Ranking Flows from Sampled Traffic

Chadi Barakat[*]

INRIA - Planète group
Sophia Antipolis, France

Chadi.Barakat@inria.fr

Gianluca Iannaccone, Christophe Diot

Intel Research
Cambridge, UK

Gianluca.Iannaccone@intel.com,
Christophe.Diot@intel.com

ABSTRACT

Most of the theoretical work on sampling has addressed the inversion of general traffic properties such as flow size distribution, average flow size, or total number of flows. In this paper, we make a step towards understanding the impact of packet sampling on individual flow properties. We study how to detect and rank the largest flows on a link. To this end, we develop an analytical model that we validate on real traces from two networks. First we study a *blind* ranking method where only the number of sampled packets from each flow is known. Then, we propose a new method, *protocol-aware* ranking, where we make use of the packet sequence number (when available in transport header) to infer the number of non-sampled packets from a flow, and hence to improve the ranking. Surprisingly, our analytical and experimental results indicate that a high sampling rate (10% and even more depending on the number of top flows to be ranked) is required for a correct blind ranking of the largest flows. The sampling rate can be reduced by an order of magnitude if one just aims at detecting these flows or by using the protocol-aware method.

Categories and Subject Descriptors

C.4 [**Computer Systems Organization**]: PERFORMANCE OF SYSTEMS—*Measurement techniques*

General Terms

Measurement, Performance, Experimentation

Keywords

Packet sampling, largest flow detection and ranking, performance evaluation, validation with real traces

1. INTRODUCTION

The list of the top users or applications is one of the most useful statistics to be extracted from network traffic.

Network operators use the knowledge of the most popular destinations to identify emerging markets and applications or to locate where to setup new Points of Presence. Content delivery networks use the popularity of sites to define caching and replication strategies. In traffic engineering, the identification of heavy hitters in the network can be used to treat and route them differently across the network [20, 17, 10]. Keeping track of the network prefixes that generate most traffic is also of great importance for anomaly detection. A variation in the pattern of the most common applications may be used as a warning sign and trigger careful inspection of the packet streams.

However, the ability to identify the top users in a packet stream is limited by the network monitoring technology. Capturing and processing all packets on high speed links still remains a challenge for today's network equipment [16, 9]. In this context, a common solution is to sample the packet stream to reduce the load on the monitoring system and to simplify the task of sorting the list of items. The underlying assumption in this approach is that the sampling process does not alter the properties of the data distribution.

Sampled traffic data is then used to infer properties of the original data (this operation is called *inversion*). The inversion of sampled traffic is, however, an error-prone procedure that often requires a deep study of the data distribution to evaluate how the sampling rate impacts the accuracy of the metric of interest. Although the inversion may be simple for aggregate link statistics (e.g., to estimate the number of packets transmitted on a link, it is usually sufficient to multiply the number of sampled packets by the inverse of the sampling rate), it is much harder for the properties of individual connections or "flows" [9, 11, 8].

For these reasons, in this paper, we address this simple, and so far unanswered, question: *which sampling rate is needed to correctly detect and rank the flows that carry the most packets?*

We define the problem as follows. Consider a traffic monitor that samples packets independently of each other with probability p (random sampling) and classifies them into *sampled flows.* At the end of the measurement period, the monitor processes the list of sampled flows, ranks them based on their size in packets, and returns an ordered list of the t largest flows.

We are interested in knowing (i) whether the ordered list contains all the actual largest flows in the original packet

[*] This work was done while the author was visiting Intel Research Cambridge.

CoNEXT'05, October 24–27, 2005, Toulouse, France.

stream (*detection*), and (*ii*) if the items in the list appear in the correct order (*ranking*).

We build an analytical model and define a performance metric that evaluates the accuracy of identification and ranking of the largest flows. We consider a flow to consist of a single TCP connection. However, our results are general and can be applied to alternative definitions of flow, as well.

We evaluate two approaches to sort the list of flows:

(*i*) *Blind*, where the sampled flows are ranked just based on their sampled size. This method can be applied to any definition of flow.

(*ii*) *Protocol-aware*, where we make use of additional information in the packet header (e.g., the sequence number in TCP packets) to infer the number of non-sampled packets between sampled ones. This method can only be applied to flow definitions that preserve the protocol level details.

The contributions of this work are the following: (1) We perform an analytical study of the problem of ranking two sampled flows and compute the probability that they are *misranked*. We propose a Gaussian approximation to make the problem numerically tractable. (2) We introduce the protocol-aware ranking method that uses protocol level information to complement the flow statistics and render the detection and ranking of the largest flows more accurate. (3) Based on the model for the ranking of two flows, we propose a general model to study the detection and ranking problem, given a generic flow size distribution. We define a performance metric and evaluate the impact of several metric's parameter on the accuracy of the ranking. (4) We validate our findings on measurement data using publicly-available packet-level traces. Our results indicate that a surprisingly high sampling rate is required to obtain a good accuracy with the blind approach (10% and even more depending on the number of flows of interest). As for the protocol-aware approach, it allows to reduce the required sampling rate by an order of magnitude compared to the blind approach.

The paper is structured as follows. Next, we discuss the related literature. In Section 3 and 4, we present our model. Section 5 analyzes the model numerically and Section 6 validates it on real packet-level traces. Section 7 concludes the paper and provides perspectives for our future research.

2. RELATED WORK

The inversion of sampled traffic has been extensively studied in the literature. The main focus has been on the inversion of aggregate flow properties such as flow size distribution [9, 11], average flow size or total number of flows [8] on a given network link. Duffield et al. [8] study the problem of flow splitting and propose estimators for the total number of flows and for the average flow size in the original traffic stream. [9, 11] study the inversion of the flow size distribution with two different methods. They both show that the major difficulty comes from the number of flows that are not sampled at all and that need to be estimated with an auxiliary method. As an auxiliary method, [8, 9] propose the use of the SYN flag in the TCP header to mark the beginning of a flow. [9] shows that periodic and random sampling provide roughly the same result on high speed links, and so random sampling can be used for mathematical analysis due to its appealing features. [4] finds the sampling rate that assures a bounded error on the estimation of the size of flows contributing to more than some predefined percentage of the traffic volume. [14] studies whether the number of sampled packets is a good estimator for the detection of large flows without considering its impact on the flow ranking.

Given the potential applications of finding the list of top users, it does not come as a surprise that there has been a significant effort in the research community to find ways to track frequent items in a data stream [5, 7, 3, 10]. However, this problem has usually been addressed from a memory requirement standpoint. All the works in the literature assume that if the algorithm and the memory size is well chosen, the largest flows can be detected and ranked with a high precision. However, in the presence of packet sampling, even if the methods rank correctly the set of sampled flows, there is no guarantee that the sampled rank corresponds to the original rank. The problem we address in this paper complements these works as it focuses on the impact of sampling on the flow ranking.

3. BASIC MODEL: RANKING TWO FLOWS

In this section, we study the probability to misrank two flows of original sizes S_1 and S_2 in packets. This probability is the basis for the general model for detecting and ranking the largest flows that we will present later. Indeed, the detection and ranking of the largest flows can be transformed into a problem of ranking over a set of flow pairs.

Without loss of generality, we assume $S_1 < S_2$. We consider a random sampling of rate p. Let s_1 and s_2 denote the sizes in packets of both flows after sampling. The two sampled flows are misranked if (*i*) s_1 is larger than s_2, or (*ii*) both flows are not sampled, i.e., their sampled sizes equal to zero. By combining (*i*) and (*ii*), one can see that the necessary condition for a good ranking is to sample at least one packet from the larger flow (i.e., the smaller of the two flows can disappear after sampling). The probability to misrank the two flows can then be written as $P_m(S_1, S_2) = \mathbb{P}\{s_1 \geq s_2\}$. For the case $S_1 = S_2$, we consider the two flows as misranked if $s_1 \neq s_2$, or if both flows are not sampled at all, i.e. $s_1 = s_2 = 0$.

We compute and study the misranking probability of two flows of given sizes in the rest of this section. First, we consider the blind ranking method where only the number of sampled packets from a flow is known. For this method, we express the misranking probability as a double sum of binomials, then we present a Gaussian approximation to make the problem tractable numerically. Second, we consider the protocol-aware ranking method for which we calculate a numerical-tractable closed-form expression of the misraking probability. Note that the misranking probability is a symmetric function, i.e., $P_m(S_1, S_2) = P_m(S_2, S_1)$.

