

Algorithms and Estimators for Summarization of Unaggregated Data Streams

Edith Cohen^{*} Nick Duffield[†] Haim Kaplan[‡] Mikkel Thorup[§] Carsten Lund[¶]

September 10, 2007

Abstract

¹ Statistical summaries of IP traffic are at the heart of network operation and are used to recover information on arbitrary subpopulations of flows. It is therefore of great importance to collect the most accurate and informative summaries given the router's resource constraints. IP packet streams consist of multiple interleaving IP flows. While queries are posed over the set of flows, the summarization algorithm is applied to the stream of packets. Aggregation of traffic into flows before summarization is often infeasible and therefore the summary has to be produced over the unaggregated stream. Cisco's sampled NetFlow, based on aggregating a sampled packet stream into flows, is the most widely deployed such system.

We observe two sources of inefficiency in current methods. Firstly, a single parameter (the sampling rate) is used to control utilization of both memory and processing/access speed, which means that it has to be set according to the bottleneck resource. Secondly, the unbiased estimators are applicable to summaries that in effect are collected through uneven use of resources during the measurement period (information from the earlier part of the measurement period is either not collected at all and fewer counters are utilized or discarded when performing a sampling rate adaptation).

We develop algorithms that collect more informative summaries through an even and more efficient use of available resources. The heart of our approach is a novel derivation of efficiently-computable *unbiased* estimators that use these more informative counts. We prove analytically that these estimators are superior (have at most the same variance on *all* packet streams and subpopulations) to previous approaches. Simulations on Pareto distributions and IP flow data show that the new summaries provide significantly more accurate estimates. We provide an implementation design that can be efficiently deployed at routers.

1 Introduction

Collection and summarization of network traffic data is necessary for many applications including billing, provisioning, anomaly detection, inferring traffic demands, and configuring packet filters and routing protocols. Traffic consists of interleaving packets of multiple flows but the summaries should support queries on statistics of subpopulations of IP flows, such as the amount of traffic that belongs to a particular protocol, originate from a particular AS, or both. These queries are posed after the sketch is produced and therefore it is critical to retain sufficient meta data information and provide estimators that facilitate such queries.

^{*}AT&T Research Labs, 180 Park Ave. Florham Park, NJ, USA. Email: edith@research.att.com.

[†]AT&T Research Labs, 180 Park Ave. Florham Park, NJ, USA. Email: duffield@research.att.com.

[‡]Computer Science Department, Tel Aviv University, Tel Aviv, Israel. Email: haimk@post.tau.ac.il

[§]AT&T Research Labs, 180 Park Ave. Florham Park, NJ, USA. Email: mthorup@research.att.com.

[¶]AT&T Research Labs, 180 Park Ave. Florham Park, NJ, USA. Email: lund@research.att.com.

¹This is a full version that expands on two conference papers [3, 2]

IP packet streams are processed in real-time at the routers by systems such as Cisco’s sampled NetFlow (NF) or processed by software tools such as Gigascope [8]. Two critical resources in the collection of data are the high-speed memory (usually expensive fast SRAM) and CPU power that are used to process the incoming packets. The available memory limits the number of cached flows that can be actively counted. The processing power limits the level of per-packet processing and the fraction of packets that can undergo higher-level processing.

The practice is to obtain periodic summaries (sketches) of traffic by applying a data stream algorithm to the raw packet stream. NF samples packets randomly at a fixed rate. Once a flow is sampled, it is cached, and a counter is created that counts subsequent sampled packets of the same flow. The number of counters is the number of distinct sampled flows. The packet-level sampling that NF performs serves two purposes: Firstly, it addresses the memory constraint by reducing the number of distinct flows that are cached (the bulk of small flows is not sampled). Without sampling, we need a counter for each distinct flow in the original stream. Secondly, the sampling reduces the processing power needed for the aggregation, since only sampled packets require the higher-level processing required in order to determine if they belong to a cached flow.

An algorithm that is able to count more packets than NF using the same number of statistics counters (memory) is *sample-and-hold* (SH) [13, 12]. With SH, as with NF, packets are sampled at a fixed rate and once a packet from a particular flow is sampled, the flow is cached. The difference is that with SH, once a flow is actively counted, all subsequent packets that belong to the same flow are counted (with NF, only sampled packets are counted). SH sketches are considerably more accurate than NF sketches [13, 12]. A disadvantage of SH over NF, however, is that the summarization module must process *every* packet in order to determine if it belongs to a cached flow. This additional processing makes it less practical for high volume routers.

NF and SH use a *fixed* packet sampling rate, as a result, the number of distinct flows that are sampled and therefore the number of statistics counters required is variable. When conditions are stable, the number of distinct flows sampled using a given sampling rate has small variance. Therefore one can manually adjust the sampling rate so that the number of counters does not exceed the memory limit and most counters are utilized [12]. Anomalies such as DDoS attacks, however, can greatly affect the number of distinct flows. A fixed-sampling-rate scheme can not react to such anomalies as its memory requirement would exceed the available memory. Therefore, anomalies would cause disruption of measurement or affect router performance. This issue is addressed by *adaptive* variants that include *adaptive sampled NetFlow* (ANF) [13, 11, 16] and *adaptive SH* (ASH) [13, 12]. These variants adaptively decrease the sampling rate and adjust the values of the statistics counters as to emulate sampling with a lower rate.

Overview

We design sketching algorithms for packet streams that obtain considerably more accurate statistics than existing approaches. We focus on subpopulation-size queries (number of packets of a subpopulation) but the sketches can be used for other queries such as bytes and flow size distribution. We use available resources in a balanced and load-sensitive way to collect more information from the packet sample. Then we design unbiased estimators that use the additional information. Our algorithms are robust to anomalies and changes in traffic patterns, and gracefully degrade performance when there is a decrease in available resources. They are supported by rigorous analysis.

Step counts for NF and SH NF, SH, and their adaptive variants do not equally utilize available resources through the measurement period: The number of cached flows increases through the measurement period and reaches its maximum only at the end. The adaptive ANF and ASH fully utilize all counters, but this utilization is in a sense “wasted” and does not translate into more accurate estimates, as each *rate adaptation*

(decrease of the sampling rate) is implemented by discarding the more informative counts obtained with the lower sampling rate [13, 11].

Step-counting NetFlow (SNF) and *step-counting sample-and-hold* (SSH) process the same packets as their adaptive counterparts, but when performing rate adaptation, they *retain* the current counts. SSH and SNF build on a simple but powerful design [20, 18] of transferring partial counts from (fast and expensive) SRAM to (slower and cheaper) DRAM, which allows us to use smaller size counters in SRAM and add the counts to larger DRAM counters when the SRAM counters are about to overflow. This design allows us to distinguish between the resources required for active counting and those required for intermediate storage. While applicable to all methods, SNF and SSH are able to make a further use of this design by transferring, after each rate adaptation, the counts into DRAM. Counting more of the processed packets in the final summary is the key for obtaining better estimates.

Hybrids of NF and SH There are multiple resource constraints for gathering statistics. At the router, the memory size that determines the number of statistics counters and the CPU processing (or size of specialized hardware) that determines the fraction of packets that can be examined against the flow cache. Other constraints are the available bandwidth and storage to transmit and store the final sketch. Previous schemes, however, use a *single parameter* (sampling rate) with these multiple constraints: NF (and ANF and SNF), must set (or adjust) the sampling rate to be low enough so that the number of counters does not overflow or over utilize the router memory. As a result, resources available for processing packets may not be fully utilized. SH (and ASH and SSH), on the other hand, do not address CPU processing constraints at all, and all packets are processed. We propose *hybrid* schemes of SH variants that use *two* packet sampling rates. The first one controls the fraction of packets that are processed in order to determine if they belong to an already-cached flow. The second, and lower, rate, determines the fraction of packets that can create new entries of cached flows.

Subpopulation weight estimators (Section 4) The sketches we produce (for all algorithms) have the form of a subset of the flows along with the flow attributes and an *adjusted weight* associated with each flow. Adjusted weights have the property that for each flow, the expectation is equal to its actual size (Adjusted weights of flows not included in the sketch is defined to be zero). Therefore, an unbiased estimate for the size of a subpopulation of flows can be obtained by summing the adjusted weights of flows in the sketch that belong to this subpopulation. The per-flow unbiasedness property is highly desirable as accuracy increases when aggregating over larger subpopulations and when combining estimates obtained from sketches of different time periods. The heart of our work is the calculation and analysis of unbiased adjusted weights.

The derivation of adjusted weights for NF, which applies fixed-rate sampling, is standard: a simple scaling of the counts by the inverse sampling rate. We derive adjusted weights assignments for ANF, ASH, SSH and SNF. These assignments can be computed efficiently. An important contribution we make is the understanding of what information to gather and how to use it to obtain correct adjusted weights. The derivation and efficient computation of correct unbiased adjusted weights for SNF and SSH is novel and highly nontrivial. The adjusted weights we derive have minimum variance among all estimators that use the same information (the counts gathered by the algorithm for the flow), and in this sense are *optimal*.

The quality of the adjusted weight assignment depends on the distribution over subsets of flows that are included in the sketch, the information collected by the algorithm for these flows, and the procedure used to calculate these weights. We show that the distribution of the subsets of flows included in the sketch produced by each of the algorithms NF, SH, hybrids, and variants, is that of drawing a weighted sample without replacement (WS) from the full set of aggregated flows. Therefore, the difference in the quality of the sketches stems only from the variance of the adjusted weights we assign. More informative counts are beneficial only if they correspond to adjusted weights with lower variance. We analyze the variance of the adjusted weight assignment and establish a strong relation between the different methods that holds for any

packet stream and any flow or subpopulation of flows: The SSH estimator has at most the variance of ASH and of SNF, and SNF has at most the variance of ANF.

FSD and other subpopulation properties (Sections 5 and 6) There are multiple aggregates of interest over subpopulations and it is important that the same summaries can support queries for these aggregates. We derive unbiased estimators for important classes of aggregates: Flow Size Distribution (FSD) estimators, that provide unbiased estimates on the number of distinct flows of certain range of sizes in a subpopulation are derived in Section 5. Estimators for other properties such as total bytes or total number of AS hops to destination are derived in Section 6.

For a cleaner exposition, lengthy or technical proofs are deferred to Section 9.

Implementation Our implementation design piggybacks on several existing ingredients. The basis is the flow counting mechanism that Cisco’s NF deploys. (Proposed improved implementation such as [18, 11] can also be integrated.) A router implementation of adaptive sampling rate for ANF was proposed in [11, 16] (rate adaptation was termed *renormalization*. This design can also be used for ASH and the step-counting and hybrid variants.

Discretized sampling rates (Section 7) The pure adaptive models perform a rate adaptation each time a flow is “evicted” from the cache. Rate adaptations, however, are intensive operations [11, 16]. We design a “router friendly” variant of the pure model with discretized sampling rates. This design drastically reduces the number of rate adaptations and also simplifies their implementation. As in [11], discretization allows us to efficiently perform rate adaptations. The discretized model, however, differs mathematically from the pure sampling schemes. We show how to apply the estimators derived for the pure schemes to the discretized schemes. More importantly, we show that these estimators are also unbiased and retain other key properties of the estimators for the pure model. Furthermore, the particular discretization we used was critical for the unbiasedness arguments to hold.

Performance study (Section 8) We evaluate the performance of these methods on IP flows data collected by *unsampled* NF running on a gateway router and on synthetic data obtained using Pareto distribution with different parameter values. On the IP data, we consider subpopulations of flows that belong to specified applications, and on the synthetic data, we consider prefixes and suffixes of the flow size distribution. We observe that the step-counting SNF and SSH provide significantly more accurate estimates than their adaptive counterparts. The SH variants significantly dominate their NF counterparts and the hybrid version provides a smooth performance curve between these two extremes. Even with low sampling rates, the hybrids are able to provide much more accurate estimates than plain NF, ANF, and SNF. We also show that our implementation design can be tuned to provide a very low number of rate adaptations.

2 Related work

An orthogonal summarization problem is summarizing *aggregated* data [14]. For example, using k -mins or bottom- k sketches [1, 4, 7, 5, 10]. Estimators developed for these summaries utilize the weight of each item, which is not readily available in our unaggregated setup. Direct application requires *pre-aggregation*, that is, obtaining an exact packet count for each flow as when running unsampled NF. This is infeasible in high volume routers as it requires processing of every packet and storing an active counter for every flow. These methods can be used, however, to trim the size of a sketch obtained using any method that obtains unbiased adjusted weights (including NF, SH, and their variants), when trimming is needed in order to address transmission bandwidth or storage constraints.

An extension of ASH that does not discard counts when a rate adaptation is performed was considered in [12] for finding “elephant flows.” While this extension attempts to provide similar benefit to step-counting, it is not adequate for estimating subpopulation sizes. The unadjusted count itself is indeed a better estimator than the reduced count for each individual flow, but this estimator is inherently biased. The bias depends on where in the measurement period the packets occurred, and an unbiased estimator can not be constructed from the counts collected. The relative bias is very large on smaller flows (of the order of the inverse sampling-rate) and if used to estimate subpopulation sizes for such flows, we can obtain a large relative error on such subpopulations (even if the subpopulation size is large).

Kumar et al [17] proposed a streaming algorithm for IP traffic that produces sketches that allow us to estimate the flow size distribution (FSD) of subpopulations. Their design executes two modules concurrently. The first is a sampled NetFlow module that collects flow statistics, along with full flow labels, over sampled packets. The second is a streaming module that is applied to the full packet stream and uses an array of counters, accessed by hashing. Estimating the flow size distribution is a more general problem than estimating the size of a subpopulation, and therefore this approach can be used to estimate the subpopulation sizes. To be accurate, however, the number of counters in the streaming module should be roughly the same as the number of flows and therefore the size of fast memory (SRAM) should be proportional to the number of distinct flows. As we see here, however, accurate estimates for subpopulation sizes can be obtained more efficiently using other approaches.

In some cases, protocol-level information such as testing for the TCP syn flag [9] on sampled packets and using TCP sequence numbers [19] can be used to obtain better estimates of the size of the flow from sampled packets. These methods can significantly increase the accuracy of estimating the flow size distribution of TCP flows from packet samples, but are not as critical for subpopulation size estimates for subpopulations with multiple flows. We focus on general-purpose all-protocols estimators, but these methods can be integrated with our sketches.

3 Sketching algorithms

We provide the underlying mathematical models of the different flow sampling schemes. These models are used in the analysis and are mimicked by the “router friendly” implementations. Our sampling schemes are data stream algorithms that are applied to a stream of packets.

Sampled NF performs fixed-rate packet sampling. Packets are sampled independently at a rate p and sampled packets are aggregated into flows. All flows with at least one sampled packet are *cached* and there is an *active counter* for each flow. The sketch includes all flows that are cached in the end of the measurement period.

SH, like NF, samples packets at a fixed rate p and maintains a cache of all flows that have at least one sampled packet. SH, however, processes all packets and not only sampled packets. If a processed packet belongs to a cached flow, it is counted.

Rank-based view of the sample space

We facilitate the analysis through a *rank-based view* of the sample space: Each point in the sample space is a *rank assignment*, where each packet is assigned a rank value that is independently drawn from $U[0, 1]$. The actions of each of the sampling algorithms that we consider are defined by the rank assignment. Implementations do not track per-packet random rank values or even draw rank values. They maintain just enough “partial” information on the rank assignment to maintain a flow cache and counts that are consistent with the rank-based view.

For each flow $f \in F$ and position in the packet stream, we define the *current rank value* $r(f)$ to be the smallest rank assigned to a packet of the flow that occurred before the current position in the packet stream.

An NF sketch with sampling rate p is equivalent to obtaining a rank assignment and counting all packets that have rank value $< p$. The set of actively counted flows at a given time is $\{f \in F | r(f) < p\}$.

An SH sketch with sampling rate p is equivalent to obtaining a rank assignment and counting all packets such that the current rank of the flow (including the current packet) at the time the packet is processed is smaller than p .

Adaptive algorithms

The adaptive algorithms ANF and ASH work with a fixed limit k on the number of cached flows and produce a sketch of k flows. When the number of cached flows exceeds k , the sampling rate is “retroactively” decreased, and counts are adjusted, so that only k flows remain cached. The current *sampling rate* is defined to be the $(k + 1)$ st smallest rank among $r(f)$ ($f \in F$) (it is defined to be 1 if there are fewer than $(k + 1)$ distinct flows.)

It is sometimes necessary to set a limit $p_{\text{start}} < 1$ on the initial sampling rate. In this case, the current sampling rate is defined to be p_{start} if there are fewer than $(k + 1)$ distinct flows and otherwise is the minimum of p_{start} and the $(k + 1)$ st smallest rank among $r(f)$ ($f \in F$).

The sampling rate is determined by the rank assignment, the prefix of processed packets, and p_{start} (it does not depend on the sampling scheme). The *effective sampling rate* is defined as the value of the sampling rate at the end of the measurement period.

The set of cached flows at a given time (by either ANF or ASH) are the flows with current rank that is below the current sampling rate. The (current) ANF count of a cached flow is the number of packets seen so far with rank that is below the current sampling rate. The (current) ASH count of a cached flow f is the number of packets seen so far such that the rank of the flow just after the processing of the packet is below the current sampling rate.

The sketch includes the set of cached flows and their counts at the end of the measurement period.

Step-counting

We refer to a decrease of the current sampling rate as *rate adaptation*. The adaptive algorithms [13, 11] implement rate adaptation by decreasing the more informative flow counts that corresponded to the higher sampling rate. The step-counting algorithms, sNF and sSH, record these counts instead of adjusting them down. The active counting of packets performed by the step-counting algorithms is just like their adaptive counterparts, but the step-counting algorithms produce a vector of counts, rather than a single count, for each flow in the sketch.

Using the rank-based view, sNF counts include all packets such that (i) the rank of the packet is below the value of the sampling rate when (just after) the packet is processed (ii) the rank of the flow remains continuously below the sampling rate since the packet was processed until the current time. sSH counts include all packets such that when (just after) the packet is processed and continuously until the current time, the rank of the flow is below the sampling rate. With both sNF and sSH, a flow is cached if and only if its rank is below the current sampling rate. Figure 1 motivates, through an example, the design of the step-counting algorithms. The figure shows two flows and the expected number of packets included in the final count by each method. The step-counting sNF and sSH count many more packets than their adaptive counterparts.

