



Inverting sampled ADSL traffic

Nadia Ben Azzouna, Fabrice Guillemin, Stéphanie Poisson, Philippe Robert,
Christine Fricker, Nelson Antunes

► **To cite this version:**

Nadia Ben Azzouna, Fabrice Guillemin, Stéphanie Poisson, Philippe Robert, Christine Fricker, et al..
Inverting sampled ADSL traffic. [Research Report] RR-5425, INRIA. 2004, pp.14. inria-00070581

HAL Id: inria-00070581

<https://hal.inria.fr/inria-00070581>

Submitted on 19 May 2006

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.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Inverting sampled ADSL traffic

Nadia Ben Azzouna — Fabrice Guillemin — Stéphanie Poisson — Philippe Robert —
Christine Fricker — Nelson Antunes

N° 5425

Décembre 2004

Thème COM



*R*apport
de recherche



Inverting sampled ADSL traffic

Nadia Ben Azzouna^{*}, Fabrice Guillemin, Stéphanie Poisson, Philippe Robert[†], Christine Fricker, Nelson Antunes[‡]

Thème COM — Systèmes communicants
Projet Rap

Rapport de recherche n° 5425 — Décembre 2004 — 14 pages

Abstract: On the basis of a reference model for ADSL traffic on an IP backbone link, established in an earlier study, we show in this paper that it is possible to infer the characteristics of long flows by performing a deterministic $1/N$ packet sampling. By using the fact that the number of active long flows can be represented by means of the number of customers in an $M/G/\infty$ queue with Weibullian service times, we derive some probabilistic properties of sampled data. These properties are then used to infer the characteristics of original flows. The method is illustrated by considering an actual traffic trace captured in the France Telecom IP backbone network. Experimental data show that the method proves quite efficient.

Key-words: ADSL, peer-to-peer, power spectrum, long range dependence

^{*} N. Ben Azzouna, F. Guillemin and S. Poisson are with France Telecom, Division R&D, 22300 Lannion, France

[†] This work has been partially supported by the CRE research contract n° 103D1022 with France Telecom and the RNRT project Metropolis

[‡] N. Antunes, C. Fricker and Ph. Robert are with INRIA Rocquencourt, RAP project, Domaine de Voluceau, 78153 Le Chesnay, France

Étude de l'échantillonnage du trafic ADSL

Résumé : Cet article étudie les caractéristiques du trafic ADSL qui peuvent être identifiées en échantillonnant les paquets IP au niveau d'un routeur de dorsale

Mots-clés : ADSL

1 Introduction

Traffic measurement is an issue of prime interest for network operators and networking researchers in order to know the nature and the characteristics of traffic supported by IP networks. The exhaustive capture of traffic traces on high speed backbone links, with rates larger than 1 Gigabit/s, however, leads to the storage and the analysis of huge amounts of data, typically several TeraBytes per day. A method of overcoming this problem is to reduce the volume of data by sampling traffic. Several sampling techniques have been proposed in the literature (see for instance [1, 2] and references therein). In this paper, we consider the deterministic $1/N$ sampling, which consists of capturing one packet every other N packets. This sampling method has notably been implemented in CISCO routers under the name of NetFlow [3], which is widely deployed nowadays in commercial IP networks.

The major issue with $1/N$ sampling is that the correlation structure of flows is severely degraded and then any digital signal processing technique turns out very delicate to apply in order to recover the characteristics of original flows [2]. An alternative approach consists of performing a statistical analysis of flow as in [1, 4]. The accuracy of such an analysis, however, greatly depends on the number of samples for each type of flows and may lead to quite inaccurate results. In fact, this approach proves efficient only in the derivation of mean values of some characteristics of interest, for instance the mean number of packets or bytes in a flow.

