

Charging from Sampled Network Usage

Nick Duffield Carsten Lund Mikkel Thorup
AT&T Labs–Research
180 Park Avenue, Florham Park, NJ 07932, USA
E-mail: {duffield,lund,mthorup}@research.att.com

Abstract—IP flows have heavy-tailed packet and byte size distributions. This makes them poor candidates for uniform sampling—i.e. selecting 1 in N flows—since omission or inclusion of a large flow can have a large effect on estimated total traffic. Flows selected in this manner are thus unsuitable for use in usage sensitive billing. We propose instead using a size-dependent sampling scheme which gives priority to the larger contributions to customer usage. This turns the heavy tails to our advantage; we can obtain accurate estimates of customer usage from a relatively small number of important samples.

The sampling scheme allows us to control error when charging is sensitive to estimated usage only above a given base level. A refinement allows us to strictly limit the chance that a customer's estimated usage will exceed their actual usage. Furthermore, we show that a secondary goal, that of controlling the rate at which samples are produced, can be fulfilled provided the billing cycle is sufficiently long. All these claims are supported by experiments on flow traces gathered from a commercial network.

I. INTRODUCTION

A. Background

Network service providers are increasingly employing usage measurements as a component in customer charges; see e.g. [11], [19]. One driver stems from the coarse granularity in the available sizes of access ports into the network. For example, in the sequence OC3 to OC12 to OC48 to OC192, each port has a factor 4 greater capacity than the next smallest. Consider a customer charged only according to the access port size. If their demand is at the upper end of the capacity of its current port, the customer will experience a sharp increase in charges on moving to the next size up. Moreover, much of the additional resources will not be used, at least initially. Usage based charging can avoid such sharp increases by charging users for the bandwidth resources that they consume. Another driver

for usage based charging comes from the fact that in IP networks, bandwidth beyond the access point is, mostly, a shared resource. Users who are aware of the charges incurred by bandwidth usage have a greater incentive to moderate that usage. Thus charging can act as a feedback mechanism that discourages users from attempting to fill the network with their own traffic to the detriment of others; see e.g. [12]. Finally, differentiated service quality will require correspondingly differentiated charges. In particular, it is expected that premium services will be charged on a per use basis, even if best effort services remain on a flat (i.e. usage insensitive) fee. We remark that in one current proposal [13], customer awareness of the differentiated charges is to be the sole mechanism maintaining differential service quality.

In this paper, we will be concerned with customer usage as determined from direct measurements of customer traffic at routers or other network elements. Specifically, we take usage to mean the total number packets or bytes in those packets (or possibly some combination of these) that are observed during some billing interval, these totals being differentiated at some accounting granularity, e.g., by customer, service class, source and/or destination IP address, application, or some combination of these.

Where and how should such measurements be performed? One possibility is to maintain byte or packet counters at a customer's access port(s). Such counters are currently very coarsely grained, giving aggregate counts in each direction across an interface over periods of a few minutes. However, even separate counters differentiated by service quality would not suffice for all charging schemes. This is because service quality may not be the sole determinant of customer charges. These could also depend, for example, on the remote (i.e. non-customer) IP address involved; see e.g. [11]. This illustrates a broader point, namely, that the determinants of a charging scheme may be both numerous and also relatively dynamic. This may preclude using counts arising from a set of traffic filters, due to the requirement to have potentially a large number of such filters, and the administrative cost of configuring or reconfiguring such filters.

A complementary approach is to measure (or at least

summarize) all traffic, then transmit the measurements to a back-office system for interpretation according to the charging policy. In principle this could be done by gathering packet headers, or by forming flow statistics. An IP flow is a sequence of IP packets that shared a common property, as source or destination IP address or port number or combinations thereof. A flow may be terminated by a timeout criterion, so that the interpacket time within the flow does not exceed some threshold, or a protocol-based criterion, e.g., by TCP FIN packet. Flow definition schemes have been developed in research environments, see e.g. [1], and continue to be the subject of standardization efforts [10], [16]. Cisco NetFlow is an operating system feature for the collection and export of flow statistics. These include the identifying property of the flow, its start and end time, the number of packets in the flow, and the total number of bytes of all packets in the flow. Information on some commercial examples of the use of flow statistics for billing purposes can be found at [4]. Other examples of flow definitions employed as part of network management and accounting systems can be found in Inmon’s sFlow [9], Qosient’s Argus [15], Riverstone’s LFAP [17] and XACCT’s Crane [20].

B. The Motivation for Sampling

One limitation to comprehensive direct measurement of traffic stems from the immense amounts of measurement data generated. For example, a single OC48 at full utilization could generate about 100GB of packet headers, or several GB of (raw) flow statistics each hour. The demands on computational resources at the measurement point, transmission bandwidth for measured data, and back-end systems for storage and analysis of data, all increase costs for the service provider.

A common approach to dealing with large data volumes is to *sample*. In the context of network measurements, this is not a new idea; see e.g. [6]. However, a common objection to sampling has been the potential for inaccuracy; customers can be expected to be resistant to being overcharged due to overestimation of the resources that they use.

In this paper we employ a sampling mechanism that specifically addresses concerns of sampling error. Total customer usage is the sum of a number of components, some large, some small. Sampling errors arise predominantly from omission of the larger components, whereas accuracy is less sensitive to estimation of the smaller components. For example, consider a simple sampling scheme in which we estimate total bytes by sampling 1 in every N flows, then add together N times the bytes reported by each sampled flow. The underlying distribution of flow

bytes sizes has been found to follow a heavy tailed distribution [8]. In this case, the estimate can be extremely sensitive to the omission or inclusion of the larger flows. Generally, such an estimator can have high variance due to the sampling procedure itself.

The work in this paper rests on the observation that the heavy-tailed distribution of flow packet and byte sizes can be turned to our advantage for sampling *provided an appropriate sampling algorithm is used*. We replace uniform sampling with *size dependent* sampling, in which an object of size x is selected with some size dependent probability $p(x)$. The probability $p(x)$ is 1 for large x . In the case of flows, all sufficiently large flows will always be selected; there is no sampling error for such flows. On the other hand we take $p(x) < 1$ for smaller flows; this reduces the number of samples, but the error involved is small since the underlying flows are small. To estimate the total usage represented in the original set of flows, we sum up the quantities $x/p(x)$ over only the sampled flows. Applying the renormalization factor $1/p(x)$ to the small flows compensates for the fact that they might have been omitted. In fact, it can be shown that this sum is an unbiased estimator of the actual total usage (i.e. its average value over all possible random samplings is equal to the actual total usage). We remark that the uniform sampling described in the previous paragraph is a special case of this scheme with $p(x)$ constant and equal to $1/N$.