Hybrids

Figure 1 shows that ASH and sSH count many more packets than ANF and sNF. This advantage, however, comes at a cost: ANF and sNF simply ignore all packets that are not sampled while ASH and sSH have to process all packets in order to determine if they belong to a cached flow. Our *hybrid* algorithms provide a

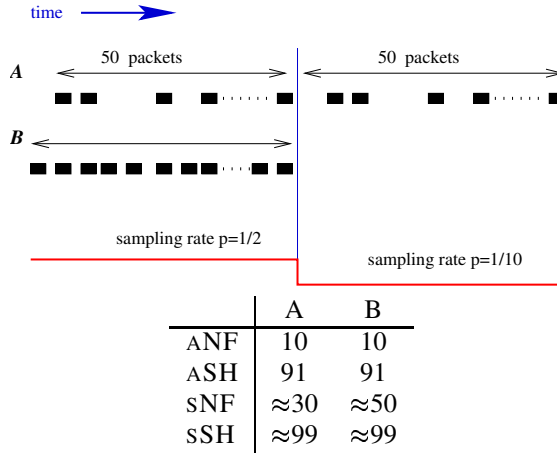


Figure 1: Expected counts of two 100-packet flows. Flow A has 50 packets with current sampling rate 0.5 and 50 packets with current rate 0.1. Flow B has 100 packets with current sampling rate 0.5. The final counts with ANF and ASH only depends on the final sampling rate. SNF and SSH benefit from the higher sampling rate at the beginning of the measurement period and the final counts account for more packets.

smooth tradeoff between the fraction of packets that are processed and the fraction that are included in the final count.

Hybrid sketching algorithms use a base sampling rate parameter p_{base} , which controls the fraction of packets that are processed by the algorithm. The initial sampling rate is $p_{\text{start}} \leq p_{\text{base}}$. A hybrid algorithm samples all packets independently at a fixed rate p_{base} and then applies a respective basic algorithm (SH, ASH, or SSH) to the p_{base} -sampled stream with initial sampling rate $p'_{\text{start}} = p_{\text{start}}/p_{\text{base}}$. (Hybrid NF variants with $p_{\text{start}} \leq p_{\text{base}}$ are equivalent to applying the underlying NF variant and therefore we only consider hybrid SH variants).

The rank-based view of the hybrid algorithms is as follows. Hybrid-ASH counts include a packet if and only if (i) the rank of the packet is below p_{base} and (ii) the rank of the flow at the time (just after) the packet is processed is below the current sampling rate. Hybrid-SSH counts include a packet if and only if (i) the rank of the packet is below p_{base} and (ii) continuously, from (just after) the time the packet is processed until the current time, the rank of the flow was below the sampling rate. A flow is cached with hybrid-ASH and hybrid-SSH if and only if its rank value is below the current sampling rate.

An equivalent rank-based view of the hybrid algorithms discards all packets with rank value above p_{base} , scales the rank values of the remaining packets and p_{start} by p_{base}^{-1} , and applies the respective basic algorithm.

The following table shows the expected number of packets that are counted and processed for the example 100-packet flows in Figure 1. The hybrids use $p_{\text{start}} = p_{\text{base}} = 0.5$ and have respective sampling rates of 1 and 0.2 on the p_{base} -sampled stream. The step-counting algorithm count more of the processed packets than their adaptive counterparts. The hybrids provides a tradeoff that preserves the higher ratios between counted and processed packets.

| | A | B |
|------------|---------|----------|
| ANF | 10/30 | 10/50 |
| ASH | ≈91/100 | ≈ 91/100 |
| hybrid-ASH | ≈46/50 | ≈46/50 |
| SNF | ≈30/30 | ≈50/50 |
| SSH | ≈99/100 | ≈99/100 |
| hybrid-SSH | ≈50/50 | ≈50/50 |

4 Adjusted weights

We use the information collected by the algorithms to compute adjusted weights that are associated with each flow that is included in the sketch (flows that are not in the sketch have an adjusted weight of zero).

We use the notation $A^L(f)$, where $L \in \{\text{ASH}, \text{ANF}, \text{SSH}, \text{SNF}, \text{NF}, \text{SH}\}$, for the random variable that is the adjusted weight assigned to the flow f by the algorithm L . The adjusted weights we assign are a function of the counts collected by the algorithm and of the sampling rate.

Clearly, a correct adjusted weight for NF counts for a flow f with count $i(f)$ is $A_p^{\text{NF}}(f) = i(f)/p$ (the number of counted packets divided by the sampling rate p). The derivation of the adjusted weight assignments for the adaptive algorithms is more subtle and based on partitioning the sample space as in [4, 5, 10]. This partitioning allows us to apply the adjusted weights expressions that are applicable to the corresponding fixed-rate variants.

It turns out that there is a unique deterministic assignment of adjusted weights for each of the algorithms considered. Since deterministic assignment has smaller variance than any randomized one, it is preferable.

Lemma 4.1. *Let $i(f)$ be the packet count collected for a flow f by ANF. $A_{p'}^{\text{ANF}}(f) = i(f)/p'$, where p' is the effective sampling rate, are correct adjusted weights.*

Proof. Consider a flow f , and the probability subspace where the k th smallest rank among $r(f')$ ($f' \in F \setminus \{f\}$) is p' . Consider the conditional distribution of the number of packets of flow f that are counted. The number of packets is just like what would have been counted with NF with rate p . Therefore, $A_{p'}^{\text{ANF}}(i(f)) = A_p^{\text{NF}}(i(f))$ are unbiased adjusted weight for f within this probability subspace. \square

Lemma 4.2. *Let $i(f)$ be the number of counted packets for a flow f with SH. The assignment $A_p^{\text{SH}}(f) = i(f) + (1 - p)/p$ if $i(f) > 0$ (the flow is sampled) are correct adjusted weights.*

Proof. We derive $A_p^{\text{SH}}(i)$, the adjusted weight assigned to a flow with count of i packets (we omit the superscript and subscript when clear from context) and show that this is the unique deterministic assignment.

We have $A_p(0) = 0$ (items that are not sampled at all obtain zero adjusted weight). In order for the assignment to be unbiased for 1-packet flows, we have to have $pA_p(1) + (1 - p)A_p(0) = 1$. Substituting $A_p(0)$, we obtain that $A_p(1) = 1/p$. To be correct for n -packet items, we must have $\sum_{i=0}^n (1 - p)^i p A_p(n - i) = n$. We can solve these for $n = 2, 3, 4, \dots$ and obtain that $A_p(n) = (1 + (n - 1)p)/p = (1 - p)/p + n$ for $n \geq 1$. \square

This assignment can also be derived by applying the Horvitz-Thompson estimator to each packet. For each packet, we look at the partition of the sample space to two parts, one where a previous packet is sampled, and the other where a previous packet is not sampled. The adjusted weight assigned to a packet is unbiased on each part: if a previous packet is sampled, the probability that a packet is counted is 1 and its adjusted weight is 1. If no previous packet is sampled, then the probability that the packet is sampled is p and the Horvitz-Thompson adjusted weight is $1/p$. The adjusted weight of the flow is the sum of the adjusted weights of sampled packets. This assignment can be interpreted as the first sampled packet of the flow representing $1/p$ unseen packets whereas subsequent counted packets of the flow represent only themselves.

Lemma 4.3. *Consider ASH and let $i(f)$ be the ASH count of a flow and p' be the effective sampling rate. The assignment $A_{p'}^{\text{ASH}}(f) = i(f) + (1 - p')/p'$ (if $i(f) > 0$) are correct adjusted weights.*

Proof. The proof is similar to that of Lemma 4.1: For each flow, we consider the probability subspace where the k th smallest rank among the flows $F \setminus \{f\}$ is fixed and assign unbiased weights within this subspace. \square

The information we collect using SNF and SSH is the step function $1 \geq p_1 > p_2 > \dots > p_r$ denoted by the vector $\mathbf{p} = (p_1, \dots, p_r)$ of the current sampling rate and for each sampled flow, the counts $\mathbf{i}(f) = (i_1(f), i_2(f), \dots, i_r(f))$ of the number of packets recorded at each step. The adjusted weight assignment for a flow f is a function of \mathbf{p} and $\mathbf{i}(f)$. We use the notation

$$A_{\mathbf{p}}^{\text{SNF}}(\mathbf{i}(f)) \equiv A_{p_1, p_2, \dots, p_r}(i_1(f), i_2(f), \dots, i_r(f))$$

for the adjusted weight assigned by SNF, and similarly, $A_{\mathbf{p}}^{\text{SSH}}(\mathbf{i}(f))$ for the adjusted weight assigned by SSH.

Adjusted weights are computed after the counting period is terminated. After they are computed, the count vectors can be discarded. Therefore, SNF and SSH produce a sketch of size k .

We will show the following:

Theorem 4.4. *The adjusted weight $A_{\mathbf{p}}^{\text{SNF}}(\mathbf{n})$ for SNF and $A_{\mathbf{p}}^{\text{SSH}}(\mathbf{n})$ for SSH can be computed using number of operations that is quadratic in the number of steps with a non-zero count.*

We derive adjusted weights for SSH and SNF sketches.

The adjusted weights for ANF and ASH are based on “plugging in” the effective (final) sampling rate in the adjusted weights expressions of the non-adaptive variant. The argument for unbiasedness is based on the fact that the adjusted weights of each flow are unbiased on each part of some partition of the sample space. For the step-counting algorithms, the adjusted weights assigned to a flow f are unbiased in the probability subspace defined by the steps of the rank value of the current k th-smallest rank of a flow among $F \setminus f$. These steps are the same as the current sampling rate when the flow is actively counted. Technically, we need to consider the k th-smallest rank of an actively counted flow on steps that precede the active counting of f . The adjusted weight function, however, has the property

$$\begin{aligned} A_{p_1, p_2, \dots, p_r}(0, \dots, 0, i_j, i_{j+1}, \dots, i_r) = \\ A_{p_j, p_{j+1}, \dots, p_r}(i_j, i_{j+1}, \dots, i_r) \end{aligned} \quad (1)$$

and therefore does not depend on the current sampling rate in the duration before the final contiguous period where the flow is actively counted. This means that it is sufficient to record the steps of the current sampling rate. Eq. (1) is an instance of the following generalization that states that the adjusted weight assignment does not depend on the values of the current sampling rate in durations when there are no counted packets.

Lemma 4.5. *Consider a correct assignment of adjusted weight $A_{\mathbf{p}}(\mathbf{n})$. For an observed count \mathbf{i} and \mathbf{p} , let $1 \leq j_1 < j_2 < \dots < j_{r'} = r$ be the coordinates such that $i_{j_k} > 0$ or $i_{j_k} = r$ (that is, r is included also if $i_r = 0$).*

$$\begin{aligned} A_{p_1, p_2, \dots, p_r}(0, \dots, 0, i_{j_1}, 0, \dots, 0, i_{j_2}, \dots) = \\ A_{p_{j_1}, p_{j_2}, \dots, p_{j_{r'}} = p_r}(i_{j_1}, i_{j_2}, \dots, i_{j_{r'}}) \end{aligned} \quad (2)$$

The lemma allows us to state the adjusted weight of a flow in terms of an equivalent flow where the number of steps is equal to the number of steps where the original flow had a nonzero count. It also allows us to assume without loss of generality in our analysis that all steps except possibly the last step have positive counts.

4.1 Adjusted weights for SSH

Let r be the number of steps and $p_1 > \dots > p_r$ the corresponding sampling rates. For a flow f , let $\mathbf{n} = (n_1, \dots, n_r)$ be the number of packets of f in each step and let $\mathbf{i} = (i_1, \dots, i_r)$ be the number of counted packets in each step. We denote by $q[\mathbf{i}|\mathbf{n}]$ the probability that a flow with \mathbf{n} packets has a count of \mathbf{i} .

We provide expressions for the adjusted weights for SSH and SNF sketches that can be computed in time quadratic in the number of steps (which is logarithmic in the number of packets). We then argue that

the number of operations can be further reduced to be quadratic in the number of steps where the flow has a non-zero count. This distinction is important since many flows, in particular bursty or small flows, can have non-zero count on a single step or very few steps.

We define the values $c_{i,j}(\mathbf{p}, \mathbf{n})$ ($1 \leq i \leq j \leq r$) as follows (the parameters (\mathbf{p}, \mathbf{n}) are omitted when clear from context, and we assume $n_1 > 0$ w.l.o.g.):

$$\begin{aligned} 1 \leq j \leq r : & \quad c_{1,j} = (1 - p_j) \\ 2 \leq j \leq r : & \quad c_{2,j} = (1 - p_j)^{n_1-1} (c_{1,j} - c_{1,1}) \\ 3 \leq i \leq j \leq r : & \quad c_{i,j} = (1 - p_j)^{n_i-1} (c_{i-1,j} - c_{i-1,i-1}) \end{aligned}$$

The following two lemmas are immediate from the definitions.

Lemma 4.6. • For $1 \leq j \leq r$, $c_{1,j}$ is the probability that the rank of the first packet of the flow is at least p_j .

• For $2 \leq i \leq j \leq r$, $c_{i,j}(\mathbf{p}, \mathbf{n})$ is the probability that the flow \mathbf{n} is fully counted by SSH until the transition into step i , and at the beginning of step i , the rank of the flow is at least p_j .

Lemma 4.7. The computation of the partial sums $\sum_{h=1}^i c_{h,h}$ for $i = 1, \dots, r$ can be performed in $O(r^2)$ operations.

By lemma 4.6, $c_{i,i}$ ($i \in \{1, \dots, r\}$) is the probability that the SSH counting of the flow progressed continuously from the start until the transition into step i , and halted in this transition (as the current rank of the flow was above p_i). So

$$q[\mathbf{n}|\mathbf{n}] = 1 - \sum_{h=1}^r c_{h,h}. \quad (3)$$

The following theorem expresses the adjusted weight $A^{\text{SSH}}(\mathbf{n})$ as a function of the diagonal sums $\sum_{h=1}^i c_{h,h}$ ($h = 1, \dots, r$). The proof is provided in Section 9.1.

Theorem 4.8.

$$A^{\text{SSH}}(\mathbf{n}) = \frac{(1 - p_1) + \sum_{i=1}^r n_i (1 - \sum_{h=1}^i c_{h,h})}{1 - \sum_{h=1}^r c_{h,h}}.$$

Lemma 4.9. The adjusted weight $A^{\text{SSH}}(\mathbf{n})$ can be computed using $O(r^2)$ operations.

Proof. The proof follows from Lemma 4.7 and Theorem 4.8. □

4.2 Adjusted weights for hybrids

Unbiased adjusted weights for hybrid-ASH and hybrid-SSH are obtained by scaling by p_{base}^{-1} the adjusted weights computed for the non-hybrid variant that is applied to the p_{base} -sampled stream.

4.3 Adjusted weights for SNF

We define $d_{i,j}(\mathbf{p}, \mathbf{n})$ ($2 \leq i \leq j \leq r$) as follows.

$$\begin{aligned} 2 \leq j \leq r : & \quad d_{2,j} = \left(\frac{p_1 - p_j}{p_1} \right)^{n_1} \prod_{h=1}^r p_h^{n_h} \\ 3 \leq i \leq j \leq r : & \quad d_{i,j} = \left(\frac{p_{i-1} - p_j}{p_{i-1}} \right)^{n_{i-1}} (d_{i-1,j} - d_{i-1,i-1}) \end{aligned}$$

For $2 \leq i \leq j \leq r$, $d_{i,j}(\mathbf{p}, \mathbf{n})$ is the probability that all packets of the flow \mathbf{n} have rank values below the sampling rate at packet arrival time, that the flow is fully counted by SNF until the transition into step i , and that at the beginning of step i , the rank of the flow is at least p_j .

The probability that all packets are counted by SNF is equal to $\prod_{h=1}^r p_h^{n_h}$ minus the probability that the counting halts at the transition into steps $2, \dots, r$:

$$q^{\text{SNF}}[\mathbf{n}|\mathbf{n}] = \prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^r d_{j,j}. \quad (4)$$

Theorem 4.10.

$$A^{\text{SNF}}[\mathbf{n}] = \frac{\sum_{j=1}^r \frac{n_j}{p_j} (\prod_{h=1}^r p_h^{n_h} - \sum_{\ell=2}^j d_{\ell,\ell})}{\prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^r d_{j,j}}$$

The proof of the Theorem is provided in Section 9.3.

The partial sums $\sum_{h=1}^i d_{h,h}$ for $(i = 1, \dots, r)$ can be computed in $O(r^2)$ operations, and therefore, using Theorem 4.10 the adjusted weight $A^{\text{SNF}}[\mathbf{n}]$ can be computed in $O(r^2)$ operations.

Theorem 4.4 is an immediate corollary of the above (for SNF) and Lemma 4.9 (for SSH), using Lemma 4.5. (According to Lemma 4.5 when we compute $q[\mathbf{n}|\mathbf{n}]$ and the $c_{i,j}$'s or $d_{i,j}$'s we can remove from \mathbf{n} and \mathbf{p} all entries which are 0 in \mathbf{n} except for the last entry in \mathbf{n} which remains even if it is 0.)

4.4 Relation between the sketching algorithms

The rank-based view shows that the distribution over subsets of flows included in the sketch is the same for ANF, ASH, SNF, SSH. The different algorithms applied with the same rank assignment result in the same set of k cached flows (or all flows if there are fewer than k distinct flows in the packet stream.). The hybrid algorithms result in “almost” the same distribution: if the p_{base} -sampled packet stream contains fewer than k distinct flows then the sketch will only include those flows, but the included flows are a subset of the flows included in a sketch generated by the non-hybrid algorithms using the same underlying assignment.

This distribution is equivalent to weighted sampling without replacement of k flows (WS) (see, e.g. [6, 7]). WS is performed as follows over the set of aggregated flows: repeatedly, k times, a flow is selected from the set of unsampled flows with probability proportional to its weight. Adjusted weights for WS can be obtained using the rank conditioning or subset conditioning methods [4, 5]. These weights are computed using the *exact packet count* of each flow and therefore can not be obtained by a stream algorithm with size- k flow cache. We include these WS sketches in our evaluation in order to understand to what extent performance deviates in comparison. The rank-conditioning adjusted weight assigned to each flow is equal to its weight (the number of packets) divided by the probability that the flow is included in the sample² in some probability subspace that includes the current sample. The probability subspace is defined as all runs that have the same effective sampling rate p' and therefore the probability is equal to $1 - (1 - p')^{|f|}$, where $|f|$ is the number of packets in the flow and p' is the effective sampling rate. Therefore, the adjusted weight is equal to $|f|/(1 - (1 - p')^{|f|})$.

