

Data Streaming Algorithms for Estimating Entropy of Network Traffic

Ashwin Lall
University of Rochester

Vyas Sekar
Carnegie Mellon University

Mitsunori Ogihara*
University of Rochester

Jun (Jim) Xu[†]
Georgia Inst. of Technology

Hui Zhang[‡]
Carnegie Mellon University

ABSTRACT

Using entropy of traffic distributions has been shown to aid a wide variety of network monitoring applications such as anomaly detection, clustering to reveal interesting patterns, and traffic classification. However, realizing this potential benefit in practice requires accurate algorithms that can operate on high-speed links, with low CPU and memory requirements. In this paper, we investigate the problem of estimating the entropy in a streaming computation model. We give lower bounds for this problem, showing that neither approximation nor randomization alone will let us compute the entropy efficiently. We present two algorithms for randomly approximating the entropy in a time and space efficient manner, applicable for use on very high speed (greater than OC-48) links. The first algorithm for entropy estimation is inspired by the structural similarity with the seminal work of Alon et al. for estimating frequency moments, and we provide strong theoretical guarantees on the error and resource usage. Our second algorithm utilizes the observation that the performance of the streaming algorithm can be enhanced by separating the high-frequency items (or elephants) from the low-frequency items (or mice). We evaluate our algorithms on traffic traces from different deployment scenarios.

Categories and Subject Descriptors

C.2.3 [Computer Systems Organization]: Computer-Communication Networks: Network Operations—Network Monitoring;

*Supported in part by grants Xerox/NYSRAT #C040130 and NSF-EIA-0205061.

[†]Supported in part by NSF grant NETS-NBD 0519745 and NSF CAREER Award ANI 0238315.

[‡]Supported in part by grants NSF CNS-0433540 and ANI-0331653 and U.S. Army Research Office contract number DAAD19-02-1-0389.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS/Performance'06, June 26–30, 2006, Saint Malo, France.
Copyright 2006 ACM 1-59593-320-4/06/0006 ...\$5.00.

General Terms

Algorithms, Measurement, Theory

Keywords

Traffic Analysis, Data Streaming

1. INTRODUCTION

In network traffic flow analysis there has been a shift of focus from simple volume-based analysis to network flow distribution-based analysis. Much work has been published for making inference about the network status from such statistics [12, 17, 24]. Intrinsically, distribution-based analysis could capture the network status more succinctly than volume-based analysis would, but it requires appropriate metrics to encapsulate and capture features of the underlying traffic distribution.

The standard quantities in assessing distributions are the moments (the mean, standard deviation, skewness, kurtosis, etc.). A number of recent empirical studies [7, 17, 23, 24] have suggested the use of *entropy* as a succinct means of summarizing traffic distributions for different applications, in particular, in anomaly detection and in fine-grained traffic analysis and classification. With respect to anomaly detection [17], the use of entropy for tracking changes in traffic distributions provides two significant benefits. First, the use of entropy can increase the sensitivity of detection to uncover anomalous incidents that may not manifest as volume anomalies. Second, using such traffic features provides additional diagnostic information into the nature of the anomalous incidents (e.g., making distinction among worms, DDoS attacks, and scans) that is not available from just volume-based anomaly detection. With respect to fine-grained traffic analysis and traffic classification [24], the entropy of traffic feature distributions offers useful information to measure distance among (traffic) clusters.

While these recent studies demonstrate that using the entropy of traffic distributions has tremendous value for network monitoring applications, realizing the potential benefit requires efficient algorithms for computing the entropy. In general, computing traffic statistics on high-speed links is a hard task, because it is infeasible for traditional methods to keep up with the line-rates, due to constraints on available processing capacity. In addition, constraints imposed on memory make it almost impossible to compute the statistics per flow, or even to maintain per-flow state. Then, the

use of sampling comes as a natural solution. Sampling based methods [5, 6] have been shown to be able to reduce the processing and memory requirements, and to be suitable for capturing some traffic statistics. However, one must trade off accuracy for efficiency—the estimates obtained from sampled data may have large errors [10].

One may then naturally wonder whether there are efficient methods for accurately estimating the entropy. In particular, we ask the following questions:

- What amount of resources (time and space) do we provably need to capture the entropy of a stream of packets on a high-speed link?
- Are there efficient algorithms for entropy computation that can operate on high-speed links which have low memory and CPU costs?

To address these questions, data streaming algorithms assume significance. Data streaming algorithms [19] for computing different statistics over input streams have recently received tremendous interest from the networking and theory communities. Data streaming algorithms have the desirable property that both the computational and memory requirements are low. This property makes them ideal for such high-speed monitoring applications. They are also guaranteed to work with *any* distribution, which makes them useful in dealing with data for which the distribution is not known.

The contribution of this paper is the investigation and application of streaming algorithms to compute the entropy over network traffic streams. The challenge is to design algorithms for estimating entropy that are lightweight in terms of both memory and computational complexity. We present two algorithms for computing the entropy in a streaming model. The first algorithm is based on the insight that estimating the entropy shares structural similarity with the well-known problem of estimating the frequency moments [2].¹ Despite the apparent structural similarity, providing theoretical approximation and resource guarantees for entropy estimation is a challenging task. Our contributions are the identification of appropriate estimator functions for calculating the entropy accurately, and providing proofs of approximation guarantees and resource usage. The theoretical guarantees hold for arbitrary streams, without making any assumptions regarding the underlying distributions and structural properties of their distribution.

Network traffic data-streams have considerable underlying structure (e.g., they may have a Zipfian or power-law distribution), which suggests that we can optimize algorithms further by leveraging this fact. Our second algorithm builds on the basic streaming algorithm, but can substantially improve the efficiency based on techniques for separating the large (elephant) flows from the small (mice) flows. We use a lightweight sampling method that enables sieving out the elephant flows from the stream, and extend the earlier algorithm to utilize this separation to achieve better performance in practice.

We evaluate our algorithms on real traffic traces collected from three different deployment scenarios. The first streaming algorithm outperforms traditional sampling based approaches, and provides much lower estimation errors while

¹This approach has also been independently proposed by Chakrabarti et al. [1]. We will discuss this and other approaches in Section 8, highlighting that while our intellectual trails cross each other on some results, our approaches and evaluations differ substantially in others.

using similar (or lesser) memory resources. Interestingly, we notice that the observed errors are an order of magnitude smaller than the theoretical error guarantees. While it has proved difficult to provide rigorous theoretical (i.e., worst-case) guarantees for the second algorithm (which makes use of the elephant-mice separation), we find that the observed errors are further reduced with this approach.

The remainder of this paper is organized as follows. We introduce the notation that we will use and formally define the problem in Section 2. In Section 3 we prove that any (deterministic) approximation algorithm or (exact) randomized algorithm must use a linear amount of space. Section 4 outlines the basic streaming algorithm and provides theoretical approximation guarantees, while Section 5 provides improvements based on the technique of separating the elephant and mice flows. We evaluate our algorithms on real-world traces in Section 6, confirming the effectiveness of our approaches. We discuss some features of our algorithms in Section 7 and related work in Section 8, before concluding in Section 9.

2. PROBLEM FORMULATION

We first outline the notation used in the remainder of the paper, and formulate the problem of estimating entropy in a streaming context. Throughout this paper we will assume that all items coming over the stream are drawn from the set $[n] = \{1, 2, 3, \dots, n\}$. For example if we are interested in measuring the entropy of packets over various application ports, then n is the number of ports (maximum of 65535 ports for each protocol). Similarly, if we are interested in measuring the entropy of packets over unique source or destination addresses in the traffic stream, then n would have a maximum value of 2^{32} for 32-bit IPv4 addresses.

