Deterministic versus probabilistic packet sampling in the Internet

Yousra Chabchoub, Christine Fricker, Fabrice Guillemin, Philippe Robert

► To cite this version:

Yousra Chabchoub, Christine Fricker, Fabrice Guillemin, Philippe Robert. Deterministic versus probabilistic packet sampling in the Internet. 20th International Teletraffic Congress - ITC 20, Jun 2007, Ottawa / Canada. inria-00130394

HAL Id: inria-00130394

<https://hal.inria.fr/inria-00130394>

Submitted on 12 Feb 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DETERMINISTIC VERSUS PROBABILISTIC PACKET SAMPLING IN THE INTERNET

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

ABSTRACT. Under the assumption that packets are sufficiently interleaved and the sampling rate is small, we show in this paper that those characteristics of flows like the number of packets, volume, etc. obtained through deterministic 1-out-of- k packet sampling is equivalent to random packet sampling with rate $p = 1/k$. In particular, it is shown that under mild assumptions, the tail distribution of the total number of packets in a given flow can be estimated from the distribution of the number of sampled packets. Explicit theoretical bounds are then derived by using technical tools relying on bounds of Poisson approximation (Le Cam's Inequality) and refinements of the central limit theorem (Berry-Essen bounds). Experimental results from an ADSL trace show a good agreement with the theoretical results established in this paper.

1. INTRODUCTION

Packet sampling is an efficient method of reducing the amount of data to retrieve and to analyze in order to study the characteristics of IP traffic (cf. the drafts of IPFIX [11] and PSAMP [12] working groups at the IETF). The simplest approach to packet sampling is certainly the so-called 1-out-of- k sampling technique, which consists of capturing and analyzing one packet every other k packets. This method will be referred to in the following as deterministic sampling, which has been implemented, for instance, in CISCO routers (NetFlow facility [6]) and is widely used in today's operational networks, even if it suffers from several shortcomings identified in [8]. In particular, recovering original flow statistics from sampled data is a difficult task (see [7] for instance). Different solutions have been introduced to overcome these limitations (e.g., the "sample and hold" technique by Estan and Varghese [9], adaptive sampling [5, 8], etc.).

Because deterministic sampling may introduce some synchronization and then some bias in sampled data, which bias is not easy to determine because it depends upon the realization of flows (i.e., the relative position of packets between each other), several studies and IETF drafts [15] recommend probabilistic sampling. In its basic version, random sampling consists of picking up a packet, independently from other packets, with a given probability p . The major advantage is that random sampling provides isolation between flows: the selection of a packet does not depend upon the relative position of flows between each other.

In this paper, it is shown that if packets are sufficiently interleaved (which is definitely the case on a transmission link of a backbone network), then 1-out-of- k deterministic sampling is equivalent to random sampling with $p = 1/k$. More precisely, an explicit estimation of the distance (for the total variation norm) between

the distributions of the numbers of packets in a flow sampled with the two sampling techniques is obtained.

On the basis of this result, bounds on the difference between the distributions of the original flow size and of the sampled flow size rescaled by the sampling factor are established. If the estimation of the size of a flow with the number of sampled packets scaled by the sampling factor is natural and frequently used in the literature, it is not always accurate and can be wrong sometimes. A bound to estimate the accuracy of this estimation is therefore important in practice. Provided that the flow size is sufficiently heavy tailed, it can be shown that the original size of a flow can be indeed estimated from the number of sampled packets.

The different theoretical results obtained in this paper are illustrated on a traffic trace from the France Telecom backbone network carrying ADSL traffic. For this purpose, we introduce a flow decomposition technique based on an ad-hoc mouse/elephant dichotomy. The theoretical results are applied to elephants. Mice appear as background noise in sampled data and their flow size distribution is of less interest, since their volume represents only a small fraction of global traffic. Experimental data show good agreement with theoretical results.

The paper is organized as follows: In Section 2, we describe the traffic analysis methodology. The comparison between deterministic sampling and random sampling is discussed in Section 3 and results on random sampling are then established. These results are compared in Section 4 against experimental results. Concluding remarks are presented in Section 5.