The size-dependent sampling scheme employed here has recently been proposed in [7]. It enjoys a number of useful properties. Firstly, the sampling probabilities $p(x)$ can be chosen to satisfy a certain optimality criterion for estimator variance; we describe this later. Secondly, a simple adaptive scheme allows dynamic tuning of p in order to keep the total number of samples within a given bound. Thus, in the context of flow measurement, the number of flow statistics transmitted to the back end system can be controlled. Thirdly, on binding the sampling parameters (i.e. $p(x)$) to the data x in constructing the rescaled size $x/p(x)$, we obviate the need to keep independent track of p , or even the original per flow usage x if only total usage estimation is required. Thus, p can vary at different times, or across different regions of the network, or even for different traffic classes, as needed, but estimation remains unbiased. Fourth, sampling is composable in the sense that the first three properties above are preserved under successive resampling. Thus, one could progressively resample at different points in the measurement system in order to limit the volume of samples. Lastly, we observe that although we have framed the discussion in terms of flow sampling, it could apply equally well to packet sampling. However, we expect the performance benefit over 1

in N sampling to be smaller in this case, since packet sizes do not have a heavy-tailed distribution.

C. Contribution

In this paper we take an approach to usage-sensitive charging that mirrors the foregoing approach to sampling. The sampling scheme determines the size of the larger flows with no error. Estimation error arises entirely from sampling smaller flows. For billing purposes we wish to measure the total usage for each billed entity (e.g. for each customer at a given service level) over each billing cycle. Larger totals have a smaller relative error due to sampling, whereas estimation of total usage for the smallest customers may be subject to greater relative error. Therefore, we set a level L on the total usage, with a fixed charge for all usage up to L , then a usage sensitive charge for all usage above L . Thus we need only tune our sampling scheme in order that estimation of usage above L be sufficiently accurate.

In Section II we describe a parameterized family of size dependent probabilities $p(x)$ for sampling, and describe their statistical properties. Each member of the family is specified by a single parameter, the *sampling threshold* z . Flows of size $x \geq z$ are always sampled, whereas smaller flows are sampled with probability x/z . We demonstrate their application to flow sampling using flow traces taken from a commercial network. We review the statistical properties of the flow packet and byte size distributions, and compare the accuracy of estimation of our proposed sampling approach with that of uniform 1 in N sampling. Our approach has strikingly greater accuracy.

In Section III we show how to choose the sampling threshold z in order to uniformly control the variance of estimates of total usage above a given level L . This in turn allows us to control the sampling variance of charges in pricing schemes that are sensitive to usage only above the level L . This control is based on a bound on the sampling variance that is robust in that it is independent of the distribution of the underlying per flow usage. Two variations of the scheme are discussed. In the first, we use the bound on the variance in order to systematically adjust the *estimated* usage downward so as to reduce the possibility of erroneous overcharging to due sampling errors. However, this is achieved at the cost of rendering unbillable a small but consistent portion of customers' usage. The second variation shows how the amount of unbillable usage can be reduced by appropriately lowering the sampling threshold z . In Section IV we investigate the efficacy of all of these approaches by examining their performance on the flow traces mentioned above.

In Section V we propose some refinements of the esti-

mation of sample variance that was used to determine the appropriate sampling threshold z . Specifically, we derive an unbiased estimator for the sampling variance of the total usage from a given customer, an estimate which can be calculated from the sampled usage alone. This would allow reduction of the number of samples needed for a given estimation accuracy. These properties and some trade-offs are further discussed.

In Section VI we establish an independent condition on the sampling threshold z in order that the mean rate at which samples are generated can be kept within a given level. This is motivated by the desire to control the use of processing, transmission and storage resources in the measurement subsystem itself. We establish conditions under which such a goal is compatible with the goal of controlling sampling error for billing. We argue, with some support from the statistics of the flow traces, that these goals can be made compatible provided that the billing cycle is sufficiently long.

We summarize the technical results of the paper in Section VII and conclude in Section VIII.

II. SIZE-DEPENDENT SAMPLING

A. Objectives

The basis of the sampling scheme is that sufficiently large objects are always sampled, while smaller objects are sampled with progressively smaller probability. Let's concentrate on the specific example of flows. Suppose a set of flows labeled by $i = 1, 2, \dots, n$ have summaries generated by through measurement in the network during some time period. Let x_i be the usage attribute of interest from the flow i , e.g., the number of packets in the flow, or the total number of bytes in the flow, or any other positive quantity of interest. Recall each packet in a flow possesses a common attribute, such as IP address (or net), port number, or Type of Service (ToS) field. We refer to each combination of interest of such attributes as a "color"; c_i will be the color of flow i . In the context of billing, a color might correspond to a customer address, or this plus a remote network, and possibly a ToS specification. The mapping that associates a particular customer with a set of packet attributes may be relatively complex; we assume this to be performed by the subsystem that collects and interprets the measurements. Our object here is to estimate the totals for for each color c of interest, i.e., $X(c) = \sum_{1 \leq i \leq n: c_i = c} x_i$. We shall abbreviate the last sum by $\sum'_c x_i$.

A.1 Size-dependent Sampling.

For each positive number z we defined the *sampling probability* function $p_z(x) = \min\{1, x/z\}$. The role of z

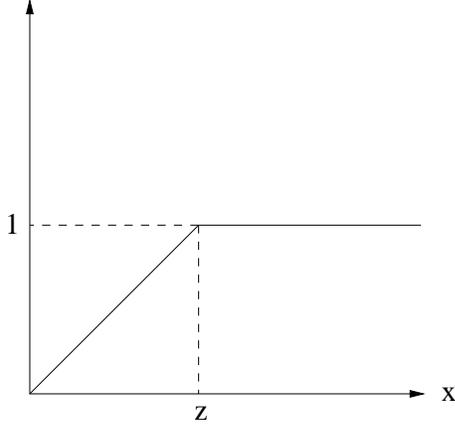


Fig. 1. SAMPLING PROBABILITY FUNCTIONS: $p_z(x) = \min\{1, x/z\}$. Flows of size less $x < z$ are sampled with probability x/z . Flows of size greater $x \geq z$ are always sampled.

will be explained a little later; for the moment let us fix it at some value. In our sampling scheme, a flow with size x is sampled with probability $p_z(x)$. The parameter z acts as a threshold: flows of size z or above are always sampled: see Figure 1. More formally, let $(w_i)_{i=1}^n$ be independent random variables, with w_i taking the values 1 with probability $p_z(x_i)$ and 0 otherwise. Thus w_i indicates whether flow i is to be sampled ($w_i = 1$) or not ($w_i = 0$). Each sampled value x_i is to be renormalized by division by $p(x_i)$. The estimate $\hat{X}(c)$ of the $X(c)$ is the random quantity

$$\hat{X}(c) = \sum_{1 \leq i \leq n: c_i=c} w_i x_i / p_z(x_i) \quad (1)$$

A.2 Statistical Properties.

We now turn to the statistical properties of the estimator $\hat{X}(c)$. In what follows we consider the underlying quantities x_i for be fixed for a given set of flows. The only randomness that enters is through the sampling operation itself, i.e., through the random quantities w_i . For example, the variance of $\hat{X}(c)$ is

$$\text{Var} \hat{X}(c) = \text{Var}(\sum'_c w_i x_i / p_z(x_i)) \quad (2)$$

$$= \sum'_c (x_i / p_z(x_i))^2 \text{Var}(w_i) \quad (3)$$

$$= \sum'_c x_i \max\{z - x_i, 0\}. \quad (4)$$

The last step follows from the fact that each w_i has variance $p(x_i)(1 - p(x_i))$. Note that flows i with $x_i > z$ make no contribution to $\text{Var}(\hat{X}(c))$ since they are always sampled.