Since these algorithms (see Table 1) share the same distribution, the difference in estimate accuracy stems from the adjusted weight assignment. The quality of the assignment depends on the information the algorithm gathers and the method we apply to derive the adjusted weights. When the adjusted weights have smaller variance, the estimates we obtain are more accurate. We explore the relation of estimate quality between the different estimators/algorithms.

²This is the Horvitz-Thompson unbiased estimator obtained by dividing the weight of the item by the probability that it is sampled.

| alg | sketch size | active counters | counts collected |
|-----|-------------|-----------------|------------------------|
| ANF | k | k | for k flows |
| ASH | k | k | for k flows |
| SNF | k | k | per-step for k flows |
| SSH | k | k | per-step for k flows |
| WS | k | $ F $ | for all flows in F |

Table 1: Methods that obtain a sketch of size k that is a weighted sample without replacement from the set of flows F .

4.5 Covariances

An important property of the algorithms we consider is *zero covariances*:

Lemma 4.11. *Consider $L \in \{\text{ASH}, \text{ANF}, \text{SSH}, \text{SNF}, \text{NF}, \text{SH}\}$ and two flows $f_1 \neq f_2$. Then $\text{COV}(A^L(f_1), A^L(f_2)) = 0$.*

(The proof is provided in Section 9.4.)

The zero covariance property is trivial for fixed-rate sampling (NF or SH), since each flow is selected independently. With the adaptive algorithms, however, the adjusted weights are not independent since inclusion of one flow makes it less likely for the other flow to be included. For WS, this property is established in [4].

The zero covariance property of the random variables $A^L(f)$ ($f \in F$) implies that

Corollary 4.12. *For any $J \subset F$ and $L \in \{\text{ASH}, \text{ANF}, \text{SSH}\}$,*

$$\text{VAR}(A^L(J)) = \text{VAR}\left(\sum_{f \in J} A^L(f)\right) = \sum_{f \in J} \text{VAR}(A^L(f)).$$

Therefore, to show that an adjusted weight assignment has lower variance than another on *all* subpopulations, it suffices to show lower variance on all individual flows.

4.6 Variance of adjusted weights

An algorithm dominates another, in terms of the information it collects on each sketched flow, if we can use its output to emulate an output of the second algorithm. It is not hard to see that SNF dominates NF, that SSH dominates both ASH and SNF (and therefore also dominates NF), that SNF and ASH are incomparable, and that they are all dominated by WS. Therefore, SSH sketches are the most powerful and ANF sketches are the least powerful. The variance of the adjusted weight assignments reflects this dominance relation, with lower variance for the methods that gather more information.

Theorem 4.13. *For any packet stream and any flow f we have the following relation between the variance of the adjusted weight assignments for f .*

$$\text{VAR}(A^{\text{WS}}(f)) \leq \text{VAR}(A^{\text{SSH}}(f)) \leq \text{VAR}(A^{\text{ASH}}(f)) \leq \text{VAR}(A^{\text{ANF}}(f)) \quad (5)$$

$$\text{VAR}(A^{\text{SSH}}(f)) \leq \text{VAR}(A^{\text{SNF}}(f)) \leq \text{VAR}(A^{\text{ANF}}(f)) \quad (6)$$

The proof is provided in Sections 9.2 and 9.3. The relation also holds to the fixed-rate variants of the algorithms (when the sampling rate or rate steps are fixed).

A variance relation also holds for the hybrids: the variance is non-increasing with the packet-processing rate p_{base} .

4.7 Estimators with negative covariances

The adjusted weights assignments $A_{\mathbf{p}}^L(\mathbf{n})$ were a function of the observed counts of the flow and the “sampling rate” (or sampling rate steps). We consider estimators that utilize different information: the counts collected for other flows in the sketch and the total packet count of the stream.

The *selectivity* of a packet is $1/w(F)$, selectivity of a flow is $\rho(f) = w(f)/w(F)$, and the selectivity of a subpopulation J is $\rho(J) = w(J)/w(F)$. When the total weight $w(F)$ is known, subpopulation weight and selectivity queries are equivalent: We can obtain an estimator for subpopulation weight by multiplying the subpopulation selectivity estimator with $w(F)$ and vice versa, by dividing the adjusted weight estimator with $w(F)$.

Adjusted selectivities, $R^L()$ are unbiased estimators for selectivity. We derive adjusted selectivity estimators for the adaptive algorithms that are based on the observed counts of all flows but do not depend on $w(F)$. We derive adjusted weight assignments $A^{+L}()$ for the adaptive algorithms that depends on $w(F)$, and the observed counts of all flows.

We derive estimators for ANF and ASH sketches. (Estimators for SSH and SSH sketches can be obtained using the same methodology.) We state the per-packet adjusted weights, as this is a more general form (proofs are provided in Section 9.6).

- $A^{+ANF}(v, n : F, w(F)) = \frac{w(F)}{n(F)}$ if the packet v is counted and 0 otherwise.
- $R^{ANF}(v, n(F)) = \frac{1}{n(F)}$ if the packet v is counted and 0 otherwise.

For ASH,

- $A^{+ASH}(v, n : F, w(F)) = \frac{w(F) - n(F) + k}{k}$, if v is the first-counted packet of a flow.
- $A^{+ASH}(v, n : F, w(F)) = 1$, if v is counted but is not the first-counted packet of a flow.
- $A^{+ASH}(v, n : F, w(F)) = 0$, if v is not counted.

Adjusted weight assignments can be incomparable. We consider an assignment $A_1(f)$ to be at least as good as $A_2(f)$ if $A_1()$ has at most the variance of $A_2()$ on *any subpopulation*. A sufficient condition is that for all $f \in F$, $\text{VAR}(A_1(f)) \leq \text{VAR}(A_2(f))$ and for all $f_1 \neq f_2$,

$$\text{COV}(A_1(f_1), A_1(f_2)) \leq \text{COV}(A_2(f_1), A_2(f_2)) .$$

The adjusted weights assignments $A^{+L}(f)$ are such that

1. For each f , $\text{VAR}(A^{+L}(f)) \leq \text{VAR}(A^L(f))$,
2. For $f_1 \neq f_2$, $\text{COV}(A^{+L}(f_1), A^{+L}(f_2)) \leq 0$, and
3. $\sum_{f_1 \in F} \sum_{f_2 \in F} \text{COV}(A^{+L}(f_1), A^{+L}(f_2)) = 0$. ($A^{+L}(F) \equiv w(F)$).

These properties imply that $A^{+L}(f)$ is better than $A^L(f)$ (at least as good on any subpopulation and better on some subpopulations). A case for these properties was made in [4, 5] for aggregated data, where there is a similar use of the total weight of the dataset to derive tighter estimators with negative covariances between subsets and zero sum of variances. These properties are also motivated by an interesting relation that shows that the variance of an “average” subset is a linear combination of the variance of the sum and the sum of variances [21].

These properties can not be obtained for fixed-rate sampling algorithms such as NF and SH: Since there is positive probability of an empty sample, we can not have (unbiased) adjusted weights such that $A^{+L}(F) \equiv w(F)$. The relation $A^{+L}(F) \equiv w(F)$ is immediate from the definitions and the variance relations are established in Section 9.6.

5 FSD estimators

A flow attribute that is lost when there is no pre-aggregation is the size of the flow (exact number of packets or bytes). This could be an important attribute for some aggregations, for example, to trace the origin of port scanning or worm activity we may want to aggregate over all flows that originate from a certain AS and have at most 10 packets. We also may want to estimate the number of flows in a subpopulation that are within a certain range of sizes.

FSD estimation is facilitated by assigning *adjusted FSD estimators* $\alpha_i^L()$ ($i \geq 1$) such that for any flow f , $E(\alpha_{|f|}^L(f)) = 1$ and for $i \neq |f|$, $E(\alpha_i^L(f)) = 0$. Similarly to adjusted weights, $\alpha_i^L() = 0$ for flows that are not included in the sketch. An unbiased estimator of the number of flows of size i in a subpopulation J is $\sum_{f \in J} \alpha_i^L(f)$. By summing the estimators $\alpha_i^L()$ over a desired range of values $i \in R$, we obtain unbiased estimates for the number of flows in this range in the subpopulation.

An important special case is the total number of flows in a subpopulation, which can be estimated using *adjusted counts* $\#^L()$ (for any flow f , $E(\#^L(f)) = 1$). We can obtain adjusted counts from adjusted FSD estimators using $\#^L(\mathbf{p}, \mathbf{n}) = \sum_{i \geq 1} \alpha_i^L(\mathbf{p}, \mathbf{n})$.

$L \in \{\text{NF}, \text{ANF}\}$ Let O_i be the random variable representing the number of flows of count i with NF (or ANF). Let p be the sampling rate (or effective sampling rate), and let C_i be the number of flows of weight i .

For $i \geq 1$, the expectation of O_i is $p^i C_i + i(1-p)p^{i+1} C_{i+1} + \dots = \sum_{j \geq i} \binom{j}{i} p^i (1-p)^{j-i} C_j$. Therefore, the inverse of the matrix with entries $\binom{j}{i} p^i (1-p)^{j-i}$, expresses each C_j as a linear combination of $E(O_i)$'s, and provides unbiased estimators [15, 9]. The entries of this inverted matrix are the FSD estimators $\alpha_j^L(p, i)$. The resulting estimators, however, are not well behaved [9]. Better estimators that use the TCP syn flag were proposed in [9].

$L \in \{\text{SH}, \text{ASH}\}$ We use a similar derivation for adjusted FSD estimators for SH and ASH. These estimators can assume negative values, but are well behaved. Let O_i be the random variable representing the number of flows of count i , p is the sampling rate, and C_i be the number of flows of weight i .

Lemma 5.1. $\hat{C}_i = O_i/p - O_{i+1}(1-p)/p$ is an unbiased estimate of the number of flows of size i .

Proof. The expectation of O_i is $C_i p + C_{i+1}(1-p)p + C_{i+2}(1-p)^2 p + \dots = \sum_{j \geq i} C_j (1-p)^{j-i} p$. Therefore, the expectation of $O_i/p - O_{i+1}(1-p)/p$ is C_i . \square

The respective nonzero coefficients are $\alpha_j^L(p, j) = 1/p$ and $\alpha_j^L(p, j+1) = -(1-p)/p$ for $j \geq 1$.

The estimator for the total number of flows is $O_1/p + \sum_{i > 1} O_i$, which corresponds to the adjusted counts $\#\text{ASH}(p, 1) = 1/p$ and $\#\text{ASH}(p, i) = 1$ for $(i > 1)$.

For ASH (proof provided in Section 9.6) we obtain tighter estimators. Let $n(f)$ be the number of counted packets for flow f .

$$\hat{C}_i = \left(1 + \frac{w(F) - n(F)}{k}\right) O_i - \frac{w(F) - n(F)}{k} O_{i+1}.$$

The respective nonzero FSD estimators are $\alpha_j^{\text{ASH}}(f, n : F, w(F)) = \frac{w(F) - n(F) + k}{k}$ if $n(f) = j$ and $\alpha_j^{\text{ASH}}(f, n : F, w(F)) = -\frac{w(F) - n(F)}{k}$ if $n(f) = j + 1$. The resulting estimator for adjusted count is $\#\text{ASH}(f) = \frac{w(F) - n(F) + k}{k}$ if $n(f) = 1$ and 1 otherwise.

SSH The SSH adjusted FSD estimators are more accurate (have smaller variance) than those obtained using SH and ASH sketches. Let $O_{\mathbf{s}}$ be the random variable representing the number of flows with observed count vector \mathbf{s} .

Lemma 5.2. Let \mathbf{n} be the observed count and $\mathbf{p} = (p_1, \dots, p_r)$ the corresponding sampling rate steps. Assume WLOG that $n_\ell > 0$ for $1 \leq \ell < r$. The following are correct adjusted FSD estimators $\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n})$ ($i \geq 1$). We state only the nonzero values.

- If $|\mathbf{n}| = 1$, $\alpha_1^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 1/p_r$. Otherwise,
- If $n_r = |\mathbf{n}| > 1$ (there is only one step and it is the last one) then $\alpha_{|\mathbf{n}|}^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 1/p_r$ and $\alpha_{|\mathbf{n}|-1}^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = -(1 - p_r)/p_r$.
- For $\ell = 3, \dots, r-1$, for $\ell = 2$ if $n_1 > 1$, and for $\ell = r$ if $n_r > 0$:

$$\alpha_{\sum_{h=\ell}^r n_h}^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{-c_{\ell, \ell}}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]}.$$

- If $n_1 = 1$, $\alpha_{\sum_{\ell=2}^r n_h}^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{-(1-p_2)}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]}$.
- If $n_1 > 1$, $\alpha_{-1+\sum_{h=1}^r n_h}^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{-(1-p_1)}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]}$
- $\alpha_{\sum_{h=1}^r n_h}^{\text{SSH}} = \frac{1}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]}$

The proof is deferred to Section 9.5.

Using the relation $\#^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \sum_{i \geq 1} \alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n})$ and Eq.(3) we obtain the following adjusted counts:

Corollary 5.3. • If $|\mathbf{n}| = 1$, $\#^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 1/p_r$.

- If $n_r \geq 1$ and $|\mathbf{n}| > 1$, $\#^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 1$.
- If $n_r = 0$, $|\mathbf{n}| > 1$, let $\ell < r$ be the last step with a positive packet count. $\#^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 1 + \frac{c_{\ell+1, r}}{q[\mathbf{n}|\mathbf{n}]}$.
(If all steps but the last have nonzero counts, then $\#^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 1 + c_{r, r}/q[\mathbf{n}|\mathbf{n}]$.)

Expressions for adjusted counts can be derived directly through the methods we developed for adjusted weight derivation (see Section 9). As is the case for adjusted weights, there is a unique adjusted counts function $\#^L(\mathbf{p}, \mathbf{n})$. For $L \in \{\text{NF}, \text{ANF}\}$, using the unbiasedness constraints, we obtain that the adjusted counts are the solution of the triangular system of linear equations: For a flow with count $|f|$,

$$\sum_{i=0}^{|f|} \binom{|f|}{i} p^i (1-p)^{|f|-i} \#^L(p, i) = 1.$$

Adjusted FSD estimators are computed based on the step-counts and therefore to facilitate FSD estimation, they need to be computed before the step-counts are discarded.

Hybrid algorithms We use as components the existing FSD estimators for NF sketches [15, 9] and ASH and SSH sketches.

For hybrid-ASH (or hybrid-SSH) we apply the respective (ASH or SSH) estimators to the counts obtained on the p_{base} -sampled stream. We obtain unbiased estimates $\alpha_j(J)$ ($j \geq 1$) for the number of flows of size j in the subpopulation J in the p_{base} -sampled stream. We then “treat” these $\alpha_j(J)$ values as observed counts with NF with sampling rate p_{base} , and “plug them” in the NF FSD estimator. Unfortunately, since the NF estimators are ill behaved for low sampling rates, these estimators are ill behaved for low values of p_{base} .

Lemma 5.4. *The resulting estimates \hat{C}'_i ($i \geq 1$) are unbiased estimates of the number of flows of size i in the original stream.*

Proof. The NF FSD estimators are derived by expressing O_j ($j \geq 1$), the expected number of flows with a certain observed counts as a linear combination of C_i ($i \geq 1$), the number of flows of size i . The matrix is then inverted, and we express each C_i as a linear combination of the expectations of the observed counts. The estimators \hat{C}'_i are linear combinations of the observed counts. Since this is a linear combination, we can replace the observed counts with any other random variables with the same expectation and still obtain unbiased estimators \hat{C}'_i . \square

6 Estimating other aggregates

The sketches support estimators for aggregates of other numeric flow properties over a queried subpopulation. We distinguish between *flow-level* and *packet-level* properties.

6.1 Flow-level properties

We classify a numeric property $h(f)$ of the flow f as flow level if it can be extracted from any packet of the flow and some external data (therefore, we know $h(f)$ for all the flows that are included in the sketch). Examples are the number of hops to the destination AS, unity (flow count), and flow identifiers (source or destination IP address and port, protocol). Flow-level properties can be aggregated *per-packet* or *per-flow*.

Per-packet aggregation. For a subpopulation $J \subset F$, the *per-packet sum* of $h()$ over J is $\sum_{f \in J} w(f)h(f)$.

The *per-packet average* is $\frac{\sum_{f \in J} w(f)h(f)}{\sum_{f \in J} w(f)}$. If $h(f)$ is the number of AS hops traveled by the flow f then the per-packet sum is the total number of AS hops traveled by packets in the subpopulation J and the per-packet average is the average number of hops traveled by a packet in J . If $h(f)$ is unity, the per-packet sum is the weight of the subpopulation. It is not hard to see that for a sketch with unbiased adjusted weights, $\sum_{f \in J} A(f)h(f)$ is an unbiased estimator of the per-packet sum of $h()$ over J . (A (possibly biased) estimator for the per-packet average is $\frac{\sum_{f \in J} A(f)h(f)}{\sum_{f \in J} A(f)}$.)

Per-flow aggregation. The *per-flow sum* of $h()$ over J is $\sum_{f \in J} h(f)$. The *per-flow average* of $h()$ over J is $\sum_{f \in J} h(f)/|J|$. If $h(f) \equiv 1$, the per-flow sum is the number of distinct flows in a subpopulation. If $h(f)$ is the number of AS hops then the per-flow average is the average “length” of a flow in J .

Our generic estimator for per-flow sums is based on adjusted counts. For each $f \in F$, $E(\#(f)h(f)) = h(f)$, therefore

$$\sum_{f \in J} \#(f)h(f) = \sum_{f \in J \cap \mathbf{S}} \#(f)h(f)$$

is an unbiased estimator of the per-flow sum of $h()$ over J .

6.2 Packet-level properties

Packet-level properties have numeric $h()$ -values that are associated with each packet. For a flow f we define the $h()$ -value of f as $h(f) = \sum_{c \in f} h(c)$. If $h(c)$ is the number of bytes in the packet c then $h(f)$ is the number of bytes in the flow. If $h(c)$ is unity, then $h(f) = w(f)$ is the number of packets of the flow. $h(f)$ is available only if all packets of f are processed and therefore is not provided for flows included in NF and SH variants sketches.

The algorithms are adapted to collect information needed to facilitate unbiased estimators. For any desired packet-level property $h()$, we produce *adjusted $h()$ -values* $H_{\mathbf{P}}^h(f)$. For any f , $H(f)$ is an unbiased estimator of $h(f)$ and $H(f) = 0$ for flows that are not included in the sketch. For any subpopulation J , $\sum_{f \in J} H(f) = \sum_{f \in J \cap \mathbf{S}} H(f)$ is an unbiased estimator for $\sum_{f \in J} h(f)$.

