

Per-Flow Traffic Measurement Through Randomized Counter Sharing

Tao Li, Shigang Chen, and Yibei Ling, *Senior Member, IEEE*

Abstract—Traffic measurement provides critical real-world data for service providers and network administrators to perform capacity planning, accounting and billing, anomaly detection, and service provision. One of the greatest challenges in designing an on-line measurement module is to minimize the per-packet processing time in order to keep up with the line speed of the modern routers. To meet this challenge, we should minimize the number of memory accesses per packet and implement the measurement module in the on-die SRAM. The small size of SRAM requires extremely compact data structures to be designed for storing per-flow information. The best existing work, called *counter braids*, requires more than 4 bits per flow and performs six or more memory accesses per packet. In this paper, we design a fast and compact measurement function that estimates the sizes of all flows. It achieves the optimal processing speed: two memory accesses per packet. In addition, it provides reasonable measurement accuracy in a tight space where the counter braids no longer work. Our design is based on a new data encoding/decoding scheme, called *randomized counter sharing*. This scheme allows us to mix per-flow information together in storage for compactness and, at the decoding time, separate the information of each flow through statistical removal of the error introduced during information mixing from other flows. The effectiveness of our online per-flow measurement approach is analyzed and confirmed through extensive experiments based on real network traffic traces. We also propose several methods to increase the estimation range of flow sizes.

Index Terms—Computer networks.

I. INTRODUCTION

A. Motivation

TRAFFIC monitoring and measurement provide critical information for capacity planning, accounting and billing, anomaly detection, and service provision in modern computer networks [1]–[5]. This paper focuses on a particularly challenging problem, the measurement of per-flow information for a high-speed link without using per-flow data structures. The goal is to estimate the *size* of each flow (in terms of number of packets). A flow is identified by a label that can be a source address, a destination address, or any combination of addresses, ports, and other fields in the packet header.

Manuscript received January 22, 2011; revised October 21, 2011; accepted January 09, 2012; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor P. Crowley. Date of publication April 20, 2012; date of current version October 11, 2012. This work was supported in part by the US National Science Foundation under Grant CNS-1115548. A preliminary version of this paper appeared in the Proceedings of the IEEE International Conference on Computer Communications (INFOCOM), Shanghai, China, April 10–15, 2011.

T. Li and S. Chen are with the Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL 32608 USA (e-mail: tali@cise.ufl.edu; sgchen@cise.ufl.edu).

Y. Ling is with Applied Research Laboratories, Telcordia Technologies, Morristown, NJ 07960 USA.

Digital Object Identifier 10.1109/TNET.2012.2192447

Most prior work investigates how to identify the elephants (i.e., flows with large sizes) [1], [2], [6], [7] or estimate the flow size distribution (i.e., the number of flows that have a certain size) [8], [9]. They do not measure the sizes of individual flows, which is important information for many applications. For example, if we use the addresses of the users as flow labels, per-flow traffic measurement provides the basis for usage-based billing and graceful service differentiation, where a user's service priority gracefully drops as he overspends his resource quota. Studying per-flow data over consecutive measurement periods may help us discover network access patterns and, together with user profiling, reveal geographic/demographic traffic distributions among users. Such information will help Internet service providers and application developers to align network resource allocation with the majority's needs. In the event of a worm attack, per-source data can be used to estimate the scanning rates of worm-infected hosts. In the event of a botnet attack where there is a sudden surge of small flows, a security administrator may analyze the change in the flow size distribution and use per-flow information to compile the list of candidate bots that contribute to the change, helping to narrow down the scope for further investigation.

B. Challenges and Prior Art

Per-flow traffic measurement is a very challenging problem. In order to monitor small flows with a few packets, it is desirable to record information for each packet because aggressive sampling, such as what's used in Cisco Netflow and others [10], [11], will not only introduce significant measurement error, but also miss some small flows altogether. DRAM cannot keep up with today's line speed (40 Gb/s for OC-768, and 16.4 Tb/s in experimental systems [12]). The recent research trend is to implement online measurement functions in high-speed but expensive on-die SRAM [5], [9], [13], [14], [38].

However, the size of SRAM is limited, and it has to be shared among many critical functions for routing, scheduling, traffic management, and security. Even for traffic measurement alone, there can be multiple functions running concurrently on the same SRAM, each measuring a different type of flow. These functions may have to take turns to access the SRAM, and the amount of memory allocated to each measurement function is likely to be small. Hence, it is extremely important to make per-flow measurement and other functions in SRAM fast and compact. In particular, to keep up with the router's forwarding rate (which can be as high as a couple of cycles per packet), we should minimize the number of memory accesses per packet performed by each measurement function to the shared SRAM.

It has been shown in [2] that maintaining per-flow counters cannot scale for high-speed links. Even for efficient counter

implementations [15]–[17], SRAM will only be able to hold a small fraction of per-flow state (including counters and indexing data structures such as pointers and flow identities for locating the counters). The *counter braids* avoid per-flow counters and achieve near-optimal memory efficiency [18], [19]. This method maps each flow to three arbitrary counters; they are all incremented by one for every packet of the flow. Many flows may be mapped to the same counter, which stores the sum of the flow sizes. Essentially, the counters represent linear equations, which can be solved for the flow sizes. Two levels of counters are used to reduce the memory overhead. The counter braids require slightly more than 4 bits per flow and are able to count the exact sizes of all flows. However, it also has two limitations. First, it performs 6 or occasionally 12 memory accesses per packet. Second, when the memory allocated to a measurement function is far less than 4 bits per flow, our experiments show that the message passing decoding algorithm of counter braids cannot converge to any meaningful results. When the available memory is just 1–2 bits per flow, the *exact measurement* of the flow sizes is no longer possible. We have to resort to *estimation methods*. The key is to efficiently utilize the limited space to improve the accuracy of the estimated flow sizes, and do so with the minimum number of memory accesses per packet. This is what our paper tries to achieve.

A related thread of research is to collect statistical information of the flows [8], [9] or identify the largest flows and devote the available memory to measure their sizes while ignoring the smaller ones [2], [6], [7], [20]. For example, RATE [21] and ACCEL-RATE [22] measure per-flow rate by maintaining per-flow state, but they use a *two-run sampling* method to filter out small-rate flows so that only high-rate flows are measured.

Another thread of research is to maintain a large number of counters to track various networking information. One possible solution [23], [24] can be statistically update a counter according to the current counter size. This approach is fit for the applications with loose measurement accuracy. In order to enhance the accuracy performance, Zhao *et al.* [25] propose a statistical method to make a DRAM-based solution practical, which uses a small cache and request queues to balance the counter values. Since DRAM is involved and wirespeed is achieved, this approach is able to achieve decent measurement accuracy.

Also related is the work [14] that measures the number of *distinct* destinations that each source has contacted. Per-flow counters cannot be used to solve this problem because they cannot remove duplicate packets. If a source sends 1000 packets to a destination, the packets contribute only one contact, but will count as 1000 when we measure the flow size. To remove duplicates, bitmaps (instead of counters) should be used [5], [26]–[28]. From the technical point of view, this represents a separate line of research, which employs a different set of data structures and analytical tools. Attempt has also been made to use bitmaps for estimating the flow sizes, which is however far less efficient than counters, as our experiments will show.

C. Our Contributions

We design a fast and compact per-flow traffic measurement function that achieves three main objectives.

- 1) It shares counters among flows to save space, and does not incur any space overhead for mapping flows to their counters. This distinguishes our work from [15]–[17].
- 2) It updates exactly one counter per packet, which is optimal. This separates our work from the counter braids that update three or more counters per packet. Updating each counter requires two memory accesses for read and then write.
- 3) It provides estimation of the flow sizes, as well as the confidence intervals that characterize the accuracy, even when the available memory is too small such that other exact-counting methods including [18] and [19] no longer work.

We believe our work is the first one that achieves all these objectives. It complements the existing work by providing additional flexibility for the practitioners to choose when other methods cannot meet the speed and space requirements.

