Inference of Flow Statistics via Packet Sampling in the Internet

Yousra Chabchoub, Christine Fricker, Fabrice Guillemin, Philippe Robert

▶ **To cite this version:**

HAL Id: inria-00347008

https://hal.inria.fr/inria-00347008

Inference of Flow Statistics via Packet Sampling in the Internet

Yousra Chabchoub, Christine Fricker, Fabrice Guillemin, and Philippe Robert

Abstract—**We show in this note that by deterministic packet sampling, the tail of the distribution of the original flow size can be obtained by rescaling that of the sampled flow size. To recover information on the flow size distribution lost through packet sampling, we propose some heuristics based on measurements from different backbone IP networks. These heuristic arguments allow us to recover the complete flow size distribution.**

Index Terms—**Packet sampling, Flow statistics, Pareto distribution.**

I. Introduction

Packet sampling is an efficient method of reducing the amount of data to analyze when performing measurements in the Internet. The simplest and the most popular packet sampling technique consists of selecting one packet every other k packets. This technique is referred to as deterministic 1-out-of-k sampling in the technical literature and has notably been implemented in CISCO routers [1]. Even if this sampling scheme suffers from several drawbacks, identified for instance in [2], it is widely used in today's operational networks.

The basic problem of packet sampling is that it is difficult to infer the original flow statistics from sampled data. Note that a flow is defined as the set of those packets sharing some common addressing information, typically the same source and destination IP addresses, the same source and destination port numbers together with the same protocol type.

Flow statistics inference from sampled data has been addressed in previous studies. Duffield *et al* [3], [4] study the accuracy of different estimators based on multiplying the sampled flow size by the sampling factor k, but their method does not apply to the complete range of the flow size. Hohn and Veitch [5] use generating function techniques to invert the flow size distribution but the proposed procedure is numerically unstable. Mori *et al* [6] use a Bayesian approach to inferring the characteristics of long flows.

In this paper, we develop a probabilistic approach to inverting sampled traffic together with some heuristic arguments. First, we note that when observing sampled traffic, we can only compute the distribution of the random variable $\tilde{v}$ describing the number of packets in sampled flows and K_s the number of sampled flows. If there are originally K flows, we have

$$\mathbb{P}(\tilde{v} = j) = \frac{1}{K_s} \sum_{i=1}^{K} I_{i,j} \qquad (1)$$

Y. Chabchoub, C. Fricker and P. Robert are with INRIA, Domaine de Voluceau, 78153 Le Chesnay, France, Email:{Yousra.Chabchoub, Christine.Fricker,Philippe.Robert@inria.fr}

F. Guillemin is with Orange Labs, 2 Avenue Pierre Marzin, 22300 Lannion, France, Email:Fabrice.Guillemin@orange-ftgroup.com

where $I_{i,j} = 1$ if the ith flow has been sampled j times and $I_{i,j} = 0$ otherwise.

Under some reasonable assumptions on the sampling process, we show in this note that the tail of the original flow size distribution can be obtained by rescaling the distribution of the sampled flow size distribution. It is however not possible to totally recover the original flow size distribution because information on small or moderate flow sizes is loss through sampling. To overcome this problem, we propose some heuristic arguments based on measurements and exploiting a priori information on flows. We consider here TCP traffic only.

The rest of this note is organized as follows: In Section II, we make some reasonable assumptions on the sampling process. In Section III, we prove that the tail of the original flow size can be obtained by rescaling that of the sampled flow size. In Section IV, we present some heuristic arguments to recover the total flow size distribution. Concluding remarks are presented in Section V.

II. Assumptions on the sampling process

When observing in a time window of length Δ traffic on a high speed link, one may reasonably assume that the packets of the different active flows are sufficiently interleaved. Hence, one may suppose the selection of packets among active flows at a sampling time is random.