We will denote the frequency of item $i \in [n]$ (e.g., the number of packets seen at port i) by m_i and the total number of items in the stream by m , i.e., $m = \sum_{i=1}^n m_i$. The j th item observed in the stream will be denoted by $a_j \in [n]$. We define n_0 to be the number of distinct items that appear in the stream, since it is possible that not all n items are present. As a simple example consider a stream drawn from a set of $n = 4$ different possible objects $\{A, B, C, D\}$. Let the stream $X = (A, A, B, B, C, A, B, A, C)$. For this stream, the total number of items $m = 4 + 3 + 2 = 9$, with the number of distinct items $n_0 = 3$. Note that all our analysis is in terms of m , rather than n , since in general $n \gg m$.

The natural definition of entropy (sometimes referred to as sample entropy) in this setting is the expression $H \equiv -\sum_{i=1}^n \frac{m_i}{m} \log \left(\frac{m_i}{m} \right)$. Intuitively, the entropy is a measure of the diversity or randomness of the data coming over the stream. The entropy attains its minimum value of zero when all the items coming over the stream are the same and its maximum value of $\log m$ when all the items in the stream are distinct. Unless otherwise specified, all logarithms in this paper are to the base 2 and we define $0 \log 0 = 0$. For our example stream X , the entropy $H(X) = -(4/9) \log (4/9) - (3/9) \log (3/9) - (2/9) \log (2/9) = 1.53$. Often it is useful to normalize this number to compare entropy estimates across different measurement epochs. For this purpose, we define the standardized entropy to be $H/\log m$. In our example, the standardized entropy is $1.53/\log 9 = 0.48$.

To compute the entropy,

$$\begin{aligned} H &= -\sum_{i=1}^n \frac{m_i}{m} \log\left(\frac{m_i}{m}\right) \\ &= \frac{-1}{m} \left[\sum_i m_i \log m_i - \sum_i m_i \log m \right] \\ &= \log(m) - \frac{1}{m} \sum_i m_i \log m_i, \end{aligned}$$

it suffices to compute $S \equiv \sum_i m_i \log m_i$, since we can keep a count of m exactly with $\log m$ bits. For the remainder of this paper we will concern ourselves with estimating the value S . The measure of accuracy we use to evaluate our estimates is the notion of relative error, which is defined to be $|S - \tilde{S}|/S$, where \tilde{S} is the estimated value and S the true value. For practical applications in traffic monitoring, we require that the relative error be low (say less than 2-3%), so that the accuracy of applications such as anomaly detection and traffic clustering is not affected.

An accurate estimate of S may not necessarily give an accurate estimate of H . In particular, when H is very small and S is close to its maximum value, a small relative error estimate of S may not correspond to a small relative error estimation of H . Let \tilde{S} be the estimated value of S and \tilde{H} the estimated value of H computed from \tilde{S} , i.e., $\tilde{H} = \log(m) - \tilde{S}/m$. Suppose we have an algorithm to compute S with relative error at most ϵ . Then, the relative error in estimating H can be bounded as follows:

$$\begin{aligned} \frac{|H - \tilde{H}|}{H} &= \frac{|\log(m) - S/m - \log(m) + \tilde{S}/m|}{H} \\ &= \frac{|S - \tilde{S}|}{Hm} \\ &\leq \epsilon \frac{S}{Hm}. \end{aligned}$$

Note that the relative error in H actually depends on the ratio $\frac{S}{Hm}$, which can theoretically become arbitrarily high if H is close to zero. However, given reasonable lower bounds for how small H can get, an algorithm that can give an approximation of S with relative error at most ϵ can be converted to one that gives an approximation of H with relative error $\epsilon' = \Theta(\epsilon)$. Specifically, since we know that $S \leq m \log m$, if we assume a lower bound of $\alpha \log m$ for H (for some constant α) then the relative error in estimating H is at most $\epsilon' = \epsilon/\alpha$. Thus any approximation scheme for S can be converted to one for H if we can assume a lower bound on the entropy. Our evaluations (Section 6.2) confirm that the errors for H and S are comparable.

3. LOWER BOUNDS

In this paper we will present a randomized approximation algorithm that uses $O(\log m)$ space for computing the value S of a stream. Before we do this, we would like to answer the first question of how much effort is required to estimate the entropy of a given traffic distribution. We will demonstrate that any exact randomized algorithm or any deterministic approximation algorithm needs at least linear (in the length of the stream) space. This motivates the need to use both randomization and approximation.

We first demonstrate that any randomized algorithm to compute S must use $\Omega(m)$ space by reducing the commu-

nication complexity problem of set intersection to it. Using communication complexity is a common way to prove lower bounds for streaming algorithms [2, 18]. We show here how to apply it to the computation of S (and hence the entropy H).

In the communication complexity model two parties (typically called Alice and Bob), who have non-overlapping but jointly complete parts of the input, wish to compute some function of the input. The communication complexity of the function at input size n is then the largest number of bits that the parties have to communicate using the *best* protocol to compute the function, for any input of size n . There are no bounds on the computational power of either party and the only resource being measured is the number of bits communicated.

For the problem of set intersection, Alice and Bob have subsets A and B of $\{1, \dots, N\}$ as input. The question is then whether the sets A and B have any elements in common. It is known that the deterministic communication complexity of this problem is $\Theta(N)$ [15]. It was shown by Kalyanasundaram and Schnitger in [11] that any communication complexity protocol for set intersection that has probability of error at most δ , for any $\delta < 1/2$, must use $\Omega(N)$ bits of communication. We make use of this result in the proof.

THEOREM 1. *Any randomized streaming algorithm to compute the exact value of S when there are at most m items must use $\Omega(m)$ bits of space.*

PROOF. Let us assume that we have a randomized streaming algorithm that computes $S = \sum_i m_i \log m_i$ for any stream exactly using s bits of space. This gives rise to a communication complexity protocol, using $\Theta(s)$ bits of communication, for computing set intersection that works as follows. Suppose that Alice and Bob have as input subsets of the set $\{1, \dots, m/2\}$. Alice simulates the algorithm using her set (in any arbitrary order) as input into the algorithm and sends the saved state of the algorithm (at most $\Theta(s)$ bits) to Bob. Bob then restarts the algorithm, starting with that saved state, and enters his entire set. At the end of this run, Bob checks the output of the algorithm—if the output is zero, he outputs “disjoint,” otherwise he outputs “not disjoint.”

The above protocol relies on the fact that any items that have frequency at most one do not count toward the sum S (since $1 \log 1 = 0 \log 0 = 0$). So, the value of S computed is exactly twice the size of the intersection. If we find that the intersection has size zero then we know that Alice and Bob’s sets are disjoint, otherwise they have something in common. Hence, even if the streaming algorithm is randomized, it must use $s = \Omega(m)$ bits. If it used fewer bits it would lead to a randomized protocol for set intersection with less than $\Omega(N)$ communication, which we know from [11] to be impossible. \square

THEOREM 2. *Any deterministic streaming algorithm to approximate S with relative error less than $1/3$ must use $\Omega(m)$ bits of space.*

PROOF. The proof that any (non-randomized) approximation algorithm is inefficient is similar to the proof of Proposition 3.7 in [2]. Let G be a family of $2^{\Theta(m)}$ subsets of $\{1, \dots, 2m\}$, such that each subset has cardinality $m/2$ and any pair of distinct subsets has at most $m/4$ elements in common. (It is possible to show such a G exists using the probabilistic method.)