In this paper, we show that with an a priori knowledge on flows, it is possible to recover with a reasonable accuracy some key characteristics. In order to illustrate this method, we shall consider traffic measurements from the France Telecom backbone network carrying ADSL traffic. In particular, we shall show that we can experimentally compute the characteristics of long flows (referred to as elephants), in particular the probability distribution of their duration and their arrival rate. These flows represent 94% of global traffic.

For this purpose, we shall use the model introduced in [5] for describing TCP traffic generated by ADSL customers on a backbone link in a commercial IP network. The most salient feature of ADSL traffic is that the most significant part (up to 80%) is due to peer-to-peer (p2p) applications. For traffic analyzed in [5], the dominating p2p application is eDonkey, but most current p2p protocols share some common properties, which tend to smooth traffic. For instance, current p2p protocols split large files (corresponding to audio CDs or DivX-compressed movies) into separate pieces (chunks) of moderate size (e.g., 9 Mbytes for eDonkey, see [6] for details). Chunks may be retrieved asynchronously and in parallel by peers. Segmentation of large files into chunks and multiple downloads are characteristic features, which are implemented by many p2p protocols (e.g., Bearshare, Xolox, KaZaA).

In addition, the bit rates of individual TCP connections associated with chunk downloads are in practice limited to a few tens of Kbit/s. This limitation may be due to different factors. For instance, the terminals of individual ADSL customers may play the role of servers and only a part of the available CPU is devoted to jobs related to p2p activity. Moreover, the bit rate of ADSL lines in the upload direction is often limited to a few hundreds of Kbit/s. Thus, multiple uploads may severely limit the bit rates of individual TCP connections.

The predominance of p2p applications in commercial traffic has a great impact on the characteristics of traffic. Early traffic measurements in local area or campus networks (see [7] and [8]) have shown that traffic exhibits long range dependence and self-similarity. These properties are usually explained by the fact that large files without severe access bandwidth restrictions are transmitted through such networks and TCP naturally give rise to oscillations in the case of packet loss (see for instance [9]). Because of access rate limitations and the predominance of p2p traffic with the characteristics recalled above, ADSL traffic exhibits quite different properties. In fact, it can be shown (see [10]) that the global bit rate can be described by means of the superposition of Gaussian processes with subexponentially decreasing correlations. In particular, the number of flows active at an arbitrary instant can be well approximated by the number of customers in an $M/G/\infty$ queue with Weibullian service times.

The organization of this paper is as follows: In Section 2, the basic characteristics of ADSL traffic are recalled. In Section 3, we investigate the probabilistic properties of a $1/N$ packet sampler with ADSL traffic. The results are then used in Section 4 to infer the traffic characteristics of elephants by analyzing sampled data. Some concluding remarks are presented in Section 5.

2 A reference model for ADSL traffic

To numerically illustrate the different definitions introduced in this paper, we consider traffic measurements from a Gigabit Ethernet link of the France Telecom backbone network. This link carries traffic towards different ADSL areas and only TCP traffic is observed. The traffic trace has been captured in November 2003 between 9:00 pm and 11:00 pm; this time period corresponds to the peak activity of ADSL customers. Different measurement experiments have been performed in 2002 and 2003; the one reported in this paper seems to be the most interesting one since the load of the link is quite large (about 43%) and the resulting global bit rate exhibits an almost stationary behavior (see [5] for details).

To establish a model for the global bit rate process, the well known mouse/elephant dichotomy has been used in [5]. Even though a clear definition does not exist for mice and elephants (see [11] for a thorough discussion on the difficulty of separating mice from elephants), the following simplified definitions are used in the analysis carried out in this paper. We recall that a flow is a group of TCP packets with the same source and destination IP addresses together with the same source and destination port numbers.

