

A High-Performance Algorithm for Identifying Frequent Items in Data Streams

Daniel Anderson
Georgetown University

Pryce Bevan
Georgetown University

Kevin Lang
Oath Research

Edo Liberty*
Amazon

Lee Rhodes
Oath

Justin Thaler
Georgetown University

ABSTRACT

Estimating frequencies of items over data streams is a common building block in streaming data measurement and analysis. Misra and Gries introduced their seminal algorithm for the problem in 1982, and the problem has since been revisited many times due to its practicality and applicability. We describe a highly optimized version of Misra and Gries’ algorithm that is suitable for deployment in industrial settings. Our code is made public via an open source library called Data Sketches that is already used by several companies and production systems.

Our algorithm improves on two theoretical and practical aspects of prior work. First, it handles *weighted* updates in amortized constant time, a common requirement in practice. Second, it uses a simple and fast method for merging summaries that asymptotically improves on prior work even for unweighted streams. We describe experiments confirming that our algorithms are more efficient than prior proposals.

CCS CONCEPTS

• **Theory of computation** → **Streaming, sublinear and near linear time algorithms**; • **Networks** → *Network measurement*;

KEYWORDS

streaming algorithms, mergeable summaries, frequent items

ACM Reference Format:

Daniel Anderson, Pryce Bevan, Kevin Lang, Edo Liberty, Lee Rhodes, and Justin Thaler. 2017. A High-Performance Algorithm for Identifying Frequent Items in Data Streams. In *Proceedings of IMC ’17*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3131365.3131407>

1 INTRODUCTION

Identifying frequent items (also known as heavy hitters) and answering point queries (i.e., queries of the form “approximately how many times did item i appear in the stream?”) are two of the most basic computational tasks performed on data streams. Due to their practical importance, streaming algorithms for these tasks have

*Research performed while at Yahoo Research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC ’17, November 1–3, 2017, London, United Kingdom

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5118-8/17/11...\$15.00

<https://doi.org/10.1145/3131365.3131407>

been studied intensely [6, 7, 9, 13, 14, 17, 21, 31–35]. These algorithms process a massive dataset in a single pass, and compute very small *summaries* of the dataset, from which it is possible to derive accurate—though approximate—answers to frequent items queries and point queries.

It may seem as though streaming frequency approximation is well-understood, with little room for further insight or improvement. However, when we set about implementing an algorithm suitable for industrial use on web-scale data, we found that existing algorithms have two significant shortcomings. First, they are not very efficient when handling *weighted* stream updates (i.e., data streams where each piece of data comes with an associated importance measure). Second, although existing algorithms allow for *merging* summaries produced from different streams to obtain a summary for the union of the streams, existing merging procedures are slow and space intensive. In this work, we describe an optimized streaming algorithm and a merging procedure that together address both of these shortcomings.

1.1 Problem Statement and Applications

Given a data stream σ of length n over a universe $[m] = \{1, \dots, m\}$, where each stream update is of the form (i_j, Δ_j) for some $i_j \in [m]$ and $\Delta_j > 0$ is a weight, let $f_i = \sum_{j: i_j=i} \Delta_j$ denote the (weighted) frequency of i in the stream and let $N = \sum_j \Delta_j$ denote the weighted stream length. The goal of an algorithm for answering point queries is the following. In a single pass over the stream, the algorithm must compute a small summary (sometimes also called a *sketch*) of the stream such that, for any item $i \in [m]$, it is possible to efficiently derive an estimate \hat{f}_i such that $|\hat{f}_i - f_i|$ is small.

Ideally, the algorithm should also be able to identify (ϕ, ϵ) -heavy hitters (in the ℓ_1 -norm). That is, given a user-specified threshold ϕ , the algorithm should be able to identify all items i with $f_i \geq \phi \cdot N$, possibly also returning a small number of “false positive” items whose frequency is slightly below the specified threshold ϕN .

The Importance of Weighted Updates. The majority of the literature on streaming approximate frequency analysis considers the case of unweighted updates, where $\Delta_j = 1$ for all j . However, there are many applications in which the ability to efficiently handle weighted updates is critical. For instance, much of the literature on streaming algorithms for identifying frequent items and answering point queries has been motivated by applications in network traffic measurement, many of which naturally involve weighted updates.

For example, it is often useful to track the total amount of data that individual users, or pairs of users, send over a network. This is essential for accounting and pricing, anomaly detection (such as

detecting hotspots, worms, and DDoS attacks), and traffic engineering [18, 19, 21–24, 28, 38, 39, 46, 48]. In this setting, the weight of each stream update (i.e., packet) is its size in bytes.

Alternatively, it may be useful to track the total duration of network flows, or the total time that users spend using a mobile app or watching online video. In these settings, the weight of a stream update is the duration of the event in question.

As another example, in spam detection, it is useful to track the total number of recipients of emails from a given address, and many emails have an enormous number of recipients. For this reason, Yahoo is using (weighted) frequent items algorithms in back-end systems for spam detection and email threading [3].

More generally, weighted updates arise whenever items from the data universe are associated with importance factors. For example, in information-retrieval and text mining, term frequency–inverse document frequency (tf-idf for short) is a standard statistic that assigns a value to each word in a document that is part of a larger corpus. The value assigned to a word increases with the number of times a word appears in the document, but decreases with the frequency of the word in the corpus, in order to account for the fact that some words appear more frequently in general.

In addition to the applications above, we will also see (cf. Section 3.2) that the ability to efficiently handle weighted updates enables extremely efficient *merging* of summaries—significantly faster than merging procedures proposed in prior work. Moreover, unlike algorithms tailored to unweighted updates, algorithms capable of handling weighted updates typically apply to real-valued weights. This will be the case for the algorithms we give in this work.

There are a number of simple ways to modify algorithms for the unweighted setting to handle the weighted setting. However, as we describe in Sections 1.3.4–1.3.5, all of the modifications that (to our knowledge) have been proposed in prior work have significant shortcomings.

The Importance of Mergeability. Ideally, streaming algorithms will produce summaries that are *mergeable*, meaning that one can process many different streams of data independently, and then the summaries computed from each stream can be quickly combined to obtain an accurate summary for the union of the datasets [2]. Mergeable summaries are useful in a wide variety of settings. For example, they enable massive datasets to be *automatically* processed in a fully distributed and parallel manner, by partitioning the data arbitrarily across many machines, summarizing each partition, and seamlessly combining the results. Another application is to power conservation and latency minimization in weak peripheral devices. For example, one of the key benefits of the Internet of Things (IoT) is that it enables the monitoring and aggregation of data from IoT devices such as sensors and appliances. Such devices are often power-limited, so it is essential to minimize the amount of data that must be sent from each device to the aggregation center. Mergeable summaries enable this: each device can itself compute a summary of its own data, and send only the summary to the aggregation center, which merges all received summaries to obtain a global summary for all devices' datasets. (See Section 3 for additional applications of mergeable summaries).

1.2 Our Contributions

Improved Streaming Algorithm for Weighted Updates. Our algorithm is similar to the popular Misra–Gries (MG) [35] and Space Saving (SS) [34] algorithms. However, it differs from them in several simple yet important ways, enabling our algorithm to process weighted updates in a highly efficient manner. Our algorithm runs in amortized constant time per stream update, and employs optimized data structures to minimize its memory footprint while maximizing throughput. We present experimental results validating the superiority of our algorithms over existing solutions.

Our experiments also clarify conventional wisdom regarding the best way to implement frequent items algorithms in practice. Prior literature strongly suggested that the best algorithm in practice is a min-heap based implementation of SS (we denote this implementation of SS by MHE—short for min-heap extension—cf. Section 1.3.5). Indeed, experimental work of Cormode and Hadjieleftheriou [14] that focused exclusively on the setting of unweighted streams identified SS as state of the art. They showed that a linked-list based implementation of SS (denoted in [14] as SSL) is noticeably faster than the min-heap based implementation, but also significantly more space intensive. Moreover, SSL does not naturally extend to weighted updates, while MHE does. Based on these findings, subsequent work (e.g., [36]) in the weighted setting used MHE as the implementation of choice. We overturn this conventional wisdom: our implementations do not use a min-heap, yet are significantly faster and more space efficient than MHE.

Improved Merging Procedure. We give a simple and efficient method for merging summaries produced by our algorithm, or more generally, any *counter-based* algorithm (cf. Section 1.3 for a definition of this class of algorithms). Our merging procedure significantly improves on prior work in terms of both speed and memory usage, even for unweighted streams.

Other Implications of Our Improvements. We note that algorithms for answering point queries and identifying heavy hitters are often used in a black-box fashion as subroutines in streaming algorithms solving more complicated problems. These include estimating the *empirical entropy* of a data stream [12], and identifying *hierarchical heavy hitters* [36], both of which have important applications in network traffic monitoring and anomaly detection [16, 25, 44–47].

For example, the Hierarchical Heavy Hitters (HHH) problem extends the notion of frequent items to hierarchical data. To motivate the problem definition, suppose that a single entity controls all IP addresses in a subnet (say, all addresses of the form 021.132.145.*, where * is a wildcard byte, to borrow an example from [36]). It is possible for the controlling entity to spread out traffic among the addresses in the subset so that no single IP address is a heavy hitter. Nonetheless, a network manager may want to know if the controlling entity itself is a heavy hitter (i.e., if the sum of the traffic of all IP addresses in the subnet exceeds a specified threshold). The HHH problem [5, 15, 16, 20, 26, 30, 36, 41, 43, 46] is defined to capture this scenario.

State of the art algorithms [5, 36] for the HHH problem maintain one instance of an algorithm for identifying heavy hitters for each level of the hierarchy. At query time, these algorithms identify

HHs by combining, in a simple manner, the information stored in all of the summaries. We omit further details for brevity.

Our optimized algorithms for answering point queries and identifying heavy hitters can be directly substituted without modification into these more complicated algorithms. While we leave this task to future work, we expect that our optimizations would yield similar speed and memory improvements in these settings.

1.3 Overview of Prior Work

In a comprehensive paper and survey, Cormode and Hadjieleftheriou [14] provide a detailed comparison of proposed algorithms for finding heavy hitters and answering point queries in the case of unit weight updates. They classified algorithms into three classes: counter-based algorithms, quantile algorithms, and (linear) sketches. They found that counter-based algorithms perform significantly better in terms of space, speed, and accuracy than quantile and sketching algorithms, a finding that we confirmed in our own initial experiments. Hence, we focus on counter-based algorithms for the remainder of this work.

1.3.1 Counter-Based Algorithms. A counter-based algorithm stores k counters, where k is a parameter of the algorithm that controls both the space usage and the error (larger k corresponds to more space and less error). Each counter stores an approximate count for some item of the data universe; the various counter-based algorithms differ in how they assign counters to items and how they determine the approximate counts.