Let us assume for a contradiction that there exists a deterministic streaming algorithm that estimates S with relative error at most $1/3$, using less than linear (in m) space. For every pair of elements $G_1, G_2 \in G$, let $A(G_1, G_2)$ be the sequence of length m consisting of the elements of G_1 in sorted order followed by the elements of G_2 in sorted order. By the pigeonhole principle, if the memory used by the algorithm has less than $\log |G| = \Omega(m)$ bits, then at least two distinct subsets $G_i, G_j \in G$ result in the same memory configuration when their contents are entered into the algorithm. Hence, the algorithm cannot distinguish between the streams $A(G_i, G_i)$ and $A(G_j, G_i)$. For the input $A(G_i, G_i)$ we have that $S = (m/2)(2 \log 2) = m$, but for $A(G_j, G_i)$, $S \leq (m/4)(2 \log 2) = m/2$. Now, if the relative error for $A(G_i, G_i)$ is less than $1/3$, its estimated value is more than $2m/3$, but if the relative error for $A(G_j, G_i)$ is less than $1/3$ its estimated value is less than $2m/3$. Therefore, the algorithm makes a relative error of at least $1/3$ on at least one of these inputs. This tells us that any non-randomized algorithm must either use $\Omega(m)$ space or have a relative error of at least $1/3$. \square

Thus, we see that if we use only randomization or only approximation we cannot hope to use a sublinear amount of space. As a result, the following algorithms that we present are both randomized and approximate. Fortunately, when we allow for these two relaxations we get algorithms that are sublinear (in particular, polylogarithmic) in space and time per item.

4. A STREAMING ALGORITHM

In this section we present our first algorithm and show guarantees on the performance and the size of the memory footprint. The basic algorithm is based on the key insight that estimating S is structurally similar to estimating the frequency moments [2]. The advantage of this technique is that it gives an unbiased estimate of the entropy, with strong theoretical guarantees on the space consumption based upon the desired accuracy of the algorithm. We then show how the assumptions and analysis of the algorithm can be further tightened.

4.1 Algorithm

As demonstrated in the previous section, randomization and approximation alone do not allow us to estimate S efficiently. Hence, we present an algorithm that is an (ϵ, δ) -approximation of S . An (ϵ, δ) -approximation algorithm is one that has a relative error of at most ϵ with probability at least $1 - \delta$, i.e., $\Pr(|X - \tilde{X}| \leq \epsilon X) \geq 1 - \delta$, where X and \tilde{X} are the real and estimated values, respectively.

This algorithm uses the idea of the celebrated Alon–Matias–Szegedy frequency moment estimation algorithm [2]. Conceptually, the algorithm can be divided into three stages. In the first stage we select random locations in the stream. These locations decide the set of counters that the algorithm tracks during the online stage. In the second stage, the online stage, we keep track of the number of occurrences of items that appear at the randomly selected locations. For each selected item, we keep an exact counter for the number of subsequent occurrences of that item. For example, if position k in the stream was selected, we keep an exact counter for the item at position k (denoted as a_k) for the remainder of the stream (i.e., between locations k and m). In the third

Algorithm 1: The streaming algorithm

- 1: **Pre-processing stage**
- 2: $z := \lceil 32 \log m / \epsilon^2 \rceil$, $g := 2 \log(1/\delta)$
- 3: choose $z * g$ locations in the stream at random
- 4: **Online stage**
- 5: **for** each item a_j in the stream **do**
- 6: **if** a_j already has one or more counters **then**
- 7: increment all of a_j 's counters
- 8: **if** j is one of the randomly chosen locations **then**
- 9: start keeping a count for a_j , initialized at 1
- 10: **Post-processing stage**
- 11: // View the $g * z$ counts as a matrix c of size $g \times z$
- 12: **for** $i := 1$ to g **do**
- 13: **for** $j := 1$ to z **do**
- 14: $X_{i,j} := m * (c_{i,j} \log c_{i,j} - (c_{i,j} - 1) \log(c_{i,j} - 1))$
- 15: **for** $i := 1$ to g **do**
- 16: $avg[i] :=$ the average of the X s in group i
- 17: return the median of $avg[1], \dots, avg[g]$

and final stage the algorithm uses the various counters it has tracked to obtain an estimator for the S value of the stream. The goal of the post-processing or estimating stage is to obtain an estimate of S that is unbiased and whose error is provably low.

We present the pseudocode for this algorithm in Algorithm 1. In the pre-processing stage we need to choose $z * g$ locations in the stream. Note that for this stage we need to know the length of the stream to both compute z and to choose the random locations. The choice of the random locations can be deferred as described in [2] and to compute z we can use a safe overestimate for $\log m$ without increasing the space too much. In the online stage, for each such position we keep a counter c for that item from that position on. We update at most one record per item during the online stage, using a data structure described in the following section.

In the post-processing stage, for each of the tracked counters we compute an unbiased estimator for S as follows:

$$X = m(c \log c - (c - 1) \log(c - 1)).$$

These $g * z$ unbiased estimators are then divided into g groups each containing z variables. First we compute the average over each of the g groups, and then obtain the median of the groups as our returned estimate for S . Intuitively, the estimator variable X provides us an unbiased estimate of S , but does not give good guarantees on the variance, and hence the relative error. By computing many such estimates, and obtaining the median over the averages of multiple groups, we can provide rigorous guarantees on the error as we will see in Section 4.3.

4.2 Implementation Details

One major advantage of this algorithm is that it is light weight. For any item in the stream, the algorithm has to update its count if the item is being counted. Checking whether the item is being counted can be done very quickly using a hash table. However, it is possible that a single item has multiple records for it. In the worst case, we would need to update every record for each item. We could greatly improve the efficiency of the algorithm by instead keeping

a single record for every unique item. This can be implemented by only updating the most recent record for that item and maintaining a pointer to the next most recent record. When the entire stream has been processed, the counts for the older records can be reconstructed from those of the newer ones.

The record data structure that we suggest is illustrated in Figure 1. Each record in our implementation would require ~ 200 bits because we would need to store the item label ITEM_LABEL (~ 100 bits), the counter for the item COUNTER (32 bits), a pointer CHAINING_PTR (32 bits) to resolve hash collisions if we use chaining and another pointer PREV_PTR (32 bits) to point to the older records for the item. We use a conservative estimate of 100 bits for each item label, assuming that we would store all 5 main IP packet header fields, i.e., $\langle \text{srcaddr}, \text{dstaddr}, \text{sreport}, \text{dstport}, \text{protocol} \rangle$.




Figure 1: The record data structure

At the end of each epoch the algorithm needs to perform the operations of averaging and finding the median of a list. However, both these operations only need to be done in the post-processing step. If we make an epoch sufficiently large, then these computations need be done relatively infrequently.

4.3 Theoretical Guarantees

We present analysis that shows we can give strong guarantees while using very little space. The proof is along the lines of the one in [2] and the main contribution here is to show how the variance of the variable X can be bounded to give such a small space requirement. The proof requires the assumption that $S \geq m$ or, equivalently, that $H \leq \log m - 1$. We show in Section 4.5 why this assumption is reasonable.

THEOREM 3. *If we assume that $S \geq m$, then Algorithm 1 is an (ϵ, δ) -approximation algorithm for S that uses $O(\log m \log(1/\delta)/\epsilon^2)$ records.*

PROOF. We will first show that the variable X is an unbiased estimator for S . We will then make use of Chebyshev's inequality to bound the probability of having a relative error greater than ϵ . Next, we show that if we average $z = \lceil 32 \log m / \epsilon^2 \rceil$ variables, this probability is at most $1/8$. We can then use Chernoff bounds to show that if we take $g = 2 \log(1/\delta)$ such averages, with probability at least $1 - \delta$ more than half of them have less than ϵ relative error. In this case, the median of the averages must have relative error less than ϵ .