Definition 1 (Mice and elephants). *A mouse is a flow with less than 20 packets. An elephant is a flow with more than 20 packets.*

The above definitions may appear very crude at first glance, but they allow a first segregation between long and short flows. These two types of flows are expected to have different networking behavior. As a matter of fact, flows with less than 20 packets does not or slightly leave the slow start period. Those flows are thus not so much sensitive to the fair bandwidth sharing principle imposed by TCP, even if they may react to loss by resuming

the slow start period. Long flows are expected to be regulated in the congestion avoidance regime and to fairly share the bandwidth of the network.

The segregation between long and short flows by using the above criterion is furthermore supported by the fact p2p protocols generate a huge amount of signalling traffic giving rise to flows with small size, typically less than 8 packets per flow; see [10] for details. Hence, when examining the probability distribution of the global flow duration, we can observe a bimodal distribution with one mode for short flows and another one for long flows.

With the above definitions, mice represent about 95% of flows but contribute only 5% to the global bit rate. Moreover, short flows have a little chance to be seen more than once when using a $1/N$ packet sampler, with N ranging from 100 to 1000. In fact, short flows give rise to a white noise in the sampled signal as it is seen by computing the power spectrum of the sampled bit rate. This is why we shall concentrate only on long flows in the following. But, we have to first draw attention to the fact that those flows may be of different types.

As a matter of fact, with the emergence of p2p applications, terminals of customers might play the role of servers. Hence, when observing traffic from an IP backbone network towards different ADSL areas, a significant proportion of long flows is mainly composed of ACK segments. Those elephants are referred to as ACK elephants. They contribute only 1% to the global bit rate but represent one third of elephants. Those elephants are in reality not so much interesting when studying the global bit rate. We shall subsequently introduce a heuristic rule in order to remove ACK elephants from sampled data.

The other elephants are referred to as data elephants. The mean number of active data elephants at an arbitrary instant during the observation period is approximately equal to 30,000. When studying the packet transmission pattern of a data elephant, one can observe that packets are not regularly transmitted but they are sent in bursts separated by low transmission phases. One may expect that this phenomenon is not seen by the packet sampler since the number of simultaneously active data elephants is large and two samples of a same elephant are separated by a rather large time period so that low transmission phases are unlikely seen by the packet sampler.

The empirical distribution of the duration S of data elephants is given in Figure 1. It appears that this distribution can be approximated by a two-parameter Weibullian distribution, which means that

$$\mathbb{P}(S > x) \approx \exp\left(-\left(\frac{x}{\eta}\right)^\beta\right). \quad (1)$$

where η and β are the scale and the skew parameters, respectively. One empirically finds, by using a Maximum Likelihood Estimation based on a Matlab routine, that $\eta = 116$ and $\beta = 0.39$. In addition, by following the same method as in [5] via computing the empirical spectral density of the stochastic process describing the number of active elephants and then comparing it with the theoretical spectral density, it can be checked that elephants arrive according to a Poisson process with rate λ ; the empirical value of the parameter λ is equal to 35.2 elephant/s. Thus, the number of active elephants at an arbitrary instant is equal to the number of customers in an $M/G/\infty$ queue with Weibullian service times.




Figure 1: Empirical distribution of the duration of elephants and Weibull approximation.

3 Theoretical analysis of sampled data

When sampling traffic on an IP backbone link, one picks up a packet among all active flows. If they were only data elephants and if all those elephants had similar bit rate, the probability of drawing a packet of a particular elephant at time t would then be $1/n_t$, where n_t is the number of elephants active at time t . Since the number of elephants is close to the number of customers in an $M/G/\infty$ queue in heavy traffic (since $\lambda\mathbb{E}[S]$ is large), one may approximate $1/n_t$ by $1/(\lambda\mathbb{E}[S])$. In addition, by using once more the fact that $\lambda\mathbb{E}[S]$ is large, everything happens as if one sampled an elephant according to a Poisson process, say, with rate μ .

Adding mice does not fundamentally change the situation. In that case, the probability of seeing an elephant is smaller since there are more flows, but the presence of mice accelerates the sampling clock. Thus, one may reasonably assume that $1/N$ sampling is equivalent to a Poisson sampling of elephants, i.e., the times between two successive samples of a same elephant are independent and exponentially distributed, with a large mean value denoted by $1/\mu$.

