Inverting Flow Durations from Sampled Traffic

Nelson Antunes
FCT University of Algarve and CEMAT
Campus de Gambelas
8005-139 Faro, Portugal
Email: nantunes@ualg.pt

Vladas Pipiras
Department of Statistics and Operations Research
University of North Carolina
CB 3260, Chapel Hill, NC 27599, USA
Email: pipiras@email.unc.edu

***Abstract*—The flow duration distribution is an important metric to the network operator for traffic prediction and accounting but also arguably from the viewpoint of the user. The inversion problem of recovering the flow duration distribution from sampled traffic is addressed here under several sampling methods. A theoretical framework for the inversion problem is developed using a probabilistic flow model. In this framework, direct equations for the distributions of sampled flow quantities are derived based on the distributions of original flow characteristics. The inversion of these equations provides estimators for the flow duration distribution. Finally, the inversion techniques under several sampling schemes are evaluated on two Internet traces.**

I. Introduction

A. Motivation and Related Work

Passive measurements are critical to understand the characteristics and the dynamics of the Internet traffic, and are used extensively by network operators and networking researchers. But the exhaustive capture of packets leads rapidly to a huge amount of data to store and to analyze. For instance, a capture of a few hours on a Gigabit/s link at a medium load level yields several hundreds of gigabytes of data. A way to reduce the volume of data is by sampling packets. In this work, we are interested in the flow level characteristics. A flow corresponds to a sequence of packets sharing common attributes (e.g. source/destination addresses, source/destination ports and protocol). The problem is to deduce the structure of original flows from sampled packets.

This topic has attracted much attention recently in the networking community, with the focus almost exclusively on inference of the original distribution of flow sizes (number of packets/bytes in a flow) (see e.g. [1], [2], [3], [4]). The flow duration distribution is another important metric and can yield useful information for traffic prediction, traffic engineering to support QoS and in accounting [5], [6], [7]. Its relations to flow sizes and rates have been studied extensively in e.g. [8]. It is also an important (arguably the most important) performance metric for the user [9]. An accurate estimation of the flow duration distribution is substantially more complex – a reconstruction of the flow duration now involves not only the flow size but also the flow interarrival times (IATs) between consecutive packets. To the best of our knowledge, only [14] recently used inversion to infer the original flow duration distribution from sampled traffic under the packet sampling (PS) method (see below).

B. Sampling Methods

A sampling method has a huge impact on the estimation accuracy of the flow characteristics. A number of sampling methods were proposed in the literature. We briefly review some of them. One of the simplest approaches is the so-called packet sampling (PS), where each packet, independently of others, is captured and analyzed with a fixed probability p_p. As demonstrated in [3], inversion of the flow size distribution is essentially impossible with PS for any small p_p useful in practice. To overcome the limitations of PS method, several authors [1], [2] proposed the use of TCP protocol level details, such as SYN and FIN flags. In the same direction, the use of TCP sequence numbers (SEQ) to infer the number of original packets between two sampled packets of a flow was explored in [10]. Several combinations of these methods are possible e.g. PS+(SYN)+(SEQ). Dual sampling (DS) introduced a hybrid approach combining the advantages of both packet sampling and flow sampling (FS) [11]. Other methods aim at identifying certain types of flows, for example large flows when using sample-and-hold (SH) [12].

C. Aims, Contributions and Outline

The aims of this paper are to provide an analytical framework for recovering the flow duration distribution under various sampling schemes (PS+(SYN)+(SEQ) and DS), and to see whether the techniques work in practice with real traces. The analytical framework is developed assuming that the packet arrivals within original flows follow a finite renewal process (proposed in [13] and also checked on real traces below). Assuming this model, the flow duration distribution can be expressed in terms of the flow size distribution and the distribution of flow IATs (between consecutive packets). The estimation of the flow size distribution under the considered sampling schemes has already been derived in [11]. In this work, we provide ways to estimate the distribution of flow IATs. The original packets of a flow which are sampled with the times between sampled packets lead to a sampled flow. We first derive direct equations that relate the distributions of sampled flow times quantities (IATs and duration) in terms of the distribution of flow IATs. The inversion of direct equations allows us to express and estimate the distribution of flow duration in terms of the *sampled flow* IATs or the *sampled flow duration* distributions.

The proposed estimation techniques are checked on two Internet traces. We find that the flow duration distribution can be recovered successfully when the *sampled flow* IATs are used in inversion under DS, supposing that the probability of sampling packets within a flow is large. This is a plausible scenario on certain links and architectures (see Sec. VI-A in [11]). When inversion is used in other cases (namely, using *sampled flow* IATs under DS with small probability or PS+SYN+SEQ, or using *sampled flow duration* instead), the flow duration distribution can be estimated well only until certain time, after which the estimators inherently diverge due to the reversion operation. Alternatively, for the sampling methods DS and PS+(SYN)+SEQ, the whole duration distribution can be recovered when using the *sampled flow* IATs from consecutive original packets of a flow (inferred from TCP sequence numbers) – though in some cases this estimator can be less robust. Also, under PS+SEQ, the flow duration estimation is poor due to the inaccurate estimation of the flow size distribution. Finally, PS+(SYN) can not estimate the flow size distribution for any useful sampling probability ($\leq 1\%$) which turns the inversion of the flow duration distribution essentially impossible.

The paper is organized as follows. Sec. II introduces the probabilistic flow model considered and discusses its adequacy based on two Internet traces. Sec. III reviews the various sampling methods and defines the variables of interest from sampled flows. Sec. IV derives direct equations (and their inversions) relating the distributions of sampled flows to those of original flows under the various sampling methods. The implementation of the inverse equations is also discussed. Sec. V applies the inversion techniques to two Internet traces. Finally, Sec. VI concludes.

II. Modeling traffic flows

A. Flows and Their Characteristics

A flow is defined as a unidirectional sequence of packets that have the same 5-tuple: IP source address, destination address, source port, destination port and protocol type, and where packet interarrival times (IATs) in a flow do not exceed a timeout of 60 seconds.