We first observe that the expected value of each variable X is an unbiased estimate of our desired quantity S :

$$\begin{aligned} E[X] &= \frac{m}{m} \sum_{i=1}^n \sum_{j=1}^{m_i} (j \log j - (j-1) \log(j-1)) \\ &= \sum_{i=1}^n m_i \log m_i \\ &= S. \end{aligned}$$

To make use of Chebyshev's inequality, we need to bound the variance of X from above, in terms of S^2 . The bound proceeds as follows:

$$\begin{aligned} \text{Var}(X) &= E(X^2) - E(X)^2 \\ &\leq E(X^2) \\ &= \frac{m^2}{m} \left[\sum_{j=1}^n \sum_{i=2}^{m_j} (i \log i - (i-1) \log(i-1))^2 \right]. \end{aligned}$$

Now we observe that

$$\begin{aligned} n \log n - (n-1) \log(n-1) &= \log \frac{n^n}{(n-1)^{n-1}} \\ &\leq \log \frac{n^n}{n^{n-2}} \\ &= 2 \log n, \end{aligned} \tag{1}$$

where the inequality comes from the facts that the logarithm function is monotonically increasing and that for all $n > 1$, $n^{n-2} \leq (n-1)^{n-1}$, which is proven as follows:

For $n = 2$ the fact can easily be checked. For all other n , $n > e$, so

$$\frac{n^{n-2}}{(n-1)^{n-1}} = \frac{1}{n} \left(\frac{n}{n-1} \right)^{n-1} = \frac{1}{n} \left(1 + \frac{1}{n-1} \right)^{n-1}.$$

This is at most $e/n \leq 1$. So, the inequality holds.

Now, substituting (1) into the bound on the variance, we get that

$$\begin{aligned} \text{Var}(X) &\leq m \sum_{i=1}^n \sum_{j=2}^{m_i} (2 \log j)^2 \\ &\leq 4m \sum_{i=1}^n m_i \log^2 m_i \\ &\leq 4m \log m \left(\sum_i m_i \log m_i \right) \\ &\leq 4S \log m \left(\sum_i m_i \log m_i \right) \\ &= 4S^2 \log m, \end{aligned}$$

where for the last inequality we make use of our assumption that $S \geq m$.

Let the average of the i th group be Y_i . We know that $\text{Var}(Y_i) = \text{Var}(X)/z$ and that it is also an unbiased estimator of S . Applying Chebyshev's inequality, we get that

for each Y_i ,

$$\begin{aligned} Pr(|Y_i - S| > \epsilon S) &\leq \frac{Var(Y_i)}{\epsilon^2 S^2} \\ &\leq \frac{4S^2 \log m}{z\epsilon^2 S^2} \\ &= \frac{4 \log m}{z\epsilon^2} \\ &\leq \frac{1}{8}. \end{aligned}$$

Now, by Chernoff bounds we get that with probability at least $1 - \delta$, at least $g/2$ of the averages have at most ϵ relative error. Hence, the median of the averages has relative error at most ϵ with probability at least $1 - \delta$. \square

Note that if we had chosen $z = \lceil \log m / (\epsilon^2 * \delta) \rceil$ we could have guaranteed an error probability of at most δ with just this one bigger group. While the analysis in the proof works well for smaller δ (i.e., $\delta \leq 1/128$), for practical applications we may want to use larger δ . Because of the independence of each run, with $\delta = 10\%$ we detect anomalous entropy values within one epoch with 90% certainty, within two epochs with 99% certainty and so on. For the case where δ is greater than $1/128 \approx 0.8\%$ we can use the average of a single group of $z = \lceil \log m / (\epsilon^2 * \delta) \rceil$ estimators for our estimate.

The total space (in bits) used by this algorithm is

$$O\left(\frac{\log m \log(1/\delta)}{\epsilon^2} (\log n + \log m)\right).$$

For fixed δ and ϵ this algorithm uses $O(\log m)$ records of size $O(\log m + \log n)$ bits.

Numerical Illustration: To put this into a practical perspective, let us consider an example where we have a stream of length $m = 2^{26} \approx 67$ million, with $n_0 = 6$ million distinct items. To compute the entropy exactly, we could have to maintain counts for each item using 6 million item labels and counters ($132 \text{ bits/record} \times 6 \text{ million records} = 94 \text{ MB}$). Using Algorithm 1 we could approximate the entropy with at most 25% relative error at least 75% of the time with 54 thousand records or 1.4 MB, using 200 bit records as discussed earlier.

4.4 Exact Space Bounds

In practical settings we want to know the exact values of the parameters of the above algorithm so that we use as little space as possible. We tighten the bound on the number of groups needed by making the observation that $\frac{j^j}{(j-1)^{j-1}} = j(1 + \frac{1}{j-1})^{j-1} < ej$. Here is a tighter (non-asymptotic) analysis for the bound on the variance:

THEOREM 4. *If we assume that $S \geq m$, then Algorithm 1 can be modified to use exactly $\lceil \frac{(16 \log m + 64) \log(1/\delta)}{\epsilon^2} \rceil$ records.*

PROOF.

$$\begin{aligned} E(X^2) &= m \sum_{i=1}^n \sum_{j=1}^{m_i} (j \log j - (j-1) \log(j-1))^2 \\ &\leq m \sum_{i=1}^n \sum_{j=1}^{m_i} \log^2(ej) \\ &= m \left(\sum_{i=1}^n \sum_{j=1}^{m_i} \log^2 j + m \log^2 e + 2 \log e \sum_{i=1}^n \sum_{j=1}^{m_i} \log j \right) \\ &\leq m \left(\sum_{i=1}^n m_i \log^2 m_i + m \log^2 e + 2S \log e \right) \\ &\leq S \left(\sum_{i=1}^n m_i \log^2 m_i + m \log^2 e + 2S \log e \right) \quad (2) \\ &\leq S (S \log m + m \log^2 e + 2S \log e) \\ &= S^2 (\log m + m \log^2 e / S + 2 \log e) \\ &\leq S^2 (\log m + \log^2 e + 2 \log e) \quad (3) \\ &\leq S^2 (\log m + 5), \end{aligned}$$

where (2) and (3) require the assumption that $S \geq m$.

Hence we have that the variance

$$Var(X) = E(X^2) - (E(X))^2 \leq S^2 (\log m + 4).$$

So, we see that $z = \lceil \frac{8 \log m + 32}{\epsilon^2} \rceil$ suffices. \square

Numerical Illustration: Returning to our example of a stream of size 67 million, the above improvements would drop the number of records for the case of at most 25% error with 75% probability to just 16 thousand (400 KB).

4.5 A Note on Assumptions

For the above analysis, we needed to make the assumption that $S \geq m$. It is not hard to see (and prove) that we need some kind of lower bound on the value of S to protect ourselves from the case that we are trying to distinguish two streams of low S value. If one stream has all unique elements (so that $S = 0$) and another has only one repeated element, then it is very hard to distinguish them. However, we must distinguish them to have less than 100% relative error.

Assuming that $S \geq m$, or that $H \leq \log m - 1$, is reasonable because H attains its maximum value at $\log m$. We now show some other conditions that give us that $S \geq m$, thereby making them reasonable assumptions to make.

THEOREM 5. *If $m \geq 2n_0$ then $S \geq m$.*

PROOF. It is easy to show using Lagrange multipliers that S attains its minimum value when all the items in the stream have the same count. Hence, a lower bound for S is

$$S \geq \sum_{i=1}^{n_0} \frac{m}{n_0} \log(m/n_0) = m \log(m/n_0).$$

Since we have assumed that $m \geq 2n_0$, this gives us that $S \geq m \log(m/n_0) \geq m \log 2 = m$. \square

Hence, we need only assume that each item in the stream appears at least twice on average. This assumption protects us from the case described earlier and in any setting where S can get arbitrarily small. We feel that in any practical setting this simple assumption is very reasonable. For example, on all the traces that we experimented on, the factor m/n_0 was in the range of 50 to 300.