The design of our measurement function is based on a new data encoding/decoding scheme, called *randomized counter sharing*. It splits the size of each flow among a number of counters that are randomly selected from a counter pool. These counters form the *storage vector* of the flow. For each packet of a flow, we randomly select a counter from the flow’s storage vector and increment the counter by one. Such a simple online operation can be implemented very efficiently. The storage vectors of different flows share counters uniformly at random; the size information of one flow in a counter is the noise to other flows that share the same counter. Fortunately, this noise can be quantitatively measured and removed through statistical methods, which allow us to estimate the size of a flow from the information in its storage vector. We propose two estimation methods whose accuracies are statistically guaranteed. They work well even when the total number of counters in the pool is by far smaller than the total number of flows that share the counters. Our experimental results based on real traffic traces demonstrate that the new methods can achieve good accuracy in a tight space. We also propose several methods to increase the range of flow sizes that the estimators can measure.

The randomized counter sharing scheme proposed in this paper for per-flow traffic measurement has applications beyond the networking field. It may be used in the data streaming applications to collect per-item information from a stream of data items.

The rest of this paper is organized as follows: Section II discusses the performance metrics. Section III gives an overview of our system design. Section IV presents the online data encoding module. Sections V and VI propose two offline data decoding modules. Section VII discusses the problem of setting counter length. Section VIII addresses the problem of collecting flow labels. Section IX presents the experimental results. Section X extends our estimators for large flow sizes. Section XI draws the conclusion.

II. PERFORMANCE METRICS

We measure the number of packets in each flow during a measurement period, which ends every time after a certain number (e.g., 10 million) of packets are processed. The design of per-flow measurement functions should consider the following three key performance metrics.

A. Processing Time

The per-packet processing time of an online measurement function determines the maximum packet throughput that the function can operate at. It should be made as small as possible in order to keep up with the line speed. This is especially true when multiple routing, security, measurement, and resource management functions share SRAM and processing circuits.

The processing time is mainly determined by the number of memory accesses and the number of hash computations (which can be efficiently implemented in hardware [29]). The counter braids [18], [19] update three counters at the first level for each packet. When a counter at the first level overflows, it needs to update three additional counters at the second level. Hence, it requires 3 hashes and 6 memory accesses to read and then write back after counter increment. However, in the worse case, it requires 6 hashes and 12 memory accesses. The multiresolution space-code Bloom filters [13], [38] probabilistically select one or more of its nine filters and set 3–6 bits in each of the selected ones. Each of those bits requires one memory access and one hash computation.

Our objective is to achieve a constant per-packet processing time of one hash computation and two memory accesses (for updating a single counter). This is the minimum processing time for any method that use hash operations to identify counters for update.

B. Storage Overhead

The need to reduce the SRAM overhead has been discussed in Section I. One may argue that because the amount of memory needed is related to the number of packets in a measurement period, we can reduce the memory requirement by shortening the measurement period. However, when the measurement period is smaller, more flows will span multiple periods, and consequently the average flow size in each period will be smaller. When we measure the flow sizes, we also need to capture the flow labels [19], e.g., a tuple of source address/port and destination address/port to identify a TCP flow. The flow labels are too large to fit in SRAM. They have to be stored in DRAM. Therefore, in a measurement period, each flow incurs at least one DRAM access to store its flow label. If the average flow size is large enough, the overhead of this DRAM access will be amortized over many packets of a flow. However, if the average flow size is too small, the DRAM access will become the performance bottleneck that seriously limits the throughput of the measurement function. This means the measurement period should not be too small. Our experiments in Section IX set a measurement period such that the average flow size is about 10.

C. Estimation Accuracy

Let s be the size of a flow and \hat{s} be the estimated size of the flow based on a measurement function. The estimation accuracy of the function can be specified by a confidence interval: The probability for s to be within $[\hat{s} \cdot (1 - \beta), \hat{s} \cdot (1 + \beta)]$ is at least a prespecified value α , e.g., 95%. A smaller value of β means that the estimated flow size is more accurate (in a probabilistic sense).

There is a tradeoff between the estimation accuracy and the storage overhead. If the storage space and the processing time




Fig. 1. Traffic distribution: Each point shows the number (y -coordinate) of flows that have a certain size (x -coordinate).

are unrestricted, we can accurately count each packet to achieve perfect accuracy. However, in practice, there will be constraints on both storage and processing speed, which make 100% accurate measurement sometimes infeasible. In this case, one has to settle with imperfect results that can be produced with the available resources. Within the bounds of the limited resources, we must explore novel measurement methods to make the estimated flow sizes as accurate as possible.

III. SYSTEM DESIGN

A. Basic Idea

We use an example to illustrate the idea behind our new measurement approach. Suppose the amount of SRAM allocated to one of the measurement functions is 2 Mb (2×2^{20} bits), and each measurement period ends after 10 million packets, which translates into about 8 s for an OC-192 link (10+ Gb/s) with an average packet size of 1000 B. The types of flows that the online functions may measure include per-source flows, per-destination flows, per-source/destination flows, TCP flows, WWW flows (with destination port 80), etc. Without losing generality, suppose the specific function under consideration in this example measures the sizes of TCP flows.

Fig. 1 shows the number of TCP flows that have a certain flow size in log scale, based on a real network trace captured by the main gateway of our campus. If we use 10 bits for each counter, there will be only 0.2 million counters. The number of concurrent flows in our trace for a typical measurement period is around 1 million. It is obvious that allocating per-flow state is not possible, and each counter has to store the information of multiple flows. However, if an “elephant” flow is mapped to a counter, that counter will overflow and lose information. On the other hand, if only a couple of “mouse” flows are mapped to a counter, the counter will be underutilized, with most of its high-order bits left as zeros.

To solve the above problems, we not only store the information of multiple flows in each counter, but also store the information of each flow in a large number of counters, such that an “elephant” is broken into many “mice” that are stored at different counters. More specifically, we map each flow to a set of l randomly selected counters and split the flow size into l roughly equal shares, each of which is stored in one counter. The value of a counter is the sum of the shares from all flows that are mapped to the counter. Because flows share counters, they introduce noise to each other’s measurement. The key to

accurately estimate the size of a flow is to measure the noise introduced by other flows in the counters to which the flow is mapped.

Fortunately, this can be done if the flows are mapped to the counters uniformly at random. Any two flows will have the same probability of sharing counters, which means that each flow will have the same probability of introducing a certain amount of noise to any other flow. If the number of flows and the number of counters are very large, the combined noise introduced by all flows will be distributed across the counter space about uniformly. The statistically uniform distribution of the noise can be measured and removed. The above scheme of information storage and recovery is called *randomized counter sharing*.

We stress that this design philosophy of “splitting” each flow among a large number of counters is very different from “replicating” each flow in three counters as the counting Bloom filter [30] or counter braids [18] do—they add the size of each flow as a whole to three randomly selected counters. Most notably, our method increments one counter for each arrival packet, while the counting Bloom filter or counter braids increment three counters. We store the information of each flow in many counters (e.g., 50), while they store the information of each flow in three counters.

B. Overall Design

Our online traffic measurement function consists of two modules. The online data encoding module stores the information of arrival packets in an array of counters. For each packet, it performs one hash function to identify a counter and then updates the counter with two memory accesses, one for reading and the other for writing. At the end of each measurement period, the counter array is stored to the disk and then reset to zeros.

The offline data decoding module answers queries for flow sizes. It is performed by a designated offline computer. We propose two methods for separating the information about the size of a flow from the noise in the counters. The first one is called the *counter sum estimation method* (CSM), which is very simple and easy to compute. The second one is called the *maximum likelihood estimation method* (MLM), which is more accurate but also more computationally intensive. The two complementary methods provide flexibility in designing a practical system, which may first use CSM for rough estimations and then apply MLM to the ones of interest.

