $\Delta = 2$ seconds for the Abilene B trace.

exploit this algorithm in the context of packet sampling in the Internet. This is a crucial issue in order to be able to perform traffic measurements in high speed backbone networks. As a matter of fact, a fundamental problem related to the computation of flow statistics from traffic crossing very high speed transmission links is the fact that, due to the enormous number of packets handled by routers, only a reduced amount of information can be available to the network operator. Packet sampling is in this context an efficient method of reducing the volume of data to analyze when performing measurements in the Internet. One popular technique consists of picking up one packet every other κ_s packets with $\kappa_s = 100$, 500, 1000 in practice. (This sampling scheme is referred to as 1-out-of-κ_s packet sampling in the technical literature.) This method is implemented for instance in CISCO routers, namely NetFlow facility [5] widely deployed in operational networks today. It suffers from different shortcomings well identified in the technical literature, see for instance Estan *et al.* [9].

We describe in this section the different assumptions made on traffic in order to develop an analytical evaluation of our method of inferring flow statistics. Throughout this paper, high speed transmission links (at least 1 Gbit/s) will be considered.

3.1. **Mixing condition.** When observing traffic, packets are assumed to be sufficiently interleaved so that those packets of a same flow are not back-to-back but mixed with packets of other flows. This introduces some randomness in the selection of packets when performing sampling. In particular, when K flows are active in a given time window and if, for $i = 1, \dots, K$, the ith flow is composed of v_i packets during that period, then the probability of selecting a packet of the ith flow is assumed to be equal $v_i/(v_1 + v_2 + \dots + v_K)$. This property will be referred to as *mixing condition* in the following and is formally defined as follows. A variant of this property is, implicitly at least, assumed in the existing literature* (see, e.g. Duffield *et al.* [8]). See also Chabchoub *et al.* [4].

Definition 1 (Mixing Condition)**.** *If K TCP flows are active during a time interval of duration Δ, traffic is said to be mixing if for all i, $1 \leq i \leq K$, the total number $\hat{v}_i$ of packets sampled from the ith flow during that time interval has the same distribution as the analog variable in the following scenario: at each sampling instant a packet of the ith flow is chosen with probability v_i/V where v_i is the number of packets of the ith flow and $V = v_1 + \dots + v_K$.*

This amounts to claim that with regard to sampling, the probability of selecting a packet of a given flow is proportional to the total number of packets of this flow. A variant of this property is generally, implicitly at least, assumed in the existing literature.

One alternative would consist of assuming that the probability of selecting a packet of the ith flow is $1/K$, the inverse of the total number of flows. This assumption, however, does not take into account the respective contributions of the different flows to the total volume and thus may be inaccurate. If all K flows had the same distribution with a small variance, then this assumption would not differ from the mixing condition. Note however that Pareto distributions may have a quite large variance. Hence, this leads us to suppose that the mixing condition holds and that the probability of selecting a packet from flow i is indeed v_i/V.

3.2. **Negligibility assumption.** We consider traffic on very high speed links and it then seems reasonable to assume that no flows contribute a significant proportion to global traffic. In other words, we suppose that the contribution of a given flow to global traffic is negligible. In the following, we go one step further by assuming that in any time window, the number of packets of a given flow is negligible when compared to the total number of packets in the observation window. By using the notation of the previous section, this amounts to assume that for any flow i, the number of packet v_i is much less than V. Furthermore, we even impose that the squared value of v_i is much less than V. We specifically formulate the above assumptions as follows.

Definition 2 (Negligibility condition)**.** *In any window of length Δ, the square of the number of packets of every flow is negligible when compared to the total number of packets V in the observation window. There specifically exists some $0 < \varepsilon \ll 1$ such that for all $i = 1, \dots, K$, $v_i^2/V \leq \varepsilon$.*

The above assumption implies that no flows are dominating when observing traffic on a high speed transmission link. There is thus no bias in the sampling process, which may be caused by the fact that some flows are oversampled because they contribute a significant part of traffic. This assumption is reasonable for commercial ADSL traffic because access links are often the bottlenecks in the network.