- $L \in \{\text{NF}, \text{ANF}\}$: Let $N(f)$ be the set of counted packets and let $n(f)$ be the packet count maintained by L for a cached flow f . To facilitate $h()$ -values estimation, the algorithm maintains the $h()$ counts $n^{(h)}(f) = h(N)$. The adjusted $h()$ -value is $H_p^L(f) = n^{(h)}(f)/p$. A subtle point is correctly updating $n^{(h)}$ for ANF when performing rate adaptation. The updated value should include the sum of the $h()$ values of resampled packets which if strictly done, requires us to store the $h()$ value of all packets in $N(f)$, which is prohibitive. We argue, however, that it is sufficient to store the total $n^{(h)}(f)$ and update it proportionally to the reduction in the packet count $n(f)$. That is, if resampling reduces the packet count to $n'(f)$, we update the $h()$ -count $n^{(h)}(f) \leftarrow n^{(h)}(f)n'(f)/n(f)$.

The updated $n^{(h)}(f)$ is the expectation of the updated $h()$ count that would have been obtained if we explicitly maintained N and per-packet $h()$ values and sampled from N a subset of size $n'(f)$. (all subsets of $N(f)$ that are of the same size have the same probability of being in the resample, regardless of the packet position or its $h()$ value.)

This consideration extends to a sequence of rate adaptations: The final $n^{(h)}(f)$ has the expectation of the $h()$ count over all resamples that resulted in the same sequence of packet count reductions. Interestingly, done this way, we obtain a lower variance estimator than if we had maintained and used per-packet $h()$ values for $N(f)$, as the $h()$ count we use is the expectation of the latter in each part of a partition of the sample space.

- $L \in \{\text{SH}, \text{ASH}\}$: we obtain the unbiased estimator $H^{\text{ASH}}(f) = h(c_0(f))\frac{1-p}{p} + h(N(f))$, where $N(f)$ is the set of counted packets of f and $c_0(f) \in N(f)$ is the first counted packet of f . To facilitate this estimator, the algorithm needs to record the $h()$ value of the first packet and the sum of $h()$ values of all subsequent packets. For ASH, the resampling makes a direct implementation infeasible, as we need to record per-packet $h()$ values for all packets in N . We can not use averaging as we did for ANF, since later packets are more likely to be counted than earlier packets. Fortunately, there is efficient implementation for SSH.
- $L \in \{\text{SNF}, \text{SSH}\}$: To facilitate estimators, the algorithms need to maintain per-step $h()$ counts,

$$h(N_1), h(N_2), \dots, h(N_r) ,$$

where N_i is the set of packets counted in step i . For SSH, we also need to maintain the $h()$ value of the first packet c_0 when the flow enters the cache. The expressions for adjusted $h()$ values are provided in Eq. (9) (for SSH) and Eq. (24) (for SNF).

- Hybrid SH variants: We obtain adjusted $h()$ values for the sampled stream and scale them by p_{base}^{-1} .

6.3 Sketches for byte counts

Byte counts can be estimated using a sketch built to estimate packet counts, but if byte counts are the main application, then SH variants can be adapted to estimate bytes directly instead of packets: The count values are applied to bytes and are captured as follows. If a packet belongs to a cached flow, the number of bytes is added to the active counter. Otherwise, we use the geometric distribution to determine what part of a packet (if at all) should be counted. For a continuous variant of this process we can use the exponential distribution.

7 Implementation design using discretized sampling rates

We design an alternative to the pure models that addresses important practical implementation issues. The first is the number of rate adaptations performed. The second is the implementation of each rate adaptation, namely, the tracking of flow ranks that determines which flows are evicted and how counts are adjusted. We

then establish that the discretized version preserves important properties of the pure model that allow for unbiased estimation and for other properties of the variance of the adjusted weights to carry over.

Our discretized algorithm uses three tunable parameters that can be set by the router manufacturer. The first is $p_{\text{base}} \leq 1$ that determines the fraction of packets that are processed. The second is $p_{\text{start}} \leq p_{\text{base}}$ that determines the initial sampling rate. The third parameter is $0 < \mu < 1$ which controls the discretization of the sampling rates.

The number of rate-adaptations is a performance factor for all adaptive algorithms. Executing each adaptation is an intensive operation and therefore it is desirable to both limit the number of rate-adaptations and to carefully implement them [11]. For the step-counting algorithms, the number of rate adaptations also affects the size of intermediate storage and the computation of the adjusted weights (which depend on the number of rate-steps in which an actively-counted flow had a nonzero count).

The pure models perform a rate adaptation to evict a single cached flow at a time. It follows from the rank-based view that all adaptive algorithms (ANF, ASH, SNF, and SSH) perform the same rate adaptations. Rate adaptation occurs when the current sampling rate (the $(k+1)$ -st-smallest rank value of a flow) decreases. This happens when the cache is full (has k flows) and a sampled packet (equivalently, packet with rank value which is smaller than the current sampling rate) does not belong to a cached flow. The number of rate adaptations depends logarithmically on the size of the stream, but linearly on the size of the flow cache:

Lemma 7.1. *Let m be the size of the packet stream. The expected number of rate adaptations is $\leq (k+1) \ln(p_{\text{start}}m)$.*

Proof. The expected number of updates to the set of cached flows for aggregated data streams is analyzed in [4, 7]. This corresponds to a situation where each flow appears as a consecutive burst and is therefore a special case of our model. The argument used in [7] for uniform weights (a stream of 1-packet flows) trivially extends to a model where there is a stream of 1-packet flows where rank values of cached flows (flows included in the current sketch) can be arbitrarily decreased at arbitrary times.

We now express an unaggregated stream of multi-packet flows in this model: Packets are processed as if they belong to a stream of 1-packet flows. Once a flow is cached, we examine the rank value of subsequent packets that belong to the flow and then delete these packets. If the rank of the deleted packet is smaller than the current rank of its flow, we simulate an arbitrary rank decrease and decrease the rank of the flow to that of the packet. An important point for the analysis is that packet deletion is independent of the rank of the deleted packet. The probability that the i th undeleted packet modifies the sketch (has rank value that is smaller than the $(k+1)$ st smallest rank) is at most $\min\{1, (k+1)/i\}$. The total number of undeleted packets is at most m . Therefore, the expected number of updates is at most $\sum_{i=1}^m \min\{1, (k+1)/i\} \leq (k+1) \ln m$. \square

This bound is nearly tight for streams that consist of 1-packet flows [7], and it is $\Omega(k \ln(p_{\text{start}}m))$ (asymptotically tight) when at least a constant fraction of packets belong to small flows. Large number of small flows is common in Zipf-like data and small flows are often introduced in DDoS attacks, port or IP address scanning, and other anomalies.

The actions of a run of a discretized variant of ANF, ASH, SNF, SSH, and hybrids, are equivalent to those of the original variant except that when the flow cache overflows (there are more than k flows with current rank value that is below the current sampling rate), the sampling rate is decreased by a factor of μ until at least one flow (but in expectation at most $(1-\mu)$ fraction of the flows) are evicted. Observe that the discretized implementation does not always produce sketches with k flows: The number of flows is at most k but can be smaller even if there are k or more distinct flows in the p_{base} -sampled stream.

We provide a discretized rank-based view of the discretized sketching algorithms. This view is equivalent to replacing sampling rates and rank values of packets and flows $x \in [0, 1]$ with $\lceil \log_{\mu}(x/p_{\text{base}}) \rceil^3$. Packets of the p_{base} -sampled stream are assigned discretized ranks using a geometric distribution with parameter $(1-\mu)$. The *discretized rank of a flow* is the largest rank of a packet of the flow. The *discretized*

³Smaller values have larger discretized values.

current sampling rate is initially set to $\lceil \log_\mu(p_{\text{start}}/p_{\text{base}}) \rceil$. After k distinct flows are cached, it is the $(k + 1)$ st largest discretized rank of a flow, which is equal to the largest discretized rank of a cached flow plus 1. The *discretized effective sampling rate* is the discretized sampling rate at the end of the measurement period. The flow counts collected over the p_{base} -sampled stream correspond to those for the pure model: A packet is included in the current discretized ANF count if and only if its discretized rank value is above the discretized current sampling rate; With discretized ASH, the packet is included if and only if the discretized rank of the flow after the packet is processed is above the discretized current sampling rate, and similarly for SNF and SSH.

The following property, that holds for the pure model, extends to the discretized model:

Lemma 7.2. *A flow f is cached in the discretized model if and only if its discretized rank is larger than the k th largest discretized rank of the flows in $F \setminus \{f\}$. (If there are fewer than k flows in $F \setminus \{f\}$ with positive discretized rank value, the flow is cached if and only if its rank is positive.)*

This property is critical for extending the analysis of unbiased estimators and variance relations to the discretized model. It allows us to simply “plug in” sampling rates and the respective flow counts into the unbiased estimators developed for the pure model. The subtle arguments do not carry over to other conceivable implementations of rate adaptations such as removing a constant fraction of (highest-ranked) cached flows [16, 11] without having to maintain additional state.

A side benefit of discretization is that fewer bits are needed to encode rank values of active flows, as the expected maximum discretized rank of a flow is $\log_\mu(mp_{\text{base}})$. Other advantages of such discretization, such as layered transmission of summaries, are provided in [11].

The number of rate adaptations is tuned using the parameters μ and p_{start} . Larger values of μ correspond to a higher number of rate adaptations but also to better memory utilization and more flows in the final sketch (the expected number of active counters and number of flows in the final sketch is about $k(1 + \mu)/2$). Lower values of p_{start} eliminate up to $\lceil \log_\mu(p_{\text{start}}/p_{\text{base}}) \rceil$ rate adaptations. The counts of the step-counting algorithms are reduced by lower p_{start} , but if p_{start} is larger than the effective sampling rate, then the sketch produced by ANF and ASH are not affected by the lower p_{start} . The discretized implementation has a considerably better bound on the number of rate adaptations than the pure model. In particular, the linear dependence on k exhibited by the bound obtained for the pure version, is eliminated.

Lemma 7.3. *The expected number of rate adaptations performed is at most $\log_\mu(p_{\text{start}} * m)$.*

We outline an implementation of the discretized algorithms. The execution is divided into *counting phases* and *rate adaptations* (a design of [11] allows them to run concurrently.)

Counting phase. Each counting phase starts with a set of statistics counters indexed by the flow attributes of cached flows and applied to the p_{base} -sampled packet stream. Each packet is labeled as “sampled” (again) with probability μ^t , where t is the current discretized sampling rate. We perform the following⁴: (i) If the flow is cached then: If the algorithm is one of ASH, SSH, and hybrids **or** the packet is labeled sampled and the algorithm is one of ANF and SNF, we increment the respective flow counter. (ii) If the packet is labeled sampled **and** the flow is not cached, we create a new entry in the flow cache. If there are k cached flows, we perform a *rate adaptation*.

Rate adaptation. The adaptive algorithms, ANF, ASH, and hybrid-ASH maintain a single packet count for each cached flow. These counts are updated during the rate adaptation until at least one flow has a zero count. For ANF, the updated count is a binomial random variable with parameters μ and the current count; for ASH, with probability μ the count remains unchanged and otherwise it is a geometric random variable with parameters that are the current count minus 1 and rate μ^{t+1} , where t is the current discretized sampling rate. After the counts of all cached flows are updated, the current discretized sampling rate is incremented.

⁴With SH variants and hybrids we need to sample the packet only if the flow is not cached.

The update process is repeated until at least one flow has a count of zero.⁵ All flows with a count of zero are then evicted from the cache.

The step-counting algorithms SNF, SSH, and hybrid-SSH store a discretized rank for each cached flow. The flow ranks are updated using the counts collected in the most recent counting phase. The set of cached flows with smallest discretized rank value are then evicted. The discretized sampling rate is updated to be the discretized rank of evicted flows.

The update process of the rank of cached flows emulates the following process that assigns ranks individually to packets counted in the recent counting step. The first packet counted (SH and hybrid variants with flow that was not cached at the beginning of the phase) and all packets counted (NF variants) obtain a random rank from a geometric distribution with parameter $(1 - \mu)$, conditioned on it being larger than the current discretized sampling rate. The rank of each flow is updated to be the maximum of its current rank and the ranks assigned to the packets of the flow counted at the recent counting phase.

These updates can be performed efficiently (computation steps proportional to the number of cached flows with non-empty counts at the recent step) using the exponential distribution to find the maximum discretized rank over a set of packets.

7.1 Adjusted weights for discretized algorithms

Unbiased adjusted weights for the discretized algorithms are obtained by recording the discretized sampling rates for each step. We then convert each discretized rate t to a corresponding sampling rate $(p_{\text{start}}/p_{\text{base}})\mu^t$ and plug them in the corresponding expressions for ANF, ASH, SNF, or SSH adjusted weights. For the hybrid versions ($p_{\text{base}} < 1$), we scale the adjusted weights by p_{base}^{-1} .

The arguments for correctness, that are based on obtaining an unbiased estimator on each part in a partition of the same space, extend to the discretized version using Lemma 7.2. If the ranks of all packets in $F \setminus \{f\}$ are fixed, the discretized sampling rate when f is counted depends only on these fixed ranks (and not on ranks assigned to previous packets of f) and is equal to the sampling rate at measurement time. Therefore, we can compute unbiased adjusted weights while treating the effective sampling rate (for ASH and ANF) or the steps of the current sampling rate (for SNF and SSH) as being fixed.

The proofs of other properties, such as the relation between the algorithms (Section 4.4), the variance relation (Theorem 4.13), and the zero covariances (Lemma 4.11) extend to the discretized variants.

8 Simulations

We use simulations in order to understand several performance parameters: The accuracy of the estimates derived and its dependence on the algorithms, parameter settings, and the consistency of the subpopulation, the tradeoffs of the hybrid approach, and the effectiveness of the parameters controlling the number of rate adaptations. The simulations were performed using the discretized variants of the algorithms, with parameters k (maximum number of counters), $0 < p_{\text{base}} \leq 1$ that determines the fraction of processed packets (hybrid approach), and μ and p_{start} , that control a tradeoff between accuracy/utilization and the number of rate adaptations.

8.1 Data

We used both synthetic and IP flows datasets. The IP flows data was collected using *unsampled* NetFlow (flow-level summary of each 10 minute time period that includes a complete packet count for each flow) deployed at a gateway router. A typical period has about 5000 distinct flows and 100K packets. The synthetic datasets were produced using Pareto distributions with parameters $\alpha = 1.1$ and $\alpha = 1.5$. We

⁵We can avoid the repeated process by storing discretized rank value for each flow as proposed next for the step-counting algorithms.

generated distributions of flow sizes by drawing 5000 flow sizes. We simulated a packet stream from each distribution of flow sizes by randomly permuting the packets of all flows.

The cumulative distributions of the weight of the top i flow sizes for each distribution is provided in Figure 2.




Figure 2: Fraction of packets belonging to the i heaviest flows.

The subsets (subpopulations) we considered for the synthetic datasets were the 2^i largest flows and the 50%, 30%, and 10% smallest flows. This selection enables us to understand how performance depends on the consistency of the subpopulation (many smaller flows or fewer larger flows) and the skew of the data. The subpopulations we used for the IP flows (gateway) data were a partition of the flows according to destination port.

8.2 Quality of sketches

We compare the accuracy of subpopulation-size estimates obtained using ANF, SNF, ASH, and SSH. We evaluate performance as a function of the size k of the sketch (and the size of the flow cache). We also include weighted sampling without replacement.

Results that show the average absolute value of the relative error as a function of the cache size k are provided in Figures 3 and 4. The averaging for each data point was performed over 200 runs. The figures reflect the relation established theoretically: SSH dominating ASH which in turn dominates ANF and that SNF dominates ANF. They also show that ASH dominates SNF on our data. For subpopulations consisting of large flows, such as top- i flows and applications with medium to large flows, the performance gaps are more pronounced. This is because on these flows, the more dominant methods count more packets and obtain adjusted weights with smaller variance. Our adjusted weight assignments have minimum variance with respect to the information they use (the counts collected for the flow). This “optimality” enables us to translate the larger counts to smaller errors.

On subpopulations of very small flows, such as bottom-50% of flows or DNS (port 53) traffic (only the latter is shown), all methods have similar performance. In particular, there is no advantage for WS (the strongest method) over ANF (the weakest) on subpopulations consisting of 1-packet flows.

The results strongly support the use of step-counting as an alternative to the adaptive variants: On subpopulations consisting of many medium to large size flows, the relative error obtained using SNF and SSH is significantly smaller than what is obtained using ANF or ASH, respectively.

8.3 Evaluation of the hybrid algorithms

We decrease the parameter p_{base} while maintaining the same flow cache size $k = 400$. Figure 5 shows that the estimate quality gracefully degrades and a smooth performance curve is obtained. Even for $p_{\text{base}} = 0.1$, the hybrid-ASH and hybrid-SSH outperform ANF and SNF, respectively.




Figure 3: Estimating subpopulations consisting of top flows $p_{\text{base}} = 1$, $\mu = 0.9$, $p_{\text{start}} = 1$.




Figure 4: Estimating subpopulations of IP flows for various port numbers, $p_{\text{base}} = 1$, $\mu = 0.9$, $p_{\text{start}} = 1$.

Figure 6 shows the number of packets counted as a function of p_{base} as a fraction of the number of packets in the p_{base} -sampled stream and as a fraction of the total number of packets. The hybrids provide a desirable smooth tradeoff that provides high counting rate of processed packets, while enabling us to fully control the fraction of total packets that are processed.




Figure 5: Sweeping p_{base} to evaluate hybrid algorithm. Estimating subpopulations of top-64 flows in Pareto 1.1 and Pareto 1.4, $k = 400$, $p_{\text{start}} = p_{\text{base}}$, $\mu = 0.9$.

8.4 Controlling the number of rate adaptations

We sweep the parameter μ , which controls the rate of decrease of the sampling rate, and through it, the total number of rate adaptations performed. We expect the (absolute value of the) relative error of the estimates to increase when μ is decreased, as fewer packets are counted and reflected in the final sketch. On the other hand, the number of rate adaptations performed and the size of intermediate temporary storage needed to




Figure 6: Sweeping p_{base} to evaluate the tradeoffs in the hybrid algorithms. Fraction of packets in the p_{base} -sampled stream that are counted (left). Fraction of total packets that are counted (right). ($k = 400$, $p_{\text{start}} = p_{\text{base}}$, $\mu = 0.9$).

store the count vectors for sNF and sSH should decrease with μ . The effectiveness of the parameter μ and the feasibility of a router implementation depends on this tradeoff.