IV. ONLINE DATA ENCODING

The flow size information is stored in an array C of m counters. The i th counter in the array is denoted as $C[i]$, $0 \leq i \leq m - 1$. The size of the counters should be set so that the chance of overflow is negligible; we will discuss this issue in detail in Section VII. Each flow is mapped to l counters that are randomly selected from C through hash functions. These counters logically form a *storage vector* of the flow, denoted as C_f , where f is the label of the flow. The i th counter of the vector, denoted as $C_f[i]$, $0 \leq i \leq l - 1$, is selected from C as follows:

$$C_f[i] = C[H_i(f)] \quad (1)$$

where $H_i(\dots)$ is a hash function whose range is $[0, m)$. We want to stress that C_f is not a separate array for flow f . It is merely a logical construction from counters in C for the purpose of simplifying the presentation. In all our formulas, one should treat the notation $C_f[i]$ simply as $C[H_i(f)]$. The hash function H_i , $0 \leq i \leq l - 1$, can be implemented from a master function $H(\dots)$ as follows: $H_i(f) = H(f|i)$ or $H_i(f) = H(f \oplus R[i])$, where “ $|$ ” is the concatenation operator, “ \oplus ” is the XOR operator, and $R[i]$ is a constant whose bits differ randomly for different indices i .

All counters are initialized to zeros at the beginning of each measurement period. The operation of online data encoding is very simple: When the router receives a packet, it extracts the flow label f from the packet header, randomly selects a counter from C_f , and increases the counter by one. More specifically, the router randomly picks a number i between 0 and $l - 1$, computes the hash $H_i(f)$, and increases the counter $C[H_i(f)]$, which is physically in the array C , but logically the i th element in the vector C_f .

V. OFFLINE COUNTER SUM ESTIMATION

A. Estimation Method

At the end of a measurement period, the router stores the counter array C to a disk for long-term storage and offline data analysis. Let n be the combined size of all flows, which is $\sum_{i=0}^{m-1} C[i]$. Let s be the true size of a flow f during the measurement period. The estimated size, \hat{s} , based on our CSM is

$$\hat{s} = \sum_{i=0}^{l-1} C_f[i] - l \frac{n}{m}. \quad (2)$$

The first item is the sum of the counters in the storage vector of flow f . It can also be interpreted as the sum of the flow size s and the noise from other flows due to counter sharing. The second item captures the expected noise. We formally derive (2) as follows.

Consider an arbitrary counter in the storage vector of flow f . We treat the value of the counter as a random variable X . Let Y be the portion of X contributed by the packets of flow f , and Z be the portion of X contributed by the packets of other flows. Obviously, $X = Y + Z$.

Each of the s packets in flow f has a probability of $\frac{1}{l}$ to increase the value of the counter by one. Hence, Y follows a binomial distribution

$$Y \sim \text{Bino} \left(s, \frac{1}{l} \right). \quad (3)$$

Each packet of another flow f' has a probability of $\frac{1}{m}$ to increase the counter by one. That is because the probability for the counter to belong to the storage vector of flow f' is $\frac{l}{m}$, and if that happens, the counter has a probability of $\frac{1}{l}$ to be selected for increment. Assume there is a large number of flows, the size of each flow is negligible when comparing with the total size of all flows, and l is large such that each flow’s size is randomly spread among many counters. We can approximately treat the

packets independently. Hence, according to the central limit theorem, Z approximately follows a binomial distribution:

$$Z \sim \text{Bino}\left(n-s, \frac{1}{m}\right) \approx \text{Bino}\left(n, \frac{1}{m}\right), \quad \text{because } s \ll n. \quad (4)$$

We must have

$$E(X) = E(Y + Z) = E(Y) + E(Z) = \frac{s}{l} + \frac{n}{m}. \quad (5)$$

That is

$$s = l \times E(X) - l \frac{n}{m}. \quad (6)$$

From the observed counter values $C_f[i]$, $E(X)$ can be measured as $\frac{\sum_{i=0}^{l-1} C_f[i]}{l}$. We have the following estimation for s :

$$\hat{s} = \sum_{i=0}^{l-1} C_f[i] - l \frac{n}{m}. \quad (7)$$

If a flow shares a counter with an ‘‘elephant’’ flow, its size estimation can be skewed. However, our experiments show that CSM works well in general because the number of ‘‘elephants’’ is typically small (as shown in Fig. 1), and thus their impact is also small, particularly when there are a very large number of counters and flows. Moreover, our next method based on maximum likelihood estimation can effectively reduce the impact of an outlier in a flow’s storage vector that is caused by an ‘‘elephant’’ flow.

B. Estimation Accuracy

The mean and variance of \hat{s} will be given in (9) and (10), respectively. They are derived as follows: Because $X = Y + Z$, we have

$$\begin{aligned} E(X^2) &= E((Y + Z)^2) = E(Y^2) + 2E(YZ) + E(Z^2) \\ &= E(Y^2) + 2E(Y)E(Z) + E(Z^2) \\ &= \frac{s^2}{l^2} - \frac{s}{l^2} + \frac{s}{l} + 2 \cdot \frac{s}{l} \cdot \frac{n}{m} + \frac{n^2}{m^2} - \frac{n}{m^2} + \frac{n}{m}. \end{aligned}$$

The following facts are used in the above mathematical process: $E(Y^2) = \frac{s^2}{l^2} - \frac{s}{l^2} + \frac{s}{l}$ because $Y \sim \text{Bino}(s, 1/l)$. $E(YZ) = E(Y)E(Z)$ since Y and Z are independent. $E(Z^2) = \frac{n^2}{m^2} - \frac{n}{m^2} + \frac{n}{m}$ because $Z \sim \text{Bino}(n, 1/m)$

$$\begin{aligned} \text{Var}(X) &= E(X^2) - (E(X))^2 \\ &= \frac{s}{l} \left(1 - \frac{1}{l}\right) + \frac{n}{m} \left(1 - \frac{1}{m}\right). \end{aligned} \quad (8)$$

In (7), $C_f[i]$, $0 \leq i \leq l-1$, are independent samples of X . We can interpret \hat{s} as a random variable in the sense that a different set of samples of X may result in a different value of \hat{s} . From (7), we have

$$\begin{aligned} E(\hat{s}) &= l \times E(X) - l \frac{n}{m} \\ &= l \left(\frac{s}{l} + \frac{n}{m}\right) - l \frac{n}{m} = s \end{aligned} \quad (9)$$




Fig. 2. Standard deviation with respect to s , when l is set to 50, 100, and 500, respectively.

which means our estimation is unbiased. The variance of \hat{s} can be written as

$$\begin{aligned} \text{Var}(\hat{s}) &= l^2 \times \text{Var}(X) = l^2 \left(\frac{s}{l} \left(1 - \frac{1}{l}\right) + \frac{n}{m} \left(1 - \frac{1}{m}\right) \right) \\ &= s(l-1) + l^2 \frac{n}{m} \left(1 - \frac{1}{m}\right). \end{aligned} \quad (10)$$

The standard deviation, divided by s to show the relative value, is

$$\text{StdDev} \left(\frac{\hat{s}}{s} \right) = \frac{\sqrt{\text{Var}(\hat{s})}}{s} \quad (11)$$

We set l to 50, 100, and 500, respectively, and illustrate the standard deviation of s in Fig. 2.

C. Confidence Interval

The confidence interval for the estimation will be given in (14), and it is derived as follows. The binomial distribution, $Z \sim \text{Bino}(n, 1/m)$, can be closely approximated as a Gaussian distribution, $\text{Norm}\left(\frac{n}{m}, \frac{n}{m} \left(1 - \frac{1}{m}\right)\right)$, when n is large. Similarly, the binomial distribution, $Y \sim \text{Bino}(s, 1/l)$, can be approximated by $\text{Norm}\left(\frac{s}{l}, \frac{s}{l} \left(1 - \frac{1}{l}\right)\right)$. Because the linear combination of two independent Gaussian random variables is also normally distributed [31], we have $X \sim \text{Norm}\left(\frac{s}{l} + \frac{n}{m}, \frac{s}{l} \left(1 - \frac{1}{l}\right) + \frac{n}{m} \left(1 - \frac{1}{m}\right)\right)$. To simplify the presentation, let $\mu = \frac{s}{l} + \frac{n}{m}$ and $\Delta = \frac{s}{l} \left(1 - \frac{1}{l}\right) + \frac{n}{m} \left(1 - \frac{1}{m}\right)$

$$X \sim \text{Norm}(\mu, \Delta) \quad (12)$$

where the mean μ and the variance Δ agree with (5) and (8), respectively.