Moreover, in a time window of length Δ, flows start and finish and some of them may be silent (for instance in the case of flows alternating between On and Off periods). In [7, Section 3.3], it is shown that these fluctuations may neglected at the first order (i.e., when computing mean values) and it can be assumed that flows are permanent. Under the two above assumptions, we suppose that the probability of selecting a packet of a given flow, say, flow i, is equal to v_i/V_i, where v_i is the size of flow i and V_i is the total number of packets arrived when flow i is active.

III. Tail of the sampled flow size

Let $W_j \stackrel{\text{def}}{=} \sum_{i=1}^{K} I_{i,j}$, the number of flows sampled j times.

Proposition 1: If K flows are active during a time window of length Δ, the mean value $\mathbb{E}(W_j)$ satisfies

$$|\mathbb{E}(W_j) - K\mathbb{Q}_j| \leq p \sum_{i=1}^{K} \mathbb{E}\left(v_i^2/V_i\right), \qquad (2)$$

where $p = 1/k$, v_i is the random number of packets in a flow, and $\mathbb{Q}$ is the probability distribution defined by

$$\mathbb{P}(\mathbb{Q} = j) \stackrel{def}{=} \mathbb{Q}_j = \mathbb{E}\left(\frac{(pv)^j}{j!}e^{-pv}\right), \qquad (3)$$

Proof: Let us condition on the values of the set $\mathcal{F} = \{v_1, \ldots, v_K, V_1, \ldots, V_K\}$. Under the assumptions of Section II, the number of times that the ith flow is sampled is equal to the sum

$$S_i = B_1^i + B_2^i + \cdots + B_{pV_i}^i,$$

where 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_i. The random variables $(B_\ell^i, \ell \geq 1)$ are i.i.d. Bernoulli random variables and Le Cam's Inequality [8] then states

$$\left\|\mathbb{P}(S_i \in \cdot) - \mathbb{P}\left(Q_{\mathbb{E}(S_i)} \in \cdot\right)\right\|_{tv} \leq \sum_{\ell=1}^{pV_i} \mathbb{P}(B_\ell^i = 1)^2,$$

where $\|.\|_{tv}$ is the total variation norm and $Q_{\mathbb{E}(S_i)}$ is a Poisson random variable with mean $\mathbb{E}(S_i)$. By deconditioning with respect to the set $\mathcal{F}$, we have by using the distribution $\mathbb{Q}$

$$\|\mathbb{P}(S_i \in \cdot) - \mathbb{Q}\|_{tv} \leq p\mathbb{E}\left(v_i^2/V_i\right). \qquad (4)$$

In particular, for $j \in \mathbb{N}$, $|\mathbb{P}(S_i = j) - \mathbb{Q}_j| \leq p\mathbb{E}\left(v_i^2/V_i\right)$. Since $\mathbb{E}(W_j) = \sum_{i=1}^{K} \mathbb{P}(\tilde{v}_i = j)$, summing on i yields Equation (2). ∎

If K is sufficiently large, we have from Equation (1) and the above proposition, we have for $j \geq 1$

$$\mathbb{P}(\tilde{v} = j) \sim \frac{1}{\nu}\frac{\mathbb{E}(W_j)}{K}, \qquad (5)$$

where $\nu = K_s/K$ is the probability of sampling a flow.

Proposition 2: If all flows have a negligible contribution to the total volume of traffic (i.e., $\mathbb{E}(v_i^2/V_i) \ll 1$ for all $i = 1, \ldots, K$), if K is sufficiently large, and if the flow size distribution has a slowly varying tail, then when $j \to \infty$

$$\mathbb{P}(\tilde{v} \geq j) \sim \mathbb{P}\left(v \geq \frac{j}{p}\right)/\nu. \qquad (6)$$

Proof: From Equation (5)

$$\mathbb{P}(\tilde{v} = j) \sim \frac{1}{\nu}\mathbb{E}\left(\frac{(pv)^j}{j!}e^{-pv}\right) = \frac{p^j}{j!\nu}\sum_{\ell=0}^{\infty} e^{j\log\ell - p\ell}\mathbb{P}(v = \ell). \qquad (7)$$