A flow is characterized by the IATs (between consecutive packets), and its size (the number of packets). Let W be the size of a flow, and D_i, $i = 1, \ldots, W-1$, be the IATs between the ith and $(i+1)$th packets of a flow. The duration of a flow is given by

$$V = \sum_{k=1}^{W-1} D_k.$$

B. Probabilistic Model of Flows

We will assume the following probabilistic model for flows. The size W of a flow has a probability mass function (p.m.f.)

$$f_W(w), \quad w \geq 1.$$

The IATs D_i, $i = 1, \ldots, W-1$, between the packets in a flow are i.i.d. random variables, independent of W. We will refer to this model as *Flow Model*. Another common name is a finite renewal process.

Trace	Date	Local start time	Duration	TCP Flows
Auckland IX	20080327	12:00:01	1:00:00	1,371,756
Waikato V	20070626	21:00:01	1:40:00	842,578

TABLE I
Traces Summary

The common cumulative distribution function (c.d.f.) of IATs will be denoted

$$F_D(t) = P(D \leq t).$$

The (conditional) c.d.f. of the total duration V will be denoted

$$F_V(t) = P(V \leq t | W \geq 2).$$

In a measurement interval, flows are supposed to be independent copies drawn from Flow Model. This is in contrast to [11] where the focus is on sizes of flows, and where randomness is induced only by a sampling procedure (on the other hand, the approach analogous to here is taken in [3], [14]). We do not believe we can start without assuming a probabilistic model for flows. In particular, our inference methods about distributions of flow times (i.e. IATs and duration) will be based on Flow Model.

Flow Model provides the simplest way of modeling flows. This model appears in [13], as part of a more general flow-based model for cumulative packet traffic.

C. Data Traces and Goodness of Fit of Flow Model

We use two recent publicly available Internet traces, Auckland IX [15] and Waikato V [16]. The binary traces were processed with Libtrace [17] and Matlab to extract IP packet headers and timestamps in order to reconstruct flows. The traces are summarized in Table I. Since most of the sampling methods use the TCP protocol information, we focus only on TCP flows, as they account for 80-90% of packets in the Internet [10].

Does Flow Model provide a good fit with real data traces? Several exploratory tools to evaluate Flow Model were suggested in [3]. We briefly illustrate one of these tools here first. Note that $V/(W-1)$, $W \geq 2$, can be thought as the average IAT between packets of a flow. Figure 1 depicts W against the average IAT for Auckland IX (top plot). The level corresponds to the flow density (number of flows in each small square). A logarithmic scale is used to accentuate the differences. Most flows have a small number of packets where the density is highest. For small flow sizes the average IAT covers a wide range of values. By independence assumptions of Flow Model, the average IAT should concentrate around the value where the density is higher as the flow size increases. In the top plot shown, however, large flows have shorter IATs (i.e. the top of the triangle is to the left of the bright area with the highest density), indicating small deviations from Flow Model. As a consequence, inference of the distribution of IATs may be susceptible to whether packets are more likely to be selected from these flows according to a particular sampling scheme. Figure 1 provides the density plot for Waikato V (bottom plot) trace which is consistent with Flow Model.

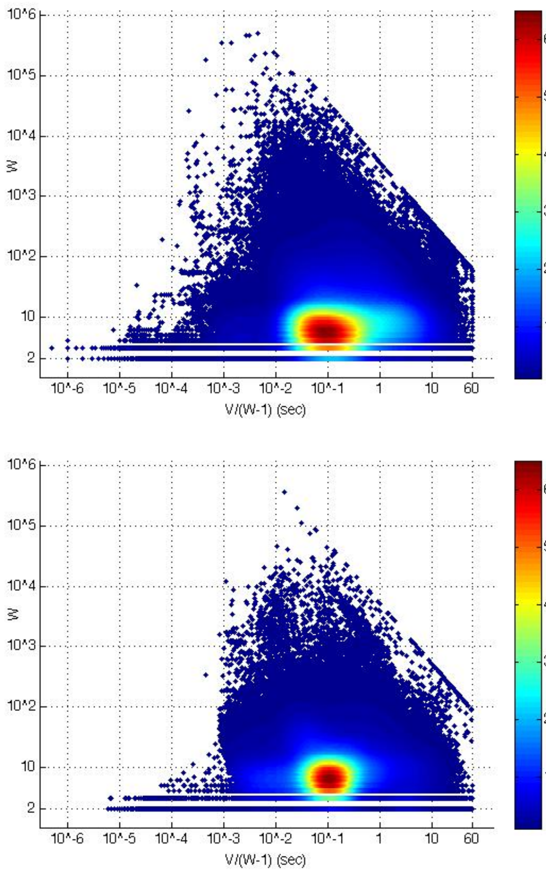

Fig. 1. Flow density for flow size against the average IAT. Auckland IX (top), Waikato V (bottom).

III. Sampling Framework

We recall here some sampling schemes (PS and some PS-based methods) introduced in the networking literature (Sec. III-A), and discuss the variables of interest related to sampled flows (Sec. III-B).

A. Sampling Schemes

Packet sampling (PS): In this sampling scheme, each packet of a flow is sampled independently of the other packets with probability p_p (see e.g. [2], [3]). Its simplicity makes it appropriate for high-speed routers.

PS+SEQ: Under this sampling method, packets are sampled as in PS. But TCP sequence numbers are now used as follows. Given two sampled packets of the same flow with sequence numbers s_l and s_h, we assume that there are $s_h - s_l$ original packets between the two sampled packets [10], [11]. In practice TCP sequence numbers count payload bytes and a mapping is needed to return the number of packet counts between two sampled packets. See [10], [11] for implementation with real data. This additional information helps inferring the characteristics of flows, at the expense of more intensive processing.

PS+SYN: The method is the same as PS except that only sampled flows with a TCP SYN (or first) packet are retained [2]. The idea is to sample flows independently of their size, thus avoiding the flow size bias characteristic to PS. The disadvantage is that many sampled packets (from flows without SYN packets) are discarded.