Because \hat{s} is a linear function of $C_f[i]$, $0 \leq i \leq l-1$, which are independent samples of X , \hat{s} must also approximately follow a Gaussian distribution. From (7) and (12), we have

$$\hat{s} \sim \text{Norm} \left(s, s(l-1) + l^2 \frac{n}{m} \left(1 - \frac{1}{m}\right) \right). \quad (13)$$

Hence, the confidence interval is

$$\hat{s} \pm Z_\alpha \sqrt{s(l-1) + l^2 \frac{n}{m} \left(1 - \frac{1}{m}\right)} \quad (14)$$

where α is the confidence level and Z_α is the α percentile for the standard Gaussian distribution. As an example, when $\alpha = 95\%$, $Z_\alpha = 1.96$.

VI. MAXIMUM LIKELIHOOD ESTIMATION

In this section, we propose the second estimation method that is more accurate, but also more computationally expensive.

A. Estimation Method

We know from Section V that any counter in the storage vector of flow f can be represented by a random variable X , which is the sum of Y and Z , where $Y \sim \text{Bino}(s, \frac{1}{l})$ and $Z \sim \text{Bino}(n, 1/m)$. For any integer $z \in [0, n]$, the probability for the event $Z = z$ to occur can be computed as follows:

$$\Pr\{Z = z\} = \binom{n}{z} \left(\frac{1}{m}\right)^z \left(1 - \frac{1}{m}\right)^{n-z}.$$

Because n and m are known, $\Pr\{Z = z\}$ is a function of a single variable z and thus denoted as $P(z)$.

Based on the probability distribution of Y and Z , the probability for the observed value of a counter, $C_f[i]$, $\forall i \in [0, l]$, to occur is

$$\begin{aligned} \Pr\{X = C_f[i]\} &= \sum_{z=0}^{C_f[i]} (\Pr\{Z = z\} \cdot \Pr\{Y = C_f[i] - z\}) \\ &= \sum_{z=0}^{C_f[i]} \binom{s}{C_f[i]-z} \left(\frac{1}{l}\right)^{C_f[i]-z} \left(1 - \frac{1}{l}\right)^{s-(C_f[i]-z)} P(z). \end{aligned} \quad (15)$$

Let $y = C_f[i] - z$ to simplify the formula. The probability for all observed values in the storage vector of flow f to occur is

$$\begin{aligned} L &= \prod_{i=0}^{l-1} \Pr\{X = C_f[i]\} \\ &= \prod_{i=0}^{l-1} \left(\sum_{z=0}^{C_f[i]} \binom{s}{y} \left(\frac{1}{l}\right)^y \left(1 - \frac{1}{l}\right)^{s-y} P(z) \right). \end{aligned} \quad (16)$$

The MLM is to find an estimated size \hat{s} of flow f that maximizes the above likelihood function. Namely, we want to find

$$\hat{s} = \arg \max_s \{L\} \quad (17)$$

To find \hat{s} , we first apply logarithm to turn the right side of the equation from product to summation

$$\ln(L) = \sum_{i=0}^{l-1} \ln \left(\sum_{z=0}^{C_f[i]} \binom{s}{y} \left(\frac{1}{l}\right)^y \left(1 - \frac{1}{l}\right)^{s-y} P(z) \right). \quad (18)$$

Because $\frac{d}{ds} \binom{s}{y} = \binom{s}{y} (\psi(s+1) - \psi(s+1-y))$, where $\psi(\dots)$ is the polygamma function [32], we have

$$\begin{aligned} &\frac{d}{ds} \left(\binom{s}{y} \left(1 - \frac{1}{l}\right)^{s-y} \right) \\ &= \binom{s}{y} \left(1 - \frac{1}{l}\right)^{s-y} \left(\psi(s+1) - \psi(s+1-y) + \ln \left(1 - \frac{1}{l}\right) \right). \end{aligned}$$

To simplify the presentation, we denote the right side of the above equation as $O(s)$. From (18), we can compute the first-order derivative of $\ln(L)$ as follows:

$$\frac{d \ln(L)}{ds} = \sum_{i=0}^{l-1} \frac{\sum_{z=0}^{C_f[i]} (O(s) \left(\frac{1}{l}\right)^y P(z))}{\sum_{z=0}^{C_f[i]} \binom{s}{y} \left(\frac{1}{l}\right)^y \left(1 - \frac{1}{l}\right)^{s-y} P(z)}. \quad (19)$$

Maximizing L is equivalent to maximizing $\ln(L)$. Hence, by setting the right side of (19) to zero, we can find the value for \hat{s} through numerical methods. Because $\frac{d \ln(L)}{ds}$ is a monotone function of s , we have used the bisection search method in all our experiments in Section IX to find the value \hat{s} that makes $\frac{d \ln(L)}{ds}$ equal to zero.

B. Estimation Accuracy

The estimation confidence interval will be given in (27), and it is derived as follows. The estimation formula is given in (17). According to the classical theory for MLM, when l is sufficiently large, the distribution of the flow size estimation \hat{s} can be approximated by

$$\text{Norm} \left(s, \frac{1}{\mathcal{I}(\hat{s})} \right) \quad (20)$$

where the *fisher information* $\mathcal{I}(\hat{s})$ [33] of L is defined as follows:

$$\mathcal{I}(\hat{s}) = -E \left(\frac{d^2 \ln(L)}{ds^2} \right). \quad (21)$$

In order to compute the second-order derivative, we begin from (12) and have the following:

$$\begin{aligned} \Pr\{X = C_f[i]\} &= \frac{1}{\sqrt{2\pi}\Delta} e^{-\frac{(C_f[i]-\mu)^2}{2\Delta}} \\ \ln(\Pr\{X = C_f[i]\}) &= -\ln(\sqrt{2\pi}\Delta) - \frac{(C_f[i]-\mu)^2}{2\Delta} \end{aligned} \quad (22)$$

where $0 \leq i \leq l-1$. Performing the second-order differentiation, we have

$$\begin{aligned} &\frac{d^2 \ln(\Pr\{X = C_f[i]\})}{ds^2} \\ &= -\frac{\mu'}{l\Delta} + \frac{\left(\frac{1}{2} \left(1 - \frac{1}{l}\right) + \mu - C_f[i]\right) \Delta'}{l\Delta^2} \\ &\quad + \frac{1}{l\Delta^3} \left(1 - \frac{1}{l}\right) \left((\mu - C_f[i])\mu' \Delta - (\mu - C_f[i])^2 \Delta' \right) \end{aligned} \quad (23)$$




Fig. 3. The standard deviation with respect to s , when l is set to 50, 100, and 500, respectively.

where $\mu' = \frac{1}{l}$ and $\Delta' = \frac{1}{l}(1 - \frac{1}{l})$. Therefore

$$\begin{aligned} & E\left(\frac{d^2 \ln(\Pr\{X = C_f[i]\})}{ds^2}\right) \\ &= -\frac{\mu'}{l\Delta} + \frac{\frac{1}{2}(1-\frac{1}{l})\Delta'}{l\Delta^2} + \frac{1}{l\Delta^3} \left(1 - \frac{1}{l}\right) E(\mu - C_f[i])^2 \Delta' \\ &= -\frac{1}{l^2\Delta} + \frac{3(1-\frac{1}{l})^2}{2l^2\Delta^2} \end{aligned} \quad (24)$$

where we have used the following facts: $E(\mu - C_f[i]) = 0$ and $E(\mu - C_f[i])^2 = \Delta$. Because $L = \prod_{i=0}^{l-1} \Pr\{X = C_f[i]\}$, we have

$$\begin{aligned} \mathcal{I}(\hat{s}) &= -E\left(\frac{d^2 \ln(L)}{ds^2}\right) = \sum_{i=0}^{l-1} E\left(\frac{d^2 \ln(\Pr\{X = C_f[i]\})}{ds^2}\right) \\ &= \frac{1}{l\Delta} - \frac{3(1-\frac{1}{l})^2}{2l\Delta^2}. \end{aligned} \quad (25)$$

From (20), the variance of \hat{s} is

$$\text{Var}(\hat{s}) = \frac{1}{\mathcal{I}(\hat{s})} = \frac{2l\Delta^2}{2\Delta - 3(1-\frac{1}{l})^2}. \quad (26)$$

Hence, the confidence interval is

$$\hat{s} \pm Z_\alpha \cdot \sqrt{\frac{2l\Delta^2}{2\Delta - 3(1-\frac{1}{l})^2}} \quad (27)$$

where Z_α is the α percentile for the standard Gaussian distribution.