4.6 A Constant-space Solution

As it turns out, if we make a stronger (but still reasonable) assumption on how large the entropy can get, we can make the space usage of the algorithm *independent of m* (assuming fixed sized records). Upper bounding the entropy is reasonable to do since even during abnormal events (e.g. worm attacks), when the randomness of the distributions are increased, there will still be a sufficiently large amount of legitimate activity to offset the increased randomness.

Recall that H attains its maximum at $\log m$, when each of the m items in the stream appears exactly once. We will assume that $H \leq \beta \log m$. This gives us the following bound on S :

$$\begin{aligned} S &= m \log m - mH \\ &\geq m \log m - \beta(m \log m) \\ &= (1 - \beta)m \log m. \end{aligned}$$

We can now apply this to decrease the space usage of our algorithm:

THEOREM 6. *If we assume that $H \leq \beta \log m$, then Algorithm 1 can be modified to use exactly $\lceil \frac{64 \log(1/\delta)}{(1-\beta)\epsilon^2} \rceil$ records.*

PROOF. We once again bound the variance:

$$\begin{aligned} \text{Var}(X) &= E(X^2) - E(X)^2 \\ &\leq E(X^2) \\ &= \frac{m^2}{m} \left[\sum_{j=1}^n \sum_{i=2}^{m_j} (i \log i - (i-1) \log(i-1))^2 \right] \\ &\leq m \sum_{i=1}^n \sum_{j=2}^{m_i} (2 \log j)^2 \\ &\leq 4m \sum_{i=1}^n m_i \log^2 m_i \\ &\leq 4m \log m \left(\sum_i m_i \log m_i \right) \\ &\leq 4S^2 / (1 - \beta). \end{aligned}$$

Hence, we need only $z = \frac{32}{(1-\beta)\epsilon^2}$ groups, which is independent of m . The desired bound on the number of records follows from this. \square

Numerical Illustration: For a stream with 67 million packets, if we make the simple assumption that the entropy never goes above 90% of its maximum value then we need 21 thousand records (525 KB), and if we assume that it never exceeds 75% of its maximum value then we only need 8,200 records (205 KB). Note that these space bounds will not increase with the size of the stream—they depend only on the error parameters. Hence, we can use a few hundred kilobytes for arbitrarily large streams, as long as we can safely make an assumption about how large its standardized entropy can get.

5. SEPARATING THE ELEPHANTS FROM THE MICE

The algorithm described in the previous section provides worst-case theoretical guarantees independent of the structure of the underlying traffic distributions. In practice, however, most network traffic streams have significant structure.

In particular a simple but useful insight [6] is that traffic distributions often have a clear demarcation between large flows (or *elephants*), and smaller flows (or *mice*). A small number of elephant flows contribute a large volume of traffic, and for many traffic monitoring applications it may often suffice to estimate the elephants accurately.

In our second algorithm (see Algorithm 2) we make use of the idea of separating the elephants from the mice in the stream. By separately estimating the contribution of the elephants and mice to the entropy we can further improve the accuracy of our results, thereby also decreasing the space usage of the algorithm. We believe that such a *sieving* idea has much broader applicability. Other streaming algorithms for estimating different traffic statistics can potentially benefit by using such an idea. Intuitively, the amount of space needed by the first algorithm is directly proportional to the variance of the estimator X (see Section 4.3), and by sieving out the high-count items we can significantly decrease the variance of the estimator and hence the space required.

For this algorithm we change the method of sampling slightly. Rather than pre-compute positions in the stream (which requires foreknowledge of the length of the stream), we sample each location with some small probability. After the item is sampled, an exact count is maintained for it, similar to the Sample and Hold algorithm described in [6]. If an item is sampled *exactly once*, then we consider it a mouse and compute the entropy of the mice using the previous algorithm. If an item is sampled *more than once*, we consider it an elephant and estimate its exact value. Note that this method is different from [9] in that we are looking for items that are sampled multiple times, not necessarily in consecutive samples.

Once an item is sampled a second time, it is considered an elephant. To estimate its exact value (i.e. to compensate for the number of times the item appeared before it was first sampled), we simply add the count between the first and second sampling. Intuitively, the number of occurrences of the item between successive samples should be equal if it is evenly distributed. This method of approximating the exact count of the elephant was empirically found to be a good estimator.

The record data structure for this sampling method is similar that used by Algorithm 1. The main difference is that we no longer need a pointer to older copies of an item since we only maintain a single count for each unique item. To be able to tell whether the item has been sampled before or not (to determine whether it should be promoted to an elephant) we require just a single additional bit. Thus, we see that this sampling method requires minimal overhead to separate the elephants from the mice.

The sieving algorithm assumes that every flow that is sampled twice is “elevated” to the status of an elephant. Rather than choose the elevation threshold, we evaluated different values of the threshold before we arrive at the number two. Figure 2 shows the relative error in estimating the S (of the destination address distribution) as a function of k , the threshold for promoting mice to elephants, for three different packet traces. The next section provides further details on the traces used in our evaluations. We observe that the lowest error is achieved with a value of $k = 2$. Intuitively, a higher strike-threshold decreases the number of elephants, and we do not achieve the desired elephant-mice separation.

Algorithm 2: The sieving algorithm

- 1: **Online stage**
 - 2: **for** each item in the stream **do**
 - 3: **if** the item is sampled **then**
 - 4: **if** the item is already being counted **then**
 - 5: promote the item to elephant status
 - 6: **else**
 - 7: allocate space for a counter for this item
 - 8: **else**
 - 9: increment the counter for this item, if there is one
 - 10: **Post-processing stage**
 - 11: $S_e := 0$
 - 12: **for** each elephant (with estimated count c) **do**
 - 13: $S_e := S_e + c \log c$
 - 14: estimate the contribution of the mice S_m from the remaining counts using Algorithm 1
 - 15: return $S_e + S_m$
-

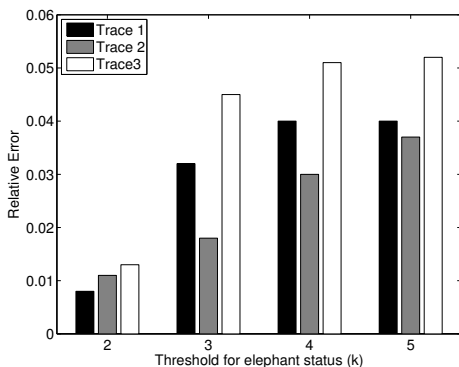


Figure 2: Selecting the threshold k for Sieving

A natural question with such a sieving algorithm is one regarding the relative weights of the two different contributing factors. Intuitively, if either the elephant or the mice flows are not substantial contributors, then we can potentially reduce the space usage further by ignoring the contribution of the insignificant one. We empirically confirmed the need for accurate estimation of both the elephant and the mice flows. Figure 3 shows the relative contribution of the elephant and mice flows to the S estimate (for the destination address distribution on Trace 1). We observe that both elephant and mice flows have substantial contributions to the overall estimation, and ignoring one of them can yield inaccurate results for estimating S , and hence H . The results across different traces and across different traffic distributions of interest were similar and are omitted for brevity.

6. EVALUATION

We first describe the datasets used in this paper. We then present a comparison of the two streaming algorithms introduced in this paper with other sampling based approaches. There are two natural metrics for characterizing the performance of the streaming algorithm for entropy computation: *resource usage* and *error*. The resource usage is related to the number of counters used by different algorithms, which directly translates into the total memory (SRAM) require-




Figure 3: Confirming that estimating both elephants and mice is necessary

ments of the algorithm, and the total CPU usage. For the following evaluations, we use the notion of relative error to determine the accuracy of different algorithms.