For instance, ADSL users may have access rates of a few Mbit/s, which are negligible when compared with backbone links of 1 to 10 Gbit/s. Moreover, the bit rate achievable by an individual flow rarely exceeds a few hundreds of Kbit/s. In the case of transit networks carrying campus traffic, the above assumption may be more questionable since bulk data transfers may take place in Ethernet local area networks and individual flows may achieve bit rates of several Mbit/s.

3.3. **The Observables.** We now introduce the different variables used to infer flow characteristics. These variables are based only upon sampled data; they can be evaluated when analyzing NetFlow records sent by routers of an IP network. For this reason, these variables are referred to as observables. Because of packet sampling, recall that the original characteristics of flows (for instance their duration or their original number of packets) cannot be directly observed.

The observables considered in this paper to infer flow characteristics are the random variables W_j, $j \geq 1$, where W_j is the number of flows sampled j times during a time interval of duration Δ. The averages of the random variables W_j are in fact the key quantities used to infer the characteristics of flows from sampled data.

The random variables W_j, $j \geq 1$ are formally defined as follows: Consider a time interval of length Δ and let K be the total number of long flows present in this time interval. Each flow $i \in \{1, \dots, K\}$ is composed of v_i packets in this time interval. Let denote by $\hat{v}_i$ the number of times that flow i is sampled. The random variable W_j is simply defined by

$$W_j = \mathbb{1}_{\{\hat{v}_1 = j\}} + \mathbb{1}_{\{\hat{v}_2 = j\}} + \cdots + \mathbb{1}_{\{\hat{v}_K = j\}}. \tag{2}$$

In practice, if Δ is not too large, the data structures used to compute the variables W_j are reasonably simple. Moreover, as it will be seen in the following, provided that Δ is appropriately chosen, the statistics of the number of packets transmitted by elephants during successive time windows with duration Δ are quite robust. Consequently, the variables W_j inherit also this property. When the number of long flows is large, the estimation of the asymptotics of their averages from the sampled traffic is easy in practice. Theoretical results on these variables are derived in the next section.

4. Mathematical Properties of the Observables

4.1. **Definitions and Le Cam's inequality.** For $j \geq 0$, the variable W_j defined by Equation (2) is a sum of Bernoulli random variables, namely

$$W_j = \mathbb{1}_{\{\hat{v}_1 = j\}} + \mathbb{1}_{\{\hat{v}_2 = j\}} + \cdots + \mathbb{1}_{\{\hat{v}_K = j\}},$$

where $\hat{v}_i$ is the number of times that the ith flow has been sampled. If these indicator functions were independent, by assuming that K is large, one could use to estimate the distribution of W_j either a Poisson approximation (in a rare event setting) or a central limit theorem (in a law of large numbers context). Since the total number of samples is known, the sum of the random variables $\hat{v}_i$ for $i = 1, \dots, K$ is known and then, the Bernoulli variables defining W_j are *not* independent.

To overcome this problem, we make use of general results on the sum of Bernoulli random variables. Let us consider a sequence (I_i) of Bernoulli random variables, i.e. $I_i \in \{0, 1\}$. The distance in total variation between the distribution of $X =$

$I_1 + \cdots + I_i + \cdots$ and a Poisson distribution with parameter $\delta > 0$ is defined by

$$\begin{aligned}\|\mathbb{P}(X \in \cdot) - \mathbb{P}(Q_\delta \in \cdot)\|_{tv} &\stackrel{\text{def.}}{=} \sup_{A \subset \mathbb{N}} |\mathbb{P}(X \in A) - \mathbb{P}(Q_\delta \in A)| \\ &= \frac{1}{2} \sum_{n \geq 0} \left| \mathbb{P}(X = n) - \frac{\delta^n}{n!} e^{-\delta} \right|.\end{aligned}$$