Cormode and Hadjieleftheriou [14] found that two counter-based algorithms are state of the art: Misra-Gries (MG) [35] and Space Saving (SS) [34]. They found that the estimates returned by **SS** tend to be more accurate than those returned by **MG** – we discuss the reasons for the improved accuracy of **SS** relative to **MG** in Section 2.3.1. However, subsequent work by Agarwal et al. [2] observed that these two algorithms are actually isomorphic, in the following sense. The estimates returned by the **SS** algorithm with $k + 1$ counters can be derived from the summary computed by the **MG** algorithm with k counters. That is, the summaries computed by both **SS** and **MG** contain essentially the same information about the stream, but the algorithms differ in how they use this information to determine a “best estimate” for any point query.

Our algorithm can be naturally viewed as an extension and modification of **MG**. Hence, we now describe **MG** in detail.

1.3.2 Description of the Misra-Gries Algorithm For Unit Weight Updates. Every time the algorithm processes a stream update $(i, +1)$, the **MG** algorithm looks to see if i is assigned a counter, and if so it increments the counter. If not, and an unassigned counter exists, the algorithm assigns the counter to i and sets the approximate count to 1. If no unassigned counter exists, the algorithm decrements all counters by 1, and marks all counters set to 0 as unassigned. When asked to provide an estimate \hat{f}_i for the frequency of item f_i , the algorithm returns 0 if i is not assigned a counter, and returns the value of the counter assigned to i otherwise.

See Algorithm 1 for pseudocode. The following standard lemma bounds the error of any estimate returned by the **MG** algorithm [35], and we provide the proof for completeness.

Algorithm 1 Misra-Gries Algorithm for Unit Weight Updates

```

1: Algorithm: Misra-Gries( $k$ ):
2:    $T \leftarrow \emptyset$  //  $T$  is set of items assigned a counter
3:   Function Update( $i, +1$ ):
4:     if  $i \in T$ :
5:        $c(i) \leftarrow c(i) + 1$ 
6:     else if  $|T| < k$ :
7:        $T = T \cup \{i\}$ 
8:        $c(i) \leftarrow 1$ 
9:     else:
10:      DecrementCounters()
11:   Function DecrementCounters():
12:     for all  $j \in T$ :
13:        $c(j) = c(j) - 1$ 
14:     if  $c(j) = 0$ :
15:        $T = T \setminus \{j\}$ 
16:   Function Estimate( $i$ ):
17:     if  $i \in T$ 
18:       return  $c(i)$ 
19:     else
20:       return 0

```

LEMMA 1.1. *The **MG** algorithm with k counters is guaranteed to return, for each $i \in [n]$, and estimate \hat{f}_i satisfying $0 \leq f_i - \hat{f}_i \leq N/(k + 1)$.*

PROOF. It is obvious from the description of the algorithm that $0 \leq f_i - \hat{f}_i$ for all i . To prove that $f_i - \hat{f}_i \leq N/k$ for all i , observe that $f_i - \hat{f}_i$ is at most the total number of decrement operations over the course of the algorithm, and there can be at most N/k decrement operations. Indeed, if there are d decrement operations, then since each decrement operation affects all k counters, the sum of the counter values at the end of the algorithm is exactly $N - d \cdot (k + 1)$. Since no counter value is ever negative, it holds that $N - d \cdot (k + 1) \geq 0$, and hence $d \leq N/(k + 1)$. \square

Using a similar proof, Berinde et al. [6] show that the **MG** algorithm also satisfies the following guarantee. This guarantee, referred to by Berinde et al. as a *tail guarantee*, is much stronger than Lemma 1.1 for streams with a very skewed frequency distribution (i.e., with a relatively small number of frequent items constituting the bulk of the stream).

LEMMA 1.2 (BERINDE ET AL. [6]). *Let $N^{\text{res}(j)}$ denote the sum of the frequencies of all but the top j most frequent items. For any $j < k$, the **MG** algorithm with k counters is guaranteed to return, for each $i \in [n]$, and estimate \hat{f}_i satisfying $0 \leq f_i - \hat{f}_i \leq N^{\text{res}(j)}/(k + 1 - j)$.*

Efficiently Implementing Misra-Gries for Unit Weight Updates. A natural way to implement the **MG** algorithm is to maintain a hash table that stores all assigned counters. The key for a counter is the item it is assigned to, and the value is the count. Whenever a stream update $(i, +1)$ arrives, the algorithm looks up key i in the hash table, and if it is present, it increments the corresponding value (i.e., count). If not, the algorithm checks to see if the hash table is storing k key-value pairs. If not, the key-value pair $(i, 1)$ is inserted into the hash table. If so, the algorithm iterates over all key-value pairs in the hash table, decrementing each value and deleting any pair if the value becomes zero. Assuming all hash table operations

Algorithm 2 Space Saving Algorithm for Unit Weight Updates

```

1: Algorithm: SpaceSaving(k):
2:    $T \leftarrow \emptyset$  //  $T$  is set of items assigned a counter
3:   Function Update( $i$ , +1):
4:     if  $i \in T$ :
5:        $c(i) \leftarrow c(i) + 1$ 
6:     else if  $|T| < k$ :
7:        $T = T \cup \{i\}$ 
8:        $c(i) \leftarrow c(i) + 1$ 
9:     else :
10:       $j \leftarrow \arg \min_{j \in T} c(j)$ 
11:       $c(i) \leftarrow c(j) + 1$ 
12:       $T = (T \setminus \{j\}) \cup \{i\}$ 
13:   Function Estimate( $i$ ):
14:     if  $i \in T$ :
15:       return  $c(i)$ 
16:     else
17:       return  $\min_{j \in T} \{c(j)\}$ 

```

(i.e., lookup, insert, and delete) take $O(1)$ amortized time, and that enumerating all key-value pairs in the hash table can be done in time $O(k)$, then this implementation runs in amortized $O(1)$ time per stream update. This amortized time bound exploits the fact that decrement operations, which require enumerating all k key-value pairs, only occur at most N/k times.

1.3.3 Description of the Space Saving Algorithm For Unit Weight Updates. When processing a stream update $(i, +1)$, the **SS** algorithm behaves identically to **MG**, except in the case where i is not assigned a counter, and no unassigned counter exists. In this case, **SS** increments the counter with the smallest value and assigns it to i . When asked to provide an estimate \hat{f}_i for item the frequency of item f_i , the algorithm returns the value of the smallest counter if i is not assigned a counter, and returns the value of the counter assigned to i otherwise. See Algorithm 2 for pseudocode.

Efficiently Implementing Space Saving for Unit Weight Updates. The **SS** algorithm can be naturally implemented by using a min-heap data structure to track the smallest counter at all times. Following [14], we refer to this implementation as **SSH**. Unfortunately, the use of a min-heap is slow (taking time $O(\log k)$ per stream update), and it also nearly doubles the space usage compared to the **MG** implementation described above, since both a hash table of capacity k and a min-heap must be stored.

In the case of unit updates, Metwally et al. [34] propose a doubly linked list based data structure they call *Stream Summary* that can process updates in $O(1)$ time. However, the need to store pointers in the linked list will more than double the space usage of the Misra-Gries implementation described above. Moreover, as mentioned again in Section 1.3.5 below, this implementation method does not naturally extend to weighted updates.

1.3.4 Extending Misra-Gries to the Weighted Case: Prior Work and Its Limitations. There are a number of simple ways to extend the **MG** algorithm to handle weighted updates. Perhaps the simplest is to treat an update (i_j, Δ_j) as Δ_j unit updates. We refer to this as the *Reduce-To-Unit-Case* (**RTUC-MG**) extension of the **MG** algorithm. However, this takes time at least Δ_j per stream update, which is

unacceptable when the weights may be large (this approach is also awkward to apply to the case of non-integer weights).

To our knowledge, the only other proposal to extend the **MG** algorithm to weighted updates is due to Berinde et al. [6], who suggest the following modification. Every time the algorithm processes a stream update (i_j, Δ_j) , the algorithm looks to see if i_j is assigned a counter, and if so it increments the counter by Δ_j . If not, and an unassigned counter exists, the algorithm assigns the counter to i and sets the approximate count to Δ_j . If no unassigned counter exists, the algorithm's behavior depends on whether $\Delta_j \leq c_{\min}$, where c_{\min} denotes the minimum value of any counter. If $\Delta_j \leq c_{\min}$, then all stored counters are reduced by Δ_j . Otherwise, all counters are reduced by c_{\min} , and some counter with zero count (there must be at least one now) is assigned to i_j and given count $\Delta_j - c_{\min}$. We refer to this as the *Reduce-By-Min-Counter* (**RBMC**) extension of the **MG** algorithm to weighted updates.

It is easy to see that the **RBMC** algorithm produces estimates identical to the **RTUC-MG** algorithm, and hence satisfies the accuracy guarantees of Lemmas 1.1 and 1.2. The advantage of the **RBMC-MG** algorithm is in its runtime, which does not grow linearly with Δ_j . The main shortcoming of the **RBMC** algorithm is that it still may not run in amortized $O(1)$ time per stream update. In fact, there are streams on which the proposal will perform expensive decrement operations on essentially every stream update: consider a stream where the first k updates increment the counts of distinct items by an arbitrarily large number M , and then the following M stream updates are unit updates to different items. Decrement operations, each requiring time $\Theta(k)$, will be performed on every one of the last M updates in the stream. While such examples may be contrived, Section 4 provides experimental results on real datasets showing that the runtime of this algorithm is significantly higher in practice than the alternatives presented in this paper.

1.3.5 Extending Space Saving to the Weighted Case: Prior Work and Its Limitations. The **SS** algorithm also has a natural *Reduce-To-Unit-Case* (**RTUC-SS**) extension to weighted updates. However, like the **RTUC-MG** algorithm, this takes time at least Δ_j per stream update, which is unacceptable when the weights may be large.

The min-heap based implementation, **SSH**, of **SS** naturally extends to the weighted case. We refer to this as the *Min-Heap-Extension* (**MHE**) extension of the **SS** algorithm to weighted updates. However, **MHE** suffers from the same poor (i.e., $O(\log k)$) update time and larger space usage (relative to **RBMC**) as in the unit weight update case. Despite these shortcomings, this implementation of **SS** for weighted updates was used in at least one prior work on computing Hierarchical Heavy Hitters [36]. The implementation of Space Saving that uses the *Stream Summary* data structure does not naturally extend to weighted updates [14].

We discuss additional prior work in Section 5.

1.4 Isomorphism Results for Weighted Streams

We mention that the **MHE** and **RBMC** algorithms are isomorphic, in the sense that the estimates returned by the **MHE** algorithm with $k + 1$ counters can be derived from the summary computed by the **RBMC** algorithm with k counters. This follows immediately from the facts that both the **MHE** and **RBMC** produce the same

estimates as **RTUC-SS** and **RTUC-MG** respectively, and the latter two algorithms were shown to be isomorphic by Agarwal et al. [2].

2 OUR ALGORITHM

For expository purposes, we first propose an initial modification of the **MG** algorithm that achieves (up to a constant factor) the same accuracy guarantees as the **RBMC** algorithm, while guaranteeing $O(1)$ amortized runtime per stream update. We then describe our final algorithm, which improves over the runtime and space usage of the initial modification by a factor close to 2.