Consider the sum $\sum_{\ell=0}^{\infty} e^{f(p,\ell)}\mathbb{P}(v = \ell)$, where $f(p,x) = j\log x - px$, which is maximum at point j/p. By assuming that the function $\ell \to \mathbb{P}(v = \ell)$ is heavy tailed, Laplace method gives for large j

$$\mathbb{E}\left(\frac{(pv)^j}{j!}e^{-pv}\right) \sim \frac{p^j}{j!}e^{f(p,j/p)}\mathbb{P}\left(v = \frac{j}{p}\right)\sum_{\ell=-\infty}^{\infty} e^{\ell^2\frac{f''(p,j/p)}{2}},$$

where $f''(p,j/p) = -p^2/j$ and $e^{f(p,j/p)} = (j/p)^j e^{-j}$. If j is sufficiently large, Stirling formula gives $j! \sim \sqrt{2\pi}j^{j+\frac{1}{2}}e^{-j}$. In addition, from [9], we have $\sum_{n=-\infty}^{\infty} e^{-an^2} \sim \sqrt{\pi/a}$ when $a \to 0$. This implies that $\sum_{\ell=-\infty}^{\infty} e^{\ell^2\frac{f''(p,j/p)}{2}} \sim \frac{\sqrt{2\pi j}}{p}$ and

then $\mathbb{P}(\tilde{v} = j) \sim \mathbb{P}(v = j/p)/p$, when j is sufficiently large. Since

$$\mathbb{P}(\tilde{v} \geq j) \sim \frac{1}{\nu p}\sum_{k=j}^{\infty} \mathbb{P}(v = k/p) \sim \frac{1}{\nu}\int_j^{\infty} d\mathbb{P}(v = k/p),$$

Equation (6) follows. ∎

IV. HEURISTICS FOR THE TOTAL FLOW SIZE DISTRIBUTION

Proposition 2 shows that the tail of the complementary cumulative distribution function (ccdf) of the original flow size can be obtained by rescaling that of the sampled flow size. We can however verify through examples that information on that distribution for small or moderate flow size values is lost.

We exemplify this phenomenon by considering a 2 hour long real traffic trace from a 1 Gbit/s transmission link of the France Telecom IP backbone network carrying ADSL traffic. The original flow size is depicted in Figure 1(a) and the deterministically sampled flow size in Figure 1(b), which exhibits good agreement with the rescaled distribution $\mathbb{P}(v = j/p)/\nu$ for sufficiently large j as predicted by Proposition 2. But all information for moderate values of the flow size is contained in a few values, in this case $\mathbb{P}(\tilde{v} \geq j)$ for $j = 2, 3$. The same phenomenon (see [10]) has been observed for an Abilene traffic trace available at http://pma.nlanr.net/Traces/Traces/long/ipls/3/.

(a) Original size.
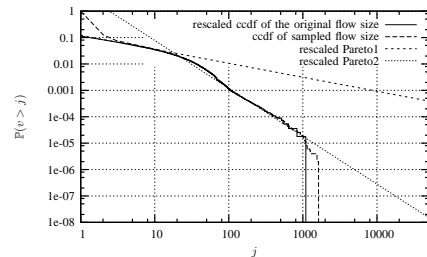
(b) Sampled size ($p = 1/100$).

Fig. 1. Flow size distribution in the France Telecom ADSL trace.

In fact, through numerous experiments with real traffic traces, it has been observed in [10] that $\mathbb{P}(v \geq j/p)$ can be approximated by $\nu\mathbb{P}(\tilde{v} \geq j)$ when $j \geq j_0$ for some $j_0 > 0$. The problem is then to estimate the quantities $\mathbb{P}(v = j)$ for $j = 1, \ldots, j_0/p - 1$.