In the following, a sampled elephant is a flow, which is seen more than once by the packet sampler. We denote by $F(x)$ the probability distribution function of the duration S of elephants, which is assumed to be Weibullian with scale parameter η and skew parameter $\beta < 1$; the function $1 - F(x)$ is given by Equation (1). The estimation of the parameters μ , β and η is discussed in Section 4. In a first step, we compute some probabilistic characteristics of the system. In a first step, we consider the distribution of the number of times an elephant is seen by the packet sampler.

Lemma 1. *The distribution of the number ν of times an elephant is seen by the packet sampler is given by*

$$n \geq 0, \quad \mathbb{P}(\nu > n) = \int_0^{+\infty} \frac{x^n}{n!} e^{-x} e^{-\left(\frac{x}{\mu\eta}\right)^\beta} dx. \quad (2)$$

Proof. Let X denote a Weibull random variable with scale parameter η and skew parameter $\beta < 1$, and let \mathcal{N} be a Poisson process with intensity μ ; the points of the Poisson process are denoted by T_n for $n \geq 0$ with $T_0 \leq 0 < T_1 < T_2 < \dots$. By definition, we have

$$\mathbb{P}(\nu > n) = \mathbb{P}(\mathcal{N}[0, X] > n) = \int_0^{+\infty} \mathbb{P}(T_{n+1} \leq x) F(dx)$$

Since T_{n+1} has an Erlang- n distribution with mean $(n+1)/\mu$ and variance $(n+1)/\mu^2$, Equation (2) directly follows via an integration by part. \square

The above lemma allows us to compute the first and second moments of the random variable ν .

Corollary 1. *The mean and the variance of the random variable ν are given by*

$$\mathbb{E}(\nu) = \mu \mathbb{E}(S) = \mu\eta\Gamma(1 + 1/\beta), \quad (3)$$

$$\text{Var}(\nu) = \mu^2 \text{Var}(S) + \mu \mathbb{E}(S), \quad (4)$$

where $\text{Var}(S)$ is the variance of the duration S of an elephant, given by

$$\text{Var}(S) = \eta^2 \Gamma\left(\frac{2}{\beta} + 1\right) - \eta^2 \Gamma\left(\frac{1}{\beta} + 1\right)^2. \quad (5)$$

Proof. Given that conditionally to the event $S = x$, the random variable ν is Poisson with mean μx , we have

$$\mathbb{E}[\nu \mid S = x] = \mu x \quad \text{and} \quad \mathbb{E}[\nu^2 \mid S = x] = (\mu x)^2 + \mu x.$$

By deconditioning, Equations (3) and (4) follow. \square

Since $\beta < 1$, it is worth noting that $\mathbb{P}(\nu > n)$ can be expanded in power series of $1/(\mu\eta)^\beta$ as

$$\mathbb{P}(\nu > n) = \frac{1}{n!} \sum_{m=0}^{\infty} \frac{\Gamma(n + m\beta + 1)}{m!} \left(\frac{-1}{(\mu\eta)^\beta}\right)^m, \quad (6)$$

which is obtained by expanding $\exp\left(-\left(\frac{x}{\mu\eta}\right)^\beta\right)$ in power series of x and by using the fact that the above series is absolutely converging when $\beta < 1$. This series can be used for numerically evaluating the quantity $\mathbb{P}(\nu > n)$. When $1/(\mu\eta)^\beta$ is small, the Wyn and Shanks

ε -algorithm might be used to accelerate the convergence of the series (see [12, 13] for the ε -algorithm).

From Lemma 1, we immediately deduce the following result for the number of packets in a sampled elephant (recalling that a sampled elephant is a flow seen more than once by the packet sampler).