2.1 An Initial Proposal

For simplicity, we assume throughout this section that k is even. Recall (cf. Section 1.3.4) that the main downside of the **RBMC** algorithm [6] is that it does not guarantee $O(1)$ amortized runtime per stream update. The reason was that **RBMC** may have to perform decrement operations essentially every stream update, and each decrement operation requires iterating over k counters.

To guarantee $O(1)$ amortized time per stream update, it is enough to ensure that decrement operations are performed at most once every (say) $k/2$ stream updates. We mention that a similar technique was used by Liberty to efficiently identify “frequent directions” in a stream of vectors [29].

To ensure that this property holds, it suffices to consider the following simple modification of **RBMC**. Every time the algorithm processes a stream update (i_j, Δ_j) , the algorithm looks to see if i_j is assigned a counter, and if so it increments the counter by Δ_j . If not, and an unassigned counter exists, the algorithm assigns the counter to i_j and sets the approximate count to Δ_j . If no unassigned counter exists, the algorithm decrements all counters by the *median* counter value c_{median} . The algorithm then marks all counters set to a non-positive value as unassigned. If Δ_j is larger than c_{median} , then i_j is then assigned a counter, which is set to $\Delta_j - c_{\text{median}}$.

We refer to this algorithm as the *Reduce-By-Median-Counter (MED)* extension of the **MG** algorithm to weighted updates. See Algorithm 3 for pseudocode, which is expressed in slightly more general form. Specifically, in Algorithm 3, c_{median} is replaced by the the k^* -th largest counter value, where k^* is a parameter of the algorithm. One recovers the description above by setting $k^* = k/2$.

2.1.1 Runtime and Accuracy Analysis.

LEMMA 2.1. *In Algorithm 3, a `DecrementCounters()` operation (Lines 15–20) is performed at most once every k^* stream updates.*

PROOF. Observe that `DecrementCounters()` operations are performed only when all k counters are assigned. Every time a `DecrementCounters()` operation is executed, Line 18 sets at least k^* counters to a non-positive value, and all such counters become unassigned in Line 20. Since any `Update(i, Δ)` operation assigns at most one counter to i , it follows that after a `DecrementCounters()` operation, at least k^* `Update()` operations must occur before all counters once again become assigned. \square

Lemma 2.1 is easily seen to imply that Algorithm 3 can be implemented in $O(1)$ amortized time per stream update. The main observation is that c_{k^*} (cf. Line 16) can be found in time $O(k)$ using the Quickselect algorithm [27]. Hence, decrement operations can be

Algorithm 3 The Reduce-By-Median-Counter (**MED**) Extension of the **MG** Algorithm to Weighted Updates

```

1: // Notation:  $k^*$  is a parameter of the algorithm
2: Algorithm: Initial-Algorithm( $k$ ):
3:    $T \leftarrow \emptyset$  //  $T$  is set of items assigned a counter
4:   Function Update( $i, \Delta$ ):
5:     if  $i \in T$ :
6:        $c(i) \leftarrow c(i) + \Delta$ 
7:     else if  $|T| < k$ :
8:        $T = T \cup \{i\}$ 
9:        $c(i) \leftarrow c(i) + \Delta$ 
10:    else :
11:      DecrementCounters()
12:      if  $\Delta \geq c_{k^*}$  // See Line 16 for definition of  $c_{k^*}$ 
13:         $T = T \cup \{i\}$  //  $|T| \leq k^* + 1 \leq k$  after this line
14:         $c(i) \leftarrow c(i) + \Delta - c_{k^*}$ 
15:   Function DecrementCounters():
16:     // Notation: Let  $c_{k^*}$  be the  $k^*$ -largest value, counting multiplicity,
17:     // in the multiset  $\{c(j) : j \in T\}$ .
18:     for all  $j \in T$ :
19:        $c(j) = c(j) - c_{k^*}$ 
20:       if  $c(j) \leq 0$ :
21:          $T = T \setminus \{j\}$ 
22:   Function Estimate( $i$ ):
23:     if  $i \in T$ :
24:       return  $c(i)$ 
25:     else
26:       return 0

```

performed in $O(k)$ time, and by Lemma 2.1, decrement operations occur at most once every $k^* = \Omega(k)$ stream updates.

THEOREM 2.2. *By maintaining all k counters in a hash table of capacity supporting (amortized) constant time insertions, deletions, and lookups, and enumeration of all counters in $O(k)$ time, Algorithm 3 can be implemented in amortized constant time as long as the parameter $k^* \geq \Omega(k)$.*

It is also not difficult to show that this implementation satisfies the same error guarantee as Lemma 1.1 up to a factor of 2. In fact, we prove the following *tail guarantee* that is significantly stronger than Lemma 1.1 for streams with skewed frequency distributions. Our analysis exploits the key ideas of earlier analyses (of the **MG** algorithm) by Agarwal et al. [2] and Berinde et al. [6].

THEOREM 2.3. *For any $j < k^*$, Algorithm 3 returns, for each $i \in [n]$, an estimate \hat{f}_i satisfying $0 \leq f_i - \hat{f}_i \leq N^{\text{res}(j)} / (k^* - j)$.*

PROOF. It is obvious from the description of the algorithm that $0 \leq f_i - \hat{f}_i$ for all i . To prove the second inequality, we will proceed by induction on the number of stream updates n . For clarity, we introduce some notation. Let $f_{i,n} = \sum_{j \leq n: i_j = i} \Delta_j$ denote the (weighted) frequency of item i in the first n stream updates, and similarly let $\hat{f}_{i,n}$ denote the output of `Estimate(i)` after the first n stream updates. Let $E_n = \max_i f_{i,n} - \hat{f}_{i,n}$ denote the maximum error of any estimate that might be returned by the summary after processing n stream updates. Let $C_n = \sum_{i \in T} c(i)$ denote the sum of all counts in the summary after the first n stream updates have been processed. Let $N_n = \sum_{i=1}^n \Delta_i$ denote the weighted stream length after n stream updates.

For expository purposes, we establish the second inequality for $j = 0$, before establishing it for any $j < k^*$.

Establishing the second inequality for $j = 0$. In the case $j = 0$ it is sufficient to show establish the following lemma.

LEMMA 2.4. $E_n \leq (N_n - C_n)/k^*$.

PROOF. We show this by induction on n . Clearly it holds for $n = 0$. Assume by way of induction that it holds for streams consisting of $n - 1$ updates.

Proof Outline for Lemma 2.4. The key observations are (a) that error is only introduced by calls to DecrementCounters() and (b) whenever DecrementCounters() is called, at least k^* counters are reduced by c_{k^*} . Hence, the error in any estimate increases by c_{k^*} , and the difference between N and C increases by at least $k^* \cdot c_{k^*}$. This ensures that the difference between N and C is always at least k^* times the error of the summary, as desired. The rest of the proof makes this argument formal.

Suppose first that the n th stream update (i_n, Δ_n) does not cause DecrementCounters() to be called. In this case, after processing the n th stream update, f_{i_n} and \hat{f}_{i_n} both increase by Δ_n , and for any $i \neq i_n$, f_i and \hat{f}_i are unchanged. Hence, $E_n = E_{n-1}$, i.e., the error of the summary is unchanged by the n th stream update. Meanwhile, N_n and C_n both increase by Δ_n . It follows that $(N_n - C_n) = (N_{n-1} - C_{n-1})$. Putting these facts together and invoking the inductive hypothesis, we conclude as desired that

$$E_n = E_{n-1} \leq (N_{n-1} - C_{n-1})/k^* = (N_n - C_n)/k^*.$$

Next, suppose that the n th stream update does cause DecrementCounters() to be called. We assume w.l.o.g. that $\Delta_n \leq c_{k^*}$; if not, then we instead treat update (i_n, Δ_n) as two separate updates (i_n, c_{k^*}) followed by $(i_n, \Delta_n - c_{k^*})$, as the algorithm behaves identically in this case.

Then the largest k^* counters are each decremented by c_{k^*} , and no counter is incremented (as Lines 13-14 are not executed due to the assumption). Hence,

$$\begin{aligned} N_n - C_n &= N_{n-1} + \Delta_n - C_n \geq N_{n-1} + \Delta_n - (C_{n-1} - k^* \cdot c_{k^*}) \\ &= N_{n-1} - C_{n-1} + k^* \cdot c_{k^*} + \Delta_n \\ &\geq N_{n-1} - C_{n-1} + k^* \cdot c_{k^*}. \end{aligned} \quad (1)$$

Meanwhile, it holds that

$$E_n \leq E_{n-1} + c_{k^*}. \quad (2)$$

To see this, first observe that, for all $i \neq i_n$, $\hat{f}_{i,n} - \hat{f}_{i,n-1} \leq c_{k^*}$, and $f_{i,n} = f_{i,n-1}$. Hence, $f_{i,n} - \hat{f}_{i,n} \leq f_{i,n-1} - \hat{f}_{i,n-1} + c_{k^*}$, i.e., the error in the estimate for i increased by at most c_{k^*} when processing the n th stream update.

To see that Equation (2) also holds for i_n itself, observe first that the fact that DecrementCounters() was called when processing the n th stream update means that i_n was not assigned a counter after $n - 1$ stream updates. Hence, $\hat{f}_{i_n, n-1} = 0$. Meanwhile, the assumption that $\Delta_n \leq c_{k^*}$ then ensures that $\hat{f}_{i_n, n} = 0$ as well. And clearly $f_{i_n, n} = f_{i_n, n-1} + \Delta_n$. Putting these facts together implies that $f_{i_n, n} - \hat{f}_{i_n, n} = f_{i_n, n-1} + \Delta_n - \hat{f}_{i_n, n-1} \leq f_{i_n, n-1} + c_{k^*} - \hat{f}_{i_n, n-1}$. That is, the error in the estimate for i_n increased by at most c_{k^*} when processing the n th stream update. Equation (2) follows.

Combining Equations (1) and (2) with the inductive hypothesis, we conclude as desired that

$$\begin{aligned} E_n &\leq E_{n-1} + c_{k^*} \leq (N_{n-1} - C_{n-1})/k^* + c_{k^*} \\ &\leq (N_n - C_n - k^* c_{k^*})/k^* + c_{k^*} = (N_n - C_n)/k^*. \end{aligned}$$

□

Establishing the second inequality for any $j < k^*$. For notational simplicity, let us assume that $f_1 \geq f_2 \geq \dots \geq f_j$. Lemma 2.4 implies that

$$C_n \leq N_n - k^* \cdot E_n. \quad (3)$$

Meanwhile, it is clear that at the end of the stream, $c(i) \geq f_i - E_n$ for all i , which implies that

$$C_n \geq \sum_{i=1}^j (f_i - E_n) \geq \left(\sum_{i=1}^j f_i \right) - j \cdot E_n. \quad (4)$$

Putting Equations (3) and (4) together, we conclude that

$$\left(\sum_{i=1}^j f_i \right) - j \cdot E_n \leq N_n - k^* \cdot E_n,$$

which in turn implies that $(k^* - j) \cdot E_n \leq N_n - \sum_{i=1}^j f_i = N^{\text{res}(j)}$. Hence, $E_n \leq N^{\text{res}(j)}/(k^* - j)$. This completes the proof. □

2.2 The Final Algorithm

Algorithm 3 has two disadvantages that are important in practice. First, an extra k words of space are required during every DecrementCounters() operation, on top of the $\Theta(k)$ words of space required for that hash table that maintains the counters. This extra space is required in order to find the k^* th largest counter, as the Quickselect algorithm (or any sorting procedure) cannot be done in place without destroying the hash table. This disadvantage nearly doubles the space usage of the algorithm.