The Poisson distribution Q_δ with mean δ is such that

$$\mathbb{P}(Q_\delta = n) = \frac{\delta^n}{n!} \exp(-\delta).$$

Note that the total variation distance is a strong distance since it is uniform with respect to all events, i.e., for all subset s A of $\mathbb{N}$,

$$|\mathbb{P}(X \in A) - \mathbb{P}(Q_\delta \in A)| \leq \|\mathbb{P}(X \in \cdot) - \mathbb{P}(Q_\delta \in \cdot)\|_{tv}.$$

The following result (see Barbour *et al.* [3]) gives a tight bound on the total variation distance between the distribution of X and the Poisson distribution with the same expected value when the Bernoulli variables are independent. In spite of the fact that this result is not directly applicable in our case, we shall show in the following how to use it to obtain information on the distributions of the observables W_j.

Theorem 1 (Le Cam's Inequality)**.** *If the random variables (I_i) are independent and if $X = \sum_i I_i$, then*

$$\|\mathbb{P}(X \in \cdot) - \mathbb{P}(Q_{\mathbb{E}(X)} \in \cdot)\|_{tv} \leq \sum_i \mathbb{P}(I_i = 1)^2 \leq \mathbb{E}(X)^2 = \mathbb{E}(X) - \mathrm{Var}(X) \tag{3}$$

If X is a Poisson distribution then $\mathrm{Var}(X) = \mathbb{E}(X)$, the above relation shows that to prove the convergence to a Poisson distribution one has only to prove that the expectation of the random variable is arbitrarily close to its variance.

4.2. **Estimation of the mean value of the observables.** We consider the 1-out-of-κ_s deterministic sampling technique, where one packet is selected every other κ_s packets. In addition, we suppose that traffic on the link observed is sufficiently mixed so that the mixing condition given by Definition 1 holds and that there are no dominating flows in traffic so that the negligibility condition (Definition 2) also pertains.

It is assumed that during a time interval of length Δ, there are K flows composed of at least B_{min} packets, where B_{min} is defined in Section 2. It has been seen that the number of packets in these flows follows a Pareto distribution defined by Relation (1) for some exponent a and parameters B_{min} and B_{max}. Let S be a random variable with this distribution. In addition, let V be the total number of packets in the observation window. Note that V is the sum of the number of packets in elephants and mice. If v_i is the number of packet in the ith elephant, then v_i has the same Pareto distribution as S (i.e., $v_i \stackrel{\text{dist.}}{=} S$) and $V \geq \sum_{i=1}^K v_i$. The difference $V - \sum_{i=1}^K v_i$ is the number of packets of mice.

Proposition 1 (Mean Value of the Observables)**.** *If K elephants are active in a time window of length Δ, the mean number $\mathbb{E}(W_j)$ of flows sampled j times, $j \geq 1$,* *satisfies the relation*

$$\left|\frac{\mathbb{E}(W_j)}{K} - \mathbb{Q}_j\right| \leq p_s \mathbb{E}\left(\frac{S^2}{V}\right), \tag{4}$$

where $\mathbb{Q}$ *is the probability distribution defined by*

$$\mathbb{P}(\mathbb{Q} = j) \stackrel{def}{=} \mathbb{Q}_j = \mathbb{E}\left(\frac{(p_s S)^j}{j!} e^{-p_s S}\right),$$

and $p_s = 1/\kappa_s$ *is the sampling rate.*

Proof. The number of times $\hat{v}_i$ that the ith flow is sampled in the time interval is given by

$$\hat{v}_i = B_1^i + B_2^i + \cdots + B_{p_s V}^i,$$

where, due to the mixing condition, B_ℓ^i is equal to one if the ℓth sampled packet is from the ith flow, which event occurs with probability v_i/V. Note that the total number of sampled packets is $p_s V$.