Corollary 2. *The distribution of the number M of packets in a sampled elephant is given by*

$$n \geq 2, \quad \mathbb{P}(M = n) = \frac{\mathbb{P}(\nu = n)}{\mathbb{P}(\nu > 1)}. \quad (7)$$

Equations (3) and (4) can be used to compute the mean and the variance of the random variable M . We now turn to the analysis of the asymptotic properties of sampled elephants, in particular the asymptotic behavior of their duration and their number of packets.

Proposition 1. *Let D denote the duration of a sampled elephant. Then, the complementary probability distribution function of the random variable D is given by*

$$\mathbb{P}(D > x) = \frac{1}{\mathbb{P}(\nu > 1)} \int_x^\infty \mu^2 (y - x) e^{-\mu(y-x)} e^{-(y/\eta)^\beta} dy. \quad (8)$$

Moreover, when x tends to infinity,

$$\mathbb{P}(D > x) \sim \frac{1}{\mathbb{P}(\nu > 1)} \exp\left(-\left(\frac{x}{\eta}\right)^\beta\right). \quad (9)$$

Proof. The duration of a sampled elephant is larger than x if the duration of the original elephant is $y > x$ and if one point of the sampling process occurs at time $u \in (0, y - x)$ and at least another point falls within the time interval $(u + x, y)$. Thus,

$$\begin{aligned} \mathbb{P}(D > x) &= \\ &= \frac{1}{\mathbb{P}(\nu > 1)} \int_x^\infty \int_0^{y-x} \mu e^{-\mu u} (1 - e^{-\mu(y-x-u)}) du F(dy) \\ &= \frac{1}{\mathbb{P}(\nu > 1)} \int_x^\infty (1 - e^{-\mu(y-x)} - \mu(y-x)e^{-\mu(y-x)}) F(dy) \end{aligned}$$

and Equation (8) follows by integrating by parts.

When x is large, we have by Equation (8) and Dominated Convergence Theorem

$$\begin{aligned} \mathbb{P}(D > x) &= \frac{e^{-(x/\eta)^\beta}}{\mathbb{P}(\nu > 1)} \int_0^{+\infty} \mu^2 z e^{-\mu z} e^{-(\frac{x}{\eta})^\beta} \left[\left(1 + \frac{z}{x}\right)^\beta - 1 \right] dz \\ &\sim \frac{1}{\mathbb{P}(\nu > 1)} e^{-(x/\eta)^\beta} \end{aligned}$$

and Equation (9) follows. □

From the above result, we can deduce via simple manipulations the mean value of the duration of sampled elephants.

Corollary 3. *The mean value of the duration of sampled elephants is given by*

$$\mathbb{E}(D) = \frac{\mathbb{E}(\nu) - \mathbb{P}(\nu > 0) - \mathbb{P}(\nu > 1)}{\mu\mathbb{P}(\nu > 1)} = \frac{1}{\mu}\mathbb{E}(M - 2), \quad (10)$$

and the second moment is given by

$$\mathbb{E}(D^2) = \frac{1}{\mu^2}\mathbb{E}((M - 2)(M - 3)), \quad (11)$$

where M is the number of packets in a sampled elephant (containing at least two packets).

Proof. We have

$$\begin{aligned} \mathbb{P}(\nu > 1)\mathbb{E}(D) &= \mathbb{P}(\nu > 1) \int_0^\infty \mathbb{P}(D > x) dx \\ &= \int_0^\infty (1 - e^{-\mu y} - \mu y e^{-\mu y}) e^{-(y/\eta)^\beta} dy. \end{aligned}$$

Equation (10) follows by using the expression of $\mathbb{P}(\nu > n)$ given by Equation (2) and the fact that $\mathbb{E}(\nu) = \mu\mathbb{E}(S)$.