PS+SYN+SEQ: Under this method, PS+SYN is performed with parameter p_p and then the sequence numbers of sampled packets of a flow are used to infer additional information as in PS+SEQ [10]. The method combines (dis)advantages of both PS+SYN and PS+SEQ methods.

Dual sampling (DS): The method consists of two PS schemes, one operating on SYN packets with sampling probability p_f and the other on non-SYN packets with probability p_p [11]. Sampled flows without SYN packets are discarded and the sequence numbers are used as before. The method includes PS+SYN+SEQ when $p_p = p_f$. It also includes the so-called flow sampling (FS) ($p_p = 1$) which is infeasible in practice and can be thought of as the ideal type of sampling. By keeping p_f small while increasing p_p, DS performance approaches that of FS.

Combinations of methods with +FIN are also possible. Here, FIN refers to retaining only those flows with their last packet (the so-called FIN packet) sampled. In practice, however and unlike SYN, a considerable proportion of flows do not terminate with FIN [11]. For this reason, we will not consider any sampling methods involving FIN and other methods proposed for some type of flows (e.g. sample-and-hold [12]).

With the exception of PS, the sampling methods can only be applied to TCP flows, or to other objects provided that suitable substitutes could be found for connection startup (SYN) and progress (sequence numbers). The use of additional information is critical since the performance of PS is notoriously poor [3]. In the remaining of this work, we focus on the following sampling schemes: PS+(SYN)+(SEQ) and DS.

B. Variables of Interest and Their Distributions

Any of the sampling procedures described above leads to a sampled flow. It is convenient to introduce the following variables for a sampled flow. Let $D_{p,i}$ be the time between the ith and $(i+1)$th sampled packets, $M_{p,i}$ be the number of original packets between the ith and $(i+1)$th sampled packets (including these two packets) and V_p be the duration of a sampled flow. Let W_p be the number of packets in a sampled flow, and we refer to it as the *sampled flow size*. For PS+(SYN) methods, W_p is just the number of sampled packets. For PS+(SYN)+SEQ and DS (i.e. SEQ methods), W_p is the number of original packets in a sampled flow, which can be deduced from the sequence numbers of the first and last packets sampled. Figure 2 illustrates the introduced notation for SYN and SEQ sampling procedures.

In terms of sampling distributions, we write

$$f_{W_p}(s) = P(W_p = s), \qquad s \geq 0,$$

for the p.m.f. of the sampled flow size. The (conditional) c.d.f. of the duration of a sampled flow will be denoted

$$F_{V_p}(t) = P(V_p \leq t | W_p \geq 2). \tag{1}$$

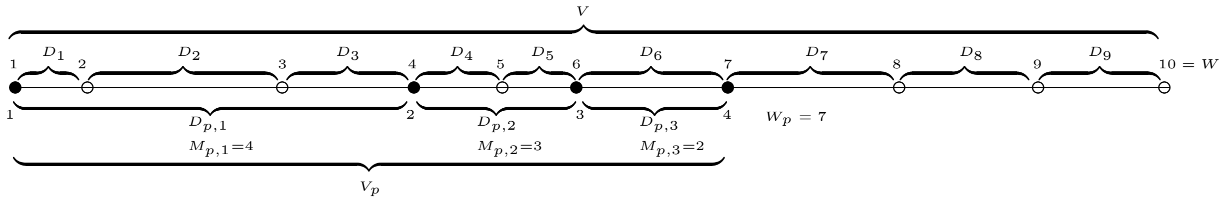

Fig. 2. Flow packets (circles) and sampled flow packets (filled circles) for a SYN and SEQ sampling method.

It can be shown that for the sampling schemes considered in this work, the c.d.f.'s of $D_{p,i}$ given $M_{p,i}$ will not depend on i (see [14] for a proof under PS). For this reason, we will also write D_p and M_p for the variables denoting the time and the number of original packets between any two consecutive sampled packets. (D_p should not be confused with the original IAT D.) We also set

$$F_{D_p}(t) = P(D_p \leq t | M_p \geq 2), \tag{2}$$

$$F_{D_p|2}(t) = P(D_p \leq t | M_p = 2) = F_D(t), \tag{3}$$

where the last equality is due to the conditioning on $M_p = 2$.

IV. Inference about Distributions of a Flow

Distributions of original flows are inferred by relating them to distributions of sampled flows. This is carried out by finding first direct equations expressing the distributions of sampled variables in terms of those of original flows (Sec. IV-B), and then inverting these equations (Sec. IV-C and IV-D). We begin by introducing the notation used below.

A. Notation

For a c.d.f. F of a non-negative random variable, we write $\widetilde{F}(v) = \int_0^t e^{-vt} F(dt)$ for its Laplace transform, and F^{*k} for its kth convolution. The indicator function of a set S will be denoted 1_S. For a sequence $A = \{A_k\}_{k\geq 1}$, we write $G_A(z) = \sum_{k=1}^{\infty} A_k z^k$ for its power series. We write G_A^{-1} for the inverse function of G_A satisfying $G_A^{-1}(G_A(z)) = z$. It is known that $G_A^{-1}(w) = G_a(w)$ where the sequence $a = \{a_k\}_{k\geq 1}$ is obtained from the sequence A by the reversion operation [18]. This is a standard operation for sequences $A = \{A_k\}_{k\geq 1}$ with $A_1 \neq 0$: the reversion sequence $a = \{a_k\}_{k\geq 1}$ of A satisfies, for example, $a_1 = 1/A_1$, $a_2 = -A_2/A_1^3$, $a_3 = (2A_2^2 - A_1 A_3)/A_1^5$, etc. In fact, the reversion operation should be used with caution: the equality $G_A^{-1}(w) = G_a(w)$ may hold only for small enough w, and

$$G_a(G_A(z)) = z \tag{4}$$

may also be true for sufficiently small z only. This is related to the radius of convergence of the power series $G_a(w)$. In our case A is a p.m.f. and hence the radius of convergence of $G_A(z)$ is at least 1. The radius of convergence of $G_a(w)$, however, will not necessarily be at least 1. Due to the reversion formula for the coefficients a_k's, one generally finds that the larger the value of A_1 ($0 < A_1 < 1$), the larger the radius of convergence of $G_a(w)$.