Figure 7 shows the dependence of the average absolute value of the relative error on the parameter μ . We can see that there is minimal performance loss in terms of estimate quality when μ is reduced from 0.9 to 0.5.

Figure 8 shows that selecting a smaller $\mu = 0.5$ is very effective. First, the number of rate adaptations is much smaller, and secondly, the size of intermediate temporary storage needed for collecting the count vectors, is much smaller with $\mu = 0.5$ than with $\mu = 0.9$.




Figure 7: Sweeping μ to evaluate how it affects the (average absolute value of the) relative error. The figure shows subpopulations of top-64 flows for pareto 1.1 and pareto 1.4 $p_{\text{start}} = p_{\text{base}} = 1$, $k = 400$.

9 Deferred proofs

The adjusted weights we assign are a function of the observed count of the flow and the sampling rate. The sampling rate (effective sampling rate or sampling rate steps) in the adaptive algorithms is treated as fixed because for any flow f , it is determined by the sampling performed on all “other” flows ($F \setminus \{f\}$). Therefore, within the probability subspace where the sampling on all other flows is fixed, the sampling rate is fixed. The adjusted weight of each flow is unbiased within each such subspace.

We deploy three different techniques to derive adjusted weights. These techniques are general tools applicable to other quantities such as adjusted counts, unbiased FSD estimators, and adjusted selectivities.

- *System of equations*: The unbiasedness constraints correspond to linear equations over the variables $A_{\mathbf{p}}^L(\mathbf{n})$. For a flow f , the expected adjusted weight over all possible observable counts $\mathbf{n} \preceq f$ of f

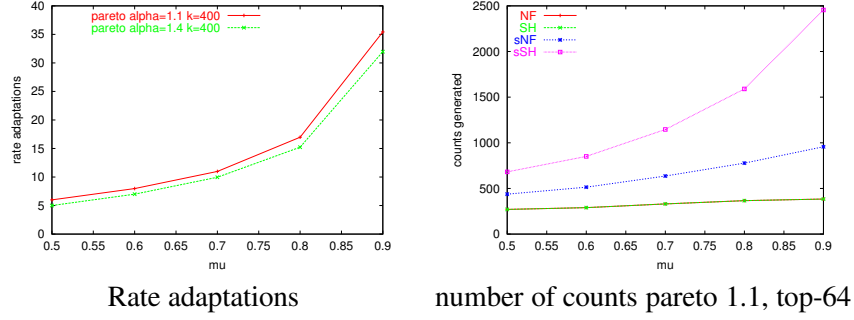


Figure 8: Sweeping μ to evaluate effectiveness decrease. Number of rate adaptations and number of counts $p_{\text{start}} = p_{\text{base}} = 1$, $k = 400$.

must be equal to $w(f)$: $w(f) = \sum_{\mathbf{n} \preceq f} q[\mathbf{n}|f] A_{\mathbf{p}}^L(\mathbf{n})$, where $q[\mathbf{n}|f]$ is the conditional probability that L obtains a count of \mathbf{n} for a flow f . The system of equations can be used to derive expressions for the adjusted weights, be solved numerically to compute adjusted weights for each instance, or establish properties of the solution such as uniqueness (A unique solution to this system implies that there is a unique deterministic assignment of adjusted weights that is a function of the observed counts of the flow and the sampling rate).

- *Dominance*: Algorithm A_1 dominates A_2 if A_1 counts are “more informative”, that is, each possible set of A_1 counts corresponds to a probability distribution over sets of A_2 counts such that applying A_1 and drawing an A_2 count from the corresponding distribution is equivalent to applying A_2 . For example, SSH dominates ASH and ASH dominates ANF. We can derive adjusted weights for A_1 from those of A_2 by taking the expectation of A_2 adjusted weights over the distribution that corresponds to the A_1 counts. If there is no closed form for A_1 adjusted weights, we can draw multiple times from the corresponding distribution and take an average of the A_2 adjusted weights. The resulting per-item adjusted weights of A_1 are unbiased and have equal or lower variance than A_2 adjusted weights (same arguments as for mimicked sketches [7]). Dominance implies that if an algorithm has unique deterministic adjusted weights, the per-item variance of these adjusted weights is at most that of any algorithm that it dominates.
- *Per-packet Horvitz-Thompson (HT) analysis*. The HT estimator is applicable when both the weight and the sampling probability of each item are provided. The weight is available for packets, but in our unaggregated setting, not for flows, and therefore we compute an adjusted weight for each packet. The adjusted weight of the flow is then the sum of the adjusted weights of packets of the flow. The second ingredient needed for HT estimators is the sampling probability, which can not be determined from the sketch. We apply HT with sample space partitioning [5] to “bypass” knowledge of the sampling probability. Per-packet analysis allows us to derive unbiased estimators for other packet-level properties.

9.1 Adjusted weights for SSH

We demonstrate all three methods by applying them to derive SSH adjusted weights (proof of Theorem 4.8). **System of equations:** Any valid observed count \mathbf{i} of a flow with \mathbf{n} packets has the form $(0, 0, i_j, n_{j+1}, \dots, n_r)$, where $1 \leq i_j \leq n_j$. We refer to vectors of this form as *suffixes* of the vector $\mathbf{n} = (n_1, \dots, n_r)$. We use the notation $\text{SUFF}_{\mathbf{n}}(j, i_j)$ for the suffix

$$(0, \dots, 0, i_j, n_{j+1}, \dots, n_r)$$

of \mathbf{n} (subscript is omitted when clear from context). We utilize the total order over the suffixes $\text{SUFF}_{\mathbf{n}}(i, i_j)$ defined by the lexicographic order on the vectors $\text{SUFF}_{\mathbf{n}}(i, i_j)$ with their coordinates reversed. That is,

$$\begin{aligned} \text{SUFF}(r, 1) = (0, \dots, 0, 1) < \text{SUFF}(r, 2) = (0, \dots, 0, 2) < \dots < \text{SUFF}(r, n_r) = (0, \dots, 0, n_r) < \text{SUFF}(r-1, 1) = (0, \dots, 0, 1, n_r) \\ < \dots < \text{SUFF}(1, n_1) = (n_1, \dots, n_r) \end{aligned}$$

A correct adjusted weight assignment must satisfy for any vector $\mathbf{n} = (n_1, \dots, n_r)$,

$$\sum_{\mathbf{s} \preceq \mathbf{n}} q[\mathbf{s}|\mathbf{n}] A^{\text{SSH}}(\mathbf{s}) = \sum_{j=1}^r n_j. \quad (7)$$

(For a flow with counts (n_1, \dots, n_r) , the expectation of the adjusted weight should be equal to the size $\sum_{j=1}^r n_j$ of the flow.)

To compute $A^{\text{SSH}}(\mathbf{n})$, we obtain a system of equations using an instance of Eq. (7) for each suffix $\mathbf{x} \preceq \mathbf{n}$ of the observed count. The system of equations includes the variables $A^{\text{SSH}}(\mathbf{x})$ for all $\mathbf{x} \preceq \mathbf{n}$ and the coefficients $q[\mathbf{s}|\mathbf{x}]$ for all $\mathbf{s} \preceq \mathbf{x}$. The matrix is triangular and can be solved by substitutions. Computation of adjusted weights that directly utilizes the system of equalities requires a number of operations is the product of the number of nonzero coordinates of \mathbf{n} and $(\sum_{h=1}^r n_h)^2$ (the square of the number of observed packets).⁶ This quadratic dependence in the number of packets makes the computation very intensive for large flows. We parametrically solve the equations to derive expressions for the adjusted weights.

We need the following lemma (proof is immediate from definitions).

Lemma 9.1. *Let $\text{SUFF}_{\mathbf{n}}(j, i_j) \preceq \text{SUFF}_{\mathbf{n}}(k, i_k) \preceq \mathbf{n}$. The following equality holds:*

$$q[\text{SUFF}_{\mathbf{n}}(j, i_j)|\mathbf{n}] = (1 - p_j)^{(n_k - i_k) + \sum_{h=1}^{k-1} n_h} q[\text{SUFF}_{\mathbf{n}}(j, i_j)|\text{SUFF}_{\mathbf{n}}(k, i_k)].$$

Proof. From Eq. (7) we have

$$q[\mathbf{n}|\mathbf{n}] A^{\text{SSH}}(\mathbf{n}) = \sum_{h=1}^r n_h - \sum_{\mathbf{s} \prec \mathbf{n}} q[\mathbf{s}|\mathbf{n}] A^{\text{SSH}}(\mathbf{s}).$$

We apply Lemma 9.1 to express the sum $\sum_{\mathbf{s} \prec \mathbf{n}} q[\mathbf{s}|\mathbf{n}] A^{\text{SSH}}(\mathbf{s})$ as a sum of sums of the form of the lhs of Eq. (7) for the vectors $\text{SUFF}_{\mathbf{n}}(1, n_1 - 1)$ and $\text{SUFF}_{\mathbf{n}}(k, n_k)$ (for $k = 2, \dots, r$). We then substitute the respective rhs.

$$\begin{aligned} \sum_{\mathbf{s} \prec \mathbf{n}} q[\mathbf{s}|\mathbf{n}] A^{\text{SSH}}(\mathbf{s}) &= \\ & (1 - p_1) \sum_{\substack{\mathbf{s} \preceq \text{SUFF}(1, n_1 - 1) \\ \text{SUFF}(2, n_2) < \mathbf{s}}} q[\mathbf{s}|\text{SUFF}(1, n_1 - 1)] A^{\text{SSH}}(\mathbf{s}) + (1 - p_2) \sum_{\substack{\mathbf{s} \preceq \text{SUFF}(2, n_2) \\ \text{SUFF}(3, n_3) < \mathbf{s}}} q[\mathbf{s}|\text{SUFF}(2, n_2)] A^{\text{SSH}}(\mathbf{s}) + \dots \\ & (1 - p_i) \sum_{\substack{\mathbf{s} \preceq \text{SUFF}(i, n_i) \\ \text{SUFF}(i+1, n_{i+1}) < \mathbf{s}}} q[\mathbf{s}|\text{SUFF}(i, n_i)] A^{\text{SSH}}(\mathbf{s}) + \dots + (1 - p_r) \sum_{\mathbf{s} \preceq \text{SUFF}(r, n_r)} q[\mathbf{s}|\text{SUFF}(r, n_r)] A^{\text{SSH}}(\mathbf{s}) = \\ &= c_{1,1} \sum_{\mathbf{s} \preceq \text{SUFF}(1, n_1 - 1)} q[\mathbf{s}|\text{SUFF}(1, n_1 - 1)] A^{\text{SSH}}(\mathbf{s}) + \sum_{h=2}^{r-1} (c_{1,h} - c_{1,1}) \sum_{\substack{\mathbf{s} \preceq \text{SUFF}(h, n_h) \\ \text{SUFF}(h+1, n_{h+1}) < \mathbf{s}}} q[\mathbf{s}|\text{SUFF}(1, n_1 - 1)] A^{\text{SSH}}(\mathbf{s}) + \\ & (c_{1,r} - c_{1,1}) \sum_{\mathbf{s} \preceq \text{SUFF}(r, n_r)} q[\mathbf{s}|\text{SUFF}(1, n_1 - 1)] A^{\text{SSH}}(\mathbf{s}) \\ &= c_{1,1} \left(\sum_{h=1}^r n_h - 1 \right) + \sum_{h=2}^{r-1} (c_{1,h} - c_{1,1}) (1 - p_h)^{n_1 - 1} \sum_{\substack{\mathbf{s} \preceq \text{SUFF}(h, n_h) \\ \text{SUFF}(h+1, n_{h+1}) < \mathbf{s}}} q[\mathbf{s}|\text{SUFF}(2, n_2)] A^{\text{SSH}}(\mathbf{s}) + \\ & (c_{1,r} - c_{1,1}) (1 - p_r)^{n_1 - 1} \sum_{\mathbf{s} \preceq \text{SUFF}(r, n_r)} q[\mathbf{s}|\text{SUFF}(2, n_2)] A^{\text{SSH}}(\mathbf{s}) \end{aligned}$$

⁶The time to solve the equations is proportional to $(\sum_{h=1}^r n_h)^2$, but we have $(\sum_{h=1}^r n_h)^2$ coefficients, and it takes time proportional to the number of nonzero coordinates of \mathbf{n} to compute each one.

$$\begin{aligned}
&= c_{1,1} \left(\sum_{h=1}^r n_h - 1 \right) + \sum_{h=2}^{r-1} c_{2,h} \sum_{\text{SUFF}(h+1, n_{h+1}) \prec \mathbf{s} \preceq \text{SUFF}(h, n_h)} q[\mathbf{s} | \text{SUFF}(2, n_2)] A^{\text{SSH}}(\mathbf{s}) + c_{2,r} \sum_{\mathbf{s} \preceq \text{SUFF}(r, n_r)} q[\mathbf{s} | \text{SUFF}(2, n_2)] A^{\text{SSH}}(\mathbf{s}) \\
&= c_{1,1} \left(\sum_{h=1}^r n_h - 1 \right) + c_{2,2} \sum_{\mathbf{s} \preceq \text{SUFF}(h, n_h)} q[\mathbf{s} | \text{SUFF}(2, n_2)] A^{\text{SSH}}(\mathbf{s}) + \\
&\quad \sum_{h=3}^{r-1} (c_{2,h} - c_{2,2}) \sum_{\text{SUFF}(h+1, n_{h+1}) \prec \mathbf{s} \preceq \text{SUFF}(h, n_h)} q[\mathbf{s} | \text{SUFF}(2, n_2)] A^{\text{SSH}}(\mathbf{s}) + (c_{2,r} - c_{2,2}) \sum_{\mathbf{s} \preceq \text{SUFF}(r, n_r)} q[\mathbf{s} | \text{SUFF}(2, n_2)] A^{\text{SSH}}(\mathbf{s}) \\
&= c_{1,1} \left(\sum_{h=1}^r n_h - 1 \right) + c_{2,2} \sum_{h=2}^r n_h + \\
&\quad \sum_{h=3}^{r-1} (c_{2,h} - c_{2,2}) (1 - p_h)^{n_2} \sum_{\substack{\mathbf{s} \preceq \text{SUFF}(h, n_h) \\ \text{SUFF}(h+1, n_{h+1}) \prec \mathbf{s}}} q[\mathbf{s} | \text{SUFF}(3, n_3)] A^{\text{SSH}}(\mathbf{s}) + (c_{2,r} - c_{2,2}) (1 - p_r)^{n_2} \sum_{\mathbf{s} \preceq \text{SUFF}(r, n_r)} q[\mathbf{s} | \text{SUFF}(3, n_3)] A^{\text{SSH}}(\mathbf{s}) \\
&= c_{1,1} \left(\sum_{h=1}^r n_h - 1 \right) + c_{2,2} \sum_{h=2}^r n_h + \sum_{h=3}^{r-1} c_{3,h} \sum_{\text{SUFF}(h+1, n_{h+1}) \prec \mathbf{s} \preceq \text{SUFF}(h, n_h)} q[\mathbf{s} | \text{SUFF}(3, n_3)] A^{\text{SSH}}(\mathbf{s}) + \\
&\quad c_{3,r} \sum_{\mathbf{s} \preceq \text{SUFF}(r, n_r)} q[\mathbf{s} | \text{SUFF}(3, n_3)] A^{\text{SSH}}(\mathbf{s}) \\
&= \dots = c_{1,1} \left(\sum_{h=1}^r n_h - 1 \right) + \sum_{j=2}^r \left(c_{j,j} \sum_{h=j}^r n_h \right).
\end{aligned}$$

Therefore,

$$q[\mathbf{n} | \mathbf{n}] A^{\text{SSH}}(\mathbf{n}) = \sum_{h=1}^r n_h - c_{1,1} \left(-1 + \sum_{h=1}^r n_h \right) - \sum_{j=2}^r c_{j,j} \sum_{h=j}^r n_h.$$

By rearranging we obtain that

$$q[\mathbf{n} | \mathbf{n}] A^{\text{SSH}}(\mathbf{n}) = c_{1,1} + \sum_{i=1}^r n_i \left(1 - \sum_{h=1}^i c_{h,h} \right).$$

The proof follows using Eq. (3). \square

Derivation based on per-packet HT estimator. Let h be a per-packet weight function and let $h(f) = \sum_{c \in f} h(c)$ be the $h()$ -value of f . ($h(c) \equiv w(c) = 1$ for packet counts but we can also use other packet-level properties such as the number of bytes in c .) We compute unbiased adjusted $h()$ values $H^{\text{SSH}}(c)$ for each packet from which we can obtain unbiased adjusted $h()$ values for each flow using $H^{\text{SSH}}(f) = \sum_{c \in f} H^{\text{SSH}}(c)$. By definition, packets that are not counted have adjusted $h()$ -values zero.

The HT estimator of $h(c)$ is the ratio of $h(c)$ and the probability that the packet c is counted in the sketch. It is clearly unbiased. This probability, however, can not be computed from the sketch. We use a partition of the sample space such that within each subspace in the partition there is a positive probability that the packet is sampled and this conditional probability can be determined from the sketch. The adjusted $h()$ value for each packet is an application of the HT estimator within this subspace.

We consider the adjusted $h()$ value $H^{\text{SSH}}(c)$ of a counted packet c of a flow f . We partition the sample space such that all rank assignment in the same subspace of the partition share the following.

1. The rank values of packets in $F \setminus \{f\}$.
2. The number of packets of f that are counted continuously up to and not including c (Note that this could be 0).

Note that the subspace that our rank assignment is mapped to also includes rank assignments where f does not appear in the sketch at all or that f appears but c is not counted. This happens if the current sampling rate drops below the current rank of the flow right before or after c is processed.

We compute the conditional probability that c is counted assuming that the rank assignment belongs to the particular subspace that it maps to. Since the ranks of packets of flows in $F \setminus \{f\}$ are fixed in this subspace then so are the steps, \mathbf{p} , of the k th smallest rank of a flow in $F \setminus \{f\}$. Furthermore, in any rank assignment in the given partition where packet c is counted, we count the same number of packets in each step. Let \mathbf{n} be the vectors of counts we obtain for f in any rank assignment where c is counted. (In rank assignments in the same subspace where c is not counted this vector could be different.)

$H^{\text{SSH}}(c)$ is computed according to one of the following cases.