We set l to 50, 100, and 500, respectively, and illustrates the standard deviation of s in Fig. 3. When comparing to Fig. 2, we find that MLM is able to achieve better estimation accuracy.

VII. SETTING COUNTER LENGTH

So far, our analysis has assumed that each counter has a sufficient number of bits such that it will not overflow. However, in order to save space, we want to set the counter length as short as possible. Suppose each measurement period ends after a pre-specified number n of packets are received. (Note that the value of n is the combined sizes of all flows during each measurement period.) The average value of all counters will be $\frac{n}{m}$. We set the number of bits in each counter, denoted as b , to be $\log_2 \frac{n}{m} + 1$. Due to the additional bit, each counter can hold at least two times

of the average before overflowing. If the allocated memory has M bits, the values of b and m can be determined from the following equations:

$$b \times m = M \quad \log_2 \frac{n}{m} + 1 = b. \quad (28)$$

Due to the randomized counter sharing design, roughly speaking, the packets are distributed in the counters at random. We observe in our experiments that the counter values approximately follow a Gaussian distribution with a mean of $\frac{n}{m}$. In this distribution, the fraction of counters that are more than four times of the mean is very small—less than 5.3% in all our experiments. Consequently, the impact of counter overflow in CSM or MLM is also very small for most flows. Though it is small, we will totally eliminate this impact in Section X-D.

VIII. FLOW LABELS

The compact online data structure introduced in Section IV only stores the flow size information. It does not store the flow labels. The labels are per-flow information, and it cannot be compressed in the same way we do for the flow sizes. In some applications, the flow labels are preknown and do not have to be collected. For example, if an ISP wants to measure the traffic from its customers, it knows their IP addresses (which are the flow labels in this case). Similarly, if the system administrator of a large enterprise network needs the information about the traffic volumes of the hosts in the network, she has the hosts' addresses.

In case the flow labels need to be collected and there is not enough SRAM to keep them, the labels have to be stored in DRAM. An efficient solution for label collection was proposed in [19]. A Bloom filter [34], [35] can be implemented in SRAM to encode the flow labels that have been by the router during a measurement period, such that each label is only stored once in DRAM when it appears for the first time in the packet stream; storing each label once is the minimum overhead if the labels must be collected.

If we use three hash functions in the Bloom filter, each packet incurs three SRAM accesses in order to check whether the flow label that carried the packet is already encoded in the Bloom filter. A recent work on one-memory-access Bloom filters [36] shows that three SRAM accesses per packet can be reduced to one. This overhead is further reduced if we only examine the UDP packets and the SYN packets (which carry the label information of TCP traffic). A recent study shows that UDP accounts for 20% of the Internet traffic [37] and the measurement of our campus traffic shows that SYN packets accounts for less than 10% of all TCP traffic. Therefore, the Bloom filter operation only needs to be carried out for less than 28% of all packets, which amortizes the overhead.

IX. EXPERIMENTS

We use experiments to evaluate our estimation methods, CSM and MLM, which are designed based on the randomized counter sharing scheme. We also compare our methods with Counter Braids (CB) [18] and Multi-Resolution Space-Code Bloom Filters (MRSCBF) [13], [38]. Our evaluation is based on the performance metrics outlined in Section II, including




Fig. 4. (left to right) *First plot*: estimation results by CSE when $M = 2$ Mb. Each flow is represented by a point in the plot, whose x -coordinate is the true flow size s and y -coordinate is the estimated flow size \hat{s} . The equality line, $y = x$, is also shown for reference. An estimation is more accurate if the point is closer to the equality line. *Second plot*: 95% confidence intervals for the estimations made by CSE when $M = 2$ Mb. The width of each vertical bar shows the size of the confidence interval at a certain flow size (which is the x -coordinate of the bar). The y -coordinate of the middle point of each bar shows the mean estimation for all flows of that size. Intuitively, the estimation is more accurate if the confidence interval is smaller, and the middle point is closer to the equality line. *Third plot*: estimation results by MLM when $M = 2$ Mb. *Fourth plot*: 95% confidence intervals for the estimations made by MLM when $M = 2$ Mb. In these experiments, $n = 10$ M.

TABLE I
NUMBER OF MEMORY ACCESSES AND NUMBER OF HASH COMPUTATIONS
PER PACKET

| | memory accesses | hash computations | constant? |
|--------|-----------------|-------------------|-----------|
| CSM | 2 | 1 | Y |
| MLM | 2 | 1 | Y |
| CB | ≥ 6 | ≥ 3 | N |
| MRSCBF | 4.47 | 4.47 | N |

per-packet processing time, memory overhead, and estimation accuracy.

The experiments use a network traffic trace obtained from the main gateway of our campus. We perform experiments on various different types of flows, such as per-source flows, per-destination flows, per-source/destination flows, and TCP flows. They all lead to the same conclusions. Without losing generality, we choose TCP flows for presentation. The trace contains about 68 million TCP flows and 750 million packets. In each measurement period, 10 million packets are processed; it typically covers slightly more than 1 million flows.

A. Processing Time

The processing time is mainly determined by the number of memory accesses and the number of hash computations per packet. Table I presents the comparison. CSM or MLM performs two memory accesses and one hash computation for each packet. CB incurs three times the overhead. It performs six memory accesses and three hash computations for each packet at the first counter level, and in the worst case makes six additional memory accesses and three additional hash computations at the second level. MRSCBF has nine filters. The i th filter uses k_i hash functions and encodes packets with a sampling probability p_i , where $k_1 = 3$, $k_2 = 4$, $k_i = 6$, $\forall i \in [3, 9]$, and $p_i = (\frac{1}{4})^{i-1}$, $\forall i \in [1, 9]$. When encoding a packet, the i th filter performs k_i hash computations and sets k_i bits. Hence, the total number of memory accesses (or hash computations) per packet for all filters is $\sum_{i=1}^9 (p_i \cdot k_i) \approx 4.47$.

B. Memory Overhead and Estimation Accuracy

We study the estimation accuracies of CSM and MLM under different levels of memory availability. In each measurement

period, 10 M packets are processed, i.e., $n = 10$ M, which translates into about 8 s for an OC-192 link (10+ Gb/s) or about 2 s for an OC-768 link (40+ Gb/s) with an average packet size of 1000 B. The memory M allocated to this particular measurement function is varied from 2 Mb (2×2^{20} bits) to 8 Mb. The counter length b and the number of counters m are determined based on (28). The size of each storage vector is 50.

When $M = 2$ Mb, the experimental results are presented in Fig. 4. The first plot from the left shows the estimation results by CSM for one measurement period; the results for other measurement periods are very similar. Each flow is represented by a point in the plot, whose x -coordinate is the true flow size s and y -coordinate is the estimated flow size \hat{s} . The equality line, $y = x$, is also shown for reference. An estimation is more accurate if the point is closer to the equality line.

The second plot presents the 95% confidence intervals for the estimations made by CSM. The width of each vertical bar shows the size of the confidence interval at a certain flow size (which is the x -coordinate of the bar). The middle point of each bar shows the mean estimation for all flows of that size. Intuitively, the estimation is more accurate if the confidence interval is smaller, and the middle point is closer to the equality line.

The third plot shows the estimation results by MLM, and the fourth plot shows the 95% confidence intervals for the estimations made by MLM. Clearly, MLM achieves better accuracy than CSM. The estimation accuracy shown in Fig. 4 is achieved with a memory of slightly less than 2 bits per flow,

We can improve the estimation accuracy of CSM or MLM by using more memory. We increase M to 4 Mb and repeat the above experiments. The results are shown in Fig. 5. We then increase M to 8 Mb and repeat the above experiments. The results are shown in Fig. 6. The accuracy clearly improves as the confidence intervals shrink when M becomes larger.