B. Direct Equations

The distributions of sampled variables (Section III-B) can be related to those of the original flows through direct equations.

Sampled flow size: A direct equation relating f_{W_p} and f_W reads [11]:

$$f_{W_p}(s) = \sum_{w=1}^{\infty} h_{sw} f_W(w), \quad s \geq 0, \tag{5}$$

where $h_{sw} = P(W_p = s | W = w)$. This is convenient to write in a matrix form as

$$\mathbf{f}_{W_p} = \mathbf{H}\mathbf{f}_W, \tag{6}$$

where $\mathbf{f}_{W_p} = (f_{W_p}(0), f_{W_p}(1), \ldots)^T$, $\mathbf{f}_W = (f_W(1), f_W(2), \ldots)^T$ and $\mathbf{H} = (h_{sw},\, s \geq 0, w \geq 1)$. Expressions for $\mathbf{H}$ under the various sampling schemes can be found in [11]. For example, for PS, $h_{sw} = \binom{w}{s} p_p^s q_p^{w-s}$ with $q_p = 1 - p_p$, for $s = 0, \ldots, w$, and $h_{sw} = 0$ for $s \geq w + 1$.

Sampled flow IATs: Turning to the distributions of flow times, consider the general distribution of the time D_p between two consecutive sampled packets. Since

$$D_p \stackrel{d}{=} \sum_{k=1}^{M_p - 1} D_k, \tag{7}$$

we obtain that

$$\widetilde{F}_{D_p}(v) = G_{A_p}(\widetilde{F}_D(v)), \quad v > 0, \tag{8}$$

where $A_p = \{A_{p,k}\}_{k\geq 1}$ is given by

$$A_{p,k} = P(M_p = k + 1 | M_p \geq 2), \quad k \geq 1. \tag{9}$$

Note that M_p is observed in PS+(SYN)+SEQ and DS through sequence numbers. For PS+(SYN), M_p is not observed and A_p is expressed in terms of the distributions of W and W_p in [14].

Alternatively, only for SEQ methods, we can write by (3)

$$\widetilde{F}_{D_p|2}(v) = \widetilde{F}_D(v), \quad v > 0. \tag{10}$$

Sampled flow duration: Recall that for SEQ methods, W_p is the number of original packets in a sampled flow. Thus, for these methods the duration distribution of a sampled flow is given by

$$V_p \stackrel{d}{=} \sum_{k=1}^{W_p - 1} D_k. \tag{11}$$

Hence,

$$\widetilde{F}_{V_p}(v) = G_{B_p}(\widetilde{F}_D(v)), \quad v > 0, \tag{12}$$

where $B_p = \{B_{p,k}\}_{k\geq 1}$ is given by

$$B_{p,k} = P(W_p = k+1 | W_p \geq 2), \quad k \geq 1. \qquad (13)$$

For PS+(SYN), W_p is just the number of packets sampled. In this case, Eq. (12) holds where B_p, the distribution of the number of original packets in a sampled flow (not observed), is expressed in terms of the distributions of W and W_p in [14].

C. Inverse Equations

The direct equations relating the sampled quantities to the original ones can be inverted.

Flow sizes: The inversion of Eq. (5) is discussed extensively in [11]. Set $\widetilde{\mathbf{f}}_{W_p} = (f_{W_p}(1), f_{W_p}(2), \ldots)^T$ and let $\widetilde{\mathbf{H}}$ be the submatrix of $\mathbf{H}$ so that

$$\widetilde{\mathbf{f}}_{W_p} = \widetilde{\mathbf{H}}\mathbf{f}_W. \qquad (14)$$

One then expects that

$$\mathbf{f}_W = \widetilde{\mathbf{H}}^{-1}\widetilde{\mathbf{f}}_{W_p}. \qquad (15)$$

Explicit forms of $\widetilde{\mathbf{H}}^{-1}$ for the various sampling schemes can be found in [11].

Flow IATs: Turning to the distributions of original flow times, Eq. (8) can be inverted, relating the original flow IATs with the sampled IATs, as

$$\widetilde{F}_D(v) = G_{a_p}(\widetilde{F}_{D_p}(v)), \quad v > 0, \qquad (16)$$

where $G_{a_p} = G_{A_p}^{-1}$ is such that

$$G_{a_p}(G_{A_p}(z)) = z \qquad (17)$$

and the sequence $a_p = \{a_{p,k}\}_{k\geq 1}$ is obtained from the sequence $A_p = \{A_{p,k}\}_{k\geq 1}$ by the reversion operation (see the discussion around Eq. (4)). No inversion is needed for Eq. (10). Similarly, Eq. (12) can be inverted, relating the original flow IATs with the sampled duration, as

$$\widetilde{F}_D(v) = G_{b_p}(\widetilde{F}_{V_p}(v)), \quad v > 0, \qquad (18)$$

where $G_{b_p} = G_{B_p}^{-1}$.

Flow duration: Arguing as for Eq. (12), note that

$$\widetilde{F}_V(v) = G_C(\widetilde{F}_D(v)), \quad v > 0, \qquad (19)$$

where $C = \{C_k\}_{k\geq 1}$ is given by

$$C_k = P(W = k+1 | W \geq 2), \quad k \geq 1. \qquad (20)$$

Substituting Eq. (18) into Eq. (19) leads to an equation relating the original flow duration with the sampled duration, as

$$\widetilde{F}_V(v) = G_C(G_{b_p}(\widetilde{F}_{V_p}(v))), \quad v > 0. \qquad (21)$$

Substituting Eq. (16) into Eq. (19) similarly leads to an expression of the original flow duration in terms of the sampled IATs, as

$$\widetilde{F}_V(v) = G_C(G_{a_p}(\widetilde{F}_{D_p}(v))), \quad v > 0. \qquad (22)$$

By Eq. (10), Eq. (19) can also be written, relating the original flow duration with the sampled IATs from consecutive original packets of a flow, as

$$\widetilde{F}_V(v) = G_C(\widetilde{F}_{D_p|2}(v)), \quad v > 0. \qquad (23)$$

D. Inference in Practice

The inverse equations discussed in Sec. IV-C can be used to estimate the distributions of original flows. For simplicity, we only discuss flow times distributions under SEQ methods (PS+(SYN) are implemented analogously using the cited results in [14]). Let $W_{p,n}$, $D_{p,i,n}$ and $M_{p,i,n}$ be the respective variables W_p, $D_{p,i}$ and $M_{p,i}$ for the sampled flows $n = 1, \ldots, N$. As in [11], N includes the number of not sampled flows and is assumed to be known. In practice, N is estimated by counting the number of SYN packets (see Sec. II-C in [11]).