3.1 Blind ranking

With this method, s_1 and s_2 represent the number of sampled packets from flows S_1 and S_2. Under our assumptions, these two variables are distributed according to a binomial distribution of probability p. Hence, we can write for $S_1 < S_2$,

$$P_m(S_1, S_2) = \mathbb{P}\{s_1 \geq s_2\} = \sum_{i=0}^{S_1} b_p(i, S_1) \sum_{j=0}^{i} b_p(j, S_2). \quad (1)$$

$b_p(i, S)$ is the probability density function of a binomial distribution of probability p, i.e., the probability of obtaining i successes out of S trials. We have $b_p(i, S) = \binom{S}{i} p^i (1-p)^{S-i}$ for $i = 0, 1, ..., S$, and $b_p(i, S) = 0$ for $i < 0$ and $i > S$. The probability to misrank two flows of equal sizes is given by $\mathbb{P}\{s_1 \neq s_2 \text{ or } s_1 = s_2 = 0\} = 1 - \mathbb{P}\{s_1 = s_2 \neq 0\}$ $= 1 - \sum_{i=1}^{S_1} b_p^2(i, S_1)$.

Unfortunately, the above expression for the misranking probability is numerically untractable since it involves two sums of binomials. For large flows of order S packets, the number of operations required to compute such a probability is on the order of $O(S^3)$, assuming that the complexity of the binomial computation is on the order of $O(S)$. The problem becomes much more complex if one has to sum over all possible flow sizes (i.e., $O(S^5)$). For this reason, we propose next a Gaussian approximation to the problem of blind ranking that is accurate and easy to compute. We use this approximation to study the ranking performance as a function of the sampling rate and the flow sizes.

3.1.1 Gaussian approximation to blind ranking

Consider a flow made of S packets and sampled at rate p. The sampled size follows a binomial distribution. However, it is well known that the binomial distribution can be approximated by a Normal (or Gaussian) distribution when p is small and when the product pS is on the order of one (flows for which, on average, at least few packets are sampled) [21, pages 108–109]. We assume that this is the case for the largest flows, and we consider the sampled size of a flow as distributed according to a Normal distribution of average pS and of variance $p(1-p)S$. Using this approximation, one can express the misranking probability for the blind ranking problem in the following simple form.

PROPOSITION 1. *For any two flows of sizes S_1 and S_2 packets ($S_1 \neq S_2$), the Gaussian approximation gives,*

$$P_m(S_1, S_2) \simeq \frac{1}{2} erfc\left(\frac{|S_2 - S_1|}{\sqrt{2(1/p-1)(S_1+S_2)}}\right), \qquad (2)$$

where $erfc(x) = (\frac{2}{\sqrt{\pi}}) \int_x^\infty e^{-u^2} \mathrm{d}u$ is the complementary error cumulative function.

Proof: Consider two flows of sizes S_1 and S_2 in packets such that $S_1 < S_2$. Their sampled versions s_1 and s_2 both follow Normal distributions of averages pS_1 and pS_2, and of variances $p(1-p)S_1$ and $p(1-p)S_2$. We know that the sum of two Normal variables is a Normal variable. So the difference $s_1 - s_2$ follows a Normal distribution of average $p(S_1 - S_2)$ and of variance $p(1-p)(S_1+S_2)$. We have then this approximation for the misranking probability:

$$\begin{aligned} P_m(S_1, S_2) &= \mathbb{P}\{s_1 - s_2 \geq 0\} \\ &\simeq \mathbb{P}\left\{V > \frac{p(S_2 - S_1)}{\sqrt{p(1-p)(S_1+S_2)}}\right\} \\ &= \frac{1}{2}\mathrm{erfc}\left(\frac{S_2 - S_1}{\sqrt{2(1/p-1)(S_1+S_2)}}\right). \qquad (3) \end{aligned}$$

V is a standard Normal random variable. Given the symmetry of the misranking probability, one can take the absolute value of $S_2 - S_1$ in (3) and get the expression stated in the proposition, which is valid for all S_1 and S_2. □

For $S_1 = S_2$, one can safely approximate the misranking probability to be equal to 1. This approximation is however of little importance given the very low probability of having two flows of equal sizes, especially when they are large.

3.2 Protocol-aware ranking

Packets can carry in their transport header an increasing sequence number. A typical example is the byte sequence number in the TCP header. Another example could be the sequence number in the header of the Real Time Protocol (RTP) [19]. One can use this sequence number, when available, to infer the number of non-sampled packets (or bytes in the case of TCP) between sampled ones, and hence to improve the accuracy of ranking. The size of the sampled flow in this case is no longer the number of packets collected, but rather the number of packets that exist between the first and last sampled packets from the flow. Although this solution is limited to flows whose packet carry a sequence number, we believe that the study of this ranking method is important given the widespread use of the TCP protocol. Our objective is to understand how the use of protocol-level information can supplement the simple, and more general, blind method and if it is worth the additional overhead it introduces (i.e., storing two sequence numbers per flow record).

In the following, we calculate the misranking probability of two flows of given sizes when using the protocol-aware method. This probability will be used later in the general ranking problem. The main contribution of this section is a closed-form expression for the misranking probability that is numerically tractable, without the need for any approximation.

Let S be the size of a flow in packets. Let s_b, $s_b = 1, 2, ..., S$, denote the (packet) sequence number carried by the first sampled packet, and let s_e, $s_e = S, S-1, ..., s_b$, denote the sequence number carried by the last sampled packet. Given s_b and s_e, one can estimate the size of the sampled flow in packets to $s = s_e - s_b + 1$. The error in this estimation comes from the non-sampled packets that are transmitted before s_b and after s_e. We give next the distribution of s, which is needed for the computation of the misranking probability, then we state our main result. Before presenting the analysis, note that this new flow size estimator only counts the packets that are transmitted with distinct sequence numbers. In the case of TCP, this corresponds to the number of bytes received at the application layer, rather then the number of bytes carried over the network. It is equivalent to assuming that the probability of sampling a retransmitted (or duplicated) packet is negligible. This is a reasonable assumption if the loss rate is low. We will address this aspect in more detail in Section 6.

Consider a flow of size $S \geq 2$ in packets. Using the above definition for s, the sampled flow has a size of i packets, $i \geq 2$, with probability:

$$\mathbb{P}\{s = i\} = \sum_{k=1}^{S-i+1} \mathbb{P}\{s_b = k\} \mathbb{P}\{s_e = k+i-1\}.$$

We have $\mathbb{P}\{s_b = k\} = (1-p)^{k-1}p$, and $\mathbb{P}\{s_e = k+i-1\} = (1-p)^{S-k-i+1}p$. This gives

$$\begin{aligned} \mathbb{P}\{s = i\} &= \sum_{k=1}^{S-i+1} (1-p)^{k-1} p (1-p)^{S-k-i+1} p \\ &= p^2 (1-p)^{S-i}(S-i+1). \qquad (4) \end{aligned}$$

As for $i = 0$, we have $\mathbb{P}\{s = 0\} = (1-p)^S$ for $S \geq 1$. And for $i = 1$, we have $\mathbb{P}\{s = 1\} = p(1-p)^{S-1}S$ for $S \geq 1$. It is easy to prove that the cumulative distribution of s is the following for all values of S:

$$\mathbb{P}\{s \le i \ne 0\} = p(1-p)^{S-i}(S-i+1) + (1-p)^{S-i+1}. \quad (5)$$

We come now to the misranking probability, which we recall is a symmetric function. For $S_1 < S_2$, we have

$$P_m(S_1, S_2) = \mathbb{P}\{s_2 \le s_1\} = \sum_{i=0}^{S_1} \mathbb{P}\{s_1 = i\} \sum_{j=0}^{i} \mathbb{P}\{s_2 = j\}. \quad (6)$$

And for $S_1 = S_2$, we have

$$P_m(S_1, S_2) = 1 - \sum_{i=1}^{S_1} \mathbb{P}\{s_1 = i\}^2. \quad (7)$$

Our main result is the following.