1. Packet c is one of the n_i packets counted in step i for some $i > 1$. In this case the conditional probability that c is counted is

$$\frac{q[\mathbf{n}|\mathbf{n}]}{1 - \sum_{h=1}^i c_{h,h}}. \quad (8)$$

To see this, fix the ranks of packets of $F \setminus \{f\}$. Then $(1 - \sum_{h=1}^i c_{h,h})$ is the probability that we count all $n_1 + \dots + n_i$ up to and including the packets of step i . $q[\mathbf{n}|\mathbf{n}]$ is the probability that we count all \mathbf{n} packets. Therefore, Eq. (8) is the conditional probability that we count \mathbf{n} given that we counted all packets up to c . It follows that $H^{\text{SSH}}(c) = h(c) \frac{1 - \sum_{h=1}^i c_{h,h}}{q[\mathbf{n}|\mathbf{n}]}$.

2. Suppose c is the first packet among the n_1 packets of step 1. In this case, the conditional probability that c is counted is $q[\mathbf{n}|\mathbf{n}]$, and $H^{\text{SSH}}(c) = h(c)/q[\mathbf{n}|\mathbf{n}]$.
3. Suppose that c is a packet of step 1 other than the first. Fixing the ranks of packets of flows in $F \setminus \{f\}$ we count the packets of step 1 with probability p_1 : That is the probability that the first packet in step 1 is counted. So the conditional probability that c is counted is $q[\mathbf{n}|\mathbf{n}]/p_1$ and $H^{\text{SSH}}(c) = h(c)p_1/q[\mathbf{n}|\mathbf{n}]$.

Let N_i be the set of packets counted in step i , and let c_0 be the first counted packet.

$$H^{\text{SSH}}(f) = \sum_{c \in f} H^{\text{SSH}}(c) = \sum_{j \geq 1} \sum_{c \in N_j} H^{\text{SSH}}(c) = \frac{h(c_0) + (h(N_1) - h(c_0))(1 - c_{1,1}) + \sum_{i=2}^r h(N_i)(1 - \sum_{h=1}^i c_{h,h})}{q[\mathbf{n}|\mathbf{n}]} = \frac{h(c_0)c_{1,1} + \sum_{i=1}^r h(N_i)(1 - \sum_{h=1}^i c_{h,h})}{q[\mathbf{n}|\mathbf{n}]}. \quad (9)$$

To facilitate this estimator, the algorithm needs to collect per-step sums $h(N_i)$ over counted packets in the step and to separately record $h(c_0)$.

Derivation based on dominance of SSH over ASH

Proof. We compute the expectation of the adjusted weight of ASH when considered nonzero only on rank assignments such that SSH fully counts the flow. We refer to this expectation as the *combined expectation*. The conditional expectation we seek is then the ratio of the combined expectation and $q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]$. Therefore, we need to show that the combined expectation is equal to

$$q^{\text{SSH}}[\mathbf{n}|\mathbf{n}] A^{\text{SSH}}(\mathbf{n}).$$

Suppose now that $|\mathbf{n}| > n_r$. If the SSH count is \mathbf{n} , the ASH count on the same rank assignment must start in a packet that occurred before the last step (for SSH to fully count the flow the rank must be at most p_r before the beginning of the step.), and therefore step r has contribution zero to the combined expectation. We compute, for each step $r - 1 \geq \ell > 1$ the ‘‘contribution’’ to the combined expectation if ASH started

counting the flow during step ℓ . The probability that at the beginning of the step the rank value was at least p_r is $(c_{\ell,r} - c_{\ell,\ell})$. Conditioned on that, the contribution to the expectation is

$$\sum_{t=0}^{n_\ell-1} (1-p_r)^t p_r \left(\sum_{h=\ell}^r n_h + (1-p_r)/p_r - t \right) = \sum_{h=\ell}^r n_h - (1-p_r)^{n_\ell} \sum_{h=\ell+1}^r n_h. \quad (10)$$

Therefore, the contribution is the product

$$(c_{\ell,r} - c_{\ell,\ell}) \left(\sum_{h=\ell}^r n_h - (1-p_r)^{n_\ell} \sum_{h=\ell+1}^r n_h \right) = (c_{\ell,r} - c_{\ell,\ell}) \sum_{h=\ell}^r n_h - c_{\ell+1,r} \sum_{h=\ell+1}^r n_h. \quad (11)$$

We first consider the case where $|\mathbf{n}| = n_r$ (all observed packets are in the last step). In this case, if the SSH count is \mathbf{n} , the ASH count must be $|\mathbf{n}|$ and $A^{\text{SSH}}(\mathbf{n}) \equiv A^{\text{ASH}}(|\mathbf{n}|)$. The combined expectation is therefore $q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]A^{\text{ASH}}(|\mathbf{n}|) = q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]A^{\text{SSH}}(\mathbf{n})$.

We need to be slightly careful with the first step by considering the rank of the first packet and then the other $n_1 - 1$ packets. With probability p_r , the first packet obtains rank value at most p_r and ASH uses the adjusted weight $\sum_{h=1}^r n_h + (1-p_r)/p_r$ and therefore the contribution is the product

$$(1-p_r) + p_r \sum_{h=1}^r n_h. \quad (12)$$

With probability $p_1 - p_r \equiv c_{1,r} - c_{1,1}$, the first packet obtains rank value in $(p_r, p_1]$, and applying a similar derivation as Eq. 10 for the remaining $n_1 - 1$ packets, we obtain $-1 + \sum_{h=1}^r n_h - (1-p_r)^{n_1-1} \sum_{h=2}^r n_h$. The contribution is therefore,

$$p_r - p_1 + (c_{1,r} - c_{1,1}) \sum_{h=1}^r n_h - c_{2,r} \sum_{h=2}^r n_h. \quad (13)$$

Summing the contributions of all the steps in Eq. 11, Eq. 12, and Eq. 13 we obtain that this expectation is

$$\begin{aligned} & (1-p_r) + p_r \sum_{h=1}^r n_h + p_r - p_1 + (c_{1,r} - c_{1,1}) \sum_{h=1}^r n_h - c_{2,r} \sum_{h=2}^r n_h + \sum_{\ell=2}^{r-1} \left((c_{\ell,r} - c_{\ell,\ell}) \sum_{h=\ell}^r n_h - c_{\ell+1,r} \sum_{h=\ell+1}^r n_h \right) \\ = & p_1 n_1 + (1-p_1) + \sum_{\ell=2}^r (1 - \sum_{h=1}^{\ell} c_{h,h}) n_\ell = q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]A^{\text{SSH}}(\mathbf{n}). \end{aligned}$$

(Using Theorem 4.8 and Eq. 3.)

□

We obtain the following:

Lemma 9.2. *Consider a flow with SSH counts \mathbf{n} . The conditional expectation of the adjusted weight assigned to the flow by ASH, given that the flow is fully counted by SSH, is equal to $A^{\text{SSH}}(\mathbf{n})$ (the adjusted weight assigned by SSH for observed count \mathbf{n}).*

9.2 Variance relation

Consider a flow f with $|f|$ packets and the probability subspace where ranks of packets belonging to all other flows $(F \setminus \{f\})$ are fixed. It is sufficient to establish the relation between the methods in this subspace. Consider such a subspace. Let \mathbf{p} be the steps of the effective sampling rate and p_r be the final effective sampling

rate. The adjusted weight assignment for all methods has expectation $|f|$ within each such subspace. We consider the variance of the different methods within such subspace. and use the notation $\text{VAR}(A^L(f)|\mathbf{p})$ for $L \in \{\text{SSH}, \text{SNF}\}$, and $\text{VAR}(A^L(f)|p_r)$ for $L \in \{\text{WS}, \text{ANF}, \text{ASH}\}$. This conditioning is equivalent to establishing the variance relation when the sampling rate p_r is fixed or when the steps \mathbf{p} are fixed (and the last step is p_r). It is the key for extending the proofs to the discretized version, since we simply condition on a different step function \mathbf{p} determined by the discretized $(k+1)$ th largest rank. It also shows that the variance relation holds for the fixed-rate and fixed-steps variants of WS NF and SH.

The variance for ANF is that of a binomial random variable and therefore is

$$\text{VAR}(A^{\text{ANF}}(f)|p_r) = |f|(1-p_r)/p_r. \quad (14)$$

For WS, the adjusted weight is $|f|/(1-(1-p_r)^{|f|})$ with probability $1-(1-p_r)^{|f|}$ and zero otherwise, and therefore, the variance is

$$\text{VAR}(A^{\text{WS}}(f)|p_r) = |f|^2(1/(1-(1-p_r)^{|f|}) - 1). \quad (15)$$

For ASH, it is

$$\text{VAR}(A^{\text{ASH}}(f)|p_r) = \sum_{i=0}^{|f|-1} (1-p_r)^i p_r (|f|-i+(1-p_r)/p_r)^2 - |f|^2 = ((1-p_r) - (1-p_r)^{|f|+1})/p_r^2. \quad (16)$$

For SSH and for SNF, the variance $\text{VAR}(A^{\text{SNF}}(f)|\mathbf{p})$ depends on the way the packets of the flow f are distributed across these steps. The variance is lowest when all packets occur when the sampling probability is highest, and the variance is highest, and equal to that of ANF, when all packets occur on the step with the lowest sampling probability. The variance relation is established using the following Lemma.

Using the explicit expressions (Eq. 14,15,16) and the inequality $(1-p)^n \geq 1-np$ for all natural n and $0 \leq p \leq 1$ it follows that $\text{VAR}(A^{\text{WS}}(f)|p_r) \leq \text{VAR}(A^{\text{ASH}}(f)|p_r) \leq \text{VAR}(A^{\text{ANF}}(f)|p_r)$. We establish the relation between the variance of the different methods via direct arguments that provide more insights and are applicable to the step-counting algorithms.

Lemma 9.3. *Consider two mappings A_1 and A_2 and suppose there exists a partition of the sample space S into subspaces such that within each subspace $S' \subset S$,*

- $\mu_{S'}(A_1) = \mu_{S'}(A_2)$ (that is, A_1 and A_2 have the same expectation on the subspace S'), and
- $\mu_{S'}^2(A_1) \geq \mu_{S'}^2(A_2)$ (the variance, or alternatively, the second raw moment, of A_1 is at least as large as A_2).

Then $\mu_S(A_1) = \mu_S(A_2)$ and $\mu_S^2(A_1) \geq \mu_S^2(A_2)$ (A_1 and A_2 have the same expectation and the variance of A_1 is at least as large as A_2 .)

Corollary 9.4. *Let A_1 be an estimator and consider a partition of the sample space. Consider the estimator A'_1 that has a value that is equal to the expectation of A_1 on the respective part of the partition. Then $E(A_1) = E(A'_1)$ and $\text{VAR}(A_1) \geq \text{VAR}(A'_1)$.*

If f is not included in the sample ($r(f) > p_r$), it obtains an adjusted weight of zero with all four methods. Therefore, it suffices to compare the methods based on the variance in the adjusted weight assignment within the probability subspace when the flow is sampled ($r(f) \leq p_r$). Since all methods are unbiased, they all have the same expectation on this subspace. We apply Lemma 9.3. With WS, f obtains an adjusted weight of $|f|/(1-(1-p_r)^{|f|})$, which is fixed, therefore the variance is zero. Therefore, this assignment is optimal among all methods that yield the same probability distributions over subsets of flows. In particular, $\text{VAR}(A^{\text{WS}}(f)|p_r) \leq \text{VAR}(A^{\text{SSH}}(f)|\mathbf{p})$.

We next compare ANF and ASH. We further partition the probability subspace by fixing the position $i \in [1, \dots, |f|]$ of the first packet in f that obtains rank that is at most p_r . The adjusted weight assigned by ASH is (a fixed value) of $(1/p_r) + (|f| - i)$, and therefore the variance is zero. ANF assignment is $(1/p_r)$ plus $(1/p_r)$ times a binomial random variable with parameters $(|f| - i)$ and p_r . This assignment has the same conditional expectation as the ASH assignment within this subspace, but also has a nonnegative, and therefore at least as large, variance (the variance is strictly positive when $i < |f|$). Therefore, using Lemma 9.3, ASH has a smaller variance overall than ANF.

We compare the variance of SSH and ASH by again applying Lemma 9.3. (Observe that we can not use the same partition that we used for comparing ASH and ANF, since SSH does not have the same expectation as ASH on each such subspace.) We partition the sample space according to the *suffix of the packets of f that are counted using SSH*. That is, each subspace contains all rank assignments to packets of f that result in this suffix being counted. The adjusted weight assigned by SSH is fixed within each partition and therefore has zero variance. The ASH adjusted weight depends on the first packet that obtains rank value below p_r , and therefore varies⁷. What remains to show is that ASH and SSH adjusted weights have the same expectation within each such subspace.

We apply the following simple observation. Consider a flow with counts \mathbf{n} and a vector $\mathbf{s} \preceq \mathbf{n}$. The conditional probability that i packets are counted using ASH given that \mathbf{s} is counted using SSH is independent of the choice of \mathbf{n} (depends only on \mathbf{s}). Therefore, the expectation of the adjusted weight assigned by ASH conditioned on SSH counting \mathbf{s} out of \mathbf{n} is equal to the expectation when SSH counts \mathbf{s} out of \mathbf{s} . This simplifies the proof, as it suffices to establish equality for flows that are fully counted by SSH, which is established in Lemma 9.2.

Lemma 9.5. $\text{VAR}(A^{\text{SSH}}(f)|\mathbf{p}) \leq \text{VAR}(A^{\text{SNF}}(f)|\mathbf{p})$

Proof. Consider a flow f and a probability space $\Omega^{(\mathbf{p})}$ containing all rank assignments such that the steps (as defined by the k th smallest rank in $F \setminus f$) are \mathbf{p} . Consider a partition of Ω to subspaces $\Omega_{\mathbf{n}}^{(\mathbf{p})}$ according to the SSH count vector \mathbf{n} obtained for f .

Consider one such subspace $\Omega_{\mathbf{n}}^{(\mathbf{p})}$. By definition, the adjusted weight assigned to the flow f in this subspace is fixed and is equal to $A_{\mathbf{p}}^{\text{SSH}}(\mathbf{n})$.

We define another SSH adjusted weight assignment, $A'_{\mathbf{p}}^{\text{SSH}}(\mathbf{n})$ as the expectation of the estimator $A_{\mathbf{p}}^{\text{SNF}}$ over rank assignments in $\Omega_{\mathbf{n}}^{(\mathbf{p})}$.

For any rank assignment, a packet is counted by SNF only if it is counted by SSH. The first counted packet by SSH, must also be counted by SNF. Therefore, \mathbf{s} is a possible SNF count of f in $\Omega_{\mathbf{n}}^{(\mathbf{p})}$ if and only if it has the form $\mathbf{s} \leq \mathbf{n}$ (component wise) and $s_1 > 0$ (we assume WLOG that $n_1 > 0$). For notation convenience, we define the vectors $\mathbf{n}' = (1, n_1 - 1, n_2, \dots, n_r)$, $\mathbf{s}' = (1, s_1 - 1, s_2, \dots, s_r)$ and $\mathbf{p}' = (p_1, p_1, p_2, \dots, p_r)$. (That is, we create a “dummy” step with probability p_1 , that precedes the first step and contains the first packet of \mathbf{n} .) This notation allows us to specify that the first “packet” of \mathbf{n} is counted.

The probability over $\Omega_{\mathbf{n}}^{(\mathbf{p})}$ of a rank assignment with corresponding SNF count \mathbf{s} is equal to

$$q_{\mathbf{p}'}^{\text{SNF}}[\mathbf{s}'|\mathbf{n}'] / q_{\mathbf{p}}^{\text{SSH}}[\mathbf{n}|\mathbf{n}].$$

Therefore,

$$A'_{\mathbf{p}}^{\text{SSH}}(\mathbf{n}) = \sum_{\mathbf{s} \leq \mathbf{n}, s_1 > 0} \frac{q_{\mathbf{p}'}^{\text{SNF}}[\mathbf{s}'|\mathbf{n}']}{q_{\mathbf{p}}^{\text{SSH}}[\mathbf{n}|\mathbf{n}]} A_{\mathbf{p}}^{\text{SNF}}(\mathbf{s}). \quad (17)$$

⁷The variance can be zero only if all suffix packets are contained in the last step. In this case, the first packet of the suffix has rank value below p_r and therefore all suffix packets are counted by ASH.

It follows from Equation (17) that $A'_{\mathbf{p}}^{\text{SSH}}(\mathbf{n})$ is a deterministic function of \mathbf{p} and \mathbf{n} .⁸ Using corollary 9.4, the estimator $A'_{\mathbf{p}}^{\text{SSH}}$ is unbiased and has variance that is at most that of $A_{\mathbf{p}}^{\text{SNF}}$ over $\Omega_{\mathbf{p}}$ (and therefore, over any probability space that consists of subspaces of the form of $\Omega_{\mathbf{p}}$.)

In Sections 4.1 and 9.1 we show that $A_{\mathbf{p}}^{\text{SSH}}(\mathbf{n})$ is the unique solution of a system of equations. Therefore, it is the only possible assignment of adjusted weights that are a deterministic function of \mathbf{p} and \mathbf{n} and are unbiased (has expectation $|f|$) for any possible f and a corresponding probability space $\Omega(\mathbf{P})$. Since the estimator $A'_{\mathbf{p}}^{\text{SSH}}(\mathbf{n})$ is also a deterministic function of \mathbf{n} and \mathbf{p} and is unbiased on $\Omega(\mathbf{P})$ it follows that $A_{\mathbf{p}}^{\text{SSH}} \equiv A'_{\mathbf{p}}^{\text{SSH}}$. □

9.3 Adjusted weights for SNF

We prove Theorem 4.10 (derivation of adjusted weights for SNF) using the dominance of SNF over ANF. This proof also establishes the variance relation $\text{VAR}(A^{\text{SNF}}(f|\mathbf{p})) \leq \text{VAR}(A^{\text{NF}}(f|p_r))$.

Proof. We fix the ranks of the packets of flows in $F \setminus f$. The the steps \mathbf{p} of the k th smallest rank of a flow in $F \setminus f$ are then fixed. Let \mathbf{n} be the number of packets of f in each of these steps. (We assume without loss of generality that $n_1 > 0$.) We consider the subspace V of rank assignments where f is fully counted. Let $A^{\text{SNF}}(\mathbf{n})$ be the adjusted weight that f obtains at any point of this subspace. We show that $A^{\text{SNF}}(\mathbf{n})$ is the average of the adjusted weight of NF in V . Note that the fraction of V is $q^{\text{SNF}}[\mathbf{n}|\mathbf{n}]$.