Flow sizes: For flow sizes, $f_{W_p}(s)$ entering $\widetilde{\mathbf{f}}_{W_p}$ in Eq. (15) is replaced by its empirical counterpart $\sum_n 1_{\{W_{p,n}=s\}}/N$.

Flow IATs: The simplest estimator for the distribution F_D of IATs is obtained by implementing Eq. (10) in the time domain (i.e. Eq. 3), namely,

$$\widehat{F}_D(t) = \frac{\sum_{n=1}^{N}\sum_{i=1}^{I} 1_{\{D_{p,i,n}\leq t, M_{p,i,n}=2\}}}{\sum_{n=1}^{N}\sum_{i=1}^{I} 1_{\{M_{p,i,n}=2\}}}. \qquad (24)$$

In applications to the real traces considered below, smaller I may perform better (see Sec. V).

The distribution F_D could also be estimated by implementing Eq. (16) in the time domain as

$$\widehat{F}_D(t) = \sum_{k=1}^{K} \widehat{a}_{p,k}\widehat{F}_{D_p}^{*k}(t), \qquad (25)$$

where hats indicate estimators and K is a truncation parameter. In view of Eq. (2), we take

$$\widehat{F}_{D_p}(t) = \frac{\sum_{n=1}^{N}\sum_{i=1}^{I} 1_{\{D_{p,i,n}\leq t, M_{p,i,n}\geq 2\}}}{\sum_{n=1}^{N}\sum_{i=1}^{I} 1_{\{M_{p,i,n}\geq 2\}}}. \qquad (26)$$

The coefficients $\widehat{a}_{p,k}$ in Eq. (25) are obtained through the reversion (see Sec. IV-A) of the coefficients $\widehat{A}_{p,k}$, which estimate $A_{p,k}$ in Eq. (9) through

$$\widehat{A}_{p,k} = \frac{\sum_{n=1}^{N}\sum_{i=1}^{I} 1_{\{M_{p,i,n}=k+1, M_{p,i,n}\geq 2\}}}{\sum_{n=1}^{N}\sum_{i=1}^{I} 1_{\{M_{p,i,n}\geq 2\}}}. \qquad (27)$$

Because of the potential convergence problems with the reversion (see Sec. IV-A), the estimator (25) may diverge for larger values of t. The larger the value of $\widehat{A}_{p,1}$, the larger the values of t where the estimator (25) diverges.

Remark: Another natural possibility to implement Eq. (16) is through the inverse Laplace transform (ILT). A number of ILT procedures are available in the literature (e.g. [19], [20]). We have implemented Eq. (16) through the ILT but generally found the results to be worse than when using Eq. (25).

Similarly Eq. (18) can be implemented as:

$$\widehat{F}_D(t) = \sum_{k=1}^{K} \widehat{b}_{p,k}\widehat{F}_{V_p}^{*k}(t), \qquad (28)$$

where $\widehat{F}_{V_p}$ is the empirical counterpart of F_{V_p} in Eq. (1) given by

$$\widehat{F}_{V_p}(t) = \frac{\sum_{n=1}^{N} 1_{\{V_{p,n}\leq t, W_{p,n}\geq 2\}}}{\sum_{n=1}^{N} 1_{\{W_{p,n}\geq 2\}}}, \qquad (29)$$

Sampling	Normalization
PS	$p_p = p$
PS + SYN	$p_p = (-1 + \sqrt{1 + 4pE[W](E[W]-1)})/(2(E[W]-1))$
DS	$p_f = pE[W]/(p_p(E[W]-1)+1)$

TABLE II
AVERAGE SAMPLING RATES AND NORMALIZATIONS.

and $\{\widehat{b}_{p,k}\}_{k\geq 1}$ is the reversion of the sequence $\{\widehat{B}_{p,k}\}_{k\geq 1}$ estimating $\{B_{p,k}\}_{k\geq 1}$ in Eq. (13) through

$$\widehat{B}_{p,k} = \frac{\sum_{n=1}^{N} 1_{\{W_{p,n}=k+1, W_{p,n}\geq 2\}}}{\sum_{n=1}^{N} 1_{\{W_{p,n}\geq 2\}}}. \qquad (30)$$

Flow duration: For flow duration, Eq. (19) is implemented as

$$\widehat{F}_V(t) = \sum_{k=1}^{K} \widehat{C}_k \widehat{F}_D^{*k}(t), \qquad (31)$$

where $\widehat{C}_k$ estimates C_k in Eq. (20) by using the estimator of f_W, and $\widehat{F}_D$ is one of the estimators of F_D discussed above, namely, (24), (25) or (28). This corresponds to implementing, respectively, Eqs. (23), (22) and (21).

V. EXPERIMENTAL INVERSION ON INTERNET DATA

We evaluate here the inversion techniques of Sec. IV-D to estimate the flow duration distribution under the different sampling methods for Auckland IX and Waikato V traces.

A. Normalization

To compare across the sampling methods, we use the Effective Sampling Rate (ESR) normalization proposed in [11]. A normalization that just equals the average sampling rate greatly disadvantages PS+SYN+(SEQ) and DS sampling methods which discard sampled packets from flows without a SYN packet. Let p be the average rate of packets which are actually used by the method in the estimation. Table II gives the normalization parameters for each sampling method given a fixed p value (see [11] for derivation). DS has two parameters with p_p chosen to be the independent one. For a fixed p, this equation gives p_f as a monotonically decreasing function of p_p. Ideally, $p_f \ll p_p$ and p_p should be as large as possible to achieve the FS performance (see [11]). Note also that the normalization parameters for X+SEQ and X (e.g. X=PS+(SYN)) are the same. In the following, we set $p = 0.01$ for all sampling methods.