The statistical properties of the estimate $\hat{X}(c)$, and the role of the parameter z are given in the following result,

from [7]. Let $N(c)$ denote the (random) number of objects of color c that are sampled, i.e., $N(c) = \sum'_c w_i$.

Theorem 1: [7] For each fixed set of sizes $(x_i)_{i=1}^n$:

(i) $\hat{X}(c)$ is an unbiased estimator of $X(c)$, i.e., $\mathbb{E}\hat{X}(c) = X(c)$.

(ii) p_z is optimal amongst the set all possible of sampling probability functions p in the sense that $\text{Var}\hat{X}(c) + z^2 \mathbb{E}N(c)$ is minimized when $p = p_z$.

Due to Theorem 1(ii) we will refer to size-dependent sampling, that uses a sampling probability p_z for some $z > 0$, as *optimal sampling*.

A.3 Trading Off Sample Volume and Variance.

So far we have interpreted z as a size threshold above which flows are always sampled. Note that the larger the value of z , the less likely we are to sample a given flow, and the greater the variance associated with sampling it. Theorem 1 furnishes us with a related interpretation, in which we think of z as the relative weight we attach to having small estimator variance vs. having a small number of samples. If z is small, then $\text{Var}\hat{X}(c) + z^2 \mathbb{E}N(c)$ is more easily minimized by making $\text{Var}\hat{X}(c)$ small, which in turns happens if we tend to sample more of the flows. Conversely, if z is large, then $\text{Var}\hat{X}(c) + z^2 \mathbb{E}N(c)$ is more easily minimized by making $\mathbb{E}N(c)$ small, which in turns happens if we tend to sample less of the flows.

A.4 Computational Issues.

It is worth remarking that there is an efficient implementation of the sampling strategy that is hardly more complex than deterministic 1 in N sampling; see [7].

B. Application to Flow Sampling

B.1 Flow Data Sources.

We illustrate the power of the optimal sampling method, including a comparison with deterministic 1 in N sampling, by applying it to IP flows. In this paper we use flow data drawn from two sources. First, we used raw NetFlow traces gathered from a router in an aggregation network serving domestic cable Internet users. The traces were gathered during one week in June 2001, and encompassed traffic generated by several thousand distinct IP address on the customer side. The second set of flows was generated from an IP header trace gathered during one week from an aggregation network serving modem banks access by dialup customers. Several hundred distinct customer IP addresses were present in this trace on the customer side. The IP header traces were used to generate flows based on the unique source and destination IP address and port numbers, and a 30 second timeout (i.e. this was the maximum

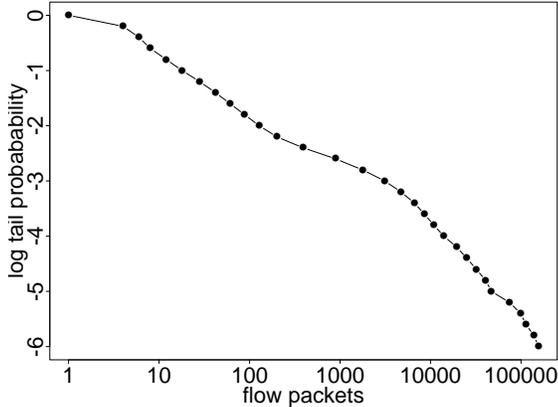


Fig. 2. CCDF OF FLOW BYTE SIZES: approximate linearity on log-log scale indicative of heavy tailed distribution.

time allowed between successive packets with matching addresses and ports). In both datasets, there is not necessarily a one-to-one mapping between IP addresses and individual customers, due to the potential for the dynamic re-assignment of IP addresses. For the purposes of this study, we confine ourselves to the problem of estimating traffic volumes generated by each customer side IP address (thus, these form the set of “colors”), regardless of whether it is actually used by more than one customer during the measurement period. We believe that the statistical effectiveness of the method would be essentially unchanged for per customer data, at least when applied to measurements over sufficiently long timescales.

B.2 Flow-size Distributions

Heavy-tailed distributions of the numbers of packets and bytes of IP flows has been previously noted; see e.g. [8]. Figure 2 displays the complementary cumulative distribution function (CCDF)—i.e. the proportion of flows with bytes greater than a given level—of the flow sizes present for one of the traces studied in the current paper. Observe the approximate linearity on the log-log scale, indicative of a heavy tailed distribution. The distribution of total bytes per customer-side IP address over a given period shares the heavy tailed property: see Figure 3.

B.3 Measures of Accuracy.

Our principle statistic for comparing estimated usage \hat{X} with its actual value is the absolute relative error $|1 - \hat{X}/X|$. We shall be interested in the distribution of this relative error over the range of colors of interest, and its dependence on the size X of the quantity to be estimated. For several experiments we shall wish to summarize this relative error over the range of colors of interests. One

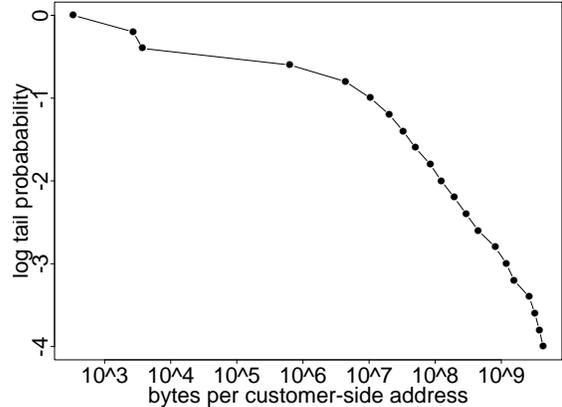


Fig. 3. CCDF OF BYTES PER CUSTOMER-SIDE IP ADDRESS: approximate linearity on log-log scale indicative of heavy tailed distribution.

measure is the Weighted Mean Relative Error (WMRE)

$$WMRE = \frac{\sum_c |\hat{X}(c) - X(c)|}{\sum_c X(c)} \quad (5)$$

(Note that \sum_c without the $'$ is just the a usual sum over all c , rather than over objects with a specific color c). The $WMRE$ averages the per-color absolute relative errors $|1 - \hat{X}(c)/X(c)|$ with a weight $X(c)/\sum_c X(c)$ in proportion to the quantity to be estimates. Thus it gives greater weight to relative errors for large volume colors than for those with small volumes.

B.4 Comparing Optimal with 1 in N Sampling.

Figure 4 compares the WMRE for 1 in N sampling and optimal sampling. For 1 in N sampling, N is called the sampling period. For optimal sampling, the reciprocal of the average probability that a flow is sampled probability the *effective* sampling period. For optimal sampling, the thresholds z took the values from 10^2 (for smaller sampling periods) up to 10^9 (for larger sampling periods).

Observe the strikingly better accuracy—i.e. smaller WMRE—of optimal sampling as compared with 1 in N sampling, over 4 orders of magnitude of the sampling period. As an example, with an effective sampling period of 100, the WMRE for optimal sampling is about only 1%, while for 1 in N sampling it is around 50%. The irregularity of the upper line reflects the sensitivity of the estimates from 1 in N sampling to random inclusion or exclusion of the largest flows during sampling. These features demonstrate the potential for inaccuracy arising from naive sampling from heavy-tailed distributions.