2. TRAFFIC ANALYSIS METHODOLOGY

Let us consider a high speed transmission link carrying Internet traffic and let us divide time into slots of length T . The constant T may range from a few seconds to several tens of minutes (say, from one to two hours).

In this paper, we are interested in the characteristics of TCP traffic since it still represents today 95 % of the total amount of traffic in IP networks, even though the proportion of UDP traffic is growing with the development of streaming applications (VoIP, video, peer-to-peer streaming, etc.). To analyze TCP traffic, we adopt a flow based approach, a flow being defined as the set of those packets with the same source and destination IP addresses together with the same source and destination port numbers (and of course the same protocol type, in this case TCP). In the literature on Internet traffic characterization, it is well known that all flows are not all equivalent: there are flows with many packets, possibly transmitted in bursts, and small flows comprising only a few packets. Many small flows are composed of single SYN segments corresponding to unsuccessful TCP connection establishments attempts.

To simplify the notation, small flows will be referred to in the following as mice and long flows as elephants. This notation corresponds more or less to the elephant/mouse dichotomy introduced by Paxson and Floyd [14], even if clear definitions for mice and elephants do not exist (see the discussion in [13]). To be more specific, we shall use the following definitions:

Definition 1 (Mouse/Elephant). *A mouse is a flow with less than b packets in a time window of length T . An elephant is a flow with at least b packets in a time window of length T .*

FIGURE 1. Ccdf of the number of packets in mice and elephants for $T=5$ seconds and $T = 3200$ seconds ($b = 20$ packets).

We do not claim that the above definitions should be *the* definitions for mice and elephants; they are introduced for convenience to split the flow population into two distinct sets. In particular, they depend upon the length T of the measurement window and the threshold b . In previous studies (see [1] for instance), a threshold $b = 20$ packets yields a neat delineation between mice and elephants when dealing with ADSL traffic even for large observation windows.

To illustrate the above definitions, we consider a traffic trace from the France Telecom IP backbone network carrying ADSL traffic. This traffic trace has been captured on a Gigabit Ethernet link in October 2003 between 9:00 pm and 11:00 pm (this time period corresponding to the peak activity by ADSL customers); the link load was equal to 43.5%. The complementary cumulative distribution function (ccdf) of the number N_{mice} of packets in mice is displayed in Figures 1(a) and 1(b) for $T = 5$ seconds and $T = 3200$ seconds, respectively. We see that for $T = 5$ seconds, the distribution of the random variable N_{mice} can reasonably be approximated by a geometric distribution (i.e., $\mathbb{P}(N_{\text{mice}} > n) \approx r_1^n$). By using a standard Maximum Likelihood Expectation (MLE) procedure, we find $r_1 = 0.75$. For $T = 3200$ seconds, only the tail of the distribution can be approximated by a geometric distribution; experimental results give $\mathbb{P}(N_{\text{mice}} > n) \approx c_2 r_2^n$ for large n , with $c_2 = .1$ and $r_2 = .6$.

The distribution of the number N_{eleph} of packets in elephants is displayed in Figure 1(c) and 1(d) for $T = 5$ and $T = 3200$ seconds, respectively. Now, we see that elephants clearly exhibit a behavior, which is significantly different from that of mice. The random variable N_{eleph} has a slowly decreasing distribution, which

can reasonably be approximated by a Pareto distribution, at least for moderate values of N_{eleph} for $T = 5$ seconds.

We specifically have $\mathbb{P}(N_{\text{eleph}} > n) \approx (b/n)^a$ for $n \geq b = 20$. For $T = 5$ seconds, we find by means of a standard MLE procedure $a = 1.95$. When $T = 3200$ seconds, the distribution of N_{eleph} is more complicated and can be approximated by two Pareto distributions, namely $\mathbb{P}(N_{\text{eleph}} > n) \approx (20/n)^{a_2}$ for $20 \leq n \leq 2000$ with $a_2 = .55$, and $\mathbb{P}(N_{\text{eleph}} > n) \approx (600/n)^{a'_2}$ for $n \geq 2000$ with $a'_2 = 1.8$.

Remark. It turns out that taking only a limited time window for the statistics of the duration of a flow gives a much more robust statistical description of the traffic. Additional works has to be done to recover the full information on the duration of the flows.