Conditionally on the values of the set $\mathcal{F} = \{v_1, \ldots, v_K\}$, the variables $(B_\ell^i, \ell \geq 1)$ are independent Bernoulli variables. For $1 \leq i \leq K$, Le Cam's Inequality (3) gives therefore the relation

$$\|\mathbb{P}(\hat{v}_i \in \cdot \mid \mathcal{F}) - Q_{p_s v_i}\|_{tv} \leq p_s \frac{v_i^2}{V}.$$

By integrating with respect to the variables $v_1, \ldots, v_K$, this gives the relation

$$\|\mathbb{P}(\hat{v}_i \in \cdot) - \mathbb{Q}\|_{tv} \leq p_s \mathbb{E}\left(\frac{v_i^2}{V}\right).$$

In particular, for $j \in \mathbb{N}$, $|\mathbb{P}(\hat{v}_i = j) - \mathbb{Q}_j| \leq p_s \mathbb{E}\left(S^2/V\right)$. Since

$$\mathbb{E}(W_j) = \sum_{i=1}^{K} \mathbb{P}(\hat{v}_i = j),$$

by summing on $i = 1, \ldots, K$, one gets

$$|\mathbb{E}(W_j) - K\mathbb{Q}_j| \leq p_s K \mathbb{E}\left(\frac{S^2}{V}\right)$$

and the result follows. □

If the number of packets of flows were constant, then $\mathbb{Q}$ would be a Poisson distribution with parameter $p_s S$, the variable S being in this case a constant. The above inequality shows that at the first order the expected value of W_j is $p_s\mathbb{E}(S)$. The expression of $\mathbb{Q}$, however, indicates that higher order moments of S play a significant role. For example, if the variable S has a significant variance, then the classical rough reduction, which consists of assuming that the size of a sampled elephant is $p_s S$, is no longer valid for estimating the original size of the elephant.

Under the negligibility condition, we deduce that

$$\left|\frac{\mathbb{E}(W_j)}{K} - \mathbb{Q}_j\right| \leq p_s \varepsilon,$$

where ε appears in Definition 2 and is assumed to much less than 1. This implies that Inequality (4) is tight and the quantity $\mathbb{E}(W_j)/K$ can accurately be approximated by the quantity $\mathbb{Q}_j$, when no flows are dominating in traffic.

It is worth noting from Equation (4) that the presence of mice accelerates the convergence of $\mathbb{E}(W_j)/K$ to $\mathbb{Q}_j$. Indeed, for a fixed total number of packets, the more numerous are the mice (i.e., the greater is V), the smaller is the quantity $\mathbb{E}(S^2/V)$ and the closer to $\mathbb{Q}_j$ is the ratio $\mathbb{E}(W_j)/K$. In fact, the presence of mice reduces the probability of sampling an elephant and the number of elephants sampled j times decreases, leading to the setting of the law of small numbers.

We are now ready to state the main result needed for estimating the number K of elephants from sampled data.

Proposition 2 (Asymptotic Mean Values)**.** *Under the same assumptions as those of Proposition 1,*

$$\lim_{K\to+\infty} \frac{\mathbb{E}(W_{j+1})}{\mathbb{E}(W_j)} \sim 1 - \frac{a+1}{j+1} \tag{5}$$

and

$$\lim_{K\to+\infty} \frac{\mathbb{E}(W_j)}{K} \sim a(p_s B_{min})^a \frac{\Gamma(j-a)}{j!}, \tag{6}$$

if $B_{max} >> 1$ *and* $p_s B_{min} << 1$*, where* Γ *is the classical Gamma function defined by*

$$\Gamma(x) = \int_0^{+\infty} u^{x-1} e^{-u}\, du, \quad x > 0.$$

Proof. For $j \geq 1$,

$$\mathbb{Q}_j = \mathbb{E}\left(\frac{(p_s S)^j}{j!} e^{-p_s S}\right) \sim a B_{min}^a \frac{p_s^{a+1}}{j!} \int_{B_{min}}^{+\infty} (p_s u)^{j-a-1} e^{-p_s u}\, du$$