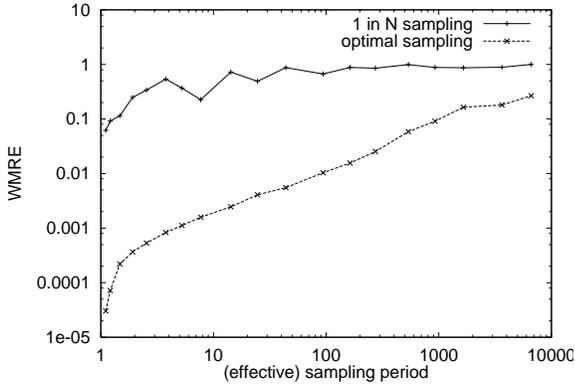


Fig. 4. WEIGHTED MEAN RELATIVE ERROR VS. EFFECTIVE SAMPLING PERIOD: for 1 in N and optimal sampling. Strikingly better accuracy from optimal sampling, over 4 orders of magnitude of the sampling period

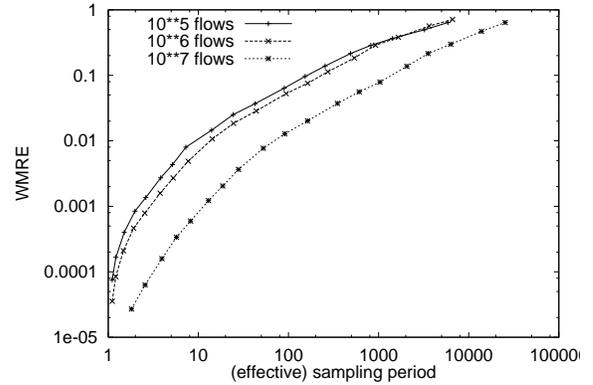


Fig. 5. WEIGHTED MEAN RELATIVE ERROR VS. EFFECTIVE SAMPLING PERIOD: for traces of 10^5 , 10^6 and 10^7 flows. Note increasing accuracy—i.e. smaller WMRE—with larger traces.

B.5 Accuracy and Trace Size.

Figure 5 displays WMRE vs. sampling period for a trace of 10^7 flows, as compared with subportions containing 10^6 and 10^5 flows. The relative error decreases as the trace length increases, since the byte total for a given IP address is composed of a greater number of contributions. This is an important observation, since, as we shall see, it may be desirable to place lower bounds on z in order to fulfill other objectives, such as limiting the rate at which samples are generated. The behavior in Figure 5 suggests that it will still be possible to fulfill simultaneously the goals of low relative error and low rate of sample production, provided that the length of the period of observation (e.g. the billing period) is sufficiently long.

III. PRICING AND THE CONTROL OF SAMPLE VARIANCE

We are now going to use the sampling approach from the last section for charging. Fair charging requires that the deviation between the traffic charged to a customer and the actual traffic is kept to a minimum. The scheme from the last section is essentially the doing best possible, in the sense that variance of \hat{X} is minimized for a given threshold z . However, the relative estimation error can be relatively large for colors with small amounts of traffic. As an extreme example, suppose the traffic associated with color c contributes total actual usage $X(c) < z$. Each such flow must therefore have size less than z , and will hence make a contribution to the estimated usage $\hat{X}(c)$ that is either 0 (if the flow is not sampled), or z (if it is). Hence $\hat{X}(c)$ will be either 0, or at least z .

A. Estimation and Small Volumes.

As a simple solution to the problem of estimating the small traffic volumes, we suggest that traffic of a given color is charged a fixed fee, plus a usage-sensitive charge only for traffic volumes that exceed a certain level L . (L could depend on the color in question, although we shall suppress this possible dependence for clarity of notation). The idea is to tune the sampling algorithms so that any usage $X(c)$ that exceeds L can be reliably estimated. Usage $X(c)$ that falls below L does not need to be reliably estimated, since the associated charge is usage-insensitive, i.e., independent of $\hat{X}(c) < L$.

Generally, then, we can consider traffic to be charged according to some function $f_c(\hat{X}(c))$ which depends on $\hat{X}(c)$ only through the quantity $\max\{\hat{X}(c), L\}$, i.e., it is independent of any usage below L . The subscript of f_c indicates that the charge could depend on the color c , e.g., through the type of service, or foreign IP address. As simple example would be do charge using

$$f_c(\hat{X}(c)) = a_c + b_c \max\{\hat{X}(c), L\}. \quad (6)$$

Here, a_c is a fixed charge, than can encompass, e.g., port charges and administrative charges, whereas b_c is a per byte charge on traffic transmitted during the billing cycle, with a minimum usage L . (6) can also express pricing models in which there is a fixed administrative charge for small customers, whose usage doesn't warrant accurate measurement. Note that both a_c and b_c are allowed to depend on the color c in question. We will assume that both are non-negative quantities.

B. Reliable Estimation of Traffic Volumes.

Reliable estimation of the volumes $X(c)$ is arranged for by choosing the sampling threshold z appropriately high for level L in question. The larger the level L and the larger the deviation of $\hat{X}(c)$ from $X(c)$ that we can tolerate, the higher a threshold z we can allow. To find the required z we first bound the variance of $\hat{X}(c)$ above, and show that the bound is achieved in the worst case.

Lemma 1: $\text{Var}\hat{X}(c) \leq zX(c)$ with $\text{Var}\hat{X}(c) \rightarrow zX(c)$ as all $x_i \rightarrow 0$.

Proof of Lemma 1: From (2), $\text{Var}\hat{X}(c) = \sum_c' x_i \max\{z - x_i, 0\} \leq \sum_c' x_i z = zX(c)$. If all $x_i < z$, then $\text{Var}\hat{X}(c) = zX(c) - \sum_c' x_i^2$, which converges to $zX(c)$ as all $x_i \rightarrow 0$. ■

Note the bound on Lemma 1 depends only on z and the total $X(c)$, *not* on the detailed distribution of the x_i .

Suppose now that we want to control the variance of all estimates $\hat{X}(c)$ greater than the level L . We express this as a condition on the standard error, i.e., the ratio of standard deviation $\sigma(\hat{X}(c)) = \sqrt{\text{Var}\hat{X}(c)}$ to the mean $X(c)$. Roughly speaking, we want the typical estimation error to be no more than about ε times X , for some target $\varepsilon > 0$. We express this as the **standard error condition**:

$$\sigma(\hat{X}(c)) < \varepsilon X(c) \text{ if } X(c) \geq L \quad (7)$$

For example, with $\varepsilon = 0.05$ we require the standard deviation to be no more than 5% of the mean.

A more detailed interpretation is as follows. Suppose that the $\hat{X}(c)$ are derived from a large number of flows of independent sizes. Then $\hat{X}(c)$ is roughly normally distributed. (7) means that the probability of overestimating $\hat{X}(c) > L$ by an amount $\delta X(c)$ —i.e., by δ/ε standard deviations—is no more than $\phi(-\delta/\varepsilon)$, where ϕ is the standard normal distribution function. Thus, with $\varepsilon = 0.05$, the probability of overestimating $\hat{X}(c)$ by more than 10% is found—since $10\% = 2 \times 5\%$ —as $\phi(-2) = 2.23\%$.