In this paper, we are interested in comparing the random variables describing the number of packets in a sampled flow, when deterministic or random sampling is performed.

3. PROPERTIES OF RANDOM AND DETERMINISTIC SAMPLING

3.1. Deterministic sampling. In the case of deterministic sampling, one packet is selected every other $1/p$ (integer) packets, where p is the sampling rate. If packets of flows are back to back, then there is little chance of seeing flows more than once if their number of packets is not significantly larger than the sampling coefficient $1/p$. Fortunately, on a high speed backbone link, the number of simultaneous flows is very large and packets of the different competing flows are highly interleaved. Hence, consecutive packets of a given flow are separated by many packets of other flows. This introduces some randomness in the selection of packets of a given flow.

More precisely, assuming that flows are permanent in a time window of length T , deterministic sampling consists of drawing $\lfloor pM(T) \rfloor$ packets out of the total number $M(T)$ of packets in the time window. If packets are sufficiently interleaved, a sampled packet belongs to a given flow f with probability $N_f/M(T)$ if flow f has originally N_f packets. Under this assumption, the number of sampled packets from flow f is $n_f = B_1^f + B_2^f + \dots + B_{pM(T)}^f$, where the quantities B_j^f are independent Bernoulli random variables equal to one if the j th sampled packet is from flow f . Note that if f and g are distinct flows, then the variables (B_j^f) and (B_j^g) are *not* independent.

The assumption of permanent flows is reasonable, when the observation window length T is small. When T is large, however, flows may be bursty and alternate between on and off periods. This phenomenon has been observed in particular when analyzing elephants in ADSL traffic [1].

3.2. Probabilistic sampling. It is assumed in this section that random sampling is performed: each packet of a given flow f with N_f packets is taken with a probability p and the number of packets in the sampled flow is exactly given by $\tilde{n}_f = \tilde{B}_1^f + \tilde{B}_2^f + \dots + \tilde{B}_{N_f}^f$, where the random variables $(\tilde{B}_i^f)$ are Bernoulli with mean p . The key property of this sampling mode is that it provides isolation between flows. Mathematically, it amounts to the fact that the Bernoulli variables $(\tilde{B}_i^f)$ and $(\tilde{B}_i^g)$ are independent for distinct flows f and g .

The comparison between the two sampling methods is done through the estimation of the *total variation distance* between the distributions of n_f and $\tilde{n}_f$,

$$\|\mathbb{P}(n_f \in \cdot) - \mathbb{P}(\tilde{n}_f \in \cdot)\|_{tv} \stackrel{\text{def.}}{=} \sup_{A \subset \mathbb{N}} |\mathbb{P}(n_f \in A) - \mathbb{P}(\tilde{n}_f \in A)|.$$

Proposition 1 (Probabilistic vs. Deterministic Sampling). *Under the above assumptions, for a flow f with N_f packets with $\mathbb{E}(N_f^2) < +\infty$, the relation*

$$(1) \quad \|\mathbb{P}(n_f \in \cdot) - \mathbb{P}(\tilde{n}_f \in \cdot)\|_{tv} \leq p \frac{\mathbb{E}(N_f^2)}{M(T)} + p^2 \mathbb{E}(N_f)$$

holds. Moreover, as $M(T)$ goes to infinity, the number of sampled packets n_f converges in distribution to $\mathbb{Q}$ defined by

$$\mathbb{Q}(k) = \frac{p^k}{k!} \mathbb{E}(N_f^k e^{-pN_f}).$$

Proof. The proof relies on Le Cam's inequality conditionally on the value of N_f , see Chapter 1 of Barbour [3]. If $\text{Pois}(\lambda)$ denotes the Poisson distribution with parameter λ , then

$$(2) \quad \|\mathbb{P}(n_f \in \cdot \mid N_f) - \text{Pois}(pN_f)\|_{tv} \leq pN_f^2/M(T).$$

By integrating this relation, we obtain $\|\mathbb{P}(n_f \in \cdot) - \mathbb{Q}\|_{tv} \leq p\mathbb{E}(N_f^2)/M(T)$. Similarly for $\tilde{n}_f$, with similar arguments, we have $\|\mathbb{P}(\tilde{n}_f \in \cdot) - \mathbb{Q}\|_{tv} \leq p^2\mathbb{E}(N_f)$. Relation (1) is proved. The convergence in distribution is a direct consequence of Inequality (2). $\square$

Equation (1) implies that when the sampling parameter p is small, the distribution of the number of sampled packets of a given flow is close to the analogue quantity obtained by probabilistic sampling.