Consider first points in V where the first packet that has rank at most p_r is a packet t of step 1.

The probability that SNF counted all packets and packet $t + 1$ was the first packet to obtain rank value at most p_r is $(p_1 - p_r)^t p_r p_1^{n_1 - t - 1} \prod_{h=2}^r p_h^{n_h}$. Conditioned on this, the adjusted weight assigned by NF is the number of counted packets divided by p_r . The expected number of counted packets from step 1 is $1 + (n_1 - t - 1)(p_r/p_1)$ and the expected number from step $2 \leq j \leq r$ is $n_j(p_r/p_j)$. Therefore, the expected adjusted weight assigned by NF is $(1/p_r + \sum_{j=1}^r n_j/p_j - (t+1)/p_1)$. We sum over $t = 0, \dots, n_1 - 1$ and divide by $q^{\text{SNF}}[\mathbf{n}|\mathbf{n}]$ to obtain the contribution to the average adjusted weight of NF in V of the points where the first packet that has rank at most p_r is of step 1.

$$\frac{\prod_{h=2}^r p_h^{n_h} \sum_{t=0}^{n_1-1} (p_1 - p_r)^t p_r p_1^{n_1 - t - 1} \left(\frac{1}{p_r} + \sum_{j=1}^r \frac{n_j}{p_j} - \frac{t+1}{p_1} \right)}{q^{\text{SNF}}[\mathbf{n}|\mathbf{n}]} . \quad (18)$$

The derivation of the contribution to the average of points where the first packet having rank at most p_r is in steps $\ell = 2, \dots, r - 1$ ⁹ is similar to that of Eq.(18), observing that $\frac{(d_{\ell,r} - d_{\ell,\ell})}{p_{\ell}^{n_{\ell}}} (p_{\ell} - p_r)^t p_r p_{\ell}^{n_{\ell} - t - 1}$ is the probability that SNF fully counted the flow and the first packet to obtain rank value at most p_r was packet $t + 1$ during step ℓ . We omit the denominator $q^{\text{SNF}}[\mathbf{n}|\mathbf{n}]$ from the following two Equations.

$$\frac{(d_{\ell,r} - d_{\ell,\ell})}{p_{\ell}^{n_{\ell}}} \sum_{t=0}^{n_{\ell}-1} (p_{\ell} - p_r)^t p_r p_{\ell}^{n_{\ell} - t - 1} \left(\frac{1}{p_r} + \sum_{j=\ell}^r \frac{n_j}{p_j} - \frac{t+1}{p_{\ell}} \right) . \quad (19)$$

⁸This also follows the fact that SSH *dominates* SNF in the sense that SNF sketches can be emulated from SSH sketches. That is, given \mathbf{p} and \mathbf{n} , we can draw an SNF sketch from $\Omega_{\mathbf{n}}(\mathbf{P})$.

⁹Observe that the contribution of the last step must be zero unless it is the only step, since if the flow is fully counted its rank must be at most p_r before the beginning of the last step.

We obtain that Eq. (19) (contribution of step $\ell > 1$) is equal to

$$\begin{aligned} & (d_{\ell,r} - d_{\ell,\ell}) \frac{p_r}{p_\ell} \sum_{t=0}^{n_\ell-1} \left(\frac{p_\ell - p_r}{p_\ell} \right)^t \left(\frac{1}{p_r} + \sum_{j=\ell}^r \frac{n_j}{p_j} - \frac{t+1}{p_\ell} \right) \\ &= (d_{\ell,r} - d_{\ell,\ell}) \frac{p_r}{p_\ell} \left\{ \left(\frac{p_\ell - p_r}{p_\ell p_r} + \sum_{j=\ell}^r \frac{n_j}{p_j} \right) \sum_{t=0}^{n_\ell-1} \left(\frac{p_\ell - p_r}{p_\ell} \right)^t \frac{1}{p_\ell} \sum_{t=0}^{n_\ell-1} \left(\frac{p_\ell - p_r}{p_\ell} \right)^t t \right\} \end{aligned} \quad (20)$$

The first sum in the expression above is geometric, and the second is of the form $\sum_{k=0}^m kq^k$ for $q = \frac{p_\ell - p_r}{p_\ell}$. Since $\sum_{k=0}^m kq^k = \sum_{k=0}^m (f(k+1) - f(k)) = f(m+1) - f(0)$ where $f(x) = \frac{1}{q-1} \left(xq^x - \frac{q^{x+1}}{q-1} \right)$, we obtain that

$$\sum_{k=0}^m kq^k = \frac{1}{q-1} \left((m+1)q^{m+1} - \frac{q^{m+2}}{q-1} + \frac{q}{q-1} \right).$$

Using these observation we obtain that Equation (20) is equal to

$$\begin{aligned} &= (d_{\ell,r} - d_{\ell,\ell}) \left\{ \left(\frac{p_\ell - p_r}{p_\ell p_r} + \sum_{j=\ell}^r \frac{n_j}{p_j} \right) \left(1 - \left(\frac{p_\ell - p_r}{p_\ell} \right)^{n_\ell} \right) + \frac{n_\ell}{p_\ell} \left(\frac{p_\ell - p_r}{p_\ell} \right)^{n_\ell} + \frac{p_\ell - p_r}{p_\ell p_r} \left(\frac{p_\ell - p_r}{p_\ell} \right)^{n_\ell} - \frac{p_\ell - p_r}{p_\ell p_r} \right\} \\ &= \left(\frac{p_\ell - p_r}{p_\ell p_r} + \sum_{j=\ell}^r \frac{n_j}{p_j} \right) (d_{\ell,r} - d_{\ell,\ell} - d_{\ell+1,r}) + \frac{n_\ell}{p_\ell} d_{\ell+1,r} + \frac{p_\ell - p_r}{p_\ell p_r} d_{\ell+1,r} - \frac{p_\ell - p_r}{p_\ell p_r} (d_{\ell,r} - d_{\ell,\ell}) \\ &= (d_{\ell,r} - d_{\ell,\ell}) \sum_{j=\ell}^r \frac{n_j}{p_j} - d_{\ell+1,r} \sum_{j=\ell+1}^r \frac{n_j}{p_j} \end{aligned} \quad (21)$$

By applying similar manipulations to Eq. (18), we obtain that the numerator of that equation is equal to

$$\prod_{h=1}^r p_h^{n_h} \sum_{j=1}^r \frac{n_j}{p_j} - d_{2,r} \sum_{j=2}^r \frac{n_j}{p_j} \quad (22)$$

We sum the contributions of steps $\ell = 1, \dots, r-1$ (Eq. (22) and Eq. (21) for $\ell = 2, \dots, r-1$) and obtain that the total contribution to the expectation is

$$\prod_{h=1}^r p_h^{n_h} \sum_{j=1}^r \frac{n_j}{p_j} - \sum_{\ell=2}^r d_{\ell,\ell} \sum_{j=\ell}^r \frac{n_j}{p_j} = \sum_{j=1}^r \frac{n_j}{p_j} \left(\prod_{h=1}^r p_h^{n_h} - \sum_{\ell=2}^j d_{\ell,\ell} \right).$$

Therefore,

$$A^{\text{SNF}}(\mathbf{n}) = \frac{\sum_{j=1}^r \frac{n_j}{p_j} \left(\prod_{h=1}^r p_h^{n_h} - \sum_{\ell=2}^j d_{\ell,\ell} \right)}{q^{\text{SNF}}[\mathbf{n}|\mathbf{n}]} \quad (23)$$

The proof follows using Eq.(4). \square

Alternate proof (sketch) based on per-packet HT estimator. We apply the HT estimator to obtain an adjusted $h(\cdot)$ -value for each observed packet. The proof methodology is similar to the one we provided for SSH, so we provide only a sketch.

Consider a rank assignment \mathbf{x} that results in an SNF sketch with steps \mathbf{p} . The sketch includes a flow f with counts \mathbf{n} . Let $N_i \subset f$ ($i = 1, \dots, r$) be the set of packets of f that are counted with \mathbf{x} at step i .

We associate each packet with a subspace of rank assignments as follows. For $c \in N_i$, the subspace is defined by the following constraints: (i) The rank values of all packets of flows in $F \setminus \{f\}$ are as in \mathbf{x} . (ii) Each packet $a \in \bigcup_{j=1}^r N_j \setminus \{c\}$ has rank value that is below the current sampling rate at the time the

packet arrives. (iii) The rank assignment is such that at the time packet c arrives, flow f is cached with a count that includes all the packets in $\bigcup_{j=1}^i N_j$ that precede c .

This is a mapping from a rank assignment \mathbf{x} and a packet c into a subspace. The subspaces partition the space of all rank assignments (including those where c is not counted in the sketch).

We next show how to compute, given the sketch and the packet c , the conditional probability, within the subspace which \mathbf{x} maps to, that c is counted. This probability is equal to the probability that all packets in $\bigcup_{j=1}^r N_j$ are counted. Consider the subspace defined by (i): The probability that (ii)+(iii) hold given (i) is: $(1/p_i)(\prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^i d_{j,j})$. The probability that all packets in $\bigcup_{j=1}^r N_j$ are counted given the constraint (i) is $\prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^r d_{j,j}$. If all packets are counted in the subspace specified by (i), then constraints (ii) and (iii) must hold. Therefore, the probability that c is counted conditioned on (i)+(ii)+(iii) is

$$\frac{\prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^r d_{j,j}}{(1/p_i)(\prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^i d_{j,j})}.$$

The adjusted $h()$ -value of the packet c is then

$$H^{\text{SNF}}(c) = \frac{(h(c)/p_i) \prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^i d_{j,j}}{\prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^r d_{j,j}}.$$

The adjusted $h()$ -value of the flow is

$$H^{\text{SNF}}(f) = \sum_{c \in f} H^{\text{SNF}}(c) = \frac{\sum_{i=1}^r \frac{h(N_i)}{p_i} \prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^i d_{j,j}}{\prod_{h=1}^r p_h^{n_h} - \sum_{j=2}^r d_{j,j}}. \quad (24)$$

9.4 Covariances

We provide the proof of Lemma 4.11. The proof is based on conditioning on the rank values of packets belonging to flows in $F \setminus \{f_1, f_2\}$, and the methodology carries over to establish this property for the discretized versions.

Proof. It suffices to show that $E(A(f_1)A(f_2)) = w(f_1)w(f_2)$ (we omit the superscript L). We adapt the proof method in [4] used to establish zero covariance for rank conditioning estimators. We partition the sample space according to the $(k-1)$ th smallest rank value among the flows in $F \setminus \{f_1, f_2\}$. Consider one part and let r_{k-1} be that rank value. The product $A(f_1)A(f_2)$ is positive only when $r(f_1) < r_{k-1}$ and $r(f_2) < r_{k-1}$ (it is zero otherwise, since at least one of f_1 or f_2 is not included in the sketch). In this case, the effective sampling rate is equal to r_{k-1} . The count obtained for f_1 using ASH or ANF only depends on r_{k-1} and not on the count of f_2 . Therefore $A(f_1)$ and $A(f_2)$ are independent conditioned on them both having ranks below r_{k-1} . The expectation of $A(f_i)$ under this conditioning is $w(f_i)/\text{PR}\{r(f_i) < r_{k-1}\}$, and the conditional expectation

$$E(A(f_1)A(f_2) | (r(f_1) < r_{k-1}) \wedge (r(f_2) < r_{k-1})) = \frac{w(f_1)w(f_2)}{\text{PR}\{r(f_1) < r_{k-1}\}\text{PR}\{r(f_2) < r_{k-1}\}}. \quad (25)$$

Therefore, on this part

$$E(A(f_1)A(f_2)) = \frac{\text{PR}\{r(f_1) < r_{k-1}\}\text{PR}\{r(f_2) < r_{k-1}\}w(f_1)w(f_2)}{\text{PR}\{r(f_1) < r_{k-1}\}\text{PR}\{r(f_2) < r_{k-1}\}} = w(f_1)w(f_2). \quad (26)$$

We next consider SSH (The proof for SNF is along the same lines and is omitted). We use the following property: fix the step function \mathbf{p} of the current sampling rate and let p_r be the effective sampling rate. The

expectation of $A^{\text{SSH}}(f_1)$ conditioned on $r(f_1) < p_r$ is equal to $w(f_1)/\text{PR}\{r(f_1) < p_r\}$. We partition the sample space according to the step functions of r_{k-1} and r_k (the k th and $(k-1)$ th smallest current rank value among the flows in $F \setminus \{f_1, f_2\}$). Consider a part in this partition. The product $A(f_1)A(f_2)$ is positive only if $r(f_1) < r_{k-1}$ and $r(f_2) < r_{k-1}$. In this case, the current sampling rate is r_{k-1} . Fix the ranks of f_2 packets. The current sampling rate is determined and it is a step function \mathbf{p} with effective sampling rate $p_r = r_{k-1}$. The conditioned expectation of $A(f_1)$ in this part after fixing f_2 ranks, and given that it has $r(f_1) < r_{k-1}$ is $w(f_1)/\text{PR}\{r(f_1) < r_{k-1}\}$. It is independent of the f_2 ranks and therefore, the adjusted weight $A(f_1)$ and $A(f_2)$ in this part, conditioned on them both having ranks below r_{k-1} , are independent. Hence, the conditioned expectation of the product is as in Eq. 25 and therefore Eq. 26 also holds. \square

9.5 FSD estimators

We provide the proof of Lemma 5.2.

Proof. The proof is along the same lines of Lemma 9.2. We compute the expectation of the ASH FSD estimator conditioned on the SSH observed counts and sampling rate steps. Therefore, we also establish that the SSH-based estimator has at most the same variance as the ASH estimator.

Denote by $\pi(\mathbf{p}, \mathbf{n}, j)$ the conditional probability of an ASH count j given SSH count \mathbf{n} and steps \mathbf{p} . For all $i \geq 1$,

$$\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \sum_{j=1}^{|\mathbf{n}|} \pi(\mathbf{p}, \mathbf{n}, j) \alpha_i^{\text{ASH}}(p_r, j) = \frac{\pi(\mathbf{p}, \mathbf{n}, i)}{p_r} - \frac{\pi(\mathbf{p}, \mathbf{n}, i+1)(1-p_r)}{p_r}. \quad (27)$$

If $|\mathbf{n}| = 1$ then $\pi(\mathbf{p}, \mathbf{n}, 1) = 1$ and $\pi(\mathbf{p}, \mathbf{n}, j) = 0$ for $j \neq 1$ (the ASH count is 1 with probability 1). Therefore $\alpha_1^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 1/p_r$ (and $\alpha_j^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 0$ for all $j \neq 1$).

If $|\mathbf{n}| = n_r$ then $\pi(\mathbf{p}, \mathbf{n}, n_r) = 1$ and $\pi(\mathbf{p}, \mathbf{n}, j) = 0$ for $j \neq n_r$ (the ASH count is n_r with probability 1). Therefore, the only nonzero $\alpha_j^{\text{SSH}}(\mathbf{p}, \mathbf{n})$ are $\alpha_{n_r}^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = 1/p_r$ and $\alpha_{n_r-1}^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = -(1-p_r)/p_r$.

Consider \mathbf{n} such that $|\mathbf{n}| > 1$ and $|\mathbf{n}| > n_r$. Assume WLOG that $n_\ell > 0$ for $1 \leq \ell < r$. The probabilities $\pi(\mathbf{p}, \mathbf{n}, x) > 0$ are positive if only if $|\mathbf{n}| \geq x \geq n_r + 1$ and are as follows:

$$\begin{aligned} \ell = 2, \dots, r-1 \text{ and } \sum_{h=\ell+1}^r n_h < x \leq \sum_{h=\ell}^r n_h : & \quad \pi(\mathbf{p}, \mathbf{n}, x) = \frac{(c_{\ell,r} - c_{\ell,\ell})(1-p_r)^{\sum_{h=\ell}^r n_h - x} p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]} . \\ \sum_{h=2}^r n_h < x < \sum_{h=1}^r n_h : & \quad \pi(\mathbf{p}, \mathbf{n}, x) = \frac{(c_{1,r} - c_{1,1})(1-p_r)^{\sum_{h=1}^r n_h - 1 - x} p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]} . \\ x = |\mathbf{n}| : & \quad \pi(\mathbf{p}, \mathbf{n}, |\mathbf{n}|) = p_r / q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]. \end{aligned}$$

Therefore, $\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n})$ can be nonzero only for $|\mathbf{n}| \geq i \geq n_r$. We compute these values case by case using Eq.(27).

- i such that $\sum_{j=\ell+1}^r n_j < i < \sum_{j=\ell}^r n_j$ for some $2 \leq \ell \leq r-1$:

$$\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{(c_{\ell,r} - c_{\ell,\ell})(1-p_r)^{\sum_{h=\ell}^r n_h - i} p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}] p_r} - \frac{(1-p_r)(c_{\ell,r} - c_{\ell,\ell})(1-p_r)^{\sum_{h=\ell}^r n_h - i - 1} p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}] p_r} = 0$$

- if $n_1 \geq 3$, i such that $\sum_{j=2}^r n_j < i < \sum_{j=1}^r n_j - 1$:

$$\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{(c_{1,r} - c_{1,1})(1-p_r)^{\sum_{h=1}^r n_h - i - 1} p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}] p_r} - \frac{(1-p_r)(c_{1,r} - c_{1,1})(1-p_r)^{\sum_{h=1}^r n_h - i - 2} p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}] p_r} = 0$$

- $i = \sum_{j=\ell}^r n_j$ **for some** $2 \leq \ell \leq r$ If $\ell = 2$ we assume $n_1 > 1$, If $\ell = r$ we assume $n_r > 0$ (because otherwise $\alpha_{n_r}()$ is not defined):

$$\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{(c_{\ell,r} - c_{\ell,\ell})p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]p_r} - \frac{(1-p_r)(c_{\ell-1,r} - c_{\ell-1,\ell-1})(1-p_r)^{n_\ell-1}p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]p_r} = \frac{-c_{\ell,\ell}}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]}$$

- $i = \sum_{j=2}^r n_j, n_1 = 1$:

$$\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{(c_{2,r} - c_{2,2})p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]p_r} - \frac{(1-p_r)p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]p_r} = \frac{-(1-p_2)}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]} \equiv \frac{-c_{1,1} - c_{2,2}}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]}$$

- $i = -1 + \sum_{j=1}^r n_j, n_1 > 1$:

$$\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{(c_{1,r} - c_{1,1})p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]p_r} - \frac{(1-p_r)p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]p_r} = \frac{-(1-p_1)}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]} = \frac{-c_{1,1}}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]}$$

- $i = \sum_{j=1}^r n_j$:

$$\alpha_i^{\text{SSH}}(\mathbf{p}, \mathbf{n}) = \frac{p_r}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]p_r} - 0 = \frac{1}{q^{\text{SSH}}[\mathbf{n}|\mathbf{n}]}$$

□

9.6 Estimators with negative covariances

ANF adjusted weights We provide a direct proof for the ANF adjusted selectivities.