The main result of this section established a condition for z under which (7) will hold, and which similarly allows us to control the standard error of the charge (6).

Theorem 2: (i) In order to guarantee that the standard error condition (7) holds for any collection (x_i) , we require that $z < \varepsilon^2 L$.

(ii) Assume $z < \varepsilon^2 L$. Then $\sigma(\max\{\hat{X}(c), L\}) \leq \varepsilon \max\{X(c), L\}$ and hence we can bound above the standard deviation for the charge f_c in (6) as

$$\sigma(f_c(\hat{X}(c))) \leq \varepsilon f_c(X(c)). \quad (8)$$

Proof of Theorem 2: (i) From Lemma 1, the standard deviation obeys the tight upper bound $\sigma(\hat{X}(c)) \leq$

$\sqrt{zX(c)}$. So keeping $\sigma(\hat{X}(c))$ below $\varepsilon X(c)$ requires that $\sqrt{zX(c)} < \varepsilon X(c)$, and hence $z < \varepsilon^2 X(c)$ for all $X(c) \geq L$, i.e., $z \leq \varepsilon^2 L$.

(ii) $\text{Var}(\max\{\hat{X}(c), L\}) \leq \text{Var}(\hat{X}(c)) \leq zX(c) \leq \varepsilon^2 LX(c) \leq \varepsilon^2 \max\{X(c), L\}^2$ where the second inequality uses Lemma 1. Hence $\text{Var}(f_c(\hat{X}(c))) \leq \varepsilon^2 b_c^2 \max\{X(c), L\}^2 \leq \varepsilon^2 f_c^2(X(c))$. ■

Note that in the event that $\varepsilon = \varepsilon_c$ and $L = L_c$ depend on color, we would require $z \leq \min_c \varepsilon_c^2 L_c$. The error bounds of Theorem 2 inherit from Lemma 1 the property of independence of the detailed distribution of the x_i ; accuracy does not depend on specific assumptions on the sizes of traffic flows.

C. Tighter Control of Potential Overcharging.

The above approach sets limits on the chance that the deviation of the estimated usage above the actual usage exceeds a given amount. A refinement allows us to set a limit on the chance that overcharging occurs. This should be more attractive from the customer's point of view since the chance of them being overbilled at all can be small. Conversely, the service provider has to accept a small persistent underbilling in order to accommodate the potential sampling error. We remark that related approaches have been taken in other industries, e.g., food manufacturers understate the weight of packaged items in order that fluctuations in the manufacturing process will rarely lead to the actual weight falling below the stated weight [18].

We exploit again that the distribution of $\hat{X}(c)$ can be well approximated by a normal distribution when it is derived from a large number of constituent samples. Then the probability of a deviation of $\hat{X}(c)$ of at least s standard deviations above the expected value $X(c)$ is

$$\Pr[\hat{X}(c) > X(c) + s\sigma(X(c))] \approx \phi(-s) \quad (9)$$

Suppose then, that the service provider makes a conservative estimate of the volume of traffic in color c as

$$\hat{X}'(c) = \hat{X}(c) - s\sqrt{z\hat{X}(c)}. \quad (10)$$

Here s is a parameter to be chosen: it is the number of standard deviations away from $X(c)$ above which we consider over-estimation sufficiently rare. The probability that $\hat{X}'(c)$ overestimates the actual traffic, i.e., $\Pr[\hat{X}'(c) > X(c)]$, is approximately $\phi(-s)$, by Lemma 1.

For example, with $s = 3$, $\phi(-s)$ is about 0.13%, i.e. about 1 in 740 traffic volumes will be overestimated. We propose that the service provider charge according to $\hat{X}'(c)$ rather than $\hat{X}(c)$, i.e., the customer is billed $f_c(\hat{X}'(c))$. Thus the chance that the customer is overbilled at all is again, in general, $\phi(-s)$.

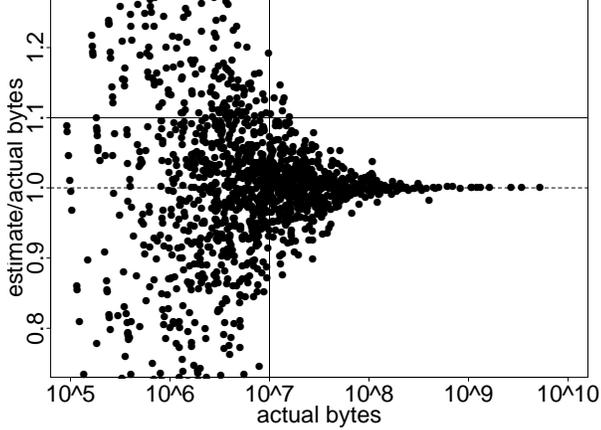


Fig. 6. ESTIMATED/ACTUAL TOTAL BYTES VS. ACTUAL BYTE TOTAL: Less than 0.1% of all samples with volumes greater than $L = z/\varepsilon^2 = 10^7$ have ratio \hat{X}/X exceed target $1 + \varepsilon$ with $\varepsilon = 0.1$ (i.e. fall in the upper right quadrant).

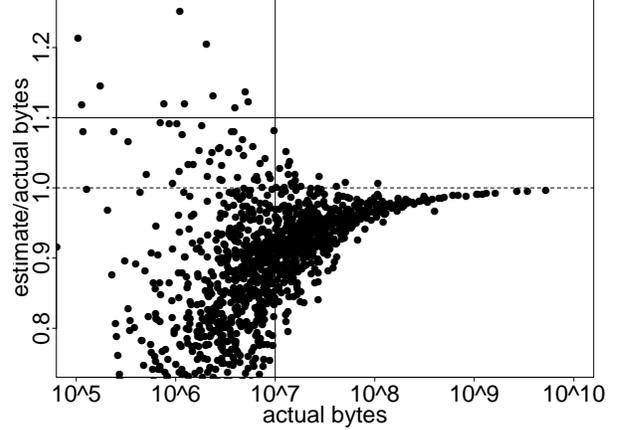


Fig. 7. ESTIMATED/ACTUAL TOTAL BYTES VS. ACTUAL BYTE TOTAL: using estimate \hat{X}' compensating with $s = 1$ standard deviations. Excursions of \hat{X}'/X above $1 + \varepsilon$ for $X > L$ have been eliminated entirely.

D. Tighter Control of Unbillable Usage.

For the service provider, the difference $\hat{X}(c) - \hat{X}'(c) = s\sqrt{z\hat{X}(c)}$ represents unbillable usage. In the charging scheme (6), this leads to underbilling by a fraction roughly $s\sqrt{z/X(c)}$. Given the minimum billed volume L , the fraction of underbilling is no more than $s\sqrt{z/L}$. (This could conceivably be systematically compensated for in the charging rate b_c). Thus in order to limit the fraction of usage that is unbillable to no more than about η , we require $sz < \eta^2 L$. Observe that, assuming $s \geq 1$, this represents at least as stringent a condition on z as that in Theorem 2, for the same proportionate error (i.e. if $\varepsilon = \eta$). In the example of $s = 3$, underbilling by a fraction of no more than $\eta = 10\%$ then requires picking z and L such that z is less than about $L/1000$.