Considering that if we deal with an elephant, the number of packets of the flow is quite large, the law of large numbers would suggest the following approximation $\tilde{B}_1^f + \tilde{B}_2^f + \dots + \tilde{B}_{N_f}^f \stackrel{\text{dist.}}{\sim} pN_f$, so that the total number of packets of a flow can be recovered from the number of sampled packets. In spite of the fact that this approximation is quite appealing and natural, it turns out that it has to be handled with care. Indeed, if N_f is geometrically distributed, then it is easily checked that the above approximation is wrong. The fact that N_f is, very likely, heavy tailed helps to establish such an approximation. This is the subject of the rest of the section. The following result is a first step in this direction.

Proposition 2. *If $h_k(x) = x^2/4p^2 \left(\sqrt{1 + 4kp/x^2} - 1 \right)^2$ $x \in \mathbb{R}$, $k > 0$, and the random variables B_i are Bernoulli with mean p , then*

$$\left| \mathbb{P} \left(\sum_{i=1}^{N_f} B_i \geq k \right) - \mathbb{P} \left[N_f \geq h_k \left(\sqrt{p(1-p)} \mathcal{G} \right) \vee k \right] \right| \leq c \mathbb{E} \left(\frac{1}{\sqrt{N_f}} \mathbb{1}_{\{N_f \geq k\}} \right),$$

where $\mathcal{G}$ is a standard Gaussian random variable, for real numbers $a \vee b = \max(a, b)$, and $c = 3(p^2 + (1-p)^2)/\sqrt{p(1-p)}$.

Proof. Let $\sigma^2 = \text{Var}(B) = p(1-p)$, $S_n = B_1 + \dots + B_n$, $\bar{S}_n = S_n/n$ and $\hat{S}_n = \sqrt{n}(\bar{S}_n - np)/\sigma$. By Berry-Essen's theorem [10], for each $n \in \mathbb{N}$ and $k > 0$,

$$\left| \mathbb{P} \left(\hat{S}_n \geq \frac{k - pn}{\sigma\sqrt{n}} \right) - \mathbb{P} \left(\mathcal{G} \geq \frac{k - pn}{\sigma\sqrt{n}} \right) \right| \leq \frac{c}{\sqrt{n}}$$

where $c = 3E((p-B)^3)/\sigma^3 = 3(p^2 + (1-p)^2)/\sqrt{p(1-p)}$. Thus, multiplying by $\mathbb{1}_{\{n \geq k\}}$, using the independence of S_n and N_f and Fubini's theorem, noticing that $\mathbb{P}(\hat{S}_N \geq (k - pN_f)/\sqrt{N_f}) = \mathbb{P}(S_{N_f} \geq k)$ and that, if $S_{N_f} \geq k$ then $N_f \geq k$, we obtain

$$\left| \mathbb{P} \left(\hat{S}_N \geq \frac{k - pN_f}{\sigma\sqrt{N_f}} \right) - \mathbb{P} \left(\mathcal{G} \geq \frac{k - pN_f}{\sigma\sqrt{N_f}}, N_f \geq k \right) \right| \leq c\mathbb{E} \left(\frac{1}{\sqrt{N_f}} \mathbb{1}_{\{N_f \geq k\}} \right).$$

Now, we prove that

$$\begin{aligned} \mathbb{P} \left(\mathcal{G} \geq \frac{k - pN_f}{\sigma\sqrt{N_f}}, N_f \geq k \right) &= \mathbb{P}(pN_f + \sqrt{N_f}\sigma\mathcal{G} \geq k, N_f \geq k) \\ &= \mathbb{P}(N_f \geq f_k(\sigma\mathcal{G}) \vee k). \end{aligned}$$

Indeed, denoting $z = \sqrt{y}$, the equation $pz^2 + zx - k = 0$ has two roots in $\mathbb{R}$, equal to $z_1 = (-x - \sqrt{x^2 + 4pk})/2p < 0$ and $z_2 = (-x + \sqrt{x^2 + 4pk})/2p > 0$. Thus, for every $x \in \mathbb{R}$, $pz^2 + zx - k \geq 0, z \geq 0$ is equivalent to $z \geq z_2$, i.e., $y \geq h_k(x)$. The result then readily follows. $\square$

From the above result, under mild assumptions on the distribution of N_f , the tail distribution of $B_1 + B_2 + \dots + B_{N_f}$ is related to the tail distribution of N_f . In particular, if N_f has a Pareto distribution, we have the following result.