and then

$$\mathbb{Q}_j \sim a B_{min}^a \frac{p_s^a}{j!} \int_{p_s B_{min}}^{+\infty} u^{j-a-1} e^{-u}\, du \sim a(p_s B_{min})^a \frac{\Gamma(j-a)}{j!},$$

since $p_s B_{min} \sim 0$. Therefore, by using the relation $\Gamma(x+1) = x\Gamma(x)$ we obtain the equivalence

$$\frac{\mathbb{Q}_{j+1}}{\mathbb{Q}_j} \sim \frac{j-a}{j+1}.$$

The proposition follows by using the fact that the upper bound of Equation (4) of Proposition 1 goes to 0 by the law of large numbers. □

As it will be seen later in the next section, Relation (5) is used to estimate the exponent a of the Pareto distribution of the number of packets of elephants, the quantities $\mathbb{E}(W_j)$ and $\mathbb{E}(W_{j+1})$ being easily derived from sampled traffic. The quantity K will be estimated from Relation (6). The estimation of the parameter B_{min} from sampled traffic as well as the correct choice of the integer j will be discussed in the next section.

5. Applications

5.1. **Traffic parameter inference algorithm.** In this section, it is assumed that only sampled traffic is available. The methods described in Section 2 to infer the statistical properties of the flows cannot be applied. An ad-hoc algorithm has to be defined in this setting. For the experiments carried out in the present section, the sampling factor $p_s = 1/\kappa_s$ has been taken equal to 1/100. To infer flow characteristics, we have to give the proper definition of the mouse and elephant dichotomy (the parameter B_{min}) and to estimate the coefficient of the corresponding Pareto distribution (the parameter a in Equation (1)).

Relation (5) gives the following equivalence, for $j \geq 1$ sufficiently large so that the impact of mice on $\mathbb{E}(W_j)$ is negligible,

$$a \sim a(j) \stackrel{\text{def.}}{=} (j+1)\left(1 - \frac{\mathbb{E}(W_{j+1})}{\mathbb{E}(W_j)}\right) - 1, \tag{7}$$

and Relation (6) yields an estimate of the number of elephants, i.e. the number of flows with a number of packets greater than or equal to B_{min}; we specifically have

$$K \sim K(j) \stackrel{\text{def.}}{=} \frac{j!\,\mathbb{E}(W_j)}{a(j)(p_s B_{min})^{a(j)}\Gamma(j - a(j))}. \tag{8}$$

These estimations greatly depend on some of the key parameters used to obtain a convenient and robust Pareto representation of the size of the flows, in particular the size of the time window Δ and the lower bound B_{min} for the elephants. The variable Δ is chosen so that

(1) the number of flows sampled twice is sufficiently large in order to obtain a significant number of samples so that the estimation of the mean values of the random variables W_j for $j \geq 2$ is accurate; this requires that Δ should not be too small,
(2) Δ is not too large in order to preserve the unimodal Pareto representation (see Section 2 for a discussion).

To count the average number of flows sampled j times, the parameter j should be chosen as large as possible in order to neglect the impact of mice (for which the Pareto representation does not hold) but not too large so that the statistics are robust to compute the mean value $\mathbb{E}(W_j)$.

In the experimental work reported below, special attention has been paid to the choice of the *universal* constants, i.e., those constants used in the analysis of sampled data, that do not depend on the traffic trace considered. In our opinion, this is a crucial in an accurate inference of traffic parameters from sampled data. These constants are defined in the algorithm given in Table 4.

Table 4. Algorithm used to identify Δ and the Pareto parameter from sampled traffic.

— Choose Δ so that $80 \leq \mathbb{E}[W_2] \leq 100$;
— Choose j so that $|a(j) - a(j+1)|$ computed with Equation (7) is minimized with for all j such that $\mathbb{E}[W_j] \geq 5$.
— B_{min} is the smallest integer so that the probability that a flow of size greater than B_{min} is sampled more than j times is greater than $p_s/10$;

5.2. **Experimental results.** Concerning the estimation of the constants B_{min}, the numerical results obtained by using the algorithm given in Table 4 are presented in Table 5, where the values of the different B_{min} estimated by the algorithm are compared against the values given in Section 2. As it can be observed, the proposed algorithm yields a rather conservative definition of elephants (i.e., flows of size greater than or equal to B_{min}).