IV. EXPERIMENTAL EVALUATION OF SAMPLING PERFORMANCE

We applied our sampling scheme to flow traces in order to determine its effectiveness, and to demonstrate the interplay between the sampling parameters, and the control they exert over sample variance, overcharging, unbillable usage, and sample volumes. Here we report the detailed behavior for one trace; the behavior for the others was similar. The trace used for this experiment comprised 10^7 flows distributed over 1,663 distinct customer-side IP addresses; these constituted the “colors” whose usage we wished to estimate. In this study we focus on usage as measured in bytes.

A. Variance and Control of Overcharging

Figure 6 is a scatter plot obtained by plotting for each color c (i.e. each IP address), the ratio $\hat{X}(c)/X(c)$ of estimated to actual usage, against the actual usage $X(c)$. With a target error ε , we aim to keep \hat{X}/X below $1 + \varepsilon$ to avoid overcharging beyond a proportion ε . The sampling threshold z was 10^5 , and the target error ε was 0.1 (i.e. 10%); we indicate $1 + \varepsilon$ by a horizontal line in the figure. We also place a dotted horizontal line at height 1; points below this line incur no overcharging. According to Theorem 2, we require a level L of at least $\varepsilon^2 z = 10^7$, indicated by the vertical line in the Figure. Note that \hat{X}/X approaches 1 for large X ; sampling is more accurate for larger total bytes X . This reflects that larger totals are typically composed of larger components (which are more accurately sampled) or more components (giving better averaging).

In order to meet the target accuracy, we require as few points as possible to fall in the upper right quadrant bounded by the horizontal $1 + \varepsilon$ and vertical L lines. Only 0.13% of all samples fall into this category. However, following Section III-C, such exceedence of the target can be reduced upon replacing the estimate \hat{X} by $\hat{X}' = \hat{X} - s\sqrt{z\hat{X}}$, from (10), compensating with some number s of standard deviations. (Recall $\hat{X} = \hat{X}'$ when $s = 0$). In the example, exceedence of the target is actually eliminated on setting $s = 1$; see Figure 7. Note that the correction $s\sqrt{z\hat{X}}$, as a proportion of \hat{X} , decreases to zero for large \hat{X} . Thus the estimation accuracy is largely unaffected for large X on increasing s above zero.

The proportion of totals $\hat{X}' > L$ that exceed X at all is reduced from about 50% at $s=0$, to 3% for $s = 1$. They

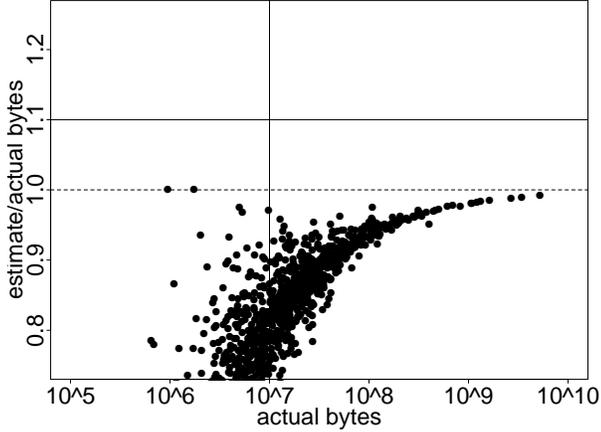


Fig. 8. ESTIMATED/ACTUAL TOTAL BYTES VS. ACTUAL BYTE TOTAL: using estimate \hat{X}' compensating with $s = 2$ standard deviations. All excursions of \hat{X}'/X above 1 have been eliminated: there is no overcharging, but unbillable usage has increased significantly.

are eliminated entirely on setting $s = 2$; see Figure 8: all points lie below the dotted horizontal line at height 1. Thus by increasing s we can eliminate overcharging of $X > L$ entirely; the potential downside is that the frequency and amount of unbillable usage increases.

We summarize these results in Table I. The important property to note is that by increasing s we reduce the number of customers whose traffic is overestimated, but increase the amount of unbillable traffic. Choosing the intermediate value $s = 1$ represents a compromise between these two effects.

The heavy tailed nature of the distribution of per color usage is striking. In this sample, some 94% of the total bytes were attributable to customer-side IP addresses with total bytes greater than L , although it took only about 10% of the addresses to generate these bytes.

s	unbillable usage	overcharged customers
0	-0.1%	50%
1	3.1%	3%
2	6.2%	0

TABLE I

TRADE-OFF BETWEEN OVERCHARGING AND UNBILLABLE TRAFFIC AS NUMBER s OF COMPENSATING STANDARD DEVIATIONS IS INCREASED IN \hat{X}'

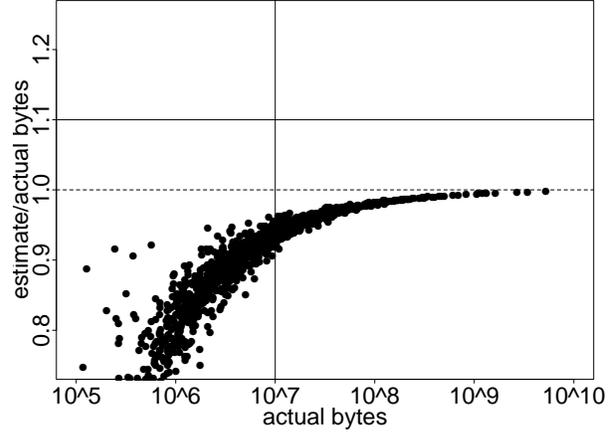


Fig. 9. ESTIMATED/ACTUAL TOTAL BYTES VS. ACTUAL BYTE TOTAL: using estimate \hat{X}' compensating with $s = 2$ standard deviations, with reduced sampling threshold z . Observe reduction in unbillable usage as compared with Figure 8.

B. Control of Unbillable Usage

The elimination of overcharging from sampling variance, achieved by passing from $s = 0$ (i.e. estimating with \hat{X}) to $s = 1$, brings its own pitfalls as described in Section III-D, namely, making a proportion of the actual usage unbillable. Let $S(L) = \sum_{c: X(c) > L} X(c)$ denote the total actual bytes from customer addresses that generated at least L bytes, with $\hat{S}(L)$ and $\hat{S}'(L)$ defined in a similar way. For $s = 0$, the ratio $\hat{S}(L)/S(L)$ is 1.001, representing a slight (0.1%) overcharging on average. For $s = 1$, the ratio $\hat{S}'(L)/S(L)$ moves down to 0.969, i.e., a 3% unbillable usage; for $s = 2$, $\hat{S}'(L)/S(L) = 0.938$, i.e., 6% unbillable usage. According to Section III-D, the underbilling can be ameliorated, at fixed L , by reducing the sampling threshold z . Suppose we now chose $z = \eta^2 L/s$, with $s = 2$ but now $\eta = 0.0447 < \varepsilon = 0.1$. In the above example, the sampling threshold z moves down from 10^5 to 10^4 . Then $\hat{S}'(L)/S(L)$ moves up to 0.980 reducing the amount of unbillable usage by roughly a factor 3. We show the effect in Figure 9. As expected, sampling variance, as exhibited by the vertical spread, is also reduced. Note that the reduction in unbillable usage comes at a cost: the number of samples taken is increased by reducing z . Decreasing z from 10^5 to 10^4 increases the sample volume by roughly a factor of 5.