PROPOSITION 2. *For $S_1 < S_2$, the misranking probability is equal to*

$$\begin{aligned} P_m(S_1, S_2) &= (1-p)^{S_1}(1-p)^{S_2} \\ &+ p(1-p)^{S_1-1}S_1[p(1-p)^{S_2-1}S_2 + (1-p)^{S_2}] \\ &+ p^3\frac{\partial^2 F(1-p,1-p)}{\partial x \partial y} + p^2\frac{\partial F(1-p,1-p)}{\partial x}, \end{aligned}$$

where

$$\begin{aligned} F(x,y) &= xy^{S_2-S_1+1} + \ldots + x^{S_1-1}y^{S_2-1} \\ &= xy^{S_2-S_1+1}(1-(xy)^{S_1-1})/(1-xy). \end{aligned}$$

For $S_1 = S_2 = S$, the misranking probability is equal to

$$P_m(S,S) = 1 - p^2(1-p)^{2(S-1)}S^2 - p^4\frac{\partial^2 G(1-p,1-p)}{\partial x \partial y},$$

where

$$G(x,y) = xy + x^2y^2 + x^{S-1}y^{S-1} = (xy - (xy)^S)/(1-xy).$$

Proof: One can validate the results by plugging (4) and (5) into (6) and (7). □

Note that the main gain of writing the misraking probability in such a condensed form is a complexity that drops from $O(S^3)$ in (6) to $O(S)$ in our final result. This gain comes from the closed-form expression for the cumulative distribution in (5), and from introducing the two functions $F(x,y)$ and $G(x,y)$. These two latter functions transform two series whose complexity is $O(S^2)$ into a closed-form expression whose complexity is $O(S)$.

We solve the derivatives in the above equations using the symbolic toolbox of matlab, which gives explicit expressions for the misranking probability. These expressions are simple to compute, but span on multiple lines, so we omit them for lack of space.

3.3 Analysis of the misranking probability

3.3.1 The blind case

We use the Gaussian approximation to study how the misranking probability varies with the sampling rate and with the sizes of both flows, in particular their difference. The study of the impact of the flow sizes is important to understand the relation between flow size distribution and ranking of the largest flows.

The misranking probability is a decreasing function of the sampling rate. It moves to zero when p moves to 1 and to 0.5 when p approaches zero[1]. Therefore, there exists one sampling rate that leads to some desired misranking probability, and any lower sampling rate results in larger error.

We study now how the misranking probability varies with the sizes of both flows. Take $S_1 = S_2 - k$, k a positive integer. From (2) and for fixed k, the misranking probability increases with S_1 and S_2 ($\mathrm{erfc}(x)$ is an increasing function in x). This indicates that it is more difficult to rank correctly two flows different by k packets as their sizes increase in absolute terms. The result is different if we take the size of one flow equal to $\alpha < 1$ times the size of the second, i.e., $S_1 = \alpha S_2$. Here, $(S_1 - S_2)/\sqrt{S_1 + S_2}$ is equal to $\sqrt{S_1}(1-\alpha)/\sqrt{1+\alpha}$, which increases with S_1. Hence, the misranking probability given in (2) decreases when S_1 increases. We conclude that, when the two flow sizes maintain the same proportion, it is easier to obtain a correct ranking when they are large in absolute terms.

We can now generalize the result above. One may think that the larger the flows, the better the ranking of their sampled versions. Our last two examples indicate that this is not always the case. The ranking accuracy depends on the relative difference of the flow sizes. In general, to have a better ranking, the difference between the two flow sizes must increase with the flow sizes and the increase must be larger than a certain threshold. This threshold is given by (2): the difference must increase at least as the square root of the flow sizes. This is an interesting finding. In the context of the general ranking problem, it can be interpreted as follows. Suppose that the flow size has a cumulative distribution function $y = F(x)$. As we move to the tail of the distribution[2], the size of the flows to be ranked increases. The ranking performance improves if the difference between flow sizes increases faster than $\sqrt{x}$. This is equivalent to saying that $\mathrm{d}x/\mathrm{d}y$ should increase with x faster than $\sqrt{x}$. All common distributions satisfy this condition, at least at their tails. For example, with the exponential distribution we have $\mathrm{d}x/\mathrm{d}y \propto e^{\lambda x}$ ($1/\lambda$ is the average), while for the Pareto distribution we have $\mathrm{d}x/\mathrm{d}y \propto x^{\beta+1}$ (β is the shape).

3.3.2 The protocol-aware case

The first difference with the blind case is in the estimation error ($S - s = s_b - 1 + S - s_e$), which can be safely assumed to be independent of the flow size for large flows (only dependent on p). This means that if two large flows keep the same distance between them while their sizes increase, their ranking maintains the same accuracy. Their ranking improves if the difference between their sizes increases as well, and it deteriorates if the difference between their sizes decreases. So in contrast to the blind case, the threshold for the ranking here to improve is that the larger flow should have its size increasing a little faster than the smaller one. In the context of the general ranking problem where flow sizes are distributed according to a cumulative distribution function $y = F(x)$, and when the top flows become larger, the protocol-aware ranking improves if the derivative $\mathrm{d}x/\mathrm{d}y$ increases with x. This is equivalent to saying that the function $F(x)$ should be concave, which is satisfied by most common distributions at their tail. For blind ranking, concavity was

[1] The Gaussian approximation does not account for the case $p = 0$ where the misranking probability should be equal to 1 based on our definition.

[2] Because we are more and more focusing on large flows or because the number of available flows for ranking increases.

not enough to obtain a better ranking; the derivative $\mathrm{d}x/\mathrm{d}y$ had to increase faster than $\sqrt{x}$. So in conclusion, the condition to have a better ranking when we move to the tail of the flow size distribution is less strict with the protocol-aware method, which is an indication of its good performance.

The second difference with the blind case is in the relation between the ranking accuracy and the sampling rate. Consider two large flows of sizes S_1 and S_2 in packets, and let s_1 and s_2 denote their sampled sizes. The coefficient of variation of the difference $s_2 - s_1$ is an indication on how well the ranking performs (a small coefficient of variation results in better ranking[3]). It is easy to prove that this coefficient of variation scales as $1/p$ for protocol-aware ranking and as $1/\sqrt{p}$ for blind ranking. This is again an important finding. It tells that when the sampling rate is very small, blind ranking could (asymptotically) perform better than protocol-aware ranking. Our numerical and experimental results will confirm this finding.

4. GENERAL MODEL: DETECTING AND RANKING THE LARGEST FLOWS

We generalize the previous model from the ranking of two flows to the detection and ranking of the top t flows, $t = 1, 2, \ldots, N$. The misranking probability $P_m(S_1, S_2)$ previously calculated is the basis for this generalization. Let $N \geq t$ denote the total number of flows available in the measurement period before sampling. We want the sampled list of top t flows to match the list of top t flows in the original traffic. Two criteria are considered to decide whether this match is accurate. First, we require the two lists to be identical. This corresponds to the *ranking* problem. The second, less constrained, criterion requires the two lists to contain the same flows regardless of their relative order within the list. This corresponds to the *detection* problem. For both problems, the quality of the result is expressed as a function of the sampling rate p, the flow size distribution, the number of flows to rank t, and the total number of flows N.

4.1 Performance metric

In order to evaluate the accuracy of detection and ranking, we need to define a performance metric that is easy to compute and that focuses on the largest flows. A flow at the top of the list can be misranked with a neighboring large flow or a distant small flow. We want our metric to differentiate between these two cases and to penalize more the latter one; a top-10 flow replaced by the 100-th flow in the sampled top list is worse than the top-10 flow being replaced by the 11-th flow. We also want our metric to be zero when the detection and ranking of the top flows are correct.