Concerning the second moment, we have

$$\begin{aligned} \mathbb{E}(D^2) &= 2 \int_0^\infty x\mathbb{P}(D > x) dx \\ &= \frac{2}{\mathbb{P}(\nu > 1)} \int_0^\infty \left(y - \frac{2}{\mu} + \left(\frac{2 + \mu y}{\mu} \right) e^{-\mu y} \right) e^{-(y/\eta)^\beta} dy \\ &= \frac{(\mu^2\mathbb{E}(S^2) - 4\mu\mathbb{E}(S) + 4\mathbb{P}(\nu > 0) + 2\mathbb{P}(\nu > 1))}{\mu^2\mathbb{P}(\nu > 1)}. \end{aligned}$$

By using Equations (3) and (4), Equation (11) follows. \square

Proposition 2. *Let ν denote the number of times a sampled elephant is seen by the packet sampler. Then, we have*

$$\log \mathbb{P}(\nu > n) \sim - \left(\frac{n}{\eta\mu} \right)^\beta \quad (12)$$

as $n \rightarrow \infty$.

Proof. To estimate the behavior of the integral appearing in Equation (2) when n is large, we use Laplace method [14]. If y_n is the root of $f'_n(x) = 0$ with

$$f_n(x) = n \log x - x - (x/\eta\mu)^\beta - \log n!,$$

then clearly (y_n) converges to infinity, since

$$n = y_n + \beta \left(\frac{y_n}{\mu\eta} \right)^\beta. \quad (13)$$

Moreover, by Stirling formula, $n! \sim n^{n+1/2} \sqrt{2\pi} e^{-n}$ when n is large. We therefore deduce that

$$\begin{aligned} e^{f_n(y_n)} &\sim \\ &\frac{1}{\sqrt{2\pi n}} \exp \left(n \log \left(1 - \frac{n - y_n}{n} \right) - (n - y_n) \frac{1 - \beta}{\beta} \right) \\ &\sim \frac{1}{\sqrt{2\pi n}} \exp \left(- \left(\frac{n}{\eta\mu} \right)^\beta \right) \end{aligned}$$

as $n \rightarrow \infty$. Finally, it is easily checked that $f_n''(y_n) < 0$ and that $\sqrt{|f_n^{(2)}(y_n)|}$ is equivalent to $1/\sqrt{n}$ for large n . A direct application of Laplace method then yields Equation (12). \square

4 Validation with actual data

In this section, we consider the traffic trace captured in November 2003 from the France Telecom backbone IP network as described in Section 2. We then perform on this trace a deterministic $1/N$ packet sampling with N set equal to 100. First, we compare the power spectra of the original and sampled bit rate processes, which are displayed in Figure 2. By convention, a bit rate process is evaluated by computing the numbers of bits in successive time intervals of length Δ and by dividing these numbers by Δ ; the power spectrum is then the spectral density of the corresponding time series. In the data reported in this paper, the parameter Δ has been set equal to 100 milliseconds.

We observe from Figure 2 that the power spectrum of the sampled bit rate has a larger noise component and less power in low frequencies when compared with the power spectrum of the original process. These observations are in line with the theory developed in [2] and [15]. In particular, it seems very difficult to extract useful information from the power spectrum of the sampled bit rate by means of digital signaling processing procedures, in particular because of the huge noise component and the low signal to noise ratio. This is why we perform in the following a probabilistic analysis based on the formulas established in the previous section.

In the remainder of this section, only traffic due to data elephants is considered since they offer about 95% of traffic. In order to distinguish between data and ACK elephants, we use the following simple criterion: If a (original or sampled) elephant has a mean packet size less than 80 bytes, this elephant is considered as an ACK elephant and is subsequently not taken into account in the analysis. This criterion is applied to the original as well as sampled trace and allows us to concentrate on the data, which are expected to be generated




Figure 2: Spectral densities of the original and sampled bit rate processes.

by data elephants. By comparing the original and the sampled traces, we observe that 47% of the original data elephants are indeed considered as data elephants in the sampled trace. The threshold of 80 bytes is arbitrary but as shown in [5], it leads to satisfying experimental results.