C. Summary of Experimental Results

These experiments demonstrate how choosing sampling $z = \varepsilon L$ is effective in keeping relative error due to sampling variance (for total bytes at least L) less than about ε . The infrequent occurrence of errors larger than ε could

be eliminated entirely by adjusting the estimate \hat{X} downwards by subtracting some number of standard deviations $s\sqrt{z\hat{X}}$. A side effect of this was the introduction of a small but consistent amount of unbillable usage. This be reduced by lowering the sampling threshold z , at a cost of increasing the volume of samples taken.

V. REFINEMENT OF VARIANCE ESTIMATION

The key idea of the previous section was that by choosing the sampling threshold z sufficiently small, one could control both overcharging and the amount of unbillable usage, by using an estimate (in fact, an upper bound) on the standard deviation $\sigma(\hat{X})$ of the estimated usage. In this section we show how this bound can be refined somewhat, if additional information from the sampling process is available. The main utility is that with a better estimate of the variance, we can use a higher value of z and hence take fewer samples.

A. Detailed Variance Estimation

One of the attractions of the bound $\text{Var}\hat{X} \leq zX(c)$ is its simplicity: no knowledge of the individual contributions to the sum $\hat{X}(c)$ is required. The bound is tight, in the sense that it is achieved in the worst case of a large number of small flows. Correspondingly, when $\hat{X}(c)$ is composed of a small number of large flows, the bound can significantly overestimate the variance. Specifically, recall from (2) the expression $\text{Var}\hat{X} = \sum_i x_i \max\{z - x_i, 0\}$. If all $x_i > 0$, this expression is zero, reflecting that all such x_i are automatically sampled. This effect is evidently at work in Figure 6; the estimates of the largest usage totals appear exact, and $\hat{X}(c)/X(c)$ is well below $1 + \varepsilon$ for many of the larger usage totals.

We now show how to obtain an unbiased estimate of $\text{Var}\hat{X}$ based only on those x_i that are sampled. The aim here is to find some function $v(x_i)$ of the sampled x_i , such that the expected value $\mathbb{E}\hat{V}$ of the sum over sampled flows $\hat{V} = \sum_i w_i v(x_i)$ is equal to the actual variance $\text{Var}\hat{X}(c)$. (Recall w_i is the random quantity indicating whether or not flow i is sampled ($w_i = 1$) or not ($w_i = 0$)).

Lemma 2: [7] $\hat{V} = \sum_i w_i z \max\{z - x_i, 0\}$ obeys $\mathbb{E}\hat{V} = \text{Var}\hat{X}$.

Proof of Lemma 2: The only random quantities in \hat{V} are the w_i , and they have expectation $\mathbb{E}w_i = p_z(x_i) = \min\{1, x_i/z\}$. Thus the only non-zero contributions to $\mathbb{E}\hat{V} = \sum_i p_z(x_i) z \max\{z - x_i, 0\}$ come from terms with $x_i < z$. These contributions take the form $x_i \max\{z - x_i, 0\}$, i.e., the same as the non-zero contributions to $\text{Var}\hat{X}$. ■

B. Using the Detailed Variance Estimator

We now look at the costs and tradeoffs involved with using the detailed variance estimator \hat{V} instead of the upper bound from Lemma 1.

Information. \hat{V} requires knowing of the x_i from sampled flows. The RHS of the bound $\sigma(\hat{X}) \leq zX$ can be estimated as $z\hat{X}$, i.e., it requires only the sum of the renormalized values $x_i/p_z(x_i)$, i.e., the original per flow usages are not required. The need not be an issue if \hat{V} can be computed at the same point as \hat{X} .

Sample Volume. In order to fulfill the standard error condition (7) we chose z so that $\hat{V} \leq \varepsilon^2 \hat{X}$. Since $\mathbb{E}\hat{V} \leq zX$, the required threshold z should be larger, on average, than the value $\varepsilon^2 L$ specified in Theorem 2. As discussed in Section II-A.3, larger thresholds z yield smaller sampling rates. To summarize, tighter control on the sampling variance allows a given level of confidence with a smaller number of samples.

Data Dependence. Since \hat{V} now depends more sensitively on the distribution of the x_i , the value of the sampling threshold z may need to be dynamically adjusted if the underlying distribution of flow sizes changes. It is worth remarking that Lemma 2 generalizes to the case that the sampling threshold z can vary across flows, i.e., on replacing z by z_i in each term involving x_i . Thus dynamic adjustment of z would not alter our ability to form an unbiased estimate of the variance.

VI. RECONCILING SAMPLING VARIANCE AND VOLUME GOALS

This paper has dealt primarily with the control of sampling variance in order to facilitate accurate charging. However, there are other potential goals for the measurement subsystem to be considered. In particular, it may be desirable to limit the rate at which samples are produced, in order to control the resulting load on the router generating the samples, the measurement collector, and the communications network that is used to transmit samples between the generator and the collector.

A. Controlling Sample Volume

Suppose that flows present themselves for sampling at a rate ρ , and that the per flow usage has a distribution function F , i.e., $F(x)$ is the proportion of flows with usage less than or equal to x . With a sampling threshold z , samples are produced at an average rate $r = \rho \int F(dx)p_z(x)$. Suppose there is a target maximum rate of samples $r^* < \rho$. Then we require the sampling threshold z to be such that $\rho \int F(dx)p_z(x) \leq r^*$. Using the fact that $p_z(x)$ is a decreasing function in z , it can be shown that this requires

$z \geq z^*$, where z^* is the unique solution z to the equation $\rho \int F(dx)p_z(x) = r^*$. We remark that an extension of this approach to *dynamically* control the sample volume—under fluctuations and systematic changes in the offered rate of flows to be sampled—has been formulated in [7].

B. Variance and the Billing Timescale

Now let z_0 denote the maximum sampling threshold allowed in order to control sampling variance, e.g., $z \leq z_0 = \varepsilon^2 L$ from Theorem 2. Clearly the goals of controlling sample volume and variance are compatible provided that $z^* \leq z_0$, for then any sampling threshold z in the interval $[z^*, z_0]$ has the property of being sufficiently small to yield small sampling variance, and sufficiently large to restrict the average sampling rate no greater than the desired rate r^* .