We introduce our performance metric using the ranking problem. The performance metric for the detection problem is a straightforward extension. Let's form all flow pairs where the first element of a pair is a flow in the top t and the second element is anywhere in the sorted list of the N original flows. The number of these pairs is equal to $N-1+N-2+\cdots+N-t = (2N-t-1)t/2$. We then count the pairs in this set that are misranked after sampling and we take the sum as our metric for ranking accuracy. This sum indicates how good the ranking is at the top of the list. It is equal to zero when the ranking is correct. When the ranking is not correct, it takes a value proportional to the original rank of the flows that have taken a slot in the top-t list. For example, if the top flow is replaced by its immediate successor in the list, the metric will return a ranking error of 1. Instead, if the same flow is replaced by a distant flow, say the 100-th, the metric will return an error of 99. Also, note that our metric does not account for any misranking of flows outside the list of top t flows. For any two flows n and m, such that $n > m > t$, the fact that n takes the position of m does not add anything to our performance metric since our metric requires at least one element of a flow pair to be in the original list of top t flows.

In the detection problem, we are no longer interested in comparing flow pairs whose both elements are in the top t list. We are only interested in the ranking between flows in the top t list and those *outside* the list. Therefore, our detection metric is defined as the number of misranked flow pairs, where the first element of a pair is in the list of top t flows and the second element is *outside* this list (non top t).

The above metrics return one value for each realization of flow sizes and of sampled packets. Given that we want to account for all realizations, we define the performance metrics as the number of misranked flow pairs *averaged* over all possible values of flow sizes in the original list of N flows and over all sampling runs. We deem the ranking/detection as acceptable when our metric takes a value below one (i.e., on average less than one flow pair is misranked).

In addition to the above, our metrics have the advantage of being easily and exactly calculable. Performance metrics based on probabilities (e.g.,[12]) require lot of assumptions that make them only suitable for computing bounds, but not exact values.

4.2 Computation of the performance metric for the ranking problem

Consider a flow of i packets belonging to the list of top t flows in the original traffic (before sampling). First, we compute the probability that this flow is misranked with another flow of general size and general position. Denote this probability by $P_{mt}(i)$, where m stands for misranking and t for top. Then, we average over all values of i to get $\bar{P}_{mt}$[4]. This latter function gives us the probability that, on average, the top t-th flow is misranked with another flow. Thus, our performance metric, which is defined as the average number of misranked flow pairs where at least one element of a pair is in the top t, is equal to $(2N - t - 1)t\bar{P}_{mt}/2$. Next, we compute the value of $\bar{P}_{mt}$.

Let p_i denote the probability that the size of a general flow is equal to i packets, and P_i denote the flow size complementary cumulative distribution, i.e., $P_i = \sum_{j=i}^{\infty} p_j$. For a large number of flows N and a high degree of multiplexing, we consider safe to assume that flow sizes are independent of each other (see [2] for a study of the flow size correlation on a OC-12 IP backbone link). A flow of size i belongs to the list of top t flows if the number of flows in the original total list, with a size larger than i, is less or equal than $t-1$. Since each flow can be larger than i with probability P_i independently of the other flows, we can write the probability that a flow of size i belongs to the list of the top t flows

[3] For $S_1 < S_2$, we are interested in $\mathbb{P}\{s_1 \geq s_2\}$. According to Tchebychev inequality, this probability can be supposed to behave like $\mathrm{VAR}[s_1 - s_2]/\mathbb{E}[s_1 - s_2]^2$, which is the square of the coefficient of variation.

[4] Note that the distribution of the size of a flow at the top of the list is different from that of a generic flow.

as $P_t(i,t,N) = \sum_{k=0}^{t-1} b_{P_i}(k, N-1)$, where $b_{P_i}(k, N-1)$ is the probability to obtain k successes out of $N-1$ trials, P_i being the probability of a success. The probability that the t-th largest flow has a size of i packets is equal to $P_t(i) = p_i P_t(i,t,N)/\bar{P}_t(t,N)$. $\bar{P}_t(t,N)$ is the probability that a flow of general size is among the top t in the original total list, which is simply equal to t/N.

Using the above notation, one can write the misranking probability between a top t flow of original size i packets and any other flow as follows

$$P_{mt}(i) = \frac{1}{P_t(i,t,N)} \left(\sum_{j=1}^{i-1} p_j P_t(i,t,N-1) P_m(j,i) + \sum_{j=i}^{\infty} p_j P_t(i,t-1,N-1) P_m(i,j) \right). \qquad (8)$$

In this expression, we sum over all possible original sizes of the other flow (the variable j) and we separate the case when this other flow is smaller than i from the case when it is larger than i [5]. $P_m(i,j)$ is the misranking probability of two flows of sizes i and j packets, which we calculated in the previous section for the two ranking methods. $\bar{P}_{mt}$ is then equal to $\sum_{i=1}^{\infty} P_t(i) P_{mt}(i)$.

For protocol-aware ranking, $P_m(i,j)$ is given explicitly in Proposition 2 and can be easily computed. For blind ranking, we use the Gaussian approximation summarized in Proposition 2, which we recall holds when at least one of the two flows to be compared is large.

4.3 Computation of the performance metric for the detection problem

Consider the probability that a flow among the top t is swapped with a flow that does belong to the top t. Let $\bar{P}^{\star}_{mt}$ denote this probability. Following the same approach described in Section 4, we can write

$$\bar{P}^{\star}_{mt} = \frac{1}{\bar{P}^{\star}_{t}} \sum_{i=1}^{\infty} \sum_{j=1}^{i-1} p_i p_j P^{\star}_t(j,i,t,N) P_m(j,i).$$

To get this expression for $\bar{P}^{\star}_{mt}$, we sum over all possible values for the size of the flow in the top t (index i) and all possible values for the size of the other flow not among the top t (index j). In this expression, p_i and p_j represent the probability that the size of a flow is equal to i or j packets, respectively. $P_m(j,i)$ is the probability that two flows of sizes i and j are misranked – it is given by the Gaussian approximation described in Proposition 1 for the blind method and the result stated in Proposition 2 for the protocol-aware method. $P^{\star}_t(j,i,t,N)$ is the joint probability that a flow of size i belongs to the list of the top t flows while another flow of size j does not belong to it (i.e., it is in the bottom $N-t$ flows). $\bar{P}^{\star}_t$ is the joint probability that a flow of any size belongs to the list of the top t flows while another flow of any size does not belong to this list. It is equal to $t(N-t)/(N(N-1))$.

We now compute $P^{\star}_t(j,i,t,N)$ for $j < i$, i.e., the probability that flow i belongs to the top list while flow j does not. The number of flows larger than i should be smaller than t, while the number of flows larger than j should be larger than t. The probability that a flow size is larger than i is $P_i = \sum_{k=i}^{\infty} p_k$. The probability that it is larger than j is $P_j = \sum_{k=j}^{\infty} p_k$. The probability that a flow size is between j and i given that it is smaller than i is $(P_j - P_i)/(1 - P_i)$. We call it $P_{j,i}$. It follows that:

$$P^{\star}_t(j,i,t,N) = \sum_{k=0}^{t-1} b_{P_i}(k, N-2) \sum_{l=t-k-1}^{N-k-2} b_{P_{j,i}}(l, N-k-2).$$

The first sum accounts for the probability to see less than t flows above i packets. The second sum accounts for the probability to see more than t flows above j given that k flows ($k < t$) were already seen above i. For $t = 1$, $P^{\star}_t(j,i,t,N)$ is no other than $P_t(i,t,N-1)$, and both $\bar{P}^{\star}_{mt}$ and $\bar{P}_{mt}$ are equal (i.e., the ranking and the detection problems are the same).

Trace	Jussieu	Abilene
Link speed	GigE (1 Gbps)	OC-48 (2.5 Gbps)
Duration	2 hours	30 minutes
TCP connections	11M	15M
Packets	112M	125M

Table 1: Summary of the traces

Once $\bar{P}^{\star}_{mt}$ is computed, we multiply it by the total number of flow pairs whose one element is in the top t and the other one is not. This total number is equal to $t(N-t)$. Our metric for the detection problem is the result of this multiplication. As for the ranking problem, we want this metric to be less than one for the detection of the top t flows to be accurate.

5. NUMERICAL RESULTS

We analyze now the accuracy of identifying and ranking the largest flows in a packet stream for both the blind and protocol-aware methods. Our metrics require the following input: p_i, the flow size distribution and N, the total number of flows observed on the link during the measurement period.

To derive realistic values for these two quantities, we consider two publicly available packet-level traces. The first trace is Abilene-I collected by NLANR [15] on an OC-48 (2.5 Gbps) link on the Abilene Network [1]. The second trace has been collected by the Metropolis project [13] on a Gigabit Ethernet access link from the Jussieu University campus in Paris to the Renater Network [18]. Table 1 summarizes the characteristics of the two traces.