From now on, we consider only sampled data, which basically correspond to the observations when NetFlow is run in a router, and we analyze sampled data elephants by considering flows with more than one packet and a mean packet size larger than 80 bytes. The sampling rate μ is estimated by using the mean values of the number M of packets and the duration D of a sampled data elephant. Experimental results give $\mathbb{E}[M] = 12.10$ and $\mathbb{E}[D] = 673$ seconds, leading to a value for the parameter μ equal to 0.018 packet/s. Note that this value leads to a theoretical ratio of sampled data to original elephants, namely $\mathbb{P}(\nu > 1)$ defined by equation (2), equal to 54.3%, numerically estimated by using the series (6). This ratio is reasonably close to the empirical value of 47% computed by considering both sampled and original traffic traces.

By using the asymptotic estimate (12), we have for large n

$$\log(-\log \mathbb{P}(\nu > n)) / \log(n) \sim \beta - \beta \log(\mu\eta) / \log(n).$$

The curve $n \rightarrow \log(-\log \mathbb{P}(\nu > n)) / \log(n)$ is displayed in Figure 3. By using a function of the form $f(x) = a + b / \log x$ and a fit algorithm, such the one offered by the Gnuplot software, we can empirically estimate the values of a and b . For large values of n , say, ranging from 95 to 150, we find $a = 0.390718$ and $b = -0.282681$. This yields a good estimate $\hat{\beta} = 0.390718$ for the value of the parameter β . For η , we have the estimate $\hat{\eta} = \exp(-b/a) / \mu = 114.66$, which is close to the actual value of the parameter $\eta = 116$.

To estimate the value of $\mathbb{P}(\nu > 1)$, we use the duration of the sampled data. Indeed, from approximation (9), we have for large x

$$\frac{\mathbb{P}(D > x)}{\exp(-(x/\eta)^\beta)} \sim \frac{1}{\mathbb{P}(\nu > 1)}. \quad (14)$$

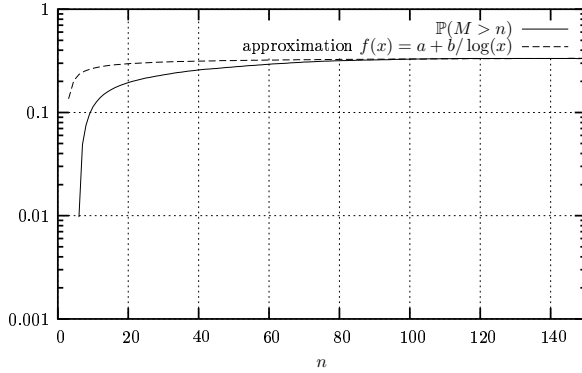


Figure 3: Logarithmic plot of the number of packets in a sampled data elephant and approximation by a function of the form $f(x) = a + b/\log(x)$.

By using the estimates $\hat{\beta}$ and $\hat{\eta}$ for the values of the parameters β and η , Figure 4 displays the ratio (14) and clearly shows that the function $x \rightarrow \mathbb{P}(D > x)/\exp(-(x/\hat{\eta})^{\hat{\beta}})$ tends to the value 2.4, which yields an estimate equal to 41% for the value of $\mathbb{P}(\nu > 1)$, which is reasonably close to the actual value equal to 47%. By using the values of $\hat{\eta}$ and $\hat{\beta}$ as well as the value of μ , a direct estimation of the series (6) yields a value of 53.4% for $\mathbb{P}(\nu > 1)$, which is again reasonably close to the empirical ratio equal to 41%.




Figure 4: Complementary probability distribution function of the duration of a sampled data elephant rescaled by the approximating function $g(x) = \exp(-(x/\hat{\eta})^{\hat{\beta}})$.