B. Goodness of Fit of Flow Model from Sampled Data

The original flows were used in Sec. II-C to check the goodness of fit of Flow Model. It is natural to ask whether the same can be done with the sampled flows. For the sampling methods that use the SEQ information, the sampled flow size W_p can be plotted against the correspondent average sampled IAT $V_p/(W_p - 1)$. The density of sampled flows over the range of values is shown in Fig. 3 for DS with $(p_f = 0.044, p_p = 0.2)$. Our conclusions about the adequacy of Flow Model from Sec. II-C are not altered, with Auckland IX (top plot) having some large flows with shorter average IATs.

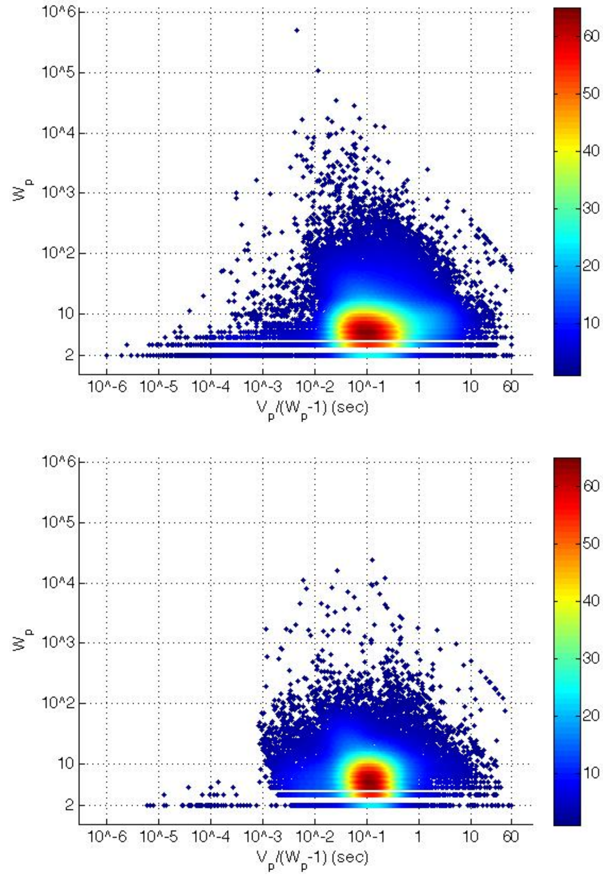

Fig. 3. Sampled flow density of the sampled flow size against average sampled IAT. Auckland IX (top). Waikato V (bottom).

C. Estimation of Flow Size Distribution

Figure 4 depicts the estimation of the flow size distribution C (with at least two packets – see Eq. (20)) obtained using the estimator of f_W for Auckland IX. The estimate $\widehat{C}$ will enter in the estimation of the flow duration distribution through Eq. (31). DS $(p_f = 0.013, p_p = 0.75)$ gives the better agreement between the estimated and original distributions. DS $(0.044, 0.2)$ shows more variation due to the decrease of p_p. PS+SYN+SEQ $(p_p = 0.084)$ also presents a reasonable fit. The performance of PS+SEQ $(p_p = 0.01)$ is worse where the discontinuous parts indicate values out of the interval $[0, 1]$. PS+SEQ tends to skip small flows and display a flow size bias. All the other methods (DS and PS+SYN+SEQ) which use SYN packets, attenuate the flow size bias by keeping flows independent of their size. PS $(p_p = 0.01)$ and PS+SYN $(p_p = 0.084)$ are not represented since all the values of the inversion are out of probability range and will not be considered below. We omit the estimate $\widehat{C}$ for Waikato V trace. Similar conclusions can be drawn.

D. Estimation of Flow Duration Distribution

Figure 5 depicts the performance of the estimator $\widehat{F}_V$ in Eq. (31) when $\widehat{F}_D$ is defined in terms of $\widehat{F}_{V_p}$ in Eq. (28) for Auckland IX. DS and PS+SYN+SEQ show a good accuracy until the estimator diverges. This is related to the convergence

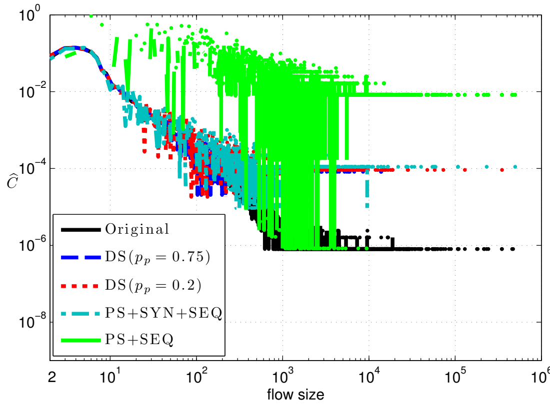

Fig. 4. Flow size distribution C (Auckland IX).

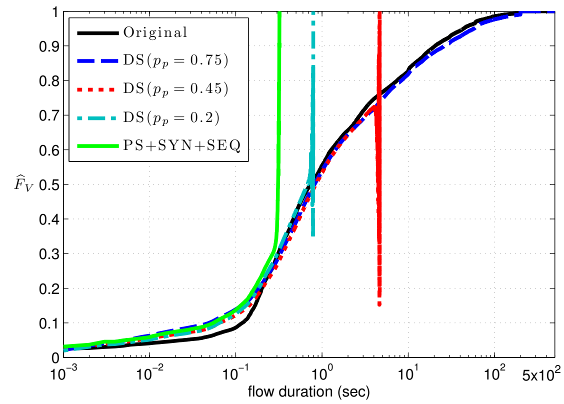

Fig. 6. Flow duration distribution using $\widehat{F}_{D_p}$ (Auckland IX).