Datasets: We use three different packet-header traces for evaluating the accuracy of our algorithms. We provide a brief description of each.²

- **University Trace (*Trace 1*):** Our first packet trace is an hour-long packet trace collected from USC’s Los Nettos collecting facility on Feb 2, 2004. We bin the trace into 1-minute epochs, with each epoch containing roughly 1.7 million TCP packets, 30267 distinct IP addresses, and 15165 ports per minute. We refer to this as *Trace 1* in the following discussion.
- **Department Trace (*Trace 2*):** We use a 5-hour long packet trace collected on Aug 5, 2003 at the gateway router of a medium sized department with approximately 1350 hosts. We observe all traffic to and from the 1350 hosts behind the access router to the commercial Internet, other non-department university hosts and servers. We bin the dataset into 5-minute epochs for our evaluation, with each epoch observing 500000 TCP packets, 2587 distinct addresses, and 4672 distinct ports on average. We refer to this as *Trace 2* in the following discussion.
- **University Trace (*Trace 3*):** The third trace we use is an hour-long trace collected at the access link of the university to the rest of the Internet, at UNC on Apr 24, 2003. We bin this trace into 1-minute epochs as with *Trace 1*. Each epoch contains on average 2.5 million packets, 25565 distinct IP addresses, and 8080 unique application ports. We refer to this as *Trace 3*.

Distributions of Interest: We focus on two main types of distributions for our evaluation. The number of distinct source and destination addresses observed in a dataset, and the distribution of traffic across destinations are typically affected by network attacks, including DDoS and worm attacks. We track the distribution of traffic across different addresses for the source and destination addresses. Understanding the application mix that traverses a network can usually be mapped into a study of the distribution of traffic on different application ports. The distribution of traffic

²The university traces are available on request from the respective universities. The department trace is a private dataset from CMU.

across different ports can also be indicative of scanning attacks or the emergence of new popular applications. In each case we are interested in the distribution of the number of packets observed at each port or address (source or destination) within the measurement epoch. Lakhina et al. [17] give an overview of different types of network events and distributions that each would affect.

For Algorithm 1 we use the assumption that $m/n_0 \geq 2$ since it is both a weak assumption (i.e., weaker than the one made in Section 4.6) and easy to check. To confirm that this assumption holds for our traces and distributions, we present the ratio m/n_0 for them here. For *Trace 1* the ratio is roughly 55 for the addresses and 115 for the ports. For *Trace 2*, $\frac{m}{n_0}$ is around 193 for the addresses and 95 for the ports. Lastly, the ratio is around 97 for the addresses and 300 for the ports in *Trace 3*. Thus we see that in all of our traces the assumption is satisfied.

6.1 Comparison with Sampling Algorithms

We first evaluate the accuracy of estimation of our streaming algorithms by comparing them against the following:

1. **Sampling:** This is the well-known uniform packet sampling approach used in most commercial router implementations [20]. Given a sampling probability p , the sampling approach will pick each packet independently with probability p . The estimation of S and H is performed over the set of sampled packets, after normalizing the counts by $1/p$.
2. **Sample and Hold:** This is the sampling approach proposed by Estan and Varghese [6]. Here given a sampling probability p , the algorithm picks each item in the stream with probability p and keeps an exact count for that item from that point on. Each sample is appropriately renormalized (incrementing by a factor $1/p$) to account for occurrences of the item before it was sampled.

The Sieving algorithm introduced in Section 5 is also conceptually a sampling algorithm, similar to *Sample and Hold* which selects a sampling probability p a priori. In order to perform a fair comparison of the performance across the different algorithms, we normalize the number of records to keep track of to be the same. For the following experiments we fix the sampling probability p , and pick the (ϵ, δ) values for *Algorithm 1*, such that the number of counters across different algorithms used is the same.

We implemented and tested our algorithms on commodity hardware (Intel Xeon 3.0 GHz desktops with 1 GB of RAM). We found that the total CPU utilization for the streaming algorithms was very low—even though we used a preliminary implementation with very few code optimizations, each measurement epoch took less than 10 seconds to process when the epoch length was an entire minute. This demonstrates that our algorithm can comfortably run in real-time. We also found that the post-processing step consumed a negligible fraction of the time of each run. Since all of the algorithms are randomized or sampling-based, for the following results we present the mean relative errors and estimates over 5 independent runs. We found that the standard deviations were very small, and do not present the deviations for clarity of presentation.

Figure 4 compares the performance of different algorithms across different traces, using a sampling rate of $p = 0.001$

for the different algorithms, and using ϵ and δ for Algorithm 1 such that the number of counters used in all four algorithms is roughly the same. The figures show the CDF of the relative error in estimating the entropy of destination addresses observed across different measurement epochs. The streaming algorithms consistently outperform the sampling based approaches. For example, on *Trace 1* we observe that the worst-case relative error with the sampling based approaches can be as high as 18%, whereas the streaming algorithms guarantee a error of at most 6% (Algorithm 1) and 4% (Sieving). We also find that the sieving algorithm provides substantially more accurate estimates for the same space usage compared to the basic streaming algorithm. The sieving algorithm has a worst-case error of at most 2-5%, which bodes well for the practical utility of the algorithms for traffic monitoring applications.

For the rest of the discussion, for brevity we only present the results from *Trace 1*, and summarize the results from *Trace 2* and *Trace 3*. Figure 5 compares the CDF of relative error across measurement epochs, for different distributions of interest from *Trace 1*. We observe a similar trend across algorithms: the sieving algorithm is consistently better than Algorithm 1, which again is substantially more accurate than the sampling based approaches. Both the streaming algorithms have a worst-case error of 7% and mean error of less than 3% across all the different traffic metrics of interest, which is a tolerable operating range for typical monitoring applications, confirming the practical utility of our approaches. We summarize the results for the other two traces in Table 1 and Table 2.

Table 1: Trace 2: Mean relative error in S estimate

Distribution	Sample	Sample&Hold	Algo. 1	Sieving
DSTADDR	0.061	0.047	0.032	0.011
SRCADDR	0.064	0.049	0.033	0.012
DSTPORT	0.077	0.067	0.037	0.016
SRCPORT	0.076	0.060	0.037	0.017

Table 2: Trace 3: Mean relative error in S estimate

Distribution	Sample	Sample&Hold	Algo. 1	Sieving
DSTADDR	0.198	0.192	0.013	0.013
SRCADDR	0.054	0.049	0.017	0.004
DSTPORT	0.116	0.109	0.016	0.015
SRCPORT	0.069	0.062	0.016	0.005

6.2 Error in estimating entropy

Recall from our discussion in Section 2, that it may be the case an accurate estimation of S does not necessarily translate into an accurate estimate of H . However, we find from our evaluations that the streaming algorithms can yield very accurate estimates of H as well. Figures 6(a) and 6(b) compare the relative error in estimating S to the relative error in estimating H , for Algorithm 1 and the sieving algorithm respectively. We observe that across different traces and distributions, that the relative error in estimating H is very low as well (less than 3% mean error with the sieving algorithm).