TABLE 5. Elephants for the France Telecom ADSL and the Abilene traffic traces.

	ADSL A	ADSL B Up	ADSL B Down	Abilene A	Abilene B
B_{min}	20	29	39	89	79
estimated B_{min}	21	45	45	77	77

The main results are gathered in Table 6 giving the quantities K and a estimated by using Equations (7) and (8) for different values of the parameters j. These values are compared against the experimental values a_{exp} and K_{exp}, referred to as the "real" a and K obtained from the complete traffic traces in Section 2. The accuracy of the estimation of K is generally quite good except for the Abilene A trace where the error is significant although not out of bound. A look at the corresponding figure in Section 2 gives a plausible explanation for this discrepancy: For this trace, the Pareto representation is not very precise.

Finally, it is worth noting from Table 6 that the estimation of the important parameter a describing the statistics of flows is also quite accurate.

TABLE 6. Estimations of the Number of Elephants from Sampled traffic

Trace	Δ	j	$\mathbb{E}(W_j)$	$\mathbb{E}(W_{j+1})$	a_{exp}	$a(j)$	K_{exp}	$K(j)$	Error
ADSL A	5s	3	12.89	3.33	1.85	1.95	943.71	1031.04	9.25%
ADSL B Do	15s	4	9.7	4.75	1.49	1.55	414.90	404.13	2.59%
ADSL B Up	15s	4	7.46	2.97	1.97	2.00	453.01	462.68	2.13%
ABILENE A	1s	5	6.04	3.21	1.38	1.81	217.44	270.79	24.53%
ABILENE B	1s	5	6.1	3.7	1.36	1.51	209.12	197.12	5.74%

6. CONCLUSION

We have developed in this paper one method of characterizing flows in IP traffic by a few parameters and another one of inferring these parameters from sampled data obtained via deterministic 1-out-of-k sampling. For this purpose, we have made some restrictive assumptions, which are in our opinion essential in order to establish an accurate characterization of flows. The basic principle we have adopted consists of describing flows in successive observation windows of limited length, which has to satisfy two contradicting requirements. On the one hand, observation windows shall not to be too large in order to preserve a description of flow statistics as simple as possible, for instance their size by means of a simple Pareto distribution.

On the other hand, a sufficiently large number of packets has to be present in each observation window in order to be able of computing flow characteristics with sufficient accuracy, in particular the tail of the distribution of the flow size. By assuming that large flows (elephants) have a size which is Pareto distributed, we have developed an algorithm to determine the optimal observation window length together with the parameters of the Pareto distribution. The location parameter

B_{min} (see Equation (1)) leads to a natural division of the total flow population into two sets: those flows with at least B_{min} packets, referred to as elephants, and those flows with less than B_{min} packets called mice. This method of characterizing flows has been tested against traffic traces from the France Telecom and Abilene networks carrying completely different types of traffic.

For interpreting sampled data, we have made assumptions on the sampling process. We have specifically supposed that flows are sufficiently interleaved in order to introduce some randomness in the packet selection process (mixing condition) and that there are no dominating flows so that there is no bias with regard to the probability of sampling a flow (negligibility condition). These two assumptions allows us to establish rigorous results for the number of times an elephant is sampled, in particular for the mean values of the random variables W_j, $j \geq 1$.

Of course, when analyzing sampled data, the original flow statistics are not known. In particular, the length of the observation window necessary to characterize the flow size by means of a unique Pareto distribution is unknown. To overcome this problem, we have proposed an algorithm to fix the observation window length and the minimal length of elephants. Then, by choosing the index j sufficiently large so as to neglect the impact of mice, the theoretical results are used to complete the flow parameter inference. This method has been tested against the Abilene and the France Telecom traffic traces and yields satisfactory results.