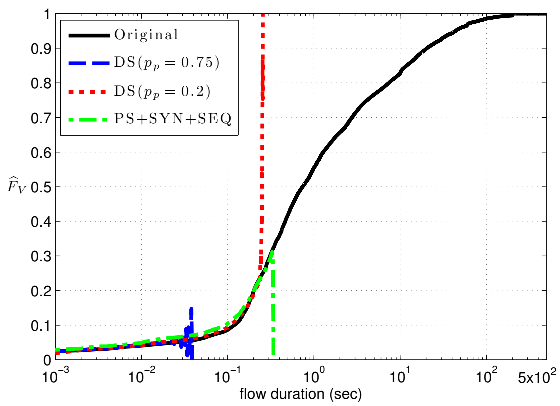

Fig. 5. Flow duration distribution using $\widehat{F}_{V_p}$ (Auckland IX).

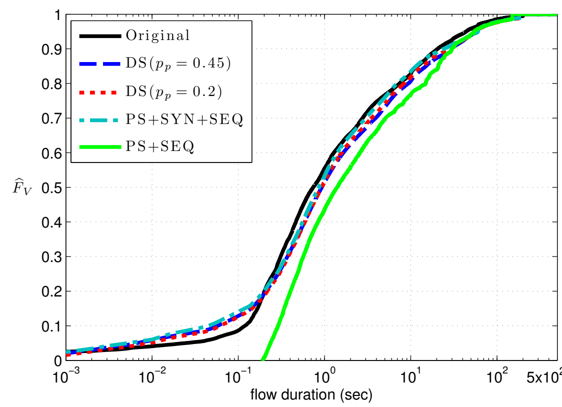

Fig. 7. Flow duration distribution using $\widehat{F}_{D_p|2}$ (Auckland IX).

of the estimator (28) where the coefficients $\widehat{b}_{p,k}$'s are obtained from the reversion of $\widehat{B}_{p,k}$'s. Recall that $\widehat{B}_{p,k}$'s estimate the distribution of the number of original packets in a sampled flow (with at least two packets). As indicated in Sec. IV-A, we generally found that the larger the first coefficient of the sequence to reverse, the larger the values of t where the estimator (28) diverges. We also found that $\widehat{B}_{p,1}$ tends to be smaller and slightly increases when p_p decreases. This explain the better performance of DS with $p_p = 0.2$ and that of PS+SYN+SEQ ($p_p = 0.084$) in the sense that the divergence takes place later.

Figure 6 illustrates the behavior of the estimator $\widehat{F}_V$ in Eq. (31) when $\widehat{F}_D$ is defined in terms of $\widehat{F}_{D_p}$ in Eq. (25) for Auckland IX. (Because of the large flows with shorter IATs found in this trace, we restrict the number of sampled packets from the same flow used in Eqs. (26) and (27) by choosing I for which $\sum_{i=1}^{I} M_{p,i} \leq 10^3$.) DS ($p_p = 0.75$) allows to recover the whole distribution of flow duration. Compared to Fig. 5, there is an improvement with DS ($p_p = 0.2$). DS $(0.0213, 0.45)$ allows to recover a larger range of values of the flow duration distribution. PS+SYN+SEQ presents a similar performance. The coefficients $\widehat{a}_{p,k}$'s in Eq. (25) are obtained through the reversion of $\widehat{A}_{p,k}$'s. Since $\widehat{A}_{p,k}$'s estimate the distribution of the number of original packets between two sampled packets (including these two), a larger probability p_p of sampling packets within a flow will result in a larger first coefficient $\widehat{A}_{p,1}$. Therefore, as p_p increases, the estimator (25) will diverge later or even converge.

Finally, Fig. 7 shows the performance of the estimator $\widehat{F}_V$ in (31) when $\widehat{F}_D$ is defined in Eq. (24) for Auckland IX. No reversion is used in Eq. (24) (though again we restrict the number of sampled packets from the same flow by choosing I as before). This is the only case where PS+SEQ works and is plotted. The poor performance is due to the inaccurate estimation of the flow size distribution $\widehat{C}$ – see Fig. 4. In contrast to Figs. 5 and 6, the whole distribution F_V can now be recovered using DS with $p_p = 0.2, 0.45$ and PS+SYN+SEQ. Note that F_D is estimated in Eq. (24) through $\widehat{F}_{D_p|2}$ (using the sampled IATs from consecutive original packets of a flow). If p_p is very small, the number of these sampled IATs can be insufficient. Moreover, the statistical mixing between them can be unbalanced (in PS+SYN+SEQ more than 50% of sampled IATs from consecutive original packets of flow are from the first and second packets). The inversion techniques behind Figs. 5 and 6 are more robust in these respects.

The performance for Waikato V is shown in Fig. 8. Given the space constraints, we only present results for DS with $(0.041, 0.2)$ and $(0.0131, 0.75)$. However, similar conclusions

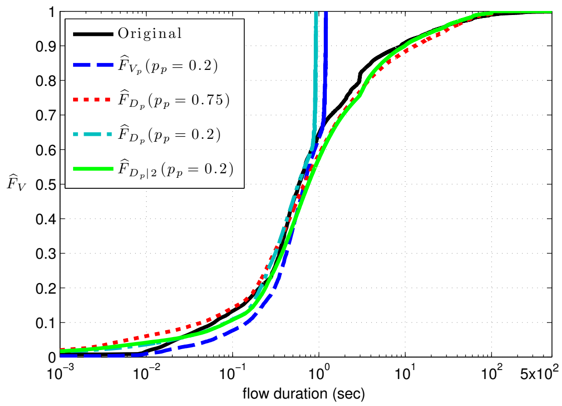

Fig. 8. Flow duration distribution under DS (Waikato V).

as for Auckland IX have been obtained for the others sampling methods. The estimation of flow duration distribution using $\widehat{F}_{V_p}$ $(p_p = 0.2)$ can now recover a larger range of values before the divergence, compared to Fig. 5. The estimation based on $\widehat{F}_{D_p}$ $(p_p = 0.75)$ presents a good accuracy with the original distribution. There is also a slight improvement in the estimation using $\widehat{F}_{D_p}$ $(p_p = 0.2)$, cf. Fig. 6. Finally, under $\widehat{F}_{D_p|2}$ $(p_p = 0.2)$, the estimation does not suffer from the limitations of sequence reversion.