Figure 7 also provides visual confirmation of the utility of the different algorithms, in tracking the standardized entropy for the destination address distribution. The sieving algorithm once again appears to have greatest accuracy,

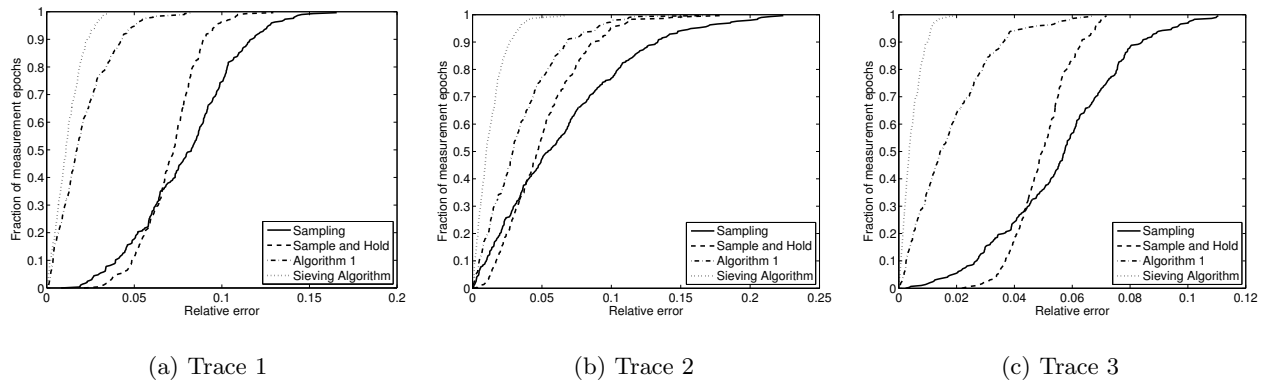


Figure 4: Comparing performance of different traces for estimating destination address entropy

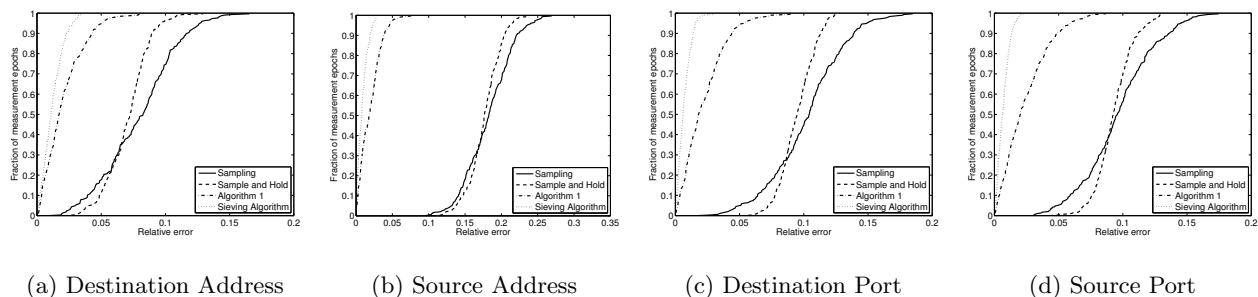


Figure 5: Comparing different distributions, using Trace 1




Figure 6: Relative error in S vs. relative error in H

which can be confirmed with visual inspection. We summarize the results for the other two traces in Table 3 and Table 4, and we observe that in each case the error in H is comparable to (or less than) the corresponding error in S (Tables 1 and 2 respectively).

Last, we vary the memory consumption of the algorithm, and show how the mean and maximum relative errors (for destination address entropy on Trace 1) vary as a function of the memory usage in Figure 8. We observe that the streaming algorithms have an order of magnitude lower error than the sampling algorithms, and can achieve very high accuracy ($< 2\%$ mean error), even with as low as 10 KB of SRAM usage. Note that even though the sampling algorithms also can give reasonably low errors at higher memory consumption (> 80 KB), the corresponding sampling rates are much

Table 3: Trace 2: Mean relative error in H estimate

Distribution	Sample	Sample&Hold	Algo. 1	Sieving
DSTADDR	0.128	0.096	0.069	0.023
SRCADDR	0.132	0.102	0.033	0.026
DSTPORT	0.154	0.131	0.076	0.032
SRCPORT	0.148	0.117	0.073	0.034

Table 4: Trace 3: Mean relative error in H estimate

Distribution	Sample	Sample&Hold	Algo. 1	Sieving
DSTADDR	0.147	0.142	0.009	0.009
SRCADDR	0.082	0.073	0.026	0.006
DSTPORT	0.121	0.114	0.017	0.015
SRCPORT	0.096	0.085	0.023	0.007

higher (> 1 in 200 packet sampling) than what is feasible for very high-speed links.

7. DISCUSSION

One interesting observation from our evaluations is that the observed errors on the traffic traces are much smaller than the theoretical guarantees for Algorithm 1. In particular, we observe that the empirical error is at least one order of magnitude smaller than the theoretical error guarantee. This is because the algorithm must guarantee the error bound for any stream with *any* distribution. Real-world packet traces have considerable underlying structure that the algorithm cannot directly take advantage of.

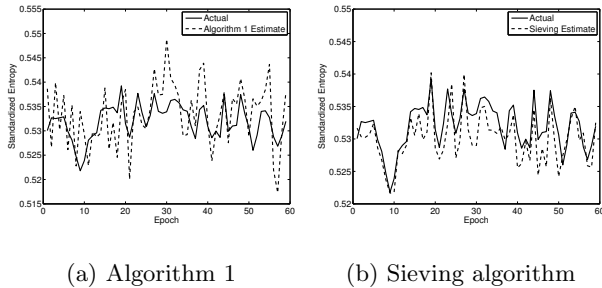


Figure 7: Verifying the accuracy in estimating the standardized entropy

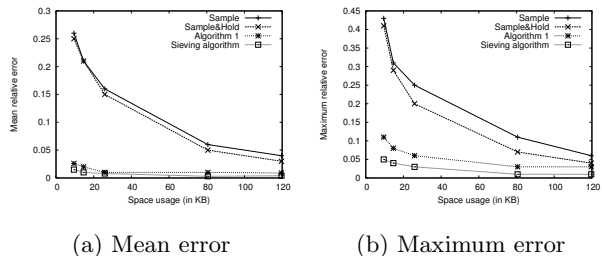


Figure 8: Relative error in estimating H vs. memory usage for different algorithms

It now follows that one way to tighten the bounds on the space/error tradeoff is to make reasonable assumptions about the distribution of the stream and have our algorithms take advantage of them. In Section 4.6 we demonstrate this by making the simple assumption that the standardized entropy of the stream never goes above some fixed constant. This gives us an algorithm that needs a fixed number of records, independent of the size of the stream. Such additional assumptions can help in tightening our space bounds. However, in order to be as general and trace-independent as possible, in our algorithms and evaluations we use very weak assumptions (i.e., $m \geq 2n_0$).

It is a common observation that network packets have a skewed Zipfian distribution. We took advantage of this fact by separating out the few high-count elephants to facilitate the estimation of the remainder more accurately. In doing so, however, we do not make any assumption about the nature of the stream. Algorithm 2 has the property that if there are no elephants in the stream, then it should perform comparably to Algorithm 1. Hence, we expect that, in general, Algorithm 2 should perform better for highly-skewed distributions, but no better than Algorithm 1 when the skew is less pronounced.

8. RELATED WORK

Many of today’s networking monitoring applications use the traffic volume, in terms of flow, packet, and byte counts as the primary metric of choice. These are especially of interest for anomaly detection mechanisms to flag incidents of interest. Some of the well-known methods include signal analysis (e.g., [3]), forecasting (e.g., [4, 21]), and other statistical approaches (e.g., [16, 25]).

There has been a recent interest in using entropy and traffic distribution features for different network monitoring applications. Lakhina et al. [17] use the entropy to augment anomaly detection and network diagnosis, within their PCA framework. Others have suggested the use of such information measures for tracking malicious network activity [7, 23]. Xu et al. [24] use the entropy as a metric to automatically cluster traffic, to infer patterns of interesting activity. For detecting specific types of attacks, researchers have suggested the use of entropy of different traffic features for worm [23] and DDoS detection [7].