We repeat the same experiments on CB, whose parameters are selected according to [18]. The results are presented in Fig. 7. The first plot shows that CB totally fails to produce any meaningful results when the available memory is too small: $M = 2$ Mb, which translates into less than 2 bits per flow. In fact, its algorithm cannot converge, but instead produce oscillating results. We have to artificially stop the algorithm after a very long time. The second plot shows that CB works well when $M = 4$ Mb. The algorithm still cannot converge by itself, even




Fig. 5. (left to right) *First plot*: estimation results by CSM when $M = 4$ Mb. *Second plot*: 95% confidence intervals for the estimations made by CSM when $M = 4$ Mb. *Third plot*: estimation results by MLM when $M = 4$ Mb. *Fourth plot*: 95% confidence intervals for the estimations made by MLM when $M = 4$ Mb. See the caption of Fig. 4 for more explanation. In these experiments, $n = 10$ M.




Fig. 6. (left to right) *First plot*: estimation results by CSM when $M = 8$ Mb. *Second plot*: 95% confidence intervals for the estimations made by CSM when $M = 8$ Mb. *Third plot*: estimation results by MLM when $M = 8$ Mb. *Fourth plot*: 95% confidence intervals for the estimations made by MLM when $M = 8$ Mb. See the caption of Fig. 4 for more explanation. In these experiments, $n = 10$ M.




Fig. 7. *Left plot*: estimation results by CB when $M = 2$ Mb. *Middle plot*: estimation results by CB when $M = 4$ Mb. *Right plot*: estimation results by CB when $M = 8$ Mb.

though it can produce very good results when we artificially stop it after a long time without observing any further improvement in the results. It can be seen that the results carry a small positive bias because most points are on one side of the equality line. The third plot shows that CB is able to return the exact sizes for most flows when the memory is $M = 8$ Mb.

Combining the results in Table I, we draw the following conclusion: 1) In practice, we should choose CSM/MLM if the requirement is to handle high measurement throughput (which means low per-packet processing time) or if the available memory is too small such that CB does not work, while relatively coarse estimation is acceptable. 2) We should choose CB if the processing time is less of a concern, sufficient memory is available, and the exact flow sizes are required.

We also run MRSCBF under different levels of memory availability. We begin with $M = 8$ Mb. CSM or MLM works very well with this memory size (Fig. 6). The performance of MRSCBF is shown in the first plot of Fig. 8. There are some very large estimated sizes. To control the scale in the vertical axis, we artificially set any estimation beyond 2800 to be

2800. The results demonstrate that MRSCBF totally fails when $M = 8$ Mb. The performance of MRSCBF improves when we increase the memory. The results when $M = 40$ Mb are shown in the second plot.¹ In the third plot, when we further increase M to 80 Mb,² no obvious improvement is observed when comparing the second plot. A final note is that the original paper of MRSCBF uses log scale in their presentation. The third plot in Fig. 8 will appear as the fourth plot in log scale.

Clearly, the bitmap-based MRSCBF performs worse than CB, CSM or MLM. To measure flow sizes, counters are superior than bitmaps.

X. EXTENSION OF ESTIMATION RANGE

We set the upper bound on the flow size that CSM and MLM can estimate in Section IX to 2500. However, in today's high-speed networks, the sizes of some flows are much larger than

¹At the end of each measurement period, about half of the bits in the filters of MRSCBF are set to ones.

²At the end of each measurement period, less than half of the bits in the filters of MRSCBF are set to ones.

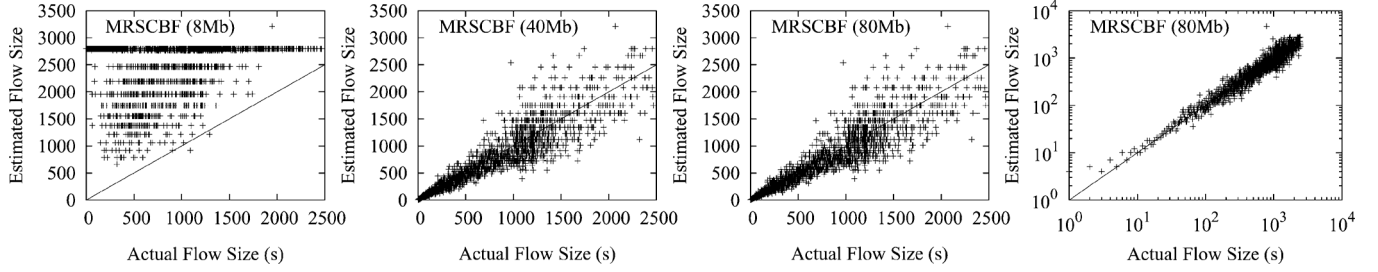


Fig. 8. (left to right) *First plot*: estimation results by MRSCBF when $M = 8$ Mb. *Second plot*: estimation results by MRSCBF when $M = 40$ Mb. *Third plot*: estimation results by MRSCBF when $M = 80$ Mb. *Fourth plot*: estimation results in logarithmic scale by MRSCBF when $M = 80$ Mb.




Fig. 9. (left to right) *First plot*: estimation results by MLM when $b = 6$. *Second plot*: estimation results by MLM when $b = 7$. *Third plot*: estimation results by MLM when $b = 8$. *Fourth plot*: estimation results by MLM when $b = 9$. In these experiments, $n = 10$ M, $M = 4$ Mb.

2500. In order to extend the estimation range to cover these large flows, we propose four approaches that increase the estimation upper bound, and we present extensive experimental results to demonstrate their effectiveness. Since MLM generally performs better than CSE, we only discuss how to extend the estimation range of MLM. CSE can be easily enhanced by similar approaches.

According to Section IV, each flow is assigned a unique storage vector. A flow's storage vector consists of l counters, and each counter has b bits. Therefore, the maximum number of packets that the storage vector can represent is $l \times (2^b - 1)$. If we increase b by one, the number of packets that the vector can represent will be doubled. Similarly, if we increase l by a certain factor, the number of packets that the vector can represent will be increased by the same factor. Based on these observations, we extend the estimation range of MLM by increasing the value of b and l , respectively. In addition, we add a sampling module to MLM and consider hybrid SRAM/DRAM implementation to extend the estimation range.

A. Increasing Counter Size b

Our first approach to extend the estimation range is to enlarge the counter size b . We repeat the same experiment on MLM presented in the third plot of Fig. 5 (Section IX-B), where $M = 4$ Mb, $l = 50$, and $n = 10$ M. This time, instead of computing b from (28), we vary its value from 6 to 9. The new experimental results are shown in Fig. 9. In the first plot, the maximum flow size that MLM can estimate is about 1400 when $b = 6$. In the second plot, where $b = 7$, the maximum flow size is about 2800, which is twice of the maximum flow size that the first plot can achieve. When b is set to 8, the third plot shows that the estimation range of MLM is further extended. The fourth plot shows that, when $b = 9$, the maximum flow size that MLM can estimate does not increase any more when




Fig. 10. *Left plot*: the estimation bias in the experimental results shown in Fig. 9. It is measured as $E(\hat{s} - s)$ with respect to s . *Right plot*: the standard deviation of the experimental results in Fig. 9. It is measured as $\frac{\sqrt{\text{Var}(\hat{s})}}{s}$.

compared to the third plot, which we will explain shortly. The estimation accuracy of the above experiments is presented in Fig. 10, where the first plot shows the estimation bias and the second plot shows the standard deviation of the experimental results in Fig. 9. Generally speaking, both bias and standard deviation increase slightly when b increases.

Since flows share counters in MLM, the size information of one flow in a counter is the noise to other flows that share the same counter. When the amount of memory allocated to MLM is fixed ($M = 4$ Mb in these experiments), a larger value for b will result in a smaller value for m , i.e., the total number of counters is reduced. Hence, each counter has to be shared by more flows, and the average number of packets stored in each counter will increase. That means heavier noise among flows, which degrades the estimation accuracy, as is demonstrated by Fig. 10. Moreover, although a counter with a larger size b can keep track of a larger number of packets, since it also carries more noise, MLM has to subtract more noise from the counter value during the estimation process. As a result, the estimation range cannot be extended indefinitely by simply increasing b , which explains the fact that the maximum flow size that MLM can estimate does not increase when b reaches 9 in Fig. 9.