We model the flow size distribution in the traces with Pareto. We opted for Pareto since it is known to be appropriate to model flow sizes in the Internet due to its heavy tailed feature [6]. Note that it is not our goal to find an accurate approximation of the distribution of flow sizes in our traces, but rather to find a general, well-known, distribution that approaches the actual flow size. In this section we analyze a wide range of parameters while Section 6 focuses on the performance we observe in the two packet-level traces.

The Pareto distribution is continuous with a complementary cumulative distribution function given by $\mathbb{P}\{S > x\} = (x/a)^{-\beta}$. $\beta > 0$ is a parameter describing the shape of the distribution and $a > 0$ is a parameter describing its scale. The Pareto random variable takes values larger than a, and has an average value equal to $a\beta/(\beta - 1)$. The tail of the Pareto distribution becomes heavier as β decreases.

We use our traces to derive an indicative value of the shape parameter β. To this end, we compute the empirical complementary cumulative distribution of flow sizes and we

[5] In the case $j \geq i$, at most $t-2$ flows can be larger than i packets if we want the flow of size i to be in the top t.

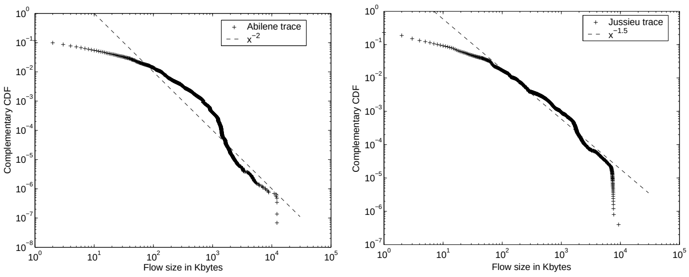

Figure 1: Empirical flow size distribution

plot it on a log-log scale. A heavy-tailed distribution of shape parameter β decays linearly on a log-log scale at rate $-\beta$. The empirical distributions are shown in Figure 1. The plots show that β equal to 2 suits the Abilene trace and β equal to 1.5 suits the Jussieu one. This means that the flow size distribution has a heavier tail in the Jussieu trace.

Then, we compute the average flow size in packets to get the starting point a for the Pareto distribution. As an average flow size we measure 5.76 Kbytes and 7.35 packets on the Abilene trace, and 9.22 Kbytes and 9.9 packets on the Jussieu trace. The total number of flows N is set by taking a measurement interval equal to one minute, then multiplying this interval by the average arrival rate of flows per second on each trace. This gives $N = 487$ Kflows for the Abilene trace and $N = 103$ Kflows for the Jussieu one.

In the rest of this section, all figures plot the ranking metric versus the packet sampling rate p on a log-log scale. We vary p from 0.1% to 50%. Each figure shows different lines that correspond to different combinations of t, β, and N. We are interested in the regions where the value of the metric is below one, indicating that the ranking is accurate on average. To ease the interpretation of results in the figures, we plot the horizontal line of ordinate 1.

5.1 Blind ranking

5.1.1 Impact of the number of flows of interest

The first parameter we study is t, the number of largest flows to rank. The purpose is to show how many flows can be detected and ranked correctly for a given sampling rate. We set β, N, and the average flow size to the values described before. The performance of blind ranking the top t flows is shown in Figure 2 for both traces. We observe that the larger the number of top flows of interest, the more difficult it is to detect and rank them correctly. In particular, with a sampling rate on the order of 1%, it is possible to rank at most the top one or two flows. As we focus at larger values of t, the required sampling rate to get a correct ranking increases well above 10%. Note that with a sampling rate on the order of 0.1%, it is almost impossible to detect even the largest flow. We also observe that the ranking on the Jussieu trace behaves slightly better than that on the Abilene trace. The Jussieu trace has a heavier tail for its flow size distribution, and so the probability to get larger flows at the top of the list is higher, which makes the ranking more accurate. This will be made clear next as we will study the impact of the shape parameter β.

5.1.2 Impact of the flow size distribution

We consider the blind ranking of the top 10 flows varying the shape parameter for the Pareto distribution among five distinct values: 3, 2.5, 2, 1.5 and 1.2. Note that for $\beta \leq 2$ the Pareto distribution is known to be heavy tailed (infinite variance). The other parameters of the model (N and the average flow size) are set as before. The values taken by our metric are shown in Figure 3 for both traces. We can make the following observations from the figure:

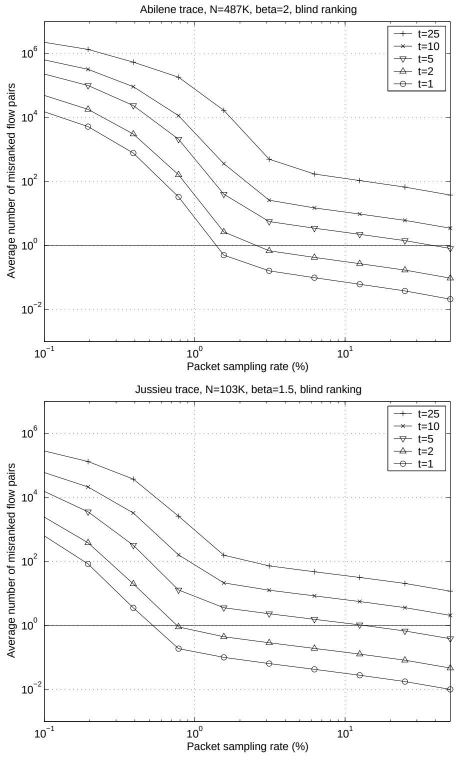

Figure 2: Performance of blind ranking varying the number t of top flows of interest

- Given a sampling rate, the ranking accuracy improves as β becomes smaller, i.e., the tail of the flow size distribution becomes heavier. Indeed, when the distribution tail becomes heavier, the probability to obtain larger flows at the top of the list increases, and since it is simpler to blindly rank larger flows (for distributions satisfying the square root condition, see Section 3.1.1), the ranking becomes more accurate.

- The ranking is never correct unless the sampling rate is very high. In our setting, one needs to sample at more than 50% to obtain an average number of misranked flow pairs below one for a value of β equal to 1.5 (i.e, heavy tailed distribution), and at more than 10% for a value of β equal to 1.2 (i.e., pronounced heavy tailed distribution). For larger values of β (i.e., lighter tail), the sampling rate needs to be as high as 100%.

5.1.3 Impact of the total number of flows

Another important parameter in the ranking problem is N, the total number of flows available during the measurement period. When N increases, the flows at the top of the list should become larger, and therefore as we saw in Section 3.1.1, the blind ranking accuracy should improve

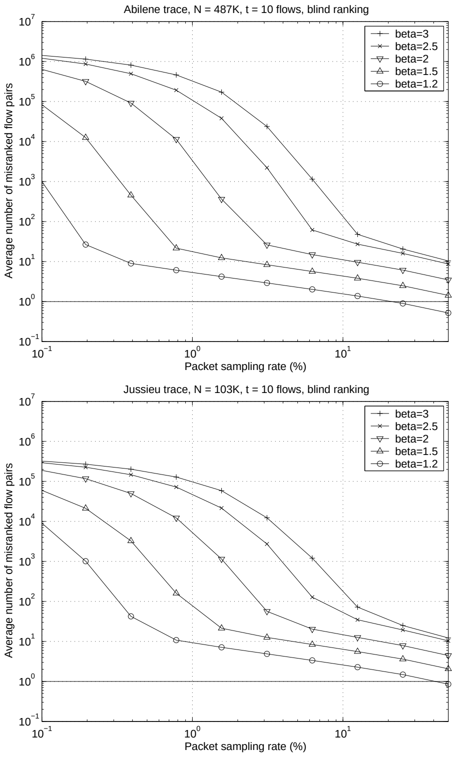

Figure 3: Performance of blind ranking varying the shape parameter of the flow size distribution