Second, during each DecrementCounters() operation, the algorithm must make an extra pass through the summary to find the k^* th largest counter, before decrementing all counters and discarding the non-positive ones. As DecrementCounters() operations is the time bottleneck in practice, this significantly slows the concrete efficiency of the algorithm.

In order to address both of these disadvantages at once, we make the following observation. In neither the proof of Theorem 2.2 (establishing constant amortized update time for Algorithm 3) nor Theorem 2.3 (establishing strong bounds on the error of Algorithm 3) is it crucial that we decrement by the k^* th largest counter value for $k^* = k/2$. The property that ensures constant amortized update time as per Theorem 2.2 is that each DecrementCounters() operation decrements by a value that is *larger than* a constant fraction of the counters. The property exploited to establish the error bounds of Theorem 2.3 is that each DecrementCounters() operations decrements by a value that is *smaller than* a constant fraction of the counters.

Both properties can be ensured by modifying the DecrementCounters() function as follows. We randomly sample ℓ counters in the hash table, where $\ell = O(\log n)$ is a suitably chosen parameter. We then compute the median of the sampled counters, and decrement all counters by this value, discarding any counters that are

non-positive after decrementing. We refer to this algorithm as the *Reduce-By-Sample-Median (SMED)* extension of the **MG** algorithm. See Algorithm 4 for pseudocode.

THEOREM 2.5. *There is an $\ell = O(\log n)$ such that Algorithm 3 can be implemented to run in amortized constant time with probability at least $1 - 1/n$.*

PROOF. Standard Chernoff bounds (e.g., [37, Theorem 4.4]) imply that if $\ell \geq c \cdot \log n$ for a suitably large constant $c > 0$, then with probability at least $1 - 1/n^2$ the median c^* of ℓ counters sampled from $\{c(i) : i \in T\}$ will satisfy $c^* \geq c(i)$ for at least $k/3$ values of i (where recall that $k = |T|$ whenever a DecrementCounters() operation occurs). By a union bound, we conclude that this event occurs for *all* DecrementCounters() operations with probability at least $1 - 1/n$. In this event, DecrementCounters() operations occur at most once every $k/3$ stream updates. The claimed time bound now follows from the same argument as in Theorem 2.2. \square

THEOREM 2.6. *For any constant $c > 2$, there is an $\ell = O(\log n)$ such that the following holds with probability at least $1 - 1/n$. For any $j < k/c$, Algorithm 4 will return, for each $i \in [n]$, and estimate \hat{f}_i satisfying $0 \leq f_i - \hat{f}_i \leq N^{\text{res}(j)} / (k/c - j)$.*

PROOF. It is easy to see that the proof of Theorem 2.3 applies even if Algorithm 4 decrements all counters by *at most* c_{k^*} rather than exactly c_{k^*} . Hence, it suffices to show that for any constant $c > 2$, there is an $\ell = O(\log n)$ such that with probability at least $1 - 1/n^2$, the median c^* of ℓ counters sampled from $\{c(i) : i \in T\}$ satisfies $c^* \leq c(i)$ for at least k/c values of $i \in T$. Indeed, this property holds by standard Chernoff bounds [37, Theorem 4.4]. By combining this property with a union bound over all DecrementCounters() operations, we conclude that with probability at least $1 - 1/n$, Algorithm 4 always decrements all counters by at most c_{k^*} where $k^* = k/c$. \square

2.3 Additional Implementation Details

A highly optimized implementation of Algorithm 4 in Java has been incorporated into an open source library [40] that is widely deployed within Yahoo and several other companies. We briefly present additional implementation details not covered above.

2.3.1 Variant of Function Estimate(). In our implementation, we modified the function Estimate() compared to the pseudocode of Algorithm 4. Specifically, our implementation keeps a variable offset that tracks the sum of all decrement values c^* over all DecrementCounters() operations. If i is assigned a counter, then our implementation's estimate \hat{f}_i for f_i is $c(i) + \text{offset}$. If i is not assigned a counter, the returned estimate is 0.

This represents a hybrid of the estimates returned by the **MG** and **SS** algorithms (cf. Algorithms 1 and 2 respectively). It behaves analogously to the **MG** algorithm on items i not assigned counters – outputting 0 in this case – and (as explained below) analogously to the **SS** algorithm on items i that are assigned counters.

Our reasoning for this choice is as follows. Before the **MG** and **SS** algorithms were realized to be isomorphic [2], Cormode and Hadjieleftheriou [14] found that the estimates returned by **SS** were superior in practice, despite the fact that the worst-case error bounds

Algorithm 4 The Reduce-By-Sample-Median (**SMED**) Extension of the **MG** Algorithm to Weighted Updates

```

1: // Notation:  $k^*$  is a parameter of the algorithm
2: Algorithm: Initial-Algorithm( $k$ ):
3:    $T \leftarrow \emptyset$  //  $T$  is set of items assigned a counter
4:   Function Update( $i, \Delta$ ):
5:     if  $i \in T$ :
6:        $c(i) \leftarrow c(i) + \Delta$ 
7:     else if  $|T| < k$ :
8:        $T = T \cup \{i\}$ 
9:        $c(i) \leftarrow c(i) + \Delta$ 
10:    else :
11:      DecrementCounters()
12:      if  $\Delta \geq c^*$  // See Line 17 for definition of  $c^*$ 
13:         $T = T \cup \{i\}$  //  $|T| \leq k^* + 1 < k$  after this line
14:         $c(i) \leftarrow c(i) + \Delta - c^*$ 
15:   Function DecrementCounters():
16:     // Notation:  $\ell$  is a parameter
17:     // Randomly sample  $\ell$  counters from  $T$ 
18:     // and let  $c^*$  be the median of the samples.
19:     for all  $j \in T$ :
20:        $c(j) = c(j) - c^*$ 
21:       if  $c(j) \leq 0$ :
22:          $T = T \setminus \{j\}$ 
23:   Function Estimate( $i$ ):
24:     if  $i \in T$ :
25:       return  $c(i)$ 
26:     else
27:       return 0

```

for both algorithms are essentially identical. The reason for this is that the **SS** estimates are “more aggressive” – in particular, whereas the **MG** algorithm will always strictly underestimate the frequency of an item (unless fewer than k items appear in the stream), the **SS** algorithm may well output *exactly correct* estimates of frequent items. The estimates returned by our actual implementation inherit this desirable property of **SS**.

On the other hand, the **SS** algorithm has the undesirable property that it will always strictly *overestimate* the frequency of an item that does not appear in the stream (unless fewer than k items appear in the stream), whereas the **MG** algorithm always outputs exactly correct estimates (namely, 0) for such items. The estimates returned by our actual implementation inherit this desirable property of **MG**.

We remark that, like **MG** and **SS** themselves, our implementation is also capable of outputting tight upper and lower bounds on the frequency of i . The upper bound is $c(i) + \text{offset}$, while the lower bound is $c(i)$ (or 0 if i is not assigned a counter).

2.3.2 Choice of ℓ . We performed numerical calculations establishing that the choice $\ell = 1024$ guarantees that the following holds. For streams of length $N \leq 10^{20}$, Algorithm 4, with probability at least $1 - 1.5 \times 10^{-8}$, returns estimates \hat{f}_i for all $i \in [m]$ satisfying the following error guarantee: $0 \leq f_i - \hat{f}_i \leq N^{\text{res}(j)} / (.33 \cdot k - j)$. Hence, we set $\ell = 1024$ in our implementation.

2.3.3 Hash Table and Exact Space Usage. We experimented with a wide variety of hash table implementations for storing counters. While many performed similarly, we discovered that a good choice

is to use linear probing. Specifically, keys and values are kept in two parallel arrays of length $L \approx 4k/3$ (the length is rounded up to the closest power of two to ensure fast modular arithmetic on indices into the arrays, in case a probe “falls off the end” of the array). We keep an additional array of state variables, where a state of 0 indicates that a cell is unoccupied, and a positive state equals the distance (plus one) of the stored key x from its preferred cell $h(x)$, where $h: [m] \rightarrow L$ is the hash function. Insertion and lookup operations are processed as in standard linear probing: the implementation successively looks at cells $h(x), h(x) + 1 \pmod L, \dots$, until the key is found or an empty cell is discovered.

The other functionality that must be supported by our hash table is decrementing all values by a specified amount c^* and deleting all that become non-positive. This is done by starting at the end (i.e., largest index) of a run (i.e., a contiguous sequence) of occupied cells; each time a non-positive value is identified, the implementation deletes it, and then works its way toward the end of the run, shifting keys and values forward as necessary to ensure that all future lookup and insert operations behave correctly.

State variables need only consist of 2 bytes with overwhelming probability (we performed numerical calculations showing that when $k \leq 2^{32}$ and $L = 4k/3$, the probability that at any given time a state variable exceeds 2^{14} is at most 10^{-250}). Hence, assuming item identifiers and approximate counts are 8 bytes each, and that $4k/3$ is a power of 2, our implementation uses $18 \cdot (4/3) \cdot k = 24 \cdot k$ bytes, plus a small constant number of additional bytes.

3 MERGING SUMMARIES

An extremely important functionality in real systems is the ability to efficiently merge the summaries of separate data sets (via an arbitrary *aggregation tree*) to obtain a summary of the union of the data sets, without increasing the size of the summary or its approximation error. This mergeability property enables a wide variety of scenarios. In addition to the settings described in Section 1.1, the following additional applications make essential use of mergeability.

Consider a system that keeps a separate summary for many subsets of a large dataset. This might occur if, e.g., a company keeps a separate summary for data obtained in each 1-hour period over the course of several years. The company may further subdivide the data by, say, geographic region or other attributes. At query time, an analyst can specify which data are of interest (e.g., the analyst may be interested in users from a given state who used the company’s service in a given 48-hour window). The summaries can then be seamlessly merged to answer approximate queries about the data of interest. In this important example, merging is occurring at query time, and millions or billions of summaries may need to be merged to answer the query. Hence, merging must be extremely efficient, justifying our focus in this section on developing a highly optimized merging procedure.

As a final example, consider multiple datasets that are geographically distributed. It may be infeasible to transmit the entirety of any one of the datasets to a central machine. If an analyst is interested in the union of the datasets, mergeable summaries enable each dataset to be summarized separately, and merely the summaries transmitted to one machine.