Lemma 9.6. *Let $n : F$ be the ANF counts. For all $f \in F$, $R^{\text{ANF}}(f) = \frac{n(f)}{n(F)}$ are unbiased selectivities, that is,*

$$E(R^{\text{ANF}}(f)) = w(f)/w(F).$$

Proof. We assign *per-packet* adjusted selectivities such that for each packet v , $R^{\text{ANF}}(v) = 1/n(F)$ if v is counted and $R^{\text{ANF}}(v) = 0$ otherwise. We have $R^{\text{ANF}}(f) = \sum_{v \in f} R^{\text{ANF}}(v)$. It suffices to show that for each packet v , $E(R^{\text{ANF}}(v)) = 1/w(F)$.

We use the rank-based view of the sample space. Consider a rank assignment and the permutation of the packets according to their order by increasing rank value. Order the flows by the position of their first packet. Then $n(F) + 1$ is the position of the first packet of the $(k+1)^{\text{st}}$ flow in this order. The first k flows are the ones included in the sketch.

Consider a packet $v \in f$ and a partition of rank assignments over $F \setminus \{v\}$ (all packets other than v) according to the induced permutation on these packets. For each permutation, we define a corresponding value of ℓ as follows. If the first k flows in the induced permutation on $F \setminus \{v\}$ include f , then ℓ is the position of the first packet of the $(k+1)^{\text{st}}$ flow. If the first k flows do not include f then ℓ is the position of the first packet of the (k) th flow.

Consider a probability subspace in this partition and the respective value of ℓ . The packet v is counted if and only if its position in the permutation is at most ℓ . The conditional probability (in this subspace) that v is counted is $\ell/w(F)$. In this case there are ℓ counted packets ($n(F) = \ell$) and the adjusted selectivity of v is $1/\ell$. So in this subspace

$$E(R^{\text{ANF}}(v)) = \frac{\ell}{w(F)} \frac{1}{\ell} + \frac{w(F) - \ell}{w(F)} 0 = 1/w(F).$$

□

The effective sampling rate distribution is determined by $w(F)$ and the observed counts $n : F = \{(f, n(f))\}$. Consider a partition of the sample space according to $n : F$, and denote by $\Omega^{(L, n:F)}$ the subspace that corresponds to $n : F$. We will show that the estimator $A^{+L}(f, n : F, w(F))$ is the expectation of the estimator $A^L(n(f))$ over $\Omega^{(L, n:F)}$ for $L \in \text{ANF}, \text{ASH}$. Therefore, the relation $\text{VAR}(A^{+L}(f)) \leq \text{VAR}(A^L(f))$ for all $f \in F$ follows using Corollary 9.4.

Lemma 9.7. *For any flow $f \in F$,*

$$\text{VAR}(A^{+\text{ANF}}) \leq \text{VAR}(A^{\text{ANF}}(f)).$$

Proof. Consider a subspace $\Omega^{(n:F)}$. In this subspace, $A^{+\text{ANF}}(f) \equiv w(F)n(f)/n(F)$ is fixed. We show that within $\Omega^{(n:F)}$,

$$E(A^{\text{ANF}}(f)) = E(w(F)R^{\text{ANF}}(f)) = w(F)n(f)/n(F).$$

$A^{\text{ANF}}(f)$ is equal to $n(f)$ divided by the effective sampling rate, which is the $n(F) + 1$ smallest rank value. If the subspace contains a rank assignment, it contains all rank assignments that result in the same permutation (rank ordering) of packets. The effective sampling rate distribution in $\Omega^{(n:F)}$ is that of the $n(F) + 1$ smallest among $w(F)$ independent random variables from $U[0, 1]$. The expectation of the inverse of the i th smallest among n independent draws from $U[0, 1]$ is known to be $n/(i - 1)$. Therefore, the expectation of the inverse of the effective sampling rate in $\Omega^{(n:F)}$ is $w(F)/n(F)$ and

$$E_{\Omega^{(n:F)}}(A^{\text{ANF}}(f)) = E_{\Omega^{(n:F)}}(n(f)/p') = w(F)n(f)/n(F).$$

□

Lemma 9.8. *For any two flows $f_1 \neq f_2$,*

$$\text{COV}(R^{\text{ANF}}(f_1), R^{\text{ANF}}(f_2)) \leq 0.$$

Proof. We need to show that

$$E(R^{\text{ANF}}(f_1)R^{\text{ANF}}(f_2)) \leq \rho(f_1)\rho(f_2). \quad (28)$$

By definition

$$\begin{aligned} R^{\text{ANF}}(f_1)R^{\text{ANF}}(f_2) &= \sum_{v_1 \in f_1} \sum_{v_2 \in f_2} R^{\text{ANF}}(v_1)R^{\text{ANF}}(v_2) \\ \rho^{\text{ANF}}(f_1)\rho^{\text{ANF}}(f_2) &= \sum_{v_1 \in f_1} \sum_{v_2 \in f_2} \rho^{\text{ANF}}(v_1)\rho^{\text{ANF}}(v_2) = \sum_{v_1 \in f_1} \sum_{v_2 \in f_2} 1/w(F)^2 \end{aligned}$$

From linearity of expectation, a sufficient condition for Eq.(28) is that for any $v_1 \in f_1$ and $v_2 \in f_2$,

$$E(R^{\text{ANF}}(v_1)R^{\text{ANF}}(v_2)) \leq 1/w(F)^2. \quad (29)$$

We consider a partition of the sample space according to the induced permutation on the packets $w(F) \setminus \{v_1, v_2\}$. We show that Eq.(29) holds within each subspace. Consider such a permutation and define ℓ as follows. If the first k flows in $w(F) \setminus \{v_1, v_2\}$ include both f_1 and f_2 , then ℓ is the position of the first packet of the $(k + 1)$ st flow. Otherwise, if the first $k - 1$ flows include exactly one of $\{f_1, f_2\}$ (observe that the other can not be the k th flow), then ℓ is the position of the first packet of the k th flow. Otherwise, (the first $k - 1$ flows do not include either of $\{f_1, f_2\}$), ℓ is the position of the first packet of the $(k - 1)$ th flow.

The packets v_1 and v_2 are both counted (and both are assigned non-zero adjusted selectivities) if and only if they both appear before the ℓ th packet of the induced permutation on $w(F) \setminus \{v_1, v_2\}$, which is

equivalent to saying that their positions in the permutation over $w(F)$ is at most $\ell + 1$. This happens with probability $\frac{\ell}{w(F)-1} \frac{\ell+1}{w(F)}$. In this case there are $\ell + 1$ counted packets and each packet is assigned adjusted selectivity $1/(\ell + 1)$. Therefore, in this subspace

$$E(R^{\text{ANF}}(v_1)R^{\text{ANF}}(v_2)) = \frac{\ell(\ell + 1)}{w(F)(w(F) - 1)} \frac{1}{(\ell + 1)^2} = \frac{1}{w(F)} \frac{\ell}{\ell + 1} \frac{1}{w(F) - 1} \leq \frac{1}{w(F)^2}.$$

(The last inequality follows using $\ell \leq w(F) - 1$.) \square

ASH adjusted weights We can use a dominance relation to define the adjusted selectivities for $L \in \{\text{ASH}, \text{SSH}, \text{SNF}\}$. The relation we used for adjusted weights emulates an ANF sketch from L 's counts and sampling rate (effective sampling rate or sampling rate steps). Correct adjusted selectivities that do not depend on $w(F)$ can be obtained by taking the expectation of the ANF adjusted selectivities over the corresponding subspace of rank assignments.

Done this way, the adjusted selectivities assignment depends on the counts and sampling rate. We obtain better estimators (lower variances) using an assignment that is based on a coarser partition of the sample space. The assignment depends on $w(F)$ and the counts but not on the sampling rate.

Consider an ASH sketch, let $n : F$ be the ASH counts and let $\Omega^{(n:F)}$ be the probability subspace of all rank assignments with ASH counts $n : F$.

Lemma 9.9. *The following is a correct adjusted selectivities assignment*

$$R^{\text{ASH}}(f) = \frac{1}{w(F)} \left(n(f) + \frac{w(F) - n(F)}{k} \right).$$

$R^{\text{ASH}}(f)$ is equal to the expectation of $R^{\text{ANF}}(f)$ over $\Omega^{(n:F)}$ and $A^{\text{ASH}}(f) = w(F)R^{\text{ASH}}(f)$ is equal to the expectation of $A^{\text{ANF}}(f)$ over $\Omega^{(n:F)}$.

Proof. We partition the packets of F into three groups $R \cup S \cup O$ as follows. R is the set of first-counted packets of the sketched flows (we have $|R| = k$), S is the set of subsequent counted packets ($|S| = n(F) - k$), and $O = F \setminus S \setminus R$ ($|O| = w(F) - n(F)$) is the set of ‘‘uncounted’’ packets. We obtain a particular ASH count (the rank assignment is in $\Omega^{(n:F)}$) if and only if the permutation over the packets has the property that the packets of R have lower ranks than the packets of O .

We consider the distribution of ANF counts in the subspace $\Omega^{(n:F)}$. We will compute the expectation of $A^{\text{ANF}}(v)$ for all ASH-counted packets. In all rank assignments in $\Omega^{(n:F)}$, the packets of R are counted by ANF and the packets of O are not. A packet from S is counted by ANF if and only if its position in the permutation over F is before all packets of O . The total number of permutations over F that lie in $\Omega^{(n:F)}$ are

$$|S|!|R|!|O|! \binom{w(F)}{|R| + |O|} = \frac{k!(w(F) - n(F))!w(F)!}{(w(F) - n(F) + k)!}. \quad (30)$$

The number of permutations over F with exactly $\ell \in \{0, \dots, n(F) - k\}$ ANF-counted packets from S (that is, total of $\ell + k$ ANF-counted packets) is

$$\binom{|S|}{\ell} (\ell + k)! |O|! (|S| - \ell + |O| - 1)! = \binom{n(F) - k}{\ell} (\ell + k)! (w(F) - n(F)) (w(F) - k - \ell - 1)! \quad (31)$$

The probability in $\Omega^{(n:F)}$ that exactly ℓ packets from S are counted is the ratio of Eq.(31) and Eq.(30):

$$p_\ell(w(F), n(F), k) = \frac{(n(F) - k)! (\ell + k)! (w(F) - k - \ell - 1)! (w(F) - n(F) + k)!}{\ell! (n(F) - k - \ell)! k! (w(F) - n(F) - 1)! w(F)!}$$

Consider the subspace of $\Omega^{(n:F)}$ that includes all permutation such that there are ℓ counted packets from S . The probability that a packet $c \in S$ is ANF-counted is $\ell/(n(F) - k)$. If counted, it is assigned adjusted selectivity $1/(\ell + k)$. Therefore, the corresponding ASH-adjusted selectivity of these packets is $R'^{\text{ASH}}(c) = \frac{\ell}{(n(F)-k)(\ell+k)}$. The probability that a packet $c \in R$ is ANF-counted is 1 and therefore $R'^{\text{ASH}}(c) = \frac{1}{\ell+k}$. Hence, for a flow f ,

$$R'^{\text{ASH}}(f) = \sum_{\ell=0}^{n(F)-k} p_{\ell}(w(F), n(F), k) \frac{n(F) - k + \ell(n(f) - 1)}{(n(F) - k)(\ell + k)}.$$

The equality $R'^{\text{ASH}}(f) \equiv R^{\text{ASH}}(f)$ follows by standard manipulations.

We next consider the expectation $A'^{\text{ASH}}(f)$ of $A^{\text{ASH}}(f)$ over this subspace. Consider the subspace of all permutation such that there are ℓ counted packets from S . The distribution of the inverse sampling rate is that of the the $(k + \ell + 1)$ smallest of $w(F)$ independent random variables from $U[0, 1]$. The expectation is therefore $w(F)/(k + \ell)$ and the expectation of $A^{\text{ASH}}(f)$ is $n(f) - 1 + w(F)/(k + \ell)$. Therefore,

$$A'^{\text{ASH}}(f) = \sum_{\ell=0}^{n(F)-k} p_{\ell}(w(F), n(F), k) \left(n(f) - 1 + \frac{w(F)}{k + \ell} \right).$$

The relation $A'^{\text{ASH}}(f) = A^{+\text{ASH}}(f)$ follows by standard manipulations. \square

The intuition for this expression is to view the process as the packets in $O \cup R$ being subjected to uniform sampling of k packets, and therefore, each packet obtains an adjusted selectivity of $1/k$ (in $R \cup O$) and adjusted weight of $|O \cup R|/k$ and selecting all packets in S and therefore each packet obtains an adjusted selectivity of $1/|S|$ (in S) and adjusted weight of 1. Therefore, $n(f) - 1 + |O \cup R|/k$ are adjusted weights for f .

It follows from Lemma 9.9 and Lemma 9.4 that $\text{VAR}(A^{+\text{ASH}}(f)) \leq \text{VAR}(A^{\text{ASH}}(f))$ and $\text{VAR}(A^{+\text{ASH}}(f)) \leq \text{VAR}(A^{+\text{ANF}}(f))$ for all $f \in F$. Using a similar proof to Lemma 9.8 we can show that for all $f_1 \neq f_2$, $\text{COV}(A^{+\text{ASH}}(f_1), A^{+\text{ASH}}(f_2)) \leq 0$.

Adjusted count and FSD ASH estimators. The expectation of the inverse sampling rate over $\Omega^{(n:F)}$ is

$$\sum_{\ell=0}^{n(F)-k} p_{\ell}(w(F), n(F), k) w(F)/(k + \ell) = \frac{w(F) - n(F) + k}{k}.$$

The ASH adjusted counts and FSD estimators depend on the counts and linearly on $1/p$. Their expectation over $\Omega^{(n:F)}$ constitutes a tighter unbiased estimator. This expectation is obtained by replacing $1/p$ with its expectation $\frac{w(F) - n(F) + k}{k}$.

10 Conclusion

Accurate summarization of IP traffic is essential for many network operations. We design summarization algorithms that generate a sketch of the packet streams that allows us to process approximate subpopulation-size queries and other aggregates. Our algorithms build on existing designs, but are yet able to obtain significantly better estimates through better utilization of available resources and careful derivation of unbiased statistical estimators that have minimum variance with respect to the information they use.

References

- [1] E. Cohen. Size-estimation framework with applications to transitive closure and reachability. *J. Comput. System Sci.*, 55:441–453, 1997.
- [2] E. Cohen, N. Duffield, H. Kaplan, C. Lund, and M. Thorup. Algorithms and estimators for accurate summarization of Internet traffic. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement (IMC)*, 2007.
- [3] E. Cohen, N. Duffield, H. Kaplan, C. Lund, and M. Thorup. Sketching unaggregated data streams for subpopulation-size queries. In *Proc. of the 2007 ACM Symp. on Principles of Database Systems (PODS 2007)*. ACM, 2007.
- [4] E. Cohen and H. Kaplan. Bottom-k sketches: Better and more efficient estimation of aggregates. In *Proceedings of the ACM SIGMETRICS'07 Conference*, 2007. poster.
- [5] E. Cohen and H. Kaplan. Sketches and estimators for subpopulation weight queries. Manuscript, 2007.
- [6] E. Cohen and H. Kaplan. Spatially-decaying aggregation over a network: model and algorithms. *J. Comput. System Sci.*, 73:265–288, 2007.
- [7] E. Cohen and H. Kaplan. Summarizing data using bottom-k sketches. In *Proceedings of the ACM PODC'07 Conference*, 2007.
- [8] C. Cranor, T. Johnson, V. Shkapenyuk, and O. Spatcheck. Gigascope: A stream database for network applications. In *Proceedings of the ACM SIGMOD*, 2003.
- [9] N. Duffield, C. Lund, and M. Thorup. Estimating flow distributions from sampled flow statistics. In *Proceedings of the ACM SIGCOMM'03 Conference*, pages 325–336, 2003.
- [10] N. Duffield, M. Thorup, and C. Lund. Flow sampling under hard resource constraints. In *Proceedings the ACM IFIP Conference on Measurement and Modeling of Computer Systems (SIGMETRICS/Performance)*, pages 85–96, 2004.
- [11] C. Estan, K. Keys, D. Moore, and G. Varghese. Building a better netflow. In *Proceedings of the ACM SIGCOMM'04 Conference*. ACM, 2004.
- [12] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proceedings of the ACM SIGCOMM'02 Conference*. ACM, 2002.
- [13] M. Gibbons and Y. Matias. New sampling-based summary statistics for improving approximate query answers. In *SIGMOD*. ACM, 1998.
- [14] J. M. Hellerstein, P. J. Haas, and H. J. Wang. Online aggregation. In *Proceedings of the ACM SIGMOD*, 1997.
- [15] N. Hohn and D. Veitch. Inverting sampled traffic. In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, pages 222–233, 2003.
- [16] K. Keys, D. Moore, and C. Estan. A robust system for accurate real-time summaries of Internet traffic. In *Proceedings of the ACM SIGMETRICS'05*. ACM, 2005.
- [17] A. Kumar, M. Sung, J. Xu, and E. W. Zegura. A data streaming algorithm for estimating subpopulation flow size distribution. *ACM SIGMETRICS Performance Evaluation Review*, 33, 2005.
- [18] S. Ramabhadran and G. Varghese. Efficient implementation of a statistics counter architecture. In *Proc. of ACM Sigmetrics 2003*, 2003.
- [19] B. Ribeiro, D. Towsley, T. Ye, and J. Bolot. Fisher information of sampled packets: an application to flow size estimation. In *Proceedings of the 2006 Internet Measurement Conference*. ACM, 2006.
- [20] D. Shah, S. Iyer, B. Prabhakar, and N. McKeown. Maintaining statistics counters in router line cards. *IEEE Micro*, 22(1):76–81, 2002.
- [21] M. Szegedy and M. Thorup. On the variance of subset sum estimation. In *Proc. 15th ESA*, 2007.