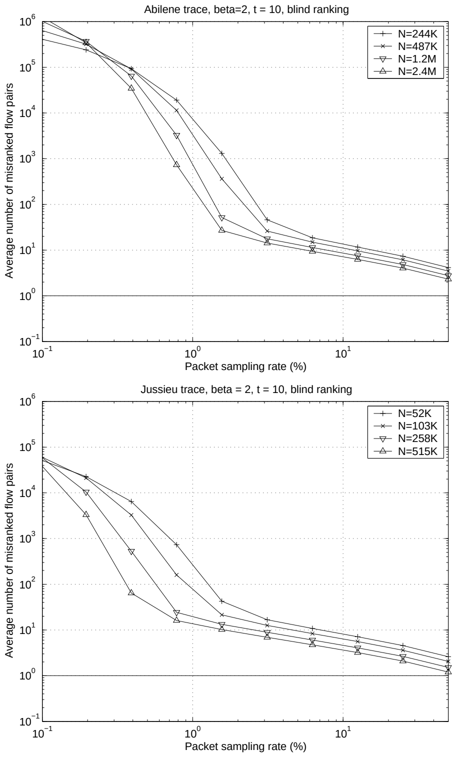

Figure 4: Performance of blind ranking varying the total number of flows

for flow size distributions satisfying the square root condition (in particular the Pareto distribution we are considering here). N varies with the utilization of the monitored link – the higher the utilization, the larger the number of flows. N can also vary with the duration of the measurement period – the longer we wait before ranking and reporting results, the larger the number of flows.

We study the impact of N on the blind ranking accuracy. We take the same value of N used in the previous sections and computed over one minute measurement period (487 Kflows for the Abilene trace and 103 Kflows for the Jussieu trace), then we multiply it by some constant factor ranging from 0.5 (2 times fewer flows) to 5 (5 times more flows). Results are shown in Figure 4. The lines in the figures correspond to a factor value equal to: 0.5, 1, 2.5, and 5. In these figures, we consider the ranking of the top 10 flows with the values of β and average flow size set from the traces. Clearly, the ranking accuracy improves as N increases. However, in our setting, this improvement is still not enough to allow a perfect ranking. One can always imagine increasing N (e.g., by increasing the measurement period) until the top t flows are extremely large and hence, perfectly detected and ranked.

5.2 Protocol-aware ranking

Protocol-aware ranking takes advantage of the information carried in the transport header of the sampled packets to infer the number of non-sampled packets of a flow. We use our model to check whether this improvement exists and to evaluate it. Remember that we are always in the context of low retransmission and duplication rates, which is necessary to remove the discrepancy between carried data volume (throughput) and application data volume (goodput).

Using the previous values for N, β and average flow size, we reproduce Figure 2, but this time for the protocol-aware case. This leads to Figure 5, which illustrates the impact of the number of largest flows to rank. For lack of space, we omit the other figures.

We compare this new figure to its counterpart in the blind case. We make the following two observations:

(i) The protocol-aware method improves the accuracy of the largest flows ranking by an order of magnitude for high sampling rates (above 1%). For example, for the Abilene trace, a sampling rate on the order of 50% was necessary to detect and rank the largest 5 flows with the blind method. Now, with the protocol-aware method, a sampling rate on the order of 5% is sufficient. The same conclusion applies to the Jussieu trace. A sampling rate on the order of 10% is needed. With the protocol-aware method, it becomes on the order of 1%.

(ii) The protocol-aware method does not improve the performance when applied at low sampling rates (above 1%). This can be clearly seen if we compare the plots between both figures for sampling rates below 1%. This results confirms our observations in Section 3.3.2.

5.3 Largest flows detection

To illustrate the difference between ranking and detection, we consider the same scenario as in Section 5.1.1. We plot the detection metric as a function of the sampling rate for different values of t (the number of top flows of interest) and for both Abilene and Jussieu traces. This gives Figure 6 for

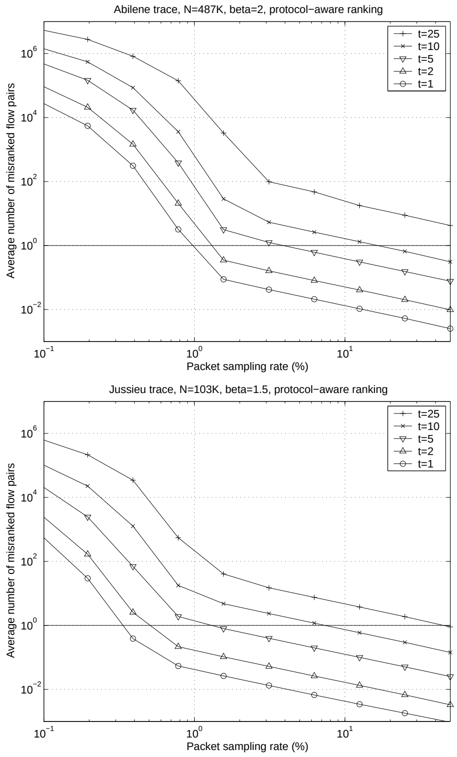

Figure 5: Performance of protocol-aware ranking varying the number t of top flows of interest

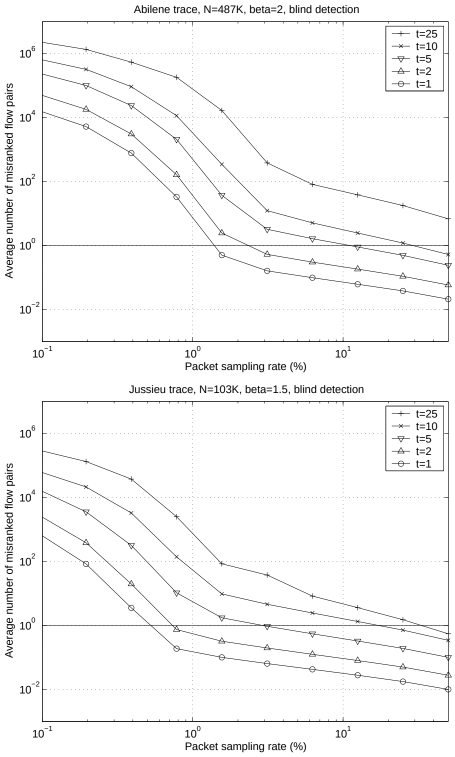

Figure 6: Only detecting the largest flows: Performance of blind ranking varying the number t of top flows of interest

blind ranking and Figure 7 for protocol-aware ranking. A comparison between these results and their counterparts in Figure 2 and 5, respectively, shows a significant improvement in the detection case for both ranking methods. All plots are shifted down by an order of magnitude. For example, in the case of blind ranking, the required sampling rate to correctly rank the top 5 flows was around 50% for the Abilene trace and 10% for the Jussieu trace. Now, with blind detection, it is around 10% and 3%, respectively. Another example is with the protocol-aware method where a sampling rate around 10% was required to rank the largest 10 flows (Figure 5), whereas now, a sampling rate around 1% is sufficient to only detect them. The same gain can be observed if we reconsider the other scenarios in Section 5.1 (not presented here for lack of space). Also, note how in the detection case the protocol aware method allows a better accuracy for high sampling rates when compared to the blind method. For low sampling rates (e.g., below 1%), the accuracy does not improve.

6. EXPERIMENTAL RESULTS

In this section we present the results of running random sampling experiments directly on the packet traces. We use the traces described in Section 5 and compute the performance metrics defined in Section 4.1.

In our traces we consider only TCP packets. Since TCP sequence numbers count bytes, we express the flow sizes in bytes instead of packets throughout this section.

Our experiments are meant to address four major issues that arise when we move from the analytical study to a real network setting: (i) how to deal with invalid TCP sequence numbers in the packet stream; (ii) the importance of flow size distributions and duration of the measurement interval; (iii) the impact of packet loss rates on individual flows – lost packets trigger retransmissions by the TCP senders; (iv) the variability of the detection/ranking performance across multiple bins and packet sampling patterns.

6.1 Implementation of protocol-aware ranking

The protocol-aware method depends on TCP sequence numbers to perform the ranking. For a given flow, it keeps track of the lowest and highest sequence number observed (taking care of packets that wrap around the sequence number space), s_b and s_e respectively.

Note that an actual implementation of this method would just require two 32 bit fields per flow to store the two sequence numbers.