VI. Conclusions

We have explored the question of recovering the distribution of flow duration from sampled data. Several sampling methods were considered: PS+(SYN)+(SEQ) and DS. We developed an exact analytical framework for inversion problem of the flow duration based on a simple flow model. Tools were provided to evaluate deviations from the model when using Internet data. We derived three inverse equations to estimate the flow duration distribution using the distributions of the *sampled flow* IATs or the *sampled flow duration*.

We have checked how successful such inversion techniques were in practice with two real Internet traces. We found that the flow duration distribution could be recovered when the *sampled flow* IATs are used in inversion under DS, given that the probability of sampling packets within a flow is large. This is a feasible parameter setting for DS on certain links and architectures [11], providing a superior statistical performance for DS method. When the *sampled flow* IATs under DS (with small packet sampling probability within a flow) or PS+SYN+SEQ are used, or the *sampled flow duration* is used instead, the flow duration can be estimated well only till certain time, after which the estimators diverge due to the reversion operation. However, under DS and PS+(SYN)+SEQ, the whole duration distribution can be recovered when using the *sampled flow* IATs from consecutive original packets of a flow – though the estimation can be less robust when the sampling probability inside the flow is very low. This was the only case where PS+SEQ worked, but the performance was poor due to inaccuracy of the flow size estimation. Finally, PS+(SYN) turn any inversion impossible for a small sampling probability.

An interesting direction for future work concerns the inference of the flow duration distribution tail without estimating the underlying flow IATs and flow size distributions. A similar study for the flow size distribution tail can be found e.g. in [21].

Acknowledgments

Research is supported in part by Fundação para a Ciência e a Tecnologia (FCT) through project PTDC/EIA-EIA/115988/2009. The authors would also like to thank Andrey Shabalin at University of North Carolina for his help with Matlab code in processing Internet traces.

References

[1] N. Duffield, C. Lund, and M. Thorup, "Properties and prediction of flow statistics from sampled packet streams," in *Proc. ACM SIGCOMM Workshop on Internet measurement*, 2002, pp. 159–171.

[2] ——, "Estimating flow distributions from sampled flow statistics," *IEEE/ACM Trans. Netw.*, vol. 13, pp. 933–946, Oct. 2005.

[3] N. Hohn and D. Veitch, "Inverting sampled traffic," *IEEE/ACM Trans. Netw.*, vol. 14, no. 1, pp. 68–80, Feb. 2006.

[4] L. Yang and G. Michailidis, "Sampled based estimation of network traffic flow characteristics," in *Proc. INFOCOM*, 2007, pp. 1775 –1783.

[5] N. Brownlee and K. Claffy, "Understanding Internet traffic streams: dragonflies and tortoises," *IEEE Commun. Mag.*, vol. 40, no. 10, pp. 110–117, Oct. 2002.

[6] N. Brownlee, "Some observations of Internet stream lifetimes," in *Proc. PAM*, 2005, pp. 265–277.

[7] L. Quan and J. Heidemann, "On the characteristics and reasons of long-lived Internet flows," in *Proc. IMC*, 2010, pp. 444–450.

[8] C. Park, F. Hernández-Campos, J. S. Marron, K. Jeffay, and F. D. Smith, "Analysis of dependence among size, rate and duration in Internet flows," *Ann. Appl. Stat.*, vol. 4, no. 1, pp. 26–52, 2010.

[9] N. Dukkipati and N. McKeown, "Why flow-completion time is the right metric for congestion control," *SIGCOMM Comput. Commun. Rev.*, vol. 36, pp. 59–62, January 2006.

[10] B. Ribeiro, D. Towsley, T. Ye, and J. C. Bolot, "Fisher information of sampled packets: an application to flow size estimation," in *Proc. ACM SIGCOMM*, 2006, pp. 15–26.

[11] P. Tune and D. Veitch, "Fisher information in flow size estimation," *IEEE Trans. Info. Theory*, vol. 57, no. 10, pp. 7011–7035, Oct. 2011.

[12] C. Estan and G. Varghese, "New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice," *ACM Trans. Comput. Syst.*, vol. 21, pp. 270–313, Aug. 2003.

[13] N. Hohn, D. Veitch, and P. Abry, "Cluster processes: a natural language for network traffic," *IEEE Trans. Signal Process.*, vol. 51, no. 8, pp. 2229–2244, Aug. 2003.

[14] N. Antunes and V. Pipiras, "Probabilistic sampling of finite renewal processes," *Bernoulli*, vol. 17, pp. 1285–1326, 2011.

[15] Auckland IX, file 20080327-080000-0. [Online]. Available: http://http://wand.net.nz/wits/auck/9/

[16] Waikato V, file 20070626-000000-0. [Online]. Available: http://wand.net.nz/wits/waikato/5/

[17] Libtrace. [Online]. Available: http://research.wand.net.nz/software/libtrace.php

[18] P. Henrici, *Applied and computational complex analysis*. New York: Wiley-Interscience [John Wiley & Sons], 1974, volume 1: Power series—integration—conformal mapping—location of zeros, Pure and Applied Mathematics.

[19] J. Abate and W. Whitt, "The Fourier-series method for inverting transforms of probability distributions," *Queueing Systems Theory Appl.*, vol. 10, no. 1-2, pp. 5–87, 1992.

[20] ——, "A unified framework for numerically inverting Laplace transforms," *INFORMS J. Comput.*, vol. 18, no. 4, pp. 408–421, 2006.

[21] P. Loiseau, P. Gonçalves, S. Girard, F. Forbes, and P. V. Primet, "Maximum likelihood estimation of the flow size distribution tail index from sampled packet data," in *Proc. ACM SIGMETRICS*, 2009, pp. 263–274.