Corollary 1. *If the random variable N_f has a Pareto distribution, i.e. for some $b > 0$ and $a > 1$, $\mathbb{P}(N_f \geq k) = (b/k)^a$, and if the random variables B_i are Bernoulli with mean p , then*

$$\lim_{k \rightarrow +\infty} \frac{\mathbb{P}(B_1 + B_2 + \dots + B_{N_f} \geq k)}{\mathbb{P}(N_f \geq k/p)} = 1.$$

Proof. We have

$$\mathbb{P}(N_f \geq h_k(\sqrt{p(1-p)}\mathcal{G}) \vee l) = \mathbb{E}((b/(h_k(\sqrt{p(1-p)}\mathcal{G}) \vee k))^a) \sim (bp/k)^a,$$

since $h_k(x) = k/p(1 + O(\frac{1}{\sqrt{k}}))$ for large x . $\square$

The above asymptotic results have been established for a random variable N_f , which has a Pareto distribution. But it is straightforwardly checked that similar results hold, when only the tail of N_f is Pareto as for the traffic trace described in Section 2. To conclude the comparison between the original flow size distribution and the rescaled sampled size distribution, let us mention that Berry-Essen bound based on the normal approximation is accurate only around the mean value. To obtain a tighter bound on the tail of the distribution, it is possible to establish the following result (see [4] for details).

Theorem 1. For $\alpha \in (1/2, 1)$, there exist positive constants C_0 and C_1 such that for any $p \in (0, 1)$ and $\ell \geq 1/p$,

$$\left| \frac{\mathbb{P}\left(\sum_{i=1}^{N_f} B_i \geq \ell\right)}{\mathbb{P}(N_f \geq \ell/p)} - 1 \right| \leq \sup_{-C_1 \leq u \leq C_1} 3 \left| \frac{\mathbb{P}\left(N_f \geq \frac{\ell}{p} + u \left(\frac{\ell}{p}\right)^\alpha\right)}{\mathbb{P}(N_f \geq \ell/p)} - 1 \right| + \frac{C_0}{\mathbb{P}(N_f \geq \ell/p)} \exp\left(-\frac{p}{4(1-p)} \ell^{2\alpha-1}\right).$$

From the above result, we see that the quantity $\mathbb{P}\left(\sum_{i=1}^{N_f} B_i \geq \ell\right)$ related to the probability that a sampled flow contains at least ℓ packets is exponentially close for sufficiently large ℓ to $\mathbb{P}(N_f \geq \ell/p)$. In Section 4, the above theoretical results are used to interpret the experimental results when performing deterministic and random sampling on the France Telecom ADSL traffic traces.

3.3. Refinements. To prove Proposition 1, it has been assumed that flows are permanent. This assumption is reasonable, when the observation window length T is small. When T is large, however, flows may be bursty and alternate between on and off periods. To take into account this phenomenon, convergence to Poisson distributions as in Proposition 1 can be proved, when flows have different transmission rates.

More precisely, let us assume that there are L classes of flows. For a class $\ell \in \{1, \dots, L\}$, $r_\ell(x)$ is the transmission rate of a flow of class ℓ after a duration of time x . The quantity $C_\ell = \frac{1}{T} \int_0^T r_\ell(u) du$ is the average transmission rate of a flow on $[0, T]$. Flows are assumed to arrive uniformly in $[0, T]$. Consequently, for each flow f in class ℓ , the number of packets transmitted up to time $t \in [0, T]$ is $N_f(t) = \int_0^t r_\ell((u - \tau_f) \bmod T) du$, where the τ_f 's are independent and uniformly distributed in $[0, T]$. It follows that the different processes $N_f(t)$ for flows f in class ℓ have the same distribution.