At the end of the measurement period, we compute the difference between the highest and lowest sequence numbers for each sampled flow, and we use the obtained values to rank flows. We then compare this ranking with the one obtained by counting all the bytes each flow transmits in the original non sampled traffic.

In order to discard invalid packets carrying incorrect sequence numbers that would corrupt the ranking, we implement a simple heuristic to update s_e and s_b. A sampled packet with sequence number $S > s_e$ causes an update $s_e \leftarrow S$ if $S - s_e < (\alpha * \mathrm{MTU})/p$. The same rule applies to the updates of s_b. This way we set a limit on the maximum distance in the sequence space between two sampled pack-

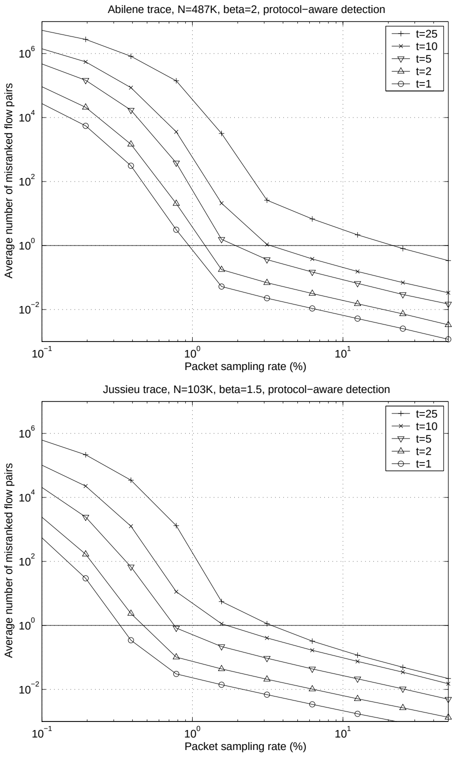

Figure 7: Only detecting the largest flows: Performance of protocol-aware ranking varying the number t of top flows of interest

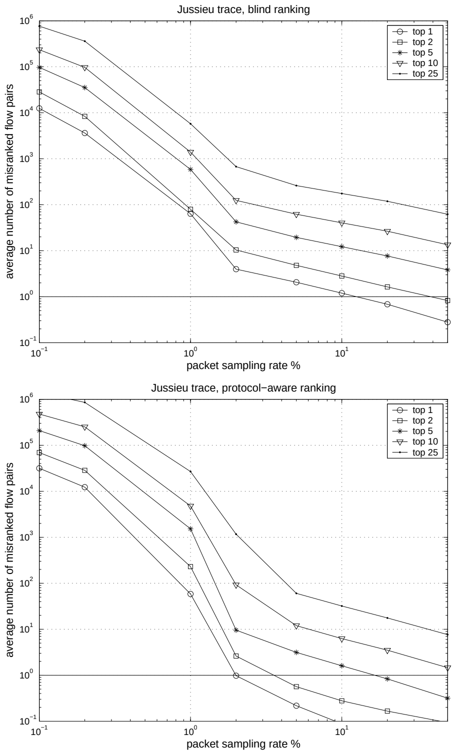

Figure 8: Performance of blind and protocol-aware ranking on Jussieu trace (60s measurement interval).

ets. This distance is inversely proportional to the sampling rate and depends on the Maximum Transmission Unit.

Furthermore, we use the parameter α that allows to make this threshold more or less "permissive" in order to account for the randomness of the sampling process and for other transport-layer events (e.g., packet retransmissions when the TCP window is large). We have run several experiments with different values of α and the results have shown little sensitivity to values of $\alpha > 10$. All the results in this Section are derived with $\alpha = 100$.

6.2 Flow size distribution and measurement interval

As shown in Figure 1, flow size distributions do not follow a perfect Pareto. Furthermore, the measurement interval itself plays a major role in shaping the distribution: it caps the size of the largest flows, that is not unbounded but now depends on the link speed. Indeed, network operators often run measurements using a "binning" method, where packets are sampled for a time interval, classified into flows, ranked, and then reported. At the end of the interval, the memory is cleared and the operation is repeated for the next measurement interval. With this binning method, all flows active at the end of the measurement interval are truncated, so that not all sampled packets of the truncated flow are considered at the same time for the ranking. The truncation may, therefore, penalize large flows and alter the tail of the flow size distribution (where flows are of large size and probably last longer than the measurement interval).

Each experiment consists of the following. We run random sampling on the packet traces and classify the sampled packets into flows. At the end of each measurement interval (set to 1 or 5 minutes), we collect the flows and rank them by the number of bytes sampled for each flow. We compare the ranking before and after sampling using our performance metric (Section 4.1). For each sampling rate we conduct 15 runs and we calculate averages.

The results of the experiments confirm the numerical results of the previous section. In the interest of space, we plot the results of two representative experiments on which we make several observations. The difference between numerical and experimental results, especially at low sampling rates, is caused by the non perfect match of the empirical flow size distribution with Pareto (Figure 1).

Figure 8 shows the performance of ranking flows on the Jussieu trace when the measurement bin is 60s. We consider a wide range of sampling rates from 0.1% to 50% and study the performance when ranking the top 1, 2, 5, 10 and 25 flows in the packet stream. The top graph in Figure 8 is derived using the blind method while the bottom graph shows the performance of the protocol-aware methods. These results are very similar to the numerical results. For sampling rates above 1%, protocol-aware ranking gives approximately an order of magnitude gain on the performance when compared to blind ranking. When the sampling rate is lower than 1%, however, the performance of the two methods is similar. Overall, the blind method requires a sampling rate of 10% to correctly identify the largest flow in the packet stream. The same sampling rate allows to correctly rank the largest 5 flows when using the protocol-aware method.

6.3 Impact of loss rate

In the analysis of the protocol-aware method in Section 3.2, we made the assumption of negligible number of retransmissions for all the flows in the packet stream.

A retransmitted packet may cause inconsistency between the blind and protocol-aware method depending on the location of the monitoring point. Indeed, the blind method counts the total number of bytes sent by the flow while the protocol-aware method considers only the data sent by the transport layer. Therefore, if the packet is lost before the monitoring point, the blind and protocol-aware method will have a consistent view of the number of bytes sent. Instead, if the packet is lost after the monitoring point, the blind method may count this packet twice.

The impact of packet losses on the detection and ranking of the largest flows depends on the metric used to estimate the size of the flows. If flow sizes are estimated according to the total number of bytes sent (i.e., the throughput), then the protocol-aware method may incur in an underestimation error that is independent of the sampling rate (it will occur even if all packets are sampled!). On the other hand, if the flow sizes are estimated according to the transport data sent (i.e., the goodput), then the blind method may incur in an overestimation error independently of the sampling rate.

To illustrate the effect of packet loss rates, we plot in Figure 9 the performance of detecting the largest flows in the Abilene trace when the measurement bin is 5 minutes and the flow sizes are measured using the total number of bytes sent over the link. The top graph shows the performance of the blind method, while the bottom graph presents the results for the protocol-aware method.

We can make the following observations:

- The protocol-aware method keeps performing better than the blind method when the sampling rate is above 1%. At lower sampling rates, the blind method performs better although it presents very large errors.

- For sampling rates above 2%, the curve relative to the detection of the top-25 flows in the protocol-aware method flattens to a value around 70. This is due to the presence of a few flows that experience a high loss rate when compared to other flows. Increasing the sampling rate does not help the protocol-aware method in detecting the largest flows when the volume of bytes sent is used to define the flow size. However, the protocol-aware method can correctly detect the top-25 flows when their size is defined in terms of transport data (see Figure 10).

In summary, the network operator has to choose the metric of interest that depends on the application. For example, for anomaly detection or traffic engineering, a metric that counts the number of bytes sent may be more appropriate. Instead, for dimensioning caches and proxies, the metric that considers the size of the objects transferred may be preferred. This latter metric suits more the protocol-aware method.

6.4 Variability of the results

A last important aspect that we need to address is the variability of the results across multiple measurement intervals and different realizations of the sampling process. Indeed, moving from one measurement interval to another, the composition of flows varies and with it the flow size distribution. Moreover, the sampling process may "get lucky" in certain cases and provide good results. The opposite is also possible.