It turns out that the condition $z^* \leq z_0$ can be realized, at least in principle. To see this, observe that the thresholds z_0 and z^* control phenomena at different timescales. z^* controls the average rate at which samples are taken. On the other hand, z_0 controls the sampling variance of the estimates $\hat{X}(c)$ of total usage over the billing timescale, potentially over days, weeks, or even months. The usage level L —under which accurate measurements are not needed—can be chosen to increase with the billing timescale. For example, we might choose L to correspond to a particular quantile of the distribution of total usage, so that only a given proportion of the total network usage is generated by customers whose usage does not exceed L during the billing cycle. Increasing the length of the billing cycle will increase the corresponding quantile L , and hence also z_0 since, as we saw in Section III, z_0 is proportional to L . Support for this approach is provided by Figure 5, which shows how the relative error in estimation decreases and the duration of collection of the flow trace increases.

VII. SUMMARY OF SAMPLING MECHANISM

Before concluding, we summarize the main technical results from this paper:

(i) **Usage Estimation.** A flow of size x is sampled with probability $p_z(x) = \min\{1, x/z\}$. If sampled, a renormalized usage $x/p(x)$ is reported. z is called the sampling threshold: flows with usage $x \geq z$ are always sampled. The sum \hat{X} of $x/p(x)$ over sampled flows is an unbiased estimator of the total usage, i.e., the sum X of x over all flows.

(ii) **Sampling Error and Pricing.** For a given usage level L , choosing $z < \varepsilon^2 L$ guarantees that the standard error $\sigma(\hat{X})/X$ of the total usage is less than ε when $X \geq L$. The price $f(\hat{X}) = a + b \max\{\hat{X}, L\}$ represents a fixed

charge for all usage up to L , and a usage sensitive charge thereafter. $f(\hat{X})$ has standard error no more than ε as compared with the corresponding price $f(X)$ if the usage had been known exactly: $\sigma(f(\hat{X})) \leq \varepsilon f(X)$.

(iii) **Control of Overcharging.** The adjusted estimate $\hat{X}' = \hat{X} - s\sqrt{z\hat{X}}$ exceeds X with probability approximately $\phi(-s)$, where ϕ is the standard normal distribution function. The chance and amount of overcharging can thus be reduced by increasing s above zero.

(iv) **Control of Unbillable Usage.** The cost of controlling the extremes of overcharging is to render a fraction of the total usage unbillable. Choosing $z < \eta^2 L/s$ then restricts this fraction to be less than about η .

(v) **Control of Sample Volume.** To keep the rate of sample production less than a target r^* requires $z > z^*$ where z^* is the root z of the equation $r^* = \rho \int F(dx)p_z(x)$. Here $\rho > r^*$ is the rate at which flows are produced, and F is the distribution of their sizes.

(vi) **Simultaneous Control of Sample Variance and Volume.** Let z_0 denote a threshold chosen to limit sample variance in (ii),(iii) or (iv) above. The average rate of sample production can be kept below r^* if $z_0 > z^*$. This can be arranged for provided the billing timescale is sufficiently long.

VIII. CONCLUSIONS

This paper was motivated by the desire to perform accurate usage sensitive billing from sampled flow statistics. We assume that it is not possible for a router to keep counters on all traffic flows of interest, either because they are too numerous, or because the set of flows of interest is itself dynamic.

The heavy-tailed nature of flow size distributions, so problematic for uniform 1 in N sampling, can be turned to our advantage when we use size-dependent flow sampling. This picks out the dominant contributions to usage. Combined with a charging scheme that is sensitive to usage only above a certain level L , this allows accurate charging of customers for their usage. The main technical results of this paper showed how to relate the sampling threshold z to the level L for a given desired level of accuracy.

Size dependent sampling also allows control of the rate at which samples are produced. Sampling rate control, which is favored by higher z , appears at first to be opposed to the goal of variance control, which if favor by higher z . However, we argued that these goals could be rendered compatible provided the billing timescale is sufficiently long. Under these conditions the goals of accurate usage-sensitive billing from sampled flow records is attainable.

REFERENCES

- [1] J. Apisdorf, K. Claffy, K. Thompson, and R. Wilder, "OC3MON: Flexible, Affordable, High Performance Statistics Collection," For further information see <http://www.nlanr.net/NA/Oc3mon>
- [2] R. Cáceres, N.G. Duffield, A. Feldmann, J. Friedmann, A. Greenberg, R. Greer, T. Johnson, C. Kalmanek, B. Krishnamurthy, D. Lavelle, P.P. Mishra, K.K. Ramakrishnan, J. Rexford, F. True, and J.E. van der Merwe, "Measurement and Analysis of IP Network Usage and Behavior", *IEEE Communications Magazine*, vol. 38, no. 5, pp. 144–151, May 2000.
- [3] Cisco Systems, Inc., "NetFlow"; for further information see: <http://www.cisco.com/warp/public/732/netflow/index.html> and http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s11/12s_sanf.htm
- [4] Cisco NetFlow Services and Applications Customer Profiles. See <http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neftct/profiles/>
- [5] K.C. Claffy, H.-W. Braun, and G.C. Polyzos. Parameterizable methodology for internet traffic flow profiling. *IEEE Journal on Selected Areas in Communications*, 13(8):1481–1494, October 1995.
- [6] Kimberly C. Claffy, George C. Polyzos, and Hans-Werner Braun. Application of Sampling Methodologies to Network Traffic Characterization. *Computer Communication Review*, 23(4):194–203, October 1993, appeared in Proceedings ACM SIGCOMM'93, San Francisco, CA, September 13–17, 1993.
- [7] N.G. Duffield, C. Lund, M. Thorup, Learn more, sample less: control of volume and variance in network measurement, Submitted for publication, 2001.
- [8] A. Feldmann, J. Rexford, and R. Cáceres, "Efficient Policies for Carrying Web Traffic over Flow-Switched Networks," *IEEE/ACM Transactions on Networking*, vol. 6, no.6, pp. 673–685, December 1998.
- [9] Inmon Corporation, "sFlow accuracy and billing", see: <http://www.inmon.com/PDF/sFlowBilling.pdf>
- [10] "Internet Protocol Flow eXport" (IPFX). IETF Working Group. See: <http://net.doit.wisc.edu/ipfx/>
- [11] Level 3 Communications, Inc., "Crossroads Internet Access", see: <http://www.level3.com/us/services/crossroads/>
- [12] J. K. MacKie-Mason and H. R. Varian. Pricing Congestible Network Resources. *IEEE Journal on Selected Areas in Communications*, 13(7):1141–1149, September 1995.
- [13] A. M. Odlyzko, Paris Metro Pricing for the Internet, Proc. ACM Conference on Electronic Commerce (EC'99), ACM, 1999, pp. 140–147.
- [14] V. Paxson, G. Almes, J. Mahdavi, M. Mathis, "Framework for IP Performance Metrics", RFC 2330, available from: <ftp://ftp.isi.edu/in-notes/rfc2330.txt>, May 1998.
- [15] Qosient, LLC, "Argus", see: <http://www.qosient.com/argus/index.htm>
- [16] Real Time Flow Measurement, see: <http://www.auckland.ac.nz/net/Internet/rtfm/>.
- [17] Riverstone Networks, Inc., see: <http://www.riverstonenet.com/technology/>
- [18] W. Shewhart, "Statistical method from the viewpoint of quality control", Dover, New York, 1939.
- [19] Sprint Corporation, "ION". See <http://ion.sprint.com>
- [20] XACCT Technologies, Inc., see: <http://www.xacct.com>