For $\ell \in \{1, \dots, L\}$, let K_ℓ be the number of flows of class ℓ in $[0, T]$ and $K = \sum_\ell K_\ell$. The total number of transmitted packets up to time u is denoted by $M(u) = \sum_{i=1}^K N_i(u)$. Let $K = \sum_\ell K_\ell$ and $pM(T)$ be the number of sampling times between 0 and T , p denoting the sampling rate. When K becomes large, assume that for every ℓ , K_ℓ/K tends to a constant α_ℓ . By the law of large numbers, $M(u)/K$ converges almost surely to $C = \sum_{\ell=1}^L \alpha_\ell C_\ell$ for all $u \in [0, T]$. The numbers of packets n_f in the sampled flows f of class ℓ have the same distribution. We have the following result, whose proof is given in Appendix A.

Proposition 3. If $pM(T)/K \rightarrow x$, the distribution of the number $n(\ell)$ of packets in a sampled flow of class ℓ converge to a Poisson distribution with parameter xC_k/C .

The above proposition shows that the distribution of the number of sampled packets of a flow in class ℓ depends only upon on the ratio of the average rate of class ℓ to the total average rate in the observation window. This indicates that we could have considered the flows permanent at the average rate in the observation window.

4. EXPERIMENTAL RESULTS

In this section, we consider the traffic trace from the France Telecom backbone network described in Section 2 and we fix the length of the observation window

equal to $T = 3200$ seconds and the sampling rate $p = 1/100$. The complementary cumulative distribution functions (ccdf) of the number of packets in original mice and elephants are displayed in Figure 1(b) and 1(d), respectively. In the original trace, there were 252,854 elephants and in the sampled trace, we found 132,352 and 132,185 of the original elephants with deterministic and random sampling, respectively.

From the above experimental results, we see that the probabilities of seeing elephants after sampling in the different cases are very close one to each other, about 0.523.

If N_f has a Pareto distribution of the form $\mathbb{P}(N_f > k) = (b/k)^a \mathbb{1}_{\{k \geq b\}}$, the probability of seeing an elephant by random sampling is

$$\mathbb{P}\left(\sum_{i=1}^{N_f} B_i > 0\right) = 1 - (1-p)^b + p \sum_{k=b}^{\infty} (1-p)^k \mathbb{P}(N_f > k) \sim bp + (bp)^a \Gamma(1-a, bp),$$

when p is small. With $a = a_2 = 0.55$, $b = 20$, and $p = 1/100$, we find that the probability of seeing an elephant is approximately equal to 55 %, which is very close to the experimental value. Hence, by estimating the exponent a of the Pareto distribution allows us to estimate the probability of seeing an elephant. This quantity is critical for the estimation of the parameters of flows. For instance, for estimating the original duration of flows, a method is presented in [2], but the estimation of ν , the probability of seeing an elephant, is critical because it relies on the tails of some probability distributions. The method based on the estimation of the exponent of the Pareto distribution is more reliable.

The major difficulty for exploiting the sampled trace comes from the fact that we do not know if a sampled flow is really an elephant or not. If we had adopted the convention that a sampled flow corresponds to an elephant as soon as it is composed of at least two packets, we would have found 143,517 and 144,000 elephants with deterministic and random sampling, respectively. We see that this convention leads to slightly overestimating the number of elephants.

Figure 2 represents the ccdf of the number of packets in elephants after probabilistic and deterministic sampling, along with the rescaled original distribution $\mathbb{P}(N > k/p)/\nu$, where ν is the probability of seeing an elephant. We can observe that the three curves coincide, which is in agreement with the results obtained in Section 3. By using Proposition 2 and Theorem 1 and assuming that random and deterministic sampling are sufficiently close one to each other, we can recover the distribution of the original elephants from the distribution of sampled elephants with known bounds.