3.1 Earlier Merging Procedures

The Procedure of Berinde et al. [6]. The merging procedure we describe in this section is highly similar to a proposal of Berinde et al. [6], who described their proposal in the context of arbitrary counter-based summaries. Berinde et al. suggested that, in order to merge many summaries, one should treat each counter in each summary as a stream update (so a counter $c(i)$ assigned to item i is treated as a stream update $(i, c(i))$) and feed all the stream updates into a single new summary. They proved that if the counter-based algorithm when run on a single stream can return estimated frequencies \hat{f}_i for each i satisfying $|f_i - \hat{f}_i| \leq N^{\text{res}(j)} / (k - j)$, then the output of the merging procedure can return estimated frequencies \hat{f}_i satisfying

$$|f_i - \hat{f}_i| \leq 3 \cdot N^{\text{res}(j)} / (k - 2j), \quad (5)$$

where the quantities f_i and $N^{\text{res}(j)}$ in Equation (5) refer to the *concatenation* of the streams that were fed into each of summaries being merged.

The proposal and accuracy analysis of Berinde et al. [6] have two significant limitations. First, the accuracy analysis only guarantees that the merged summary has at most *triple* the error of the constituent summaries. Second, as pointed out in [2], when merging many summaries, the proposal does not support an arbitrary aggregation tree (i.e., it is not possible to repeatedly merge pairs of summaries in an arbitrary manner until a single summary is obtained). Rather, each summary must be merged “into” a single output summary. This second limitation is really a different manifestation of the first: the reason that Berinde et al. cannot support an arbitrary aggregation tree is that their error analysis cannot rule out the possibility that each edge in the tree causes a tripling of the error. Hence, the final error bound for an arbitrary aggregation tree would be *exponentially* large in the number of summaries being merged.

The merging procedure we propose in this section is closely related to that of Berinde et al. However, we address both of the limitations above by giving a much tighter error analysis. Moreover, since our merging procedure uses our algorithms’ $\text{Update}(i, \Delta)$ procedure as a black box, it inherits the benefits of our $\text{Update}(i, \Delta)$ procedure (e.g., amortized constant runtime) relative to prior algorithms for weighted streams.

Our procedure has the additional advantage of not requiring the allocation of a new output summary for a merge; rather, when merging two summary via our procedure, we seamlessly update the state of one of the two summaries and output the result (discarding the other summary once merging is complete).

We mention in passing that other merging procedures for frequent items algorithms have been proposed [31, 32], but they provide weaker error guarantees than Berinde et al. (i.e., the error increases with each merge step).

The Procedure of Agarwal et al. [2]. Agarwal et al. [2] provided a merging procedure for the **MG** and **SS** algorithms that does support arbitrary aggregation trees. In this discussion, we focus on the case of the **MG** algorithm for simplicity. To merge two **MG** summaries with k counters each, their merging procedure works as follows. The counters from each summary are first added together (so if an item i is assigned a counter in both summaries, its count is set to

the sum of the two counters). The counters are then sorted, and all but the top k counters are discarded.

Suppose that for $j = \{1, 2\}$, the j 'th **MG** summary is a summary of a stream σ_j of weighted length N_j , and let $N = N_1 + N_2$ denote the weighted length of the concatenated stream $\sigma := \sigma_1 \circ \sigma_2$. Let f_i denote the frequency of i in σ . Agarwal et al. proved that the resulting summary returns estimates \hat{f}_i for all i such that:

$$0 \leq f_i - \hat{f}_i \leq (N - C)/(k + 1), \quad (6)$$

where C is the sum of the counters in the merged summary.

Implementing the Merging Procedure. The natural way to implement the merging procedure of Agarwal et al. is the following. First, allocate a new hash table capable of storing up to $2k$ counters. Second, iterate through the counters $(i, c(i))$ in both summaries, adding all $(i, c(i))$ pairs into the new hash table (if key i already resides in the hash table, then add $c(i)$ to its value). This has the effect of summing counters. Third, sort the $(i, c(i))$ pairs in the hash table by $c(i)$. Fourth, create a new summary of capacity k and insert the top k $(i, c(i))$ pairs into the summary, discarding earlier summaries and the intermediate hash table used for merging.

We observe that an alternative implementation of the third and fourth steps above is to use Quickselect [27] to identify the k th largest counter c_k in the hash table. One then makes an additional pass through the counters, identifying all that are at least as large as c_k , and feeding each into the new summary. This implementation runs in time $O(k)$.

Disadvantages. This merging procedure has a number of disadvantages in practice. First, it requires allocating a hash table of capacity $2k$ to add the counters, which potentially doubles the combined space usage of the two constituent **MG** summaries. Second, the original sorting-based proposal of Agarwal et al. requires $\Omega(k \log k)$ time. As we observe above, one can use Quickselect to reduce the runtime to $O(k)$, but the Big-Oh notation hides a substantial constant factor, and running Quickselect on up to $2k$ counters is a runtime bottleneck in practice. Third, while this merging procedure easily extends to the case where the two summaries have differing numbers of counters (say k_1 and k_2 counters respectively), it requires $\Omega((k_1 + k_2) \cdot \log(k_1 + k_2))$ time. One can hope to achieve $O(\min(k_1, k_2))$ time.

In this section, we give a different merging procedure that mitigates these disadvantages. Our procedure applies generically to any counter-based algorithm that can efficiently handle weighted updates. We now describe the procedure in this level of generality.

The Procedure of Cafaro and Pulimeno [10]. Building on earlier work of Cafaro and Tempesta [11], Cafaro and Pulimeno propose a variant of the merging procedure of Agarwal et al. [2], and prove that it satisfies a slightly stronger error guarantee than that of Agarwal et al. [2]. However, Cafaro and Tempesta's proposal suffers from the similar downsides to that of Agarwal et al.: it requires allocating a hash table of capacity $2k$ to add counters, and it inherently requires sorting counters (although [10] states that the procedure runs in time $O(k)$, this appears to assume that the counters are already sorted by count).

Algorithm 5 Our Merge Procedure

```

1: // Notation:  $T_1, T_2$  are both counter-based summaries,
2: // i.e.,  $T_2$  is a set of items  $i$  and approximate counts  $c(i)$ , and
3: //  $T_1$  comes with an  $\text{Update}_{T_1}()$  procedure.
4: Function:  $\text{Merge}(T_1, T_2)$ :
5:   for each  $i \in T_2$ :
6:      $\text{Update}_{T_1}(i, c(i))$ 
7:   return  $T_1$ 

```

3.2 Our Merging Procedure

Description of our Merging Procedure. For simplicity, let us assume that both summaries to be merged are configured to store k counters. Let k' denote the number of (non-zero) counters in the second summary to be merged. In our merging procedure, the k' non-zero counters from the second summary are treated as k' stream updates, and each stream update is fed into the first summary using the summaries $\text{Update}()$ procedure. That is, for each item i assigned a counter $c(i)$ in the second summary, then the function $\text{Update}(i, c(i))$ is called on the first summary. See Algorithm 5 for pseudocode.

Space Usage. Clearly, this merging procedure avoids the $2k$ extra words of space (on top of the space to store the two summaries being merged) required by the merging procedure of Agarwal et al. [2]. It uses no more space than that already used by the two summaries being merged, seamlessly updating the state of one of the summaries, outputting the result, and discarding the other summary upon completion.

Speed. We can guarantee that our merging procedure runs in time $O(k)$. This follows from the fact that the $\text{Update}()$ operation for our algorithms take $O(1)$ time unless a $\text{DecrementCounters}()$ operation is triggered, and a $\text{DecrementCounters}()$ operation takes $O(k)$ time and is triggered at most once every $\Omega(k)$ stream updates.

One might hope that if $k' \ll k$ (which will be the case, e.g., if the second summary is significantly smaller than the first), then the merging procedure runs in $O(k')$ time. Unfortunately, we cannot quite provide this guarantee. For example, if the second summary contains only a single nonzero counter $c(i)$, and calling $\text{Update}(i, c(i))$ on the first summary triggers a $\text{DecrementCounters}()$ operation, then the merging procedure will take $\Theta(k)$ time. However, we can guarantee the following. If $\Omega(k/k')$ summaries of size at most k' are merged into a single summary of size k , then the amortized time per merge is $O(k')$.

Error. We now show that our merging procedure, when applied to our **MED** algorithm (Algorithm 3) satisfies an error property analogous to that established by Agarwal et al. (cf. Equation (6)). Using similar ideas, it can be shown that our merging procedure satisfies the analogous guarantees when applied to Algorithm 4, but we omit this result for brevity. While our merging procedure satisfies many desirable properties compared to that of Agarwal et al. [2], our error analysis exploits the same key ideas as theirs.

THEOREM 3.1. *Suppose that for $j = \{1, 2\}$, we run Algorithm 3 with k counters on a stream σ_j of weighted length N_j . Let $N = N_1 + N_2$ denote the weighted length of the concatenated stream $\sigma := \sigma_1 \circ \sigma_2$. Let f_i denote the frequency of i in σ . Then after applying the merging*

procedure of Algorithm 5, the resulting summary can, for any i , return an estimate \hat{f}_i satisfying

$$0 \leq f_i - \hat{f}_i \leq (N - C)/k^*, \quad (7)$$

where C is the sum of the counter values in the merged summary. In fact, for any $j < k^*$, it holds that

$$0 \leq f_i - \hat{f}_i \leq N^{\text{res}(j)}/k^*. \quad (8)$$

PROOF. Clearly $0 \leq f_i - \hat{f}_i$ for all i . We now turn to establishing that $f_i - \hat{f}_i \leq (N - C)/k^*$.

For $j \in \{1, 2\}$, let $\hat{f}_{i,j}$ denote the estimate for i returned by summary j after being run on σ_j , and let $f_{i,j}$ denote the true frequency of item i in σ_j . Let $E_j := \max_i f_{i,j} - \hat{f}_{i,j}$, and let $c(j)$ denote the sum of the counter values in summary j after being run on σ_j . Let $E = \max_i f_i - \hat{f}_i$ denote the maximum error in any estimate returned by the merged summary. Finally, let S denote the sum of all values c_{k^*} that are computed during DecrementCounters() operations called during the merge procedure.

It is clear that $E \leq E_1 + E_2 + S$. By Lemma 2.4,

$$E \leq (N_1 - C_1)/k^* + (N_2 - C_2)/k^* + S. \quad (9)$$

Recall from the proof of Lemma 2.4 that any DecrementCounters() operation decreases at least k^* counters by c_{k^*} . Hence, $S \leq (C_1 + C_2 - C)/k^*$. Combining this with Equation (9), we conclude that $E \leq (N_1 - C_1)/k^* + (N_2 - C_2)/k^* + (C_1 + C_2 - C)/k^* = (N_1 + N_2 - C)/k^* = (N - C)/k^*$, establishing Equation (7). To see that Equation (8) follows from Equation (7), use the same argument from the proof of Theorem 2.3 (as the proof of Theorem 2.3 showed that the desired tail guarantee follows in a black box manner from Lemma 2.4). \square

We remark that a final advantage of our merging procedure is that it does not require significant additional code on top of the streaming algorithm itself: it simply invokes the streaming algorithm's Update() procedure in a black-box manner.