To estimate the arrival rate λ of elephants, we use the fact that the number of elephants is large and slowly fluctuating; recall that the number of active elephants is approximated by the number of customers in an $M/G/\infty$ queue with Weibullian service times. Hence,

the probability of picking up an elephant is close to $1/(\lambda\mathbb{E}[S])$. If τ denotes the mean value of the time between two packets of sampled data elephants, then the sampling rate μ of an original elephant is $1/(\tau\lambda\mathbb{E}[S])$. Therefore, we have

$$\lambda = \frac{1}{\mu\tau\eta\Gamma(1 + 1/\beta)}.$$

Experimental data yields $\tau = 0.0041$ seconds. Therefore, by using the estimates $\hat{\eta}$ and $\hat{\beta}$, the arrival rate λ can be estimated by $\hat{\lambda} = 33.29$ elephants/s, which is close to the actual value equal to 35.2 elephants/s.

5 Conclusion

We have shown in this paper that with an a priori knowledge of the form of traffic, it is possible to infer with a reasonable accuracy the characteristics of original flows from $1/N$ sampled data. The method presented in this paper proves efficient for TCP flows composed of more than 20 packets. These flows contribute 95% to the global TCP bit rate and a precise knowledge of their characteristics is necessary for charging issues. Such a method is all the more needed that $1/N$ sampling is performed by NetFlow widely deployed in commercial networks.

The next step is to find estimators for small flows, in particular the detection and the characteristics in terms of IP addresses of mice composed of one or two packets. Such mice typically arise in the case of Distributed Denial of Service attacks. The detection of such mice is then a crucial issue for security purposed. This point will be addressed in further studies.

References

- [1] Nick Duffield, Carsten Lund, and Mikkel Thorup. Properties and prediction of flow statistics properties and prediction of flow statistics. *ACM SIGCOMM Internet Measurement Workshop*, pages 6–8, November 2002.
- [2] Nicolas Hohn and Darryl Veitch. Inverting sampled traffic. *IMC*, pages 27–29, October 2003.
- [3] CISCO. *Netflow services and applications, white paper*. Specification available at <http://www.cisco.com/warp/public/732/netflow/index.html>.
- [4] Nick Duffield, Carsten Lund, and Mikkel Thorup. Estimating flow distributions from sampled flow statistics. *SIGCOMM*, pages 25–29, 2003.
- [5] 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*, To appear 2004.
- [6] K. Tutschku and P. Tran-Gia. A traffic profile of the eDonkey filesharing service. Technical report, COST 279TD(03)049, 2003.
- [7] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of Ethernet traffic. *IEEE/ACM Trans. Net.*, pages 1–15, 1994.
- [8] H.J. Fowler and W.E. Leland. Local area network traffic characteristics, with implications for broadband network congestion management. *IEEE J. Sel. Areas in Commun.*, 9(7):1139–1149, 1994.

-
- [9] M. Crovella and A. Bestavros. Self-similarity in world wide web. Evidence and possible causes. *IEEE/ACM Trans. on Networking*, pages 835–846, December 1997.
 - [10] N. Ben Azzouna and F. Guillemin. Experimental analysis of the impact of peer-to-peer applications on traffic in commercial IP networks. *European Transactions on Telecommunications, Special issues on P2P networking and P2P services*, 15(6), November-December 2004.
 - [11] 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.
 - [12] D. Shanks. Non linear transformations of divergent and slowly converging series. *J. Math. Phys.*, 34:1–42, 1955.
 - [13] P. Wyn. On a device for computing the $e_m(s_n)$ transformation. *MTAC*, 10:91–96, 1956.
 - [14] J. Dieudonné. *Calcul infinitésimal*. Hermann, 1980.
 - [15] P. Brémaud and L. Massoulié. Power spectra of general shot noises and Hawkes point process with random excitation. *Adv. Appl. Prob.*, 34(1):205–222, June 2002.



Unité de recherche INRIA Rocquencourt
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399