We have from Equation (7)

$$\mathbb{P}(\tilde{v} = j) \sim \frac{1}{\nu} \sum_{\ell=1}^{\infty} \frac{(p\ell)^j}{j!} e^{-p\ell} \mathbb{P}(v = \ell) \qquad (8)$$

and we know by Proposition 2 that for $j \geq j_0$, this equation is equivalent to $\mathbb{P}(\tilde{v} = j_0) = \mathbb{P}(v = j/p)/(\nu p)$. It follows that for determining the $(j_0/p - 1)$ quantities $\mathbb{P}(v = \ell)$ for $\ell = 1, \ldots, j_0/p - 1$, we have only j_0 equations. The problem is hence clearly under-determined. Some heuristics are needed to recover the complete flow size distribution.

It has been observed in [10] that depending on the size of the observation window Δ, the sampled flow size distribution can locally be approximated by means of Pareto distributions. This leads us to make the following assumption.

Assumption 1: There exist some $m > 0$ and some integers $j_0 < j_1 < ... < j_m = \infty$ such that for $\ell = 1, \ldots, m$ and $j \in [j_{\ell-1}, j_\ell]$, $\tilde{v}$ has a Pareto distribution of the form

$$\mathbb{P}(\tilde{v} \geq j) = \mathbb{P}(\tilde{v} \geq j_{\ell-1}) \, (j_{\ell-1}/j)^{a_\ell}$$

for some shape parameter $a_\ell > 0$.

When Δ is adequately chosen, the tail may be uni-modular (i.e., $m = 1$), but when Δ is too large, we can have $m > 1$. For the above France Telecom trace ($\Delta = 2$ hours), $m = 2$ as shown in Figure 1(b).

By using Proposition 2, we deduce that for $\frac{j_{m-1}}{p} \leq j \leq \frac{j_m}{p}$

$$\mathbb{P}(v \geq j) \sim \nu \mathbb{P}(\tilde{v} \geq j_{m-1}) \, (j_{m-1}/(pj))^{a_m}, \qquad (9)$$

The above equation implies that $\mathbb{P}(v \geq j)$ can locally be approximated by a Pareto distribution with shape parameter a_m, as shown in Figure 1(a).

For inferring the quantities $\mathbb{P}(v = j)$ for $j = 1, \ldots, j_0/p-1$, we need more assumptions. Numerous experiments [10] have shown that when $j < b_0$ for some $b_0 > 0$, $\mathbb{P}(v = j)$ follows a geometric distribution.

Assumption 2: There exists some $b_0 > 0$ such that for $1 \leq j < b_0$, $\mathbb{P}(v = j) = (1 - r)r^j$ for some $r > 0$.

The above assumption is supported by experiments, as shown in Figure 2(a) and 2(b) for the France Telecom and Abilene traffic traces, respectively. The value $b_0 = 20$ has been successfully tested in numerous experiments.


(a) ADSL trace (b) Abilene trace

Fig. 2. Ccdf of the number of packets in flows with less than $b_0 = 20$ packets.

By using Equation (9) and Assumption 2, we have the form of the distribution for $j \leq b_0$ and $j \geq j_0/p$. To fill the gap, we use the following heuristic: $\mathbb{P}(v \geq j)$ for $b_0 \leq j \leq j_1/p$

has the same form as in Equation (9), namely $\mathbb{P}(v \geq j) = \mathbb{P}(v \geq b_0)(b_0/j)^{a_1}$. Equation (8) can then be rewritten as

$$\mathbb{P}(\tilde{v} = j) \sim \frac{\mathbb{P}(v < b_0)}{\nu} \sum_{\ell=1}^{\infty} (1 - r)r^\ell \frac{(p\ell)^j}{j!} e^{-p\ell}$$
$$+ \frac{1}{\nu} \sum_{\ell=b_0}^{\infty} \frac{(p\ell)^j}{j!} e^{-p\ell} \mathbb{P}(v = \ell). \qquad (10)$$