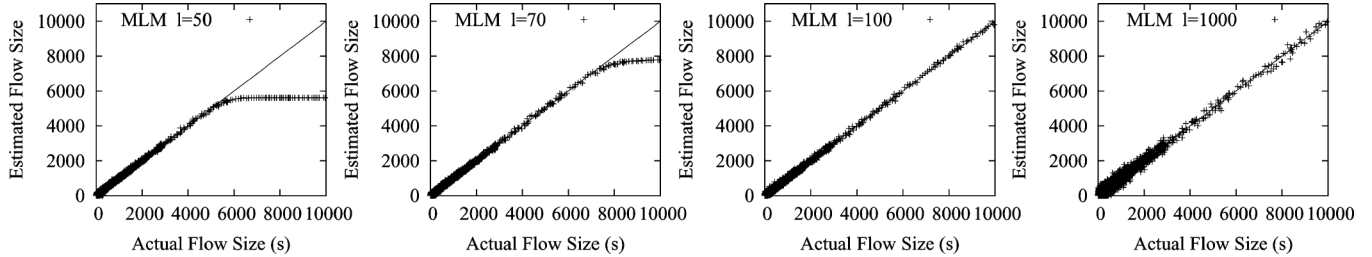


Fig. 11. (left to right) *First plot*: estimation results by MLM when $l = 50$. *Second plot*: estimation results by MLM when $l = 70$. *Third plot*: estimation results by MLM when $l = 100$. *Fourth plot*: estimation results by MLM when $l = 1000$. In these experiments, $n = 10$ M, $M = 4$ Mb.




Fig. 12. *Left plot*: the estimation bias in the experimental results shown in Fig. 11. It is measured as $E(\hat{s} - s)$ with respect to s . *Right plot*: the standard deviation of the experimental results in Fig. 9. It is measured as $\frac{\sqrt{\text{Var}(\hat{s})}}{s}$.

B. Increasing Storage Vector Size l

Our second approach for extending the estimation range is to increase the storage vector size l . We repeat the experiments in Section X-A for MLM with $M = 4$ Mb, $b = 7$, and $n = 10$ M. We vary l from 50 to 1000. Fig. 11 presents the experimental results. The first plot shows that the maximum flow size that MLM can estimate is about 5800 when $l = 50$. As we increase the value of l , MLM can estimate increasingly larger flow sizes. However, when l becomes too large, estimation accuracy will degrade, which is evident in the fourth plot. The reason is that each flow shares too many counters with others, which results in excessive noise in the counters and consequently introduce inaccuracy in the estimation process.

The estimation accuracy of the above experiments is presented in Fig. 12, where the first plot shows the estimation bias and the second plot shows the standard deviation of the experimental results in Fig. 11. Generally speaking, both bias and standard deviation increase slightly when l increases. Clearly, the value of l should not be chosen too large (such as $l = 1000$) in order to prevent estimation accuracy to degrade significantly.

C. Employing Sampling Module

In our third approach, we add a sampling module to MLM to enlarge the estimation range. The sampling technique has been widely used in network measurement [1], [8], [10], [11], [13], [27], [28], [38]. We show that it also works for MLM. Let p be the sampling probability. For each packet that the router receives in the data encoding phase, the router generates a random number r in a range $[0, N]$. If $r < p \times N$, the router processes the packet as we describe in Section IV. Otherwise, it ignores the packet without encoding it in the counter array. In the data decoding phase, the estimated flow size should be $\frac{\hat{s}}{p}$, where \hat{s}

is computed from (19). The estimation range is expanded by a factor of $\frac{1}{p}$.

We again repeat the experiments in the previous sections for MLM with $M = 4$ Mb, $l = 50$, and $n = 10$ M. The value of b is computed from (28). This time, we introduce a sampling probability p and varies its value. Fig. 13 presents the experimental results of MLM with $p = 75\%$, 50% , 25% , and 2% , respectively. It demonstrates that when the sampling probability decreases, the estimation range increases. However, it comes with a penalty on estimation accuracy. Fig. 14 shows the estimation bias and standard deviation of the estimation results in Fig. 13. If the sampling probability is not decreased too small, e.g., when $p \geq 25\%$, the increase in bias and standard deviation is insignificant. However, if the sampling probability becomes too small such as 2% , the degradation in estimation accuracy also becomes noticeable.

D. Hybrid SRAM/DRAM Design

Can we extend the estimation range without any limitation and do so without any degradation in estimation accuracy? This will require a hybrid SRAM/DRAM design. In SRAM, we still choose the value of b based on (28). The limited size of each counter means that a counter may be overflowed during the data encoding phase even though the chance for this to happen is very small (Section VII). To totally eliminate the impact of counter overflow, we keep another array of counters in DRAM, each of which has a sufficient number of bits. The counters in DRAM are one-to-one mapped to the counters in SRAM. Initially, all counters in SRAM and DRAM are set to zero. When a counter in SRAM is overflowed, it is reset to zero and the corresponding counter in DRAM is incremented by one. During offline data analysis, the counter values are set based on both SRAM and DRAM data. For example, when $b = 7$, a counter in SRAM can only store the information of 127 ($2^7 - 1$) packets, when the 128th packet comes, it increases the corresponding counter in DRAM by 1 and reset the current counter to zero. Because overflow happens only to a small fraction of SRAM counters and a DRAM access is made only after an overflowed SRAM counter is accessed 2^b times, the overall overhead of DRAM access is very small.

XI. CONCLUSION

Per-flow traffic measurement provides real-world data for a variety of applications on accounting and billing, anomaly detection, and traffic engineering. Current online data collection methods cannot meet the requirements of being both fast and




Fig. 13. (left to right) First plot: estimation results by MLM when $p = 75\%$. Second plot: estimation results by MLM when $p = 50\%$. Third plot: estimation results by MLM when $p = 25\%$. Fourth plot: estimation results by MLM when $p = 2\%$. In these experiments, $n = 10\text{ M}$, $M = 4\text{ Mb}$.




Fig. 14. Left plot: the estimation bias in the experimental results shown in Fig. 13. It is measured as $E(\hat{s} - s)$ with respect to s . Right plot: the standard deviation of the experimental results in Fig. 13. It is measured as $\frac{\sqrt{\text{Var}(\hat{s})}}{s}$.

compact. This paper proposes a novel data encoding/decoding scheme, which mixes per-flow information randomly in a tight SRAM space for compactness. Its online operation only incurs a small overhead of one hash computation and one counter update per packet. Two offline statistical methods—the counter sum estimation and the maximum likelihood estimation—are used to extract per-flow sizes from the mixed data structures with good accuracy. Due to its fundamentally different design philosophy, the new measurement function is able to work in a tight space where exact measurement is no longer possible, and it does so with the minimal number of memory accesses per packet.


ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their constructive comments.