Note. A subtlety regarding implementations of our proposed merging procedure is as follows. If counters are stored in hash tables, then the performance guarantees of the hash table assume that hash values are (roughly) uniformly distributed throughout the table. If our merge procedure is used to merge two summaries whose hash tables use *the same* hash function, then care must be taken to ensure that this property continues to hold. For example, if one iterates through the second summary's hash table from front to back, calling Update($i, c(i)$) on the first summary for each counter ($i, c(i)$) encountered, then one may "overpopulate" the front of the first summary's hash table.

One avoids this issue entirely if the summaries choose their hash functions independently of each other. Even for summaries that do use the same hash function, this issue can be addressed by, say, iterating through the second summary's counters in a random order.

4 EXPERIMENTS

The goal of this section is to investigate three issues. First, we compare our main algorithm **SMED** to alternatives: the proposals **RBMC** and **MHE** of prior work, and a variant of **SMED** that we call **SMIN**. The difference between **SMIN** and **SMED** is that the

DecrementCounters() operation of **SMIN** decrements all counters by the sample minimum counter value, instead of the sample median value. Notice that **SMIN** is identical to **RBMC**, except that the DecrementCounters() operation of **RBMC** decrements by the global minimum counter value, while **SMIN** decrements by the minimum counter value in a sample of size $\ell = 1024$.

We find that **SMED** is significantly faster than the the algorithms of prior work (5.0x-8.3x faster than **MHE** and up to 32.8x faster than **RBMC**). While it is not quite as accurate as **MHE** or **RBMC** for a given amount of memory, the maximum error of **SMED** is at most 30% larger than that of **MHE**, and less than 2.5x that of **RBMC** and **SMIN** (the latter two algorithms had nearly identical error on the datasets used in our experiments). Moreover, the increased error of **SMED** relative to alternatives can be overcome, while simultaneously decreasing runtime, by increasing the number of counters by a small factor (less than 2x). This is entirely expected, as our error analysis (cf. Theorems 2.3 and 2.6) indicate that for sufficiently large settings of the parameter ℓ , **SMED** with k counters has error that is no worse than **RBMC** with $\approx k/2$ counters. Hence, we find that our novel **SMED** algorithm is the preferable option, except possibly in settings where space and accuracy are paramount.

Even in settings in which space and accuracy are paramount, we find that on the datasets with which we experimented, the preferable solution is our novel algorithm **SMIN**. For a given space bound, **SMIN** achieves error nearly identical to **RBMC** (and much better than **MHE**), with far superior speed.

Second, our main algorithm **SMED** (cf. Algorithm 4) naturally allows for a smooth tradeoff between error and speed as follows. While **SMED** itself implements the DecrementCounters() procedure by sampling many counters and decrementing by the sample median, one could instead decrease by a different sample quantile (at the extreme end is **SMIN**, which decrements by the sample minimum). A larger decrement amount increases the error, but ensures that DecrementCounters() operations occur less frequently, and hence speeds up the algorithm. We investigate this tradeoff on real and synthetic datasets, and find that decrementing by sample medians is indeed an attractive point on this tradeoff curve.

Third, we compare our merging procedure to prior work [2], and find that ours is significantly faster (up to 12x faster than a straightforward implementation of prior work [2]) while using less than half the space.

4.1 Details

All experiments were run on a machine running Windows 10 with an Intel Core i5-6600K 3.5GHz Quad-Core Processor, (2 x 8GB) DDR4-2666 Memory, and a 500GB SSD. Test data was generated from the CAIDA Anonymized Internet Traces 2016 Dataset [1]. The first eighteen packet capture files in the dataset were chosen. These were then preprocessed into updates of the form (a_i, Δ_i) , where a_i is the source IP with decimal points excluded, and Δ_i is the packet size in bits. These eighteen preprocessed streams were then concatenated into one larger stream. The stream length was $n \approx 585.3$ million and the weighted stream length was $N := \sum_{i=1}^n \Delta_i \approx 32.4 \times 10^{10}$. The total size on the packet captures was 34.0 GB. Since the data came from an IPv4 packet capture, the universe size is $m = 2^{32}$. Although IPv4 addresses can be stored with 32 bits, our implementations, for purposes of generality, use a long long data type (64 bits) to store

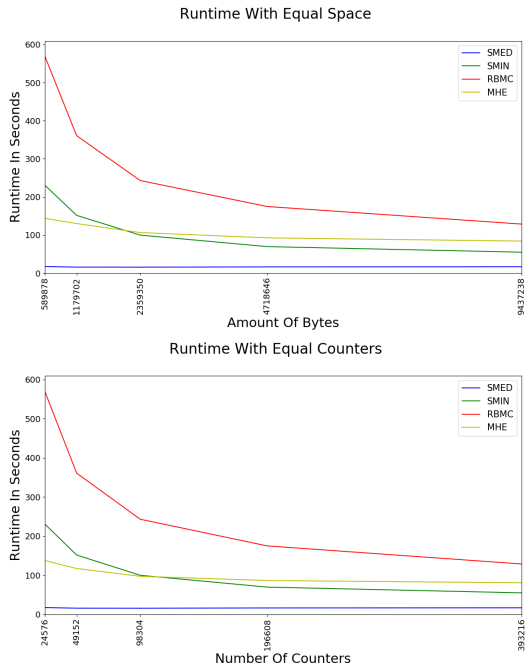


Figure 1: Runtime Comparison of Four Algorithms

each. The implementations can be trivially modified to only use 32 bits per identifier if desired.

The number of unique identifiers in the packet stream was approximately 24.7 million. Given this many unique identifiers, the trivial (exact) algorithm that keeps a hash table storing an exact count for each unique IP address in the stream will use hundreds of megabytes of space. Our new algorithms use less than 1% of the space of this trivial solution so long as they are run with less than about 247,000 counters.

We also ran experiments on synthetic data generated by a Zipfian distribution with various skewness parameters. The algorithms performed entirely similarly on these datasets and the packet trace data. Hence, for brevity we only present our experimental results on the packet trace data.

4.2 Properties Exhibited By All Algorithms

For any fixed stream, as the number of counters k used in a counter-based algorithm increases, the algorithm grows closer and closer to being exact, and the number of DecrementCount() operations relative to the number of stream updates decreases. Hence, the algorithms' performances (in both speed and error) converge as the number of counters grows. Consequently, the biggest differences between the counter-based algorithms explored in our experiments are observed when the number of counters is relatively small.

4.3 Comparison to Baselines

We compared our new algorithms, SMED and SMIN, to the algorithms for weighted streams given in prior work, RBMC and MHE. While RBMC, SMED, and SMIN all use the same amount of space (in bytes) for a given number of counters k , MHE uses additional

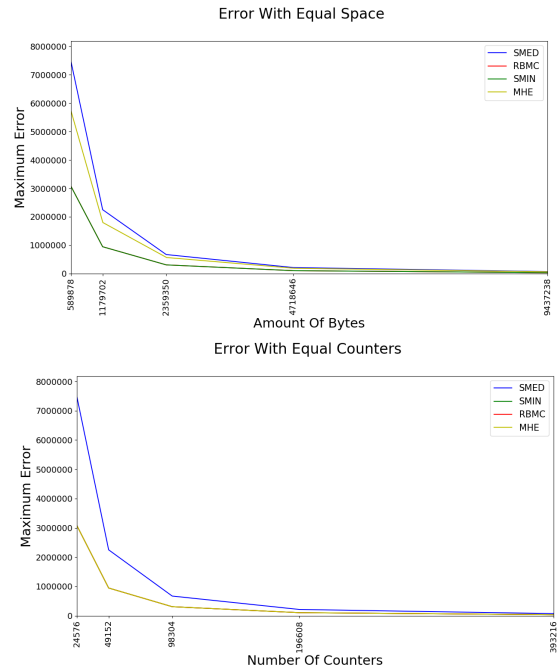


Figure 2: Maximum Error Comparison of Four Algorithms. RBMC is omitted from both plots, as its error is indistinguishable from SMIN. For an equal number of counters k (bottom plot), RBMC, MHE, and SMIN all have indistinguishable maximum error, and hence only one is displayed. Indeed, RBMC with k counters is isomorphic to MHE with $k + 1$ counters, cf. Section 1.4.

space owing to the need to maintain a min-heap data structure in addition to a hash table storing all counters. Figures 1 and 2 present both equal-space and equal-counters comparisons of the algorithms.

For each of our five tested values of k , we ran each algorithm 10 times on the same packet stream, and all numbers reported are averages of all 10 runs (the error and space usage for the MHE and RBMC algorithms are deterministic functions of the stream, so only the timings varied from run to run for these algorithms).

Figures 1 and 2 show runtime and maximum error comparisons for the four algorithms. Only the runtime graphs contain the results for RBMC, because the error of RBMC and SMIN were so close as to be indistinguishable.

Comparison of SMED to Alternatives. Figure 1 demonstrates that SMED significantly outperforms all three alternatives in terms of runtime. For an equal amount of space, SMED was faster than MHE by a factor of 5x-8.3x. Compared to SMIN, SMED was 3.3x-13.3x faster, and compared to RBMC, SMED was 7.7x-32.8x times faster. As indicated in Section 4.2, the smaller the number of counters k , the faster SMED is relative to the alternatives.

Figure 2 demonstrates that for, an equal amount of space, the maximum error of SMED is only 13%-30% larger than that of MHE. RBMC and SMIN both have significantly less error than SMED

and **MHE** for a given amount of space. However, the maximum error of **SMED** is never more than 2.5x that of **RBMC** and **SMIN**. Moreover, the figure demonstrates that this increased error of **SMED** relative to alternatives can be overcome (while simultaneously decreasing runtime) by increasing the number of counters (and hence space consumption) by a factor of less than 2x.

In conclusion, **SMED** appears to be the most attractive option, except in settings where space and error are significantly more important than speed. In these settings, our new **SMIN** algorithm is more attractive than the alternatives from prior work, namely **MHE** and **RBMC** (we note that Section 4.4 below demonstrates that it is possible to smoothly interpolate between **SMED** and **SMIN**, obtaining different tradeoffs between speed and error. Exactly which tradeoff is preferable depends on the relative importance of speed vs. error in the application domain). Figures 1 and 2 demonstrate that for an equal amount of space, **SMIN** is about 2.3-2.5x faster than **RBMC**, and their error behaviors are indistinguishable. For an equal amount of space, **SMIN** is never more than 1.6x slower than **MHE** and is actually faster than **MHE** for large k ; meanwhile, the error of **MHE** is 1.8x-1.9x that of **SMIN**.

4.4 Tradeoffs Between Speed and Error

Recall that **SMED** takes a sample of the keys currently being stored in the hash table, uses Quickselect [27] to compute the median of the sample, then decrements all counters by this quantity. We ran tests on a large number of variations of this novel algorithm, where a variation might decrement by other sample quantiles rather than the sample median. Overall, we tested fifty total variations, ranging from the 0th quantile to the 98th quantile. Note that decrementing by the 0th quantile yields the algorithm **SMIN**.