Streaming algorithms have received a lot of interest in the algorithms and networking community. The seminal work is that of Alon et al. [2] who provide a framework for estimating frequency moments. Since then, there has been a huge body of literature produced on streaming algorithms, and this is well surveyed in [19]. Kumar et al. use a combination of counting algorithms and Bayesian estimation for accurate estimation of flow size distributions [13, 14]. Streaming algorithms have also been used for identifying heavy-hitters in streams [22, 26].

While the entropy can theoretically be estimated from the flow size distribution [13], computing the flow size distribution conceptually provides much greater functionality than that required for an accurate estimate of the entropy. The complexity of estimating the flow size distribution is significantly higher than the complexity of estimating the entropy, requiring significantly more memory and effort in post-processing.

We are aware of two concurrent efforts in the streaming algorithms community for estimating entropy. Chakrabarti et al. [1] independently proposed an algorithm to estimate S that is similar to Algorithm 1. In this paper we show how this algorithm can be modified such that the memory usage is independent of the size of the stream if we make a simple assumption on how large the standardized entropy can get. McGregor et al. [8] outline algorithms estimating entropy and other information-theoretic measures in the streaming context. However, our algorithms provide unbiased estimates of the entropy, and do not make strong assumptions regarding the underlying distribution. We also provide extensive empirical validation of the utility and accuracy of our algorithms on real datasets, and observe that our sieving approach actually outperforms Algorithm 1.

9. CONCLUSIONS

In this paper, we addressed the need for efficient algorithms for estimating the entropy of network traffic streams, for enabling several real-time traffic monitoring capabilities. We presented lower bounds for the problem of estimating the entropy of a stream, demonstrating that for space-efficient estimation of entropy, both randomization and approximation are necessary. We provide two streaming algorithms for the problem of estimating the entropy. The first algorithm is based on the key insight that the problem shares structural similarity with the problem of estimating frequency moments over streams. By virtue of the strong bounds that we obtain on the variance of the estimator variable, we are able to limit the space usage of the algorithm to polylogarithmic in the length of the stream. Under some practical assumptions of the size of the entropy, we also give an algorithm that samples a number of flows that is independent of the length of the stream. We also identified a method for

increasing the accuracy of entropy estimation by separating the elephants from the mice. Our evaluations on multiple packet traces demonstrate that our techniques produce very accurate estimates, with very low CPU and memory requirements, making them suitable for deployment on routers with multi-gigabit per second links.

Acknowledgments

We would like to thank A. Chakrabarti, K. Do Ba, and S. Muthukrishnan for their useful discussion and for kindly sharing the most recent version of their paper [1] with us. We thank Minh Sung for helping us with the datasets used in this paper. We would also like to acknowledge the useful feedback provided by the anonymous reviewers.

10. REFERENCES

- [1] A. Chakrabarti, K. Do Ba, and S. Muthukrishnan. Estimating entropy and entropy norm on data streams. In *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science (STACS)*, 2006.
- [2] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of ACM Symposium on Theory of Computing (STOC)*, 1996.
- [3] P. Barford, J. Kline, D. Plonka, and A. Ron. A Signal Analysis of Network Traffic Anomalies. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop (IMW)*, 2002.
- [4] J. D. Brutlag. Aberrant behavior detection in time series for network monitoring. In *Proceedings of USENIX Large Installation System Administration Conference (LISA)*, 2000.
- [5] N. Duffield, C. Lund, and M. Thorup. Estimating flow distributions from sampled flow statistics. In *Proceedings of ACM SIGCOMM*, 2003.
- [6] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proceedings of ACM SIGCOMM*, 2002.
- [7] L. Feinstein, D. Schnackenberg, R. Balupari, and D. Kindred. Statistical approaches to DDoS attack detection and response. In *Proceedings of the DARPA Information Survivability Conference and Exposition*, 2003.
- [8] S. Guha, A. McGregor, and S. Venkatasubramanian. Streaming and sublinear approximation of entropy and information distances. In *Proceedings of ACM Symposium on Discrete Algorithms (SODA)*, 2006.
- [9] F. Hao, M. Kodialam, and T. V. Lakshman. ACCEL-RATE: a faster mechanism for memory efficient per-flow traffic estimation. In *Proceedings of ACM SIGMETRICS*, 2004.
- [10] N. Hohn and D. Veitch. Inverting sampled traffic. In *Proceedings of ACM/USENIX Internet Measurement Conference (IMC)*, 2003.
- [11] B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992.
- [12] V. Karamcheti, D. Geiger, Z. Kedem, and S. Muthukrishnan. Detecting malicious network traffic using inverse distributions of packet contents. In *Proceedings of ACM SIGCOMM Workshop on Mining Network Data (MineNet)*, 2005.
- [13] A. Kumar, M. Sung, J. Xu, and J. Wang. Data streaming algorithms for efficient and accurate estimation of flow distribution. In *Proceedings of ACM SIGMETRICS/IFIP WG 7.3 Performance*, 2004.
- [14] A. Kumar, M. Sung, J. Xu, and E. Zegura. A data streaming algorithm for estimating subpopulation flow size distribution. In *Proceedings of ACM SIGMETRICS*, 2005.
- [15] E. Kushilevitz and N. Nisan. *Communication complexity*. Cambridge University Press, New York, NY, USA, 1997.
- [16] A. Lakhina, M. Crovella, and C. Diot. Diagnosing network-wide traffic anomalies. In *Proceedings of ACM SIGCOMM*, 2004.
- [17] A. Lakhina, M. Crovella, and C. Diot. Mining anomalies using traffic feature distributions. In *Proceedings of ACM SIGCOMM*, 2005.
- [18] K. Levchenko, R. Paturi, and G. Varghese. On the difficulty of scalably detecting network attacks. In *Proceedings of ACM Conference on Computer and Communications Security (CCS)*, 2004.
- [19] S. Muthukrishnan. Data streams: algorithms and applications. <http://athos.rutgers.edu/~muthu/stream-1-1.ps>.
- [20] Cisco Netflow. <http://www.cisco.com/warp/public/732/Tech/nmp/netflow/index.shtml>.
- [21] M. Roughan, A. Greenberg, C. Kalmanek, M. Rumsewicz, J. Yates, and Y. Zhang. Experience in measuring internet backbone traffic variability: Models, metrics, measurements and meaning. In *Proceedings of International Teletraffic Congress (ITC)*, 2003.
- [22] S. Venkataraman, D. Song, P. B. Gibbons, and A. Blum. New Streaming Algorithms for Fast Detection of Superspreaders . In *Proceedings of Network and Distributed System Security Symposium (NDSS)*, 2005.
- [23] A. Wagner and B. Plattner. Entropy Based Worm and Anomaly Detection in Fast IP Networks. In *Proceedings of IEEE International Workshop on Enabling Technologies, Infrastructures for Collaborative Enterprises*, 2005.
- [24] K. Xu, Z.-L. Zhang, and S. Bhattacharya. Profiling internet backbone traffic: Behavior models and applications. In *Proceedings of ACM SIGCOMM*, 2005.
- [25] Y. Zhang, Z. Ge, M. Roughan, and A. Greenberg. Network anomography. In *Proceedings of ACM/USENIX Internet Measurement Conference (IMC)*, 2005.
- [26] Y. Zhang, S. Singh, S. Sen, N. Duffield, and C. Lund. Online identification of hierarchical heavy hitters: algorithms, evaluations, and applications. In *Proceedings of ACM/USENIX Internet Measurement Conference (IMC)*, 2004.