The shape parameters a_ℓ for $1 \leq \ell \leq m$ are determined from the sampled flow size distribution by using standard Maximum Likelihood Expectation (MLE) procedures. The parameter b_0 is set equal to 20; this choice is purely phenomenological but corresponds to the number of packets needed to leave the slow start regime with a maximum window size of 32 Kbytes. The parameter $\mathbb{P}(v \geq b_0)/\nu$ is obtained by using Proposition 2, namely by computing the ratio $\eta \stackrel{def}{=} \mathbb{P}(\tilde{v} \geq j)/(b_0 p/j)^{a_1}$ for $j \in \{j_0, \ldots, j_1\}$, which is by assumption independent of j. The number of flows with at least b_0 packets is $K_0^+ = \eta K_s$. Equation (10) multiplied by K_s for $j = 1, 2$ is then used to compute the parameter r and the number K_0^- of flows with less than b_0 packets. The total number of flows is then $K = K_0^+ + K_0^-$ and the probability of sampling a flow is estimated by the ratio K_s/K.

By using the above method for the France Telecom ADSL trace with $p = 1/100$, we find $j_0 = 3$ and the estimated shape parameters $\hat{a}_1 = .54$ and $\hat{a}_2 = 1.81$, which are close to the experimental values $a_1 = .52$ and $a_2 = 1.81$ for the original flow size. We then find $\mathbb{P}(v \geq b_0)/\nu = .3$ and since $K_s = 1,120,546$, we obtain the estimate $\hat{K}_0^+ = 336,163$, while the actual value is $K_0^+ = 343,004$. By neglecting the term due to flows with at least 20 packets in Equation (10), we then find the estimate $\hat{r} = 0.84$ while the actual experimental value is $r = .75$. This yields a number of flows with less than b_0 packets $\hat{K}_0^- \sim 20.1e6$ while the actual value is $K_0^- \approx 19.8e6$. Finally, the estimated total number of flows is $\hat{K} = 20.4e6$ while the actual value is $K = 20.1e6$ and we find the estimate $\hat{\nu} = .054$ for the probability of sampling a flow while the experimental value is $\nu = 0.057$.

V. CONCLUSION

We have shown in this paper by using probabilistic arguments that the original size distribution of large flows can be recovered from that of the sampled flow size. A critical parameter is nevertheless the flow sampling probability, which can be estimated only when the size of small flows is known. To overcome this problem, we argue that it is necessary to exploit a priori information on flows. By using this principle, we have shown that it is possible to recover the complete flow size distribution together with the number of flows.

REFERENCES

[1] CISCO, http://www.cisco.com/warp/public/netflow/index.html.
[2] 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.
[3] N. Duffield, C. Lund, and M. Thorup, "Properties and prediction of flow statistics properties and prediction of flow statistics," in *ACM SIGCOMM Internet Measurement Workshop*, November 2002, pp. 6–8.

[4] ——, "Estimating flow distributions from sampled flow statistics," in *SIGCOMM'03*, vol. August, 2003, pp. 25–29.

[5] N. Hohn and D. Veitch, "Inverting sampled traffic," in *Internet Measurement Conference*, October 2003, pp. 27–29.

[6] T. Mori, M. Uchida, and R. Kawahara, "Identifying elephant flows through periodically sampled packets," in *Internet Measurement Conference*, Taormina, Italy, 2004.

[7] Y. Chabchoub, C. Fricker, F. Guillemin, and P. Robert, "Deterministic versus probabilistic packet sampling in the Internet," in *Proceedings of ITC'20*, June 2007.

[8] A. D. Barbour, L. Holst, and S. Janson, *Poisson approximation*.   New York: The Clarendon Press Oxford University Press, 1992, oxford Science Publications.

[9] M. Abramowitz and I. Stegun, *Handbook of mathematical functions*. National Bureau of Standards, Applied Mathematics Series 55, 1972.

[10] Y. Chabchoub, C. Fricker, F. Guillemin, and P. Robert, "A robust statistical estimation of internet traffic," June 2008, preprint. http://www-rocq.inria.fr/˜robert/src/papers/2008-1.pdf.