The top plot in Figure 3 shows the runtime in seconds for each setting of k , and how it varies with the quantile used to determine decrements. The runtimes for **SMIN** and a few other low quantiles are not displayed so as not to distort the graph, as those runtimes are significantly larger than the rest. This is expected, as the less aggressively `DecrementCounters()` decreases counter values, the greater the number of `DecrementCounters()` operations that occur. From the 0th quantile (**SMIN**) to the 50th quantile (**SMED**), runtime drops heavily (specifically, as indicated previously in Section 4.3, **SMED** is 3.3x-13.3x faster than **SMIN**, depending on number of counters k .) Trivially the best runtimes are achieved by using very high quantiles in the `DecrementCounters()` operation, but the “returns” are diminishing: using the 98th quantile yields an algorithm that is only 20-30% faster than using the 20th quantile.

The middle and bottom plot in Figure 3 displays the relationship between quantile and maximum error. As expected, the trend here is that as quantile increases, error increases. However, the error increases relatively slowly as the quantile increases from 0 to about 70, before shooting up rapidly thereafter. We conclude that, on the datasets tested, any quantile less than roughly 70 represents a reasonable choice in the tradeoff curve between speed and error; exactly which point on the curve is most desirable will depend on the relative importance of speed vs. error in an application.

We believe, however, that it is wise to always use a quantile of roughly 10 or higher (i.e., we believe this is preferable to using **SMIN**). Such a choice of quantile achieves almost all of the accuracy

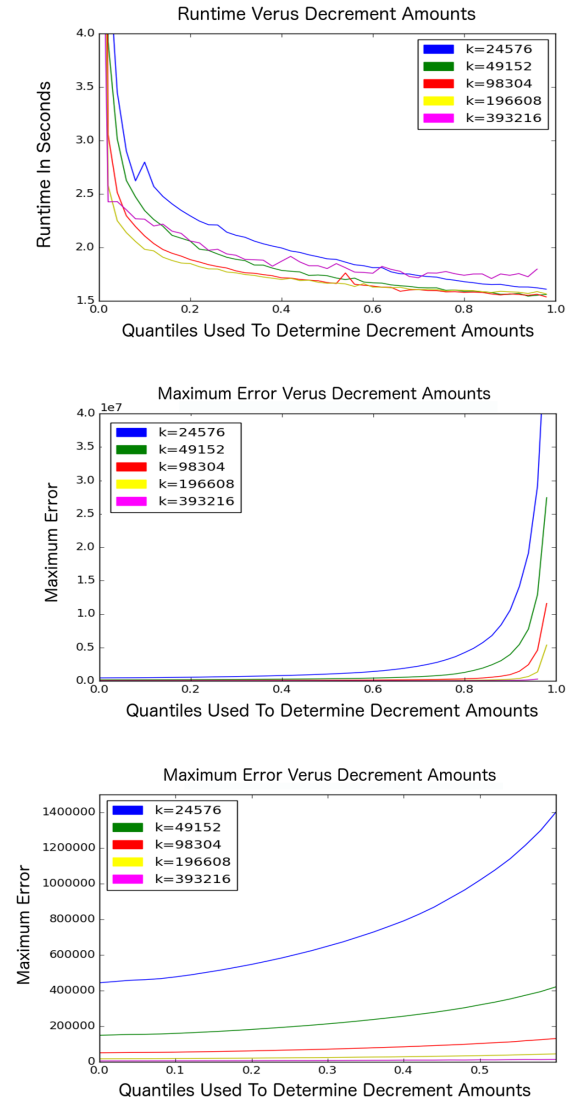


Figure 3: Time and error of our algorithms as a function of the quantile Used by `DecrementCounters()`. The middle and bottom plots are identical, except that the bottom plot cuts off large quantiles in order to better illustrate the error for small quantiles.

advantages of **SMIN**, while coming with significantly improved worst-case runtime bounds compared to **SMIN**. In particular, this provides runtime guarantees even for data distributions that differ substantially from any that we experimented on.

4.5 Merging

In addition to the prior tests, we also ran experiments to evaluate our merging procedure (cf. Section 3). We compared our proposed procedure against the proposal of Agarwal et al. [2], as well as the variant that we proposed (cf. Section 3.1) that uses Quickselect [27]

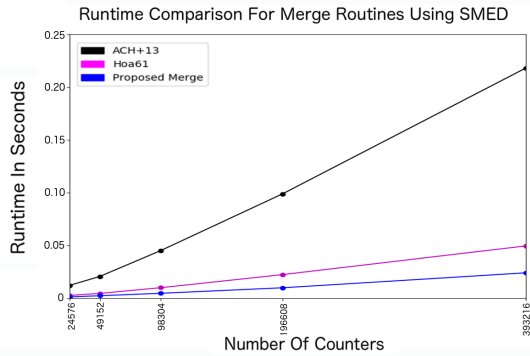


Figure 4: Speed of Our Merge Procedure Compared to Prior Work. ACH+13 denotes a direct implementation of the merging procedure of Agarwal et al. [2]. Hoa61 denotes our proposed Quickselect-based implementation of the procedure of [2] (cf. Section 3.1).

rather than sorting to identify the k th largest counter after “adding” the counters from the two sketches being merged. Recall that our merging procedure uses the Update() operation of a counter-based algorithm as a black box; we tested our merging procedure using the Update() procedure of SMED.

In each of our experiments, we merged 50 pairs of sketches, with each sketch having a capacity of k counters. These experiments require 50x the amount of raw data that is required to test a single sketch. Hence, we “filled up” the sketches using synthetic data streams before merging them: the streams had item identifiers drawn from a Zipfian distribution with parameter $\alpha = 1.05$, and item weights were generated from a uniformly random distribution from 1 to 10,000 (cf. [6, Section 5] for a definition of this distribution).

As illustrated in Figure 4, when using SMED, our proposed merging procedure runs up to 8.6x-10x faster than the procedure of [2] (the bigger the sketch, the faster our proposed procedure is relative to that of [2]). Our proposed merging procedure is 1.9x-2.26 faster than the Quickselect-based variant. Moreover, the error of our merge procedure when using SMED was always within 2.5% of the error of the procedure of [2] and the equivalent Quickselect-based variant; owing to the small magnitude of this difference, we do not display this error graphically.

It is worth mentioning that our proposed merge algorithm also uses less space than the two alternatives. Both alternatives require allocating an additional hash table of capacity $2k$ in order to store all counters from both sketches, as well as an extra hash table of capacity k to store the final merged sketch. Hence, they consume 2.5x more space than our procedure, which uses no additional space beyond that already used by the two sketches to be merged.

5 OTHER PRIOR WORK

Work of Sivaraman et al. Recent independent work of Sivaraman et al. [42] proposes modifying the SS algorithm (cf. Section 1.3.3) as follows. When processing a stream update (i, Δ) , if i is not assigned a counter and all counters are in use, sample ℓ counters at random

(where ℓ is a design parameter), and then reassign this counter to i and increment it by Δ . The primary motivation for this proposal is to minimize memory accesses on each stream update, as memory accesses are a bottleneck in network switching hardware. For $\ell = O(1)$, this proposal runs in constant time per stream update, but may have larger error than our proposals. We leave a detailed experimental comparison of our algorithms with the proposal of Sivaraman et al. [42] to future work.

Work of Ben-Basat et al. Recent independent work of Ben-Basat et al. [4] also revisits the problem of identifying frequent items in data streams, with the goal of designing algorithms that can efficiently handle weighted updates. They give two algorithms. The first achieves amortized constant update time. The second modifies the first to achieve worst-case constant update time, at the expense of constant-factor increases in space usage and update time. (A similar transformation can be applied to our algorithms to achieve worst-case constant update time). Ben-Basat et al. [4] do not consider mergability of summaries.

At a very high level, the algorithms of [4] maintain two tables of counters, an Active Table and a Passive Table. Like all counter-based algorithms, the data structures periodically “fill up” with counters, and small counters are then purged from the structure. Roughly speaking, their Active table attempts to accurately track counts of all items that have been updated since the previous purge operation, while the Passive Table tracks counters that survived the previous purge operation but have not been updated since. We refer to their paper for additional details. While the algorithms of [4] and our own work achieve similar asymptotic performance, the full version of our paper will contain an experimental comparison showing that the constant factors in the time and space complexity of our implementation are smaller than the algorithms of [4].

6 CONCLUSION AND FUTURE DIRECTIONS

We described a highly optimized version of Misra and Gries’ seminal algorithm for estimating frequencies of items over data streams. Our algorithm improves on two theoretical and practical aspects of prior work. First, it handles *weighted* updates in amortized constant time. Second, it uses a simple and fast method for merging summaries that asymptotically and concretely improves on prior work even for unweighted streams.

One direction for future work is to incorporate our optimized algorithms into more complicated streaming algorithms that use heavy hitter algorithms as subroutines. As described in Section 1.2, these include estimating the empirical entropy of a data stream [12], and identifying hierarchical heavy hitters (HHHs) [5, 36].

Another direction is to explore the possibility of obtaining further speedups by judiciously *skipping* some stream updates (i.e., dropping some stream updates altogether). This class of optimizations has previously been applied to other algorithms for computing heavy hitters [8]. In addition, Ben-Basat et al. [5] recently explored a related approach to obtain algorithms for identifying HHHs with constant update time, showing that judiciously skipping some updates offers impressive speedups with little loss in accuracy. Whether these approaches would have a similar effect on our algorithms (which already have constant amortized update time) is unclear.