For the volume V (expressed in bytes) of elephants, we can first compute the mean number of bytes in packets. For instance, for the traffic trace considered in this paper, the mean number of bytes in packets of elephants is equal to $\bar{V} = 1000$. Then, we can verify that multiplying the number of packets in elephants by the mean number $\bar{V}$ of bytes in packets give a fair estimate of the volume of elephants, as illustrated in Figure 3(a). From the results established for the number of packets in elephants and under the assumption that random sampling is sufficiently close to deterministic sampling, we can estimate the volume of original elephants with known bounds; Figure 3(b) shows that the rescaled distribution $\mathbb{P}(V > x/p)/\nu$ is close to the distribution of the volume of sampled elephants.

FIGURE 2. Number of packets in elephants after sampling and comparison with the rescaled original size $\mathbb{P}(N > k/p)/\nu$ along with the Pareto approximation.

FIGURE 3. Volume (in bytes) of elephants after deterministic and probabilistic sampling and comparison with the rescaled original volume $\mathbb{P}(V > x/p)/\nu$.

5. CONCLUSION

We have shown in this paper that as far as the volume and the number of packets in elephants are concerned, random and deterministic sampling are very close to each other, when the sampling rate becomes small. Several results for the number of packets contained in randomly sampled flows have been established. In particular, bounds between the distribution of the number of packets in a randomly sampled elephant and the rescaled original distribution have been established. Experimental results obtained by using a traffic trace from the France Telecom IP backbone network show good agreement with theoretical results.

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(11-12):1260–1299, November-December 2004.
- [2] N. Ben Azzouna, F. Guillemin, S. Poisson, P. Robert, C. Fricker, and N. Antunes. Inverting sampled ADSL traffic. In *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] Y. Chabchoub, C. Fricker, F. Guillemin, and P. Robert. Bounds for packet sampling in the Internet. In Preparation.

- [5] B.Y. Choi, J. Park, and Z.L. Zhang. Adaptive packet sampling for accurate and scalable flow measurement. In *Proc. Globecom'04*, Dallas, TX, December 2004.
- [6] CISCO. <http://www.cisco.com/warp/public/netflow/index.html>.
- [7] Nick Duffield, Carsten Lund, and Mikkel Thorup. Properties and prediction of flow statistics. In *ACM SIGCOMM Internet Measurement Workshop*, pages 6–8, November 2002.
- [8] C. Estan, K. Keys, D. Moore, and G. Varghese. Building a better NetFlow. In *Proc. ACM Sigcomm'04*, Portland, Oregon, USA, August 30 – September 3 2004.
- [9] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proc. Sigcomm'02*, Pittsburgh, Pennsylvania, USA, August 19-23 2002.
- [10] W. Feller. *An introduction to probability theory and its applications*. John Wiley and Sons, 1996.
- [11] IETF, IPFIX Working Group. IP flow information export. For information, see the url <http://www.ietf.org/html.charters/ipfix-charter.html>.
- [12] IETF, PSAMP Working Group. Packet sampling working group. See the url <https://ops.ietf.org/lists/psamp>.
- [13] K. Papagiannaki, N. Taft, S. Bhattachayya, P. Thiran, K. Salamatian, and C. Diot. On the feasibility of identifying elephants in Internet backbone traffic. Technical Report TR01-ATL-110918, Sprint Labs, Sprint ATL, November 2001.
- [14] V. Paxson and S. Floyd. Wide area traffic: The failure of the Poisson assumption. *IEEE/ACM Trans. on Networking*, pages 226–244, 1995.
- [15] T. Zseby, M. Molina, N. Duffield, S. Niccolini, and F. Raspall. Sampling and filtering techniques for IP packet selection, January 2006.

APPENDIX A. APPENDIX: PROOF OF PROPOSITION 3

Let $(t_j)_{1 \leq j \leq pM(T)}$ be the sequence of the $pM(T)$ sampling times in $[0, T]$. We have for any flow in class ℓ , say, flow i

$$\mathbb{P}(n_i = 0) = \mathbb{E} \left(\prod_{j=1}^{pM(T)} \left(1 - \frac{N_i(t_j)}{M(t_j)} \right) \right) = \mathbb{E} \left(e^{\sum_{j=1}^{pM(T)} \log(1 - N_i(t_j)/M(t_j))} \right),$$

where n_i is the number of packets in the sampled flow i . First, note that

$$\sum_{j=1}^{pM(T)} \log \left(1 - \frac{N_i(t_j)}{M(t_j)} \right) = - \sum_{j=1}^{pM(T)} \frac{N_i(t_j)}{M(t_j)} + O \left(\frac{1}{K} \right).$$