REFERENCES

- [1] X. Dimitropoulos, P. Hurley, and A. Kind, "Probabilistic lossy counting: An efficient algorithm for finding heavy hitters," *Comput. Commun. Rev.*, vol. 38, no. 1, pp. 7–16, 2008.
- [2] C. Estan and G. Varghese, "New directions in traffic measurement and accounting," in *Proc. ACM SIGCOMM*, Aug. 2002, pp. 323–336.
- [3] B. Krishnamurthy, S. Sen, Y. Zhang, and Y. Chen, "Sketch-based change detection: Methods, evaluation, and applications," in *Proc. ACM SIGCOMM Internet Meas. Conf.*, Oct. 2003, pp. 234–247.
- [4] X. Li, F. Bian, M. Crovella, C. Diot, R. Govindan, G. Iannaccone, and A. Lakhina, "Detection and identification of network anomalies using sketch subspaces," in *Proc. ACM SIGCOMM Internet Meas. Conf.*, Oct. 2006, pp. 147–152.
- [5] M. Yoon, T. Li, S. Chen, and J. Peir, "Fit a spread estimator in small memory," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 504–512.
- [6] N. Kamiyama and T. Mori, "Simple and accurate identification of high-rate flows by packet sampling," in *Proc. IEEE INFOCOM*, Apr. 2006, pp. 1–13.
- [7] R. Karp, S. Shenker, and C. Papadimitriou, "A simple algorithm for finding frequent elements in streams and bags," *Trans. Database Syst.*, vol. 28, no. 1, pp. 51–55, 2003.
- [8] N. Duffield, C. Lund, and M. Thorup, "Estimating flow distributions from sampled flow statistics," in *Proc. ACM SIGCOMM*, Oct. 2003, pp. 325–336.
- [9] A. Kumar, M. Sung, J. Xu, and J. Wang, "Data streaming algorithms for efficient and accurate estimation of flow size distribution," in *Proc. ACM SIGMETRICS*, Jun. 2004, pp. 177–188.
- [10] N. Duffield, C. Lund, and M. Thorup, "Flow sampling under hard resource constraints," in *Proc. ACM SIGMETRICS/Performance*, Jun. 2004, pp. 85–96.
- [11] N. Duffield, C. Lund, and M. Thorup, "Learn more, sample less: Control of volume and variance in network measurement," *IEEE Trans. Inf. Theory*, vol. 51, no. 5, pp. 1756–1775, May 2005.
- [12] W. D. Gardner, "Researchers transmit optical data at 16.4 Tbps," *InformationWeek*, Feb. 2008 [Online]. Available: <http://www.informationweek.com/news/206900603>
- [13] A. Kumar, J. Xu, J. Wang, O. Spatschek, and L. Li, "Space-code bloom filter for efficient per-flow traffic measurement," in *Proc. IEEE INFOCOM*, Mar. 2004, vol. 3, pp. 1762–1773.
- [14] S. Venkatataman, D. Song, P. Gibbons, and A. Blum, "New streaming algorithms for fast detection of superspreaders," in *Proc. NDSS*, Feb. 2005, pp. 149–166.
- [15] D. Shah, S. Iyer, B. Prabhakar, and N. McKeown, "Maintaining statistics counters in router line cards," *IEEE Micro*, vol. 22, no. 1, pp. 76–81, Jan.–Feb. 2002.
- [16] S. Ramabhadran and G. Varghese, "Efficient implementation of a statistics counter architecture," in *Proc. ACM SIGMETRICS*, Jun. 2003, pp. 261–271.
- [17] Q. Zhao, J. Xu, and Z. Liu, "Design of a novel statistics counter architecture with optimal space and time efficiency," in *Proc. ACM SIGMETRICS/Performance*, 2006, pp. 323–334.
- [18] Y. Lu, A. Montanari, B. Prabhakar, S. Dharmapurikar, and A. Kabbani, "Counter braids: A novel counter architecture for per-flow measurement," in *Proc. ACM SIGMETRICS*, Jun. 2008, pp. 121–132.
- [19] Y. Lu and B. Prabhakar, "Robust counting via counter braids: An error-resilient network measurement architecture," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 522–530.
- [20] E. Demaine, A. Lopez-Ortiz, and J. Munro, "Frequency estimation of internet packet streams with limited space," in *Proc. 10th ESA*, Sep. 2002, pp. 348–360.
- [21] M. Kodialam, T. V. Lakshman, and S. Mohanty, "Runs based traffic estimator (RATE): A simple, memory efficient scheme for per-flow rate estimation," in *Proc. IEEE INFOCOM*, Mar. 2004, vol. 3, pp. 1808–1818.
- [22] F. Hao, M. Kodialam, and T. V. Lakshman, "ACCEL-RATE: A faster mechanism for memory efficient per-flow traffic estimation," in *Proc. ACM SIGMETRICS/Performance*, Jun. 2004, pp. 155–166.
- [23] A. Cvetkovski, "An algorithm for approximate counting using limited memory resources," in *Proc. ACM SIGMETRICS*, 2007, pp. 181–190.
- [24] R. Stanojevic, "Small active counters," in *Proc. IEEE INFOCOM*, 2007, pp. 2153–2161.
- [25] H. Zhao, H. Wang, B. Lin, and J. Xu, "Design and performance analysis of a DRAM-based statistics counter array architecture," in *Proc. ACM/IEEE ANCS*, 2009, pp. 84–93.
- [26] C. Estan, G. Varghese, and M. Fish, "Bitmap algorithms for counting active flows on high-speed links," *IEEE/ACM Trans. Netw.*, vol. 14, no. 5, pp. 925–937, Oct. 2006.

- [27] Q. Zhao, A. Kumar, and J. Xu, "Joint data streaming and sampling techniques for detection of super sources and destinations," in *Proc. USENIX/ACM Internet Meas. Conf.*, Oct. 2005, p. 7.
- [28] J. Cao, Y. Jin, A. Chen, T. Bu, and Z. Zhang, "Identifying high cardinality Internet hosts," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 810–818.
- [29] M. Ramakrishna, E. Fu, and E. Bahcekapili, "Efficient hardware hashing functions for high performance computers," *IEEE Trans. Comput.*, vol. 46, no. 12, pp. 1378–1381, Dec. 1997.
- [30] S. Cohen and Y. Matias, "Spectral bloom filters," in *Proc. ACM SIGMOD*, Jun. 2003, pp. 241–252.
- [31] G. Casella and R. Berger, *Statistical Inference*. Pacific Grove, CA: Duxbury, Jun. 2001.
- [32] M. Abramowitz and I. Stegun, *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*. New York: Dover, Jun. 1964.
- [33] E. Lehmann and G. Casella, *Theory of Point Estimation*. New York: Springer, 1998.
- [34] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Commun. ACM*, vol. 13, no. 7, pp. 422–426, 1970.
- [35] A. Broder and M. Mitzenmacher, "Network applications of bloom filters: A survey," *Internet Math.*, vol. 1, no. 4, pp. 485–509, 2002.
- [36] Y. Qiao, S. Chen, and T. Li, "One memory access bloom filters and their generalization," in *Proc. IEEE INFOCOM*, 2011, pp. 1745–1753.
- [37] CAIDA, La Jolla, CA, "Analyzing UDP usage in Internet traffic," May 2009 [Online]. Available: <http://www.caida.org/research/traffic-analysis/tcpudpratio/>
- [38] A. Kumar, J. Xu, J. Wang, O. Spatschek, and L. Li, "Space-code bloom filter for efficient per-flow traffic measurement," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 12, pp. 2327–2339, Dec. 2006.



Tao Li received the B.S. degree in computer science from the University of Science and Technology of China, Hefei, China, in 2007, and is currently pursuing the Ph.D. degree in computer and information science and engineering at the University of Florida, Gainesville.

His advisor is Dr. Shigang Chen, and his research interests include network traffic measurement and RFID technologies.



Shigang Chen received the B.S. degree from the University of Science and Technology of China, Hefei, China, in 1993, and the M.S. and Ph.D. degrees from the University of Illinois at Urbana–Champaign in 1996 and 1999, respectively, all in computer science.

After graduation, he worked with Cisco Systems for three years before joining the University of Florida, Gainesville, in 2002, where he is currently an Associate Professor with the Department of Computer and Information Science and Engineering. His research interests include wireless networks, network protocols and algorithms, and distributed computing.

Dr. Chen is an Associate Editor of the IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY and *Computer Networks*. He served as a TPC Co-Chair for IEEE IWQoS 2009 and the Computer and Network Security Symposium of IEEE IWCC 2006, a Vice TPC Chair for IEEE MASS 2005, a Vice General Chair for QShine 2005, a TPC Co-Chair for QShine 2004, and a TPC member for many conferences including IEEE ICNP, IEEE INFOCOM, IEEE ICDCS, IEEE ICC, IEEE GLOBECOM, etc. He received the IEEE Communications Society Best Tutorial Paper Award in 1999 and the NSF CAREER Award in 2007.



Yibei Ling (M'00–SM'06) received the B.S. degree in electrical engineering from Zhejiang University, Hangzhou, China, in 1982, the M.S. degree in biostatistics from Shanghai Medical University, Shanghai, China, in 1985, and the Ph.D. degree in computer science from Florida State University at Miami in 1995.

Currently, he is a Senior Research Scientist with Applied Research Laboratories, Telcordia Technologies (formerly Bellcore), Morristown, NJ, where he is developing wireless middleware systems and wireless messaging systems. He has published several papers in the IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING, *SIGMOD*, *Data Engineering, Information Systems*, and *Operating System Review*. His research interests include distributed systems, system performance evaluation, query optimization in database management systems, and biological modeling.