REFERENCES

- [1] 2016. The CAIDA UCSD Anonymized Internet Traces 2016 Dataset. (2016). https://www.caida.org/data/passive/passive_2016_dataset.xml. Specific files used: equinix-chicago.dirA.20160121-130000.UTC.anon.pcap.gz to equinix-chicago.dirA.20160121-131800.UTC.anon.pcap.gz.
- [2] Pankaj K. Agarwal, Graham Cormode, Zengfeng Huang, Jeff M. Phillips, Zhewei Wei, and Ke Yi. 2013. Mergeable summaries. *ACM Trans. Database Syst.* 38, 4 (2013), 26. <https://doi.org/10.1145/2500128>
- [3] Nir Ailon, Zohar Shay Karnin, Edo Liberty, and Yoelle Maarek. 2013. Threading machine generated email. In *Sixth ACM International Conference on Web Search and Data Mining, WSDM 2013, Rome, Italy, February 4-8, 2013*. 405–414. <https://doi.org/10.1145/2433396.2433447>
- [4] Ran Ben Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. 2017. Optimal elephant flow detection. *arXiv preprint arXiv:1701.04021* (2017). To Appear in IEEE INFOCOM 2017.
- [5] Ran Ben Basat, Gil Einziger, Roy Friedman, Marcelo Caggiani Luizelli, and Erez Waisbard. 2017. Constant Time Updates in Hierarchical Heavy Hitters. *arXiv preprint arXiv:1707.06778* (2017). To Appear in ACM SIGCOMM 2017.
- [6] Radu Berinde, Piotr Indyk, Graham Cormode, and Martin J. Strauss. 2010. Space-optimal heavy hitters with strong error bounds. *ACM Trans. Database Syst.* 35, 4 (2010), 26. <https://doi.org/10.1145/1862919.1862923>
- [7] Arnab Bhattacharyya, Palash Dey, and David P. Woodruff. 2016. An Optimal Algorithm for ℓ_1 -Heavy Hitters in Insertion Streams and Related Problems. In *Proceedings of PODS*. 385–400. <https://doi.org/10.1145/2902251.2902284>
- [8] Supratik Bhattacharyya, Andre Madeira, S Muthukrishnan, and Tao Ye. 2007. How to scalably and accurately skip past streams. In *Data Engineering Workshop, 2007 IEEE 23rd International Conference on*. IEEE, 654–663.
- [9] Vladimir Braverman, Stephen R. Chestnut, Nikita Iykin, and David P. Woodruff. 2016. Beating CountSketch for heavy hitters in insertion streams. In *Proceedings of STOC*. 740–753. <https://doi.org/10.1145/2897518.2897558>
- [10] Massimo Cafaro and Marco Pulimeno. 2016. Merging Frequent Summaries.. In *ICTCS*. 280–285.
- [11] Massimo Cafaro and Piergiulio Tempesta. 2011. Finding frequent items in parallel. *Concurrency and Computation: Practice and Experience* 23, 15 (2011), 1774–1788.
- [12] Amit Chakrabarti, Graham Cormode, and Andrew McGregor. 2010. A near-optimal algorithm for estimating the entropy of a stream. *ACM Trans. Algorithms* 6, 3 (2010). <https://doi.org/10.1145/1798596.1798604>
- [13] Moses Charikar, Kevin Chen, and Martin Farach-Colton. 2002. Finding Frequent Items in Data Streams. In *Proceedings of ICALP*. 693–703.
- [14] Graham Cormode and Marios Hadjieleftheriou. 2010. Methods for finding frequent items in data streams. *VLDB J.* 19, 1 (2010), 3–20. <https://doi.org/10.1007/s00778-009-0172-z>
- [15] Graham Cormode, Flip Korn, S Muthukrishnan, and Divesh Srivastava. 2003. Finding hierarchical heavy hitters in data streams. In *Proceedings of the 29th international conference on Very large data bases-Volume 29*. VLDB Endowment, 464–475.
- [16] Graham Cormode, Flip Korn, S. Muthukrishnan, and Divesh Srivastava. 2004. Diamond in the Rough: Finding Hierarchical Heavy Hitters in Multi-Dimensional Data. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, Paris, France, June 13-18, 2004*, Gerhard Weikum, Arnd Christian König, and Stefan Deßloch (Eds.). ACM, 155–166. <https://doi.org/10.1145/1007568.1007588>
- [17] Graham Cormode and S. Muthukrishnan. 2004. An Improved Data Stream Summary: The Count-Min Sketch and Its Applications. In *Proceedings of LATIN*. 29–38. https://doi.org/10.1007/978-3-540-24698-5_7
- [18] Erik D Demaine, Alejandro López-Ortiz, and J Ian Munro. 2002. Frequency estimation of internet packet streams with limited space. In *European Symposium on Algorithms*. Springer, 348–360.
- [19] Nick Duffield, Carsten Lund, and Mikkel Thorup. 2001. Charging from sampled network usage. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*. ACM, 245–256.
- [20] Cristian Estan, Stefan Savage, and George Varghese. 2003. Automatically inferring patterns of resource consumption in network traffic. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 137–148.
- [21] Cristian Estan and George Varghese. 2002. New Directions in Traffic Measurement and Accounting. In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '02)*. ACM, New York, NY, USA, 323–336. <https://doi.org/10.1145/633025.633056>
- [22] Wenjia Fang and Larry Peterson. 1999. Inter-AS traffic patterns and their implications. In *Global Telecommunications Conference, 1999. GLOBECOM '99*, Vol. 3. IEEE, 1859–1868.
- [23] Anja Feldmann, Albert Greenberg, Carsten Lund, Nick Reingold, Jennifer Rexford, and Fred True. 2001. Deriving traffic demands for operational IP networks: Methodology and experience. *IEEE/ACM Transactions on Networking (ToN)* 9, 3 (2001), 265–280.
- [24] Łukasz Golab, David DeHaan, Erik D. Demaine, Alejandro Lopez-Ortiz, and J. Ian Munro. 2003. Identifying Frequent Items in Sliding Windows over Online Packet Streams. In *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement (IMC '03)*. ACM, New York, NY, USA, 173–178. <https://doi.org/10.1145/948205.948227>
- [25] Yu Gu, Andrew McCallum, and Donald F. Towsley. 2005. Detecting Anomalies in Network Traffic Using Maximum Entropy Estimation. In *Proceedings of the 5th Internet Measurement Conference, IMC 2005, Berkeley, California, USA, October 19-21, 2005*. USENIX Association, 345–350. <http://www.usenix.org/events/imc05/tech/gu.html>
- [26] John Hershberger, Nisheeth Shrivastava, Subhash Suri, and Csaba D Tóth. 2005. Space complexity of hierarchical heavy hitters in multi-dimensional data streams. In *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, 338–347.
- [27] C. A. R. Hoare. 1961. Algorithm 65: Find. *Commun. ACM* 4, 7 (July 1961), 321–322. <https://doi.org/10.1145/366622.366647>
- [28] Anukool Lakhina, Mark Crovella, and Christophe Diot. 2005. Mining anomalies using traffic feature distributions. *ACM SIGCOMM Computer Communication Review* 35, 4 (2005), 217–228.
- [29] Edo Liberty. 2013. Simple and deterministic matrix sketching. In *Proceedings of KDD*. 581–588. <https://doi.org/10.1145/2487575.2487623>
- [30] Yuan Lin and Hongyan Liu. 2007. Separator: sifting hierarchical heavy hitters accurately from data streams. *Advanced Data Mining and Applications* (2007), 170–182.
- [31] Amit Manjhi, Suman Nath, and Phillip B. Gibbons. 2005. Tributaries and Deltas: Efficient and Robust Aggregation in Sensor Network Streams. In *Proceedings of SIGMOD*. 287–298. <https://doi.org/10.1145/1066157.1066191>
- [32] Amit Manjhi, Vladislav Shkapenyuk, Kedar Dhamdhere, and Christopher Olston. 2005. Finding (Recently) Frequent Items in Distributed Data Streams. In *Proceedings of ICDE*. 767–778. <https://doi.org/10.1109/ICDE.2005.68>
- [33] Gurmeet Singh Manku and Rajeev Motwani. 2002. Approximate Frequency Counts over Data Streams. In *Proceedings of VLDB*. VLDB Endowment, 346–357. <http://dl.acm.org/citation.cfm?id=1287369.1287400>
- [34] Ahmed Metwally, Divyakant Agrawal, and Amr El Abbadi. 2005. Efficient Computation of Frequent and Top-k Elements in Data Streams. In *Proceedings of ICDT*. 398–412. https://doi.org/10.1007/978-3-540-30570-5_27
- [35] J. Misra and David Gries. 1982. Finding repeated elements. *Science of Computer Programming* 2, 2 (1982), 143–152. [https://doi.org/10.1016/0167-6423\(82\)90012-0](https://doi.org/10.1016/0167-6423(82)90012-0)
- [36] Michael Mitzenmacher, Thomas Steinke, and Justin Thaler. 2012. Hierarchical Heavy Hitters with the Space Saving Algorithm. In *Proceedings of ALENEX*. 160–174. <https://doi.org/10.1137/1.9781611972924.16>
- [37] Michael Mitzenmacher and Eli Upfal. 2005. *Probability and computing - randomized algorithms and probabilistic analysis*. Cambridge University Press.
- [38] Rong Pan, Lee Breslau, Balaji Prabhakar, and Scott Shenker. 2003. Approximate fairness through differential dropping. *ACM SIGCOMM Computer Communication Review* 33, 2 (2003), 23–39.
- [39] Frederic Raspall, Sebastia Sallent, and Josep Yufera. 2006. Shared-state Sampling. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement (IMC '06)*. ACM, New York, NY, USA, 1–14. <https://doi.org/10.1145/1177080.1177082>
- [40] Lee Rhodes, Kevin Lang, Alexander Saydakov, Justin Thaler, Edo Liberty, and Jon Malkin. 2015. DataSketches: A Java software library for streaming data algorithms. Apache License, Version 2.0. (2015). <https://datasketches.github.io>.
- [41] Vyas Sekar, Nick G Duffield, Oliver Spatscheck, Jacobus E van der Merwe, and Hui Zhang. 2006. LADS: Large-scale Automated DDoS Detection System.. In *USENIX Annual Technical Conference, General Track*. 171–184.
- [42] Vibhaalakshmi Sivaraman, Srinivas Narayana, Ori Rottenstreich, S. Muthukrishnan, and Jennifer Rexford. 2016. Smoking Out the Heavy-Hitter Flows with HashPipe. *CoRR abs/1611.04825* (2016). <http://arxiv.org/abs/1611.04825> To appear in *SDN 2017*.
- [43] Patrick Truong and Fabrice Guillemin. 2009. Identification of heavyweight address prefix pairs in IP traffic. In *Teletraffic Congress, 2009. ITC 21 2009. 21st International*. IEEE, 1–8.
- [44] Arno Wagner and Bernhard Plattner. 2005. Entropy Based Worm and Anomaly Detection in Fast IP Networks. In *14th IEEE International Workshops on Enabling Technologies (WETICE 2005), 13-15 June 2005, Linköping, Sweden*. IEEE Computer Society, 172–177. <https://doi.org/10.1109/WETICE.2005.35>
- [45] Kuai Xu, Zhi-Li Zhang, and Supratik Bhattacharyya. 2005. Profiling internet backbone traffic: behavior models and applications. In *Proceedings of the ACM SIGCOMM 2005 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Philadelphia, Pennsylvania, USA, August 22-26, 2005*, Roch Guérin, Ramesh Govindan, and Greg Minshall (Eds.). ACM, 169–180. <https://doi.org/10.1145/1080091.1080112>
- [46] Yin Zhang, Sumeet Singh, Subhabrata Sen, Nick G. Duffield, and Carsten Lund. 2004. Online identification of hierarchical heavy hitters: algorithms, evaluation, and applications. In *Proceedings of the 4th ACM SIGCOMM Internet Measurement Conference, IMC 2004, Taormina, Sicily, Italy, October 25-27, 2004*, Alfio Lombardo and James F. Kurose (Eds.). ACM, 101–114. <https://doi.org/10.1145/1028788.1028802>
- [47] Haiquan (Chuck) Zhao, Ashwin Lall, Mitsunori Ogihara, Oliver Spatscheck, Jia Wang, and Jun Xu. 2007. A Data Streaming Algorithm for Estimating Entropies

of Od Flows. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement (IMC '07)*. ACM, New York, NY, USA, 279–290. <https://doi.org/10.1145/1298306.1298345>

- [48] Qi Zhao, Abhishek Kumar, and Jun Xu. 2005. Joint Data Streaming and Sampling Techniques for Detection of Super Sources and Destinations. In *Proceedings of*

the 5th ACM SIGCOMM Conference on Internet Measurement (IMC '05). USENIX Association, Berkeley, CA, USA, 7–7. <http://dl.acm.org/citation.cfm?id=1251086.1251093>