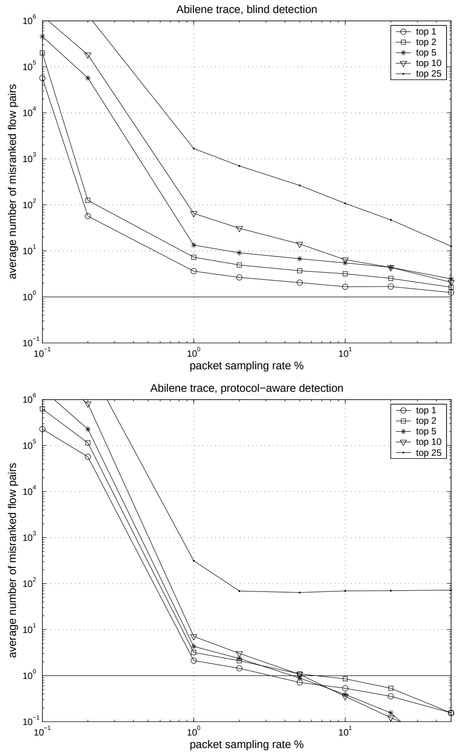

Figure 9: Performance of blind (top) and protocol-aware (bottom) detection on Abilene trace (300s measurement interval).

Figure 11 shows the average performance over 15 sampling experiments of the detection of the top-10 flows in the Abilene trace over the 5-minute measurement intervals. The error bars indicate the standard deviation across the 15 experiments. As usual, the top graph refers to the blind method, while the bottom graph presents the protocol-aware method results.

As we can see the average performance shows limited variability. A sampling rate of 0.1% gives poor results for all bins, while increasing the sampling rates consistently helps. With a sampling rate of 10% the performance metric (i.e., average number of misranked flow pairs) for the blind method is always below 100 while the protocol-aware method is always below 1.

Looking at the standard deviation, we observe large values for the blind method and much smaller values for the protocol-aware method. This indicates that the blind method is more sensitive to the sampling process than the protocol-aware method. The explanation is given in Section 3.3.2 where we showed that that the blind method presents a larger error for large flow sizes (expect when the sampling rate is very low).

7. CONCLUSIONS

We study the problem of detection and ranking the largest flows from a traffic sampled at the packet level. The study is

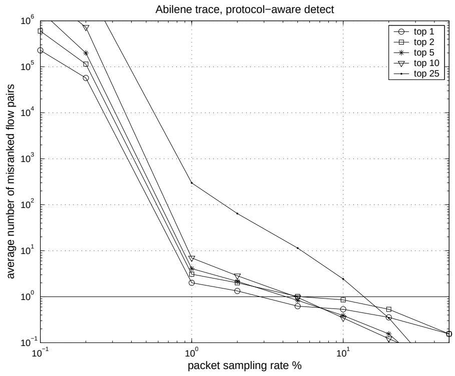

Figure 10: Performance of protocol-aware detection on Abilene trace (300s measurement interval) when using actual amount of data sent by the transport layer application.

done with stochastic tools and real packet-level traces. We find that the ranking accuracy is strongly dependent on the sampling rate, the flow size distribution, the total number of flows and the number of largest flows to be detected and ranked. By changing all these parameters, we conclude that ranking the largest flows requires a high sampling rate (10% and even more). One can reduce the required sampling rate by only detecting the largest flows without considering their relative order.

We also introduce a new method for flow ranking that exploits the information carried in transport header. By analysis and experimentation, we demonstrate that this new technique allows to reduce the required sampling rate by an order of magnitude.

We are currently exploring two possible future directions for this work. First, we want to study the accuracy of the ranking when the sampled traffic is fed into one of the mechanisms proposed in [10, 12] for sorting flows with reduced memory requirements. Second, we are exploring the use of adaptive schemes that set the sampling rate based on the characteristics of the observed traffic.

Acknowledgements

We wish to thank NLANR [15], Abilene/Internet2 [1] and the Metropolis project [13] for making available the packet traces used in this work.

8. REFERENCES

[1] Abilene: Advanced networking for leading-edge research and education. http://abilene.internet2.edu.

[2] C. Barakat, P. Thiran, G. Iannaccone, C. Diot, and P. Owezarski. Modeling Internet backbone traffic at the flow level. *IEEE Transactions on Signal Processing (Special Issue on Signal Processing in Networking)*, 51(8):2111–2124, Aug. 2003.

[3] M. Charikar, K. Chen, and M. Farach-Colton. Finding frequent items in data streams. In *Proceedings of ICALP*, 2002.

[4] B. Y. Choi, J. Park, and Z. Zhang. Adaptive packet sampling for flow volume measurement. Technical Report TR-02-040, University of Minnesota, 2002.

[5] G. Cormode and S. Muthukrishnan. What's hot and what's not: Tracking most frequent items dynamically. In *Proceedings of ACM PODS*, June 2003.

[6] M. Crovella and A. Bestravos. Self-similarity in the World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, Dec. 1997.

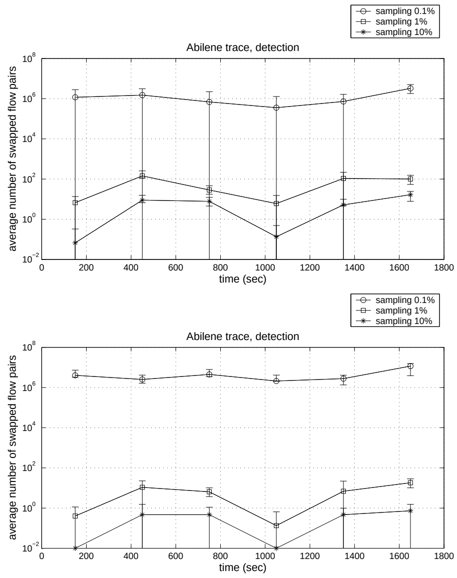

Figure 11: Performance of blind (top) and protocol-aware (bottom) detection over multiple 300s intervals (Abilene trace). Vertical bars show the standard deviation over multiple experiments.

[7] E. Demaine, A. Lopez-Ortiz, and I. Munro. Frequency estimation of internet packet streams with limited space. In *Proceedings of 10th Annual European Symposium on Algorithms*, 2002.

[8] N. G. Duffield, C. Lund, and M. Thorup. Properties and prediction of flow statistics from sampled packet streams. In *Proceedings of ACM Sigcomm Internet Measurement Workshop*, Nov. 2002.

[9] N. G. Duffield, C. Lund, and M. Thorup. Estimating flow distributions from sampled flow statistics. In *Proceedings of ACM Sigcomm*, Aug. 2003.

[10] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proceedings of ACM Sigcomm*, Aug. 2002.

[11] N. Hohn and D. Veitch. Inverting sampled traffic. In *Proceedings of ACM Sigcomm Internet Measurement Conference*, Oct. 2003.

[12] J. Jedwab, P. Phaal, and B. Pinna. Traffic estimation for the largest sources on a network, using packet sampling with limited storage. Technical Report HPL-92-35, HP Laboratories, Mar. 1992.

[13] Metropolis: METROlogie Pour l'Internet et ses services. http://www.laas.fr/ owe/METROPOLIS/metropolis_eng.html.

[14] T. Mori, M. Uchida, R. Kawahara, J. Pan, and S. Goto. Identifying elephant flows through periodically sampled packets. In *Proceedings of ACM Sigcomm Internet Measurement Conference*, Oct. 2004.

[15] NLANR: National Laboratory for Applied Network Research. http://www.nlanr.net.

[16] Packet Sampling Working Group. Internet Engineering Task Force. http://www.ietf.org/html.charters/psamp-charter.html.

[17] K. Papagiannaki, N. Taft, and C. Diot. Impact of flow dynamics on traffic engineering design principles. In *Proceedings of IEEE Infocom*, Hong Kong, China, Mar. 2004.

[18] Renater. http://www.renater.fr.

[19] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A transport protocol for real-time applications. RFC 1889, Jan. 1996.

[20] A. Shaikh, J. Rexford, and K. G. Shin. Load-sensitive routing of long-lived IP flows. In *Proceedings of ACM Sigcomm*, Sept. 1999.

[21] M. Spiegel. *Theory and Problems of Probability and Statistics*. McGraw-Hill, 1992.