Once voluminous flows are characterized in time windows of limited length, the next step is to "glue" this information in successive time windows in order to establish a complete characterization of elephants, which can span over several time windows. This point will be addressed in further studies.

References

[1] N. Ben Azzouna, F. Clérot, C. Fricker, and F. Guillemin, *A flow-based approach to modeling ADSL traffic on an IP backbone link*, Annals of Telecommunications **59** (2004), no. 11-12, 1260–1299.

[2] N. Ben Azzouna, F. Guillemin, S. Poisson, P. Robert, C. Fricker, and N. Antunes, *Inverting sampled ADSL traffic*, Proc. ICC 2005 (Seoul, Korea), May 2005.

[3] A. D. Barbour, Lars Holst, and Svante Janson, *Poisson approximation*, The Clarendon Press Oxford University Press, New York, 1992, Oxford Science Publications.

[4] Yousra Chabchoub, Christine Fricker, Fabrice Guillemin, and Philippe Robert, *Deterministic versus probabilistic packet sampling in the Internet*, Proceedings of ITC'20, June 2007.

[5] CISCO, http://www.cisco.com/warp/public/netflow/index.html.

[6] M. Crovella and A. Bestravos, *Self-similarity in world wide web. Evidence and possible causes*, IEEE/ACM Trans. on Networking (1997), 835–846.

[7] Nick Duffield, Carsten Lund, and Mikkel Thorup, *Properties and prediction of flow statistics from sampled packet streams*, IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment (New York, NY, USA), ACM Press, 2002, pp. 159–171.

[8] ———, *Properties and prediction of flow statistics properties and prediction of flow statistics*, ACM SIGCOMM Internet Measurement Workshop, November 2002, pp. 6–8.

[9] C. Estan, K. Keys, D. Moore, and G. Varghese, *Building a better NetFlow*, Proc. ACM Sigcomm'04 (Portland, Oregon, USA), August 30 – September 3 2004.

[10] Cristian Estan and George Varghese, *New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice*, ACM Trans. Comput. Syst. **21** (2003), no. 3, 270–313.

[11] A. Feldmann, A. C. Gilbert, W. Willinger, and T. G. Kurtz, *The changing nature of network traffic: scaling phenomena*, SIGCOMM Comput. Commun. Rev. **28** (1998), no. 2, 5–29.

[12] Anja Feldmann, Jennifer Rexford, and Ramón Cáceres, *Efficient policies for carrying web traffic over flow-switched networks*, IEEE/ACM Trans. Netw. **6** (1998), no. 6, 673–685.

[13] Nicolas Hohn and Darryl Veitch, *Inverting sampled traffic*, IEEE/ACM Trans. Netw. **14** (2006), no. 1, 68–80.

[14] M.M. Krunz and A.M. Makowski, *Modeling video traffic using m/g/∞ input processes: acompromise between markovian and lrd models*, IEEE Journal on Selected Areas in Communications **16** (1998), no. 5, 733–748.

[15] Konstantina Papagiannaki, Nina Taft, Supratik Bhattacharyya, Patrick Thiran, Kavé Salamatian, and Christophe Diot, *A pragmatic definition of elephants in internet backbone traffic*, Internet Measurement Workshop, ACM, 2002, pp. 175–176.

(Y. Chabchoub, C. Fricker, Ph. Robert) INRIA Rocquencourt, RAP project, Domaine de Voluceau, 78153 Le Chesnay, France.
E-mail address: Yousra.Chabchoub@inria.fr

E-mail address: Christine.Fricker@inria.fr

(F. Guillemin) Orange Labs, 2, Avenue Pierre Marzin, F-22300 Lannion
E-mail address: Fabrice.Guillemin@orange-ftgroup.com

E-mail address: Philippe.Robert@inria.fr