Second, if f is a twice continuously differentiable function in $[0, T]$, we have

$$\sum_{j=1}^{pM(T)} f(t_j) = \frac{pM(T)}{T} \int_0^T f(u) du + \frac{f(T) - f(0)}{2} + O \left(\frac{1}{pM(T)} \right),$$

since the points (t_j) are distributed more or less uniformly in $[0, T]$. Hence, we have

$$\sum_{j=1}^{pM(T)} \log \left(1 - \frac{N_i(t_j)}{M(t_j)} \right) = - \frac{pM(T)}{T} \int_0^T \frac{N_i(u)}{M(u)} du + \frac{1}{2} \frac{N_i(T)}{M(T)} + O \left(\frac{1}{K} \right).$$

The first term of the right-hand side is equal to $-\frac{x}{T} \int_0^T \frac{r_\ell(u - \tau_i)}{M(u)/K} du$, which converges a.s. to $-\frac{x}{CT} \int_0^T r_\ell(u - \tau_i) du = -xC_\ell/C$, when K tends to $+\infty$. It follows that, when K tends to $+\infty$,

$$(3) \quad \mathbb{P}(n_i = 0) \rightarrow \exp \left(\frac{-xC_\ell}{C} \right).$$

For $k \in \mathbb{N}$,

$$\begin{aligned}\mathbb{P}(n_i = k) &= \mathbb{E} \left(\sum_{i_1 < \dots < i_k} \prod_{m=1}^k \frac{N_i(t_{i_m})}{M(t_{i_m})} \prod_{j \notin \{i_1 < \dots < i_k\}} \left(1 - \frac{N_i(t_j)}{M(t_j)} \right) \right) \\ &= \mathbb{E} \left(\prod_{j=1}^{pM(T)} \left(1 - \frac{N_i(t_j)}{M(t_j)} \right) \Sigma_k(g_i(t_1), \dots, g_i(t_{pM(T)})) \right),\end{aligned}$$

where $g_i(u) = N_i(u)/(M(u) - N_i(u))$ and $\Sigma_k = \sum_{i_1 < \dots < i_k} \prod_{j=1}^k X_{i_j}$ is the symmetric homogeneous polynomial of degree k . Denoting $S_i = \sum_{j=1}^{pM(T)} X_j^i$ for $i > 1$, Newton's formula

$$(-1)^k k \Sigma_k + \sum_{p=0}^{k-1} (-1)^p \Sigma_p S_{k-p} = 0 \quad (1 \leq k \leq pM(T))$$

establishes that Σ_k can be expressed as a function of $S_1, \dots, S_k$. It is clear that $S_q(g_i(t_1), \dots, g_i(t_{pM(T)}))$ is the Riemann sum with $pM(T)$ terms associated to g_i^q on $[0, T]$. Using Newton's formula, it can be proved that when K tends to $+\infty$,

$$\Sigma_k(g_i(t_1), \dots, g_i(t_{pM(T)})) \sim \frac{\left(\sum_{j=1}^{pM(T)} g_i(t_j) \right)^k}{k!}.$$

Taking into account approximation (3), we obtain that, for flow i of class ℓ ,

$$\mathbb{P}(n_i = k) \rightarrow e^{-x \frac{C_\ell}{C}} \frac{1}{k!} \mathbb{E} \left(\left(\frac{x}{CT} \int_0^T r_l(u - \tau_i) du \right)^k \right) = \frac{e^{-x \frac{C_\ell}{C}}}{k!} \left(\frac{x C_\ell}{C} \right)^k$$

and Proposition 3 follows.

E-mail address, Y. Chabchoub: `Yousra.Chabchoub@inria.fr`

(Y. Chabchoub, C. Fricker, Ph. Robert) INRIA, DOMAINE DE VOLUCEAU, B.P. 105, 78153 LE CHESNAY CEDEX, FRANCE

E-mail address, C. Fricker: `Christine.Fricker@inria.fr`

(F. Guillemin) FRANCE TELECOM R&D, F-22300 LANNION

E-mail address: `Fabrice.Guillemin@orange-ft.com`

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