

Separations for Estimating Large Frequency Moments on Data Streams

David P. Woodruff 

Carnegie Mellon University, Pittsburgh, PA, USA

Samson Zhou  

Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

We study the classical problem of moment estimation of an underlying vector whose n coordinates are implicitly defined through a series of updates in a data stream. We show that if the updates to the vector arrive in the random-order insertion-only model, then there exist space efficient algorithms with improved dependencies on the approximation parameter ε . In particular, for any real $p > 2$, we first obtain an algorithm for F_p moment estimation using $\tilde{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p}\right)$ bits of memory. Our techniques also give algorithms for F_p moment estimation with $p > 2$ on arbitrary order insertion-only and turnstile streams, using $\tilde{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p}\right)$ bits of space and two passes, which is the first optimal multi-pass F_p estimation algorithm up to $\log n$ factors. Finally, we give an improved lower bound of $\Omega\left(\frac{1}{\varepsilon^2} \cdot n^{1-2/p}\right)$ for one-pass insertion-only streams. Our results separate the complexity of this problem both between random and non-random orders, as well as one-pass and multi-pass streams.

2012 ACM Subject Classification Theory of computation \rightarrow Streaming, sublinear and near linear time algorithms

Keywords and phrases streaming algorithms, frequency moments, random order, lower bounds

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.112

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: <https://arxiv.org/abs/2105.03773>

Funding The authors would like to thank support from NSF grant No. CCF-181584 and a Simons Investigator Award.

1 Introduction

The efficient computation of statistics has emerged as an increasingly important goal for large databases storing information generated from financial markets, internet traffic, IoT sensors, scientific observations, etc. The one-pass streaming model formally defines an implicit and underlying dataset through sequential updates that arrive one at a time and describe the evolution of the dataset over time. The goal is to aggregate or approximate some statistic of the input data using space that is sublinear in the size of the input.

The frequency moment estimation problem is fundamental to data streams. Since the celebrated paper of Alon, Matias, and Szegedy [1], the frequency moment estimation problem has been a central problem in the streaming model; more than two decades of research [1, 14, 4, 33, 24, 23, 29, 27, 26, 17, 10, 9, 6, 11, 19, 35] has studied the space or time complexity of this problem. We first consider the *insertion-only* model, where updates take the form u_1, \dots, u_m in a stream of length m and each update u_t is in $[n] = \{1, 2, \dots, n\}$ for $t \in [m]$. We assume for simplicity that $m \leq \text{poly}(n)$. The updates implicitly define a frequency vector $f \in \mathbb{R}^n$ so that each update effectively increases a coordinate of f in the sense that $f_i = |\{t : u_t = i\}|$ for each $i \in [n]$. Given $p, \varepsilon > 0$, the F_p moment estimation problem is to approximate $F_p = \sum_{i \in [n]} (f_i)^p$ within a $(1 \pm \varepsilon)$ factor. The complexity of



© David P. Woodruff and Samson Zhou;

licensed under Creative Commons License CC-BY 4.0

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).

Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 112; pp. 112:1–112:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



this problem differs greatly for the range of p . For $p > 2$, [4, 14] showed that even for the streaming model, the space usage for F_p -estimation requires polynomial factors in n and m , whereas polylogarithmic factors are achievable for $p \leq 2$ [1, 23, 29, 27, 26, 6].

We initially focus on the frequency moment estimation problem in the *random-order* model, where the set of stream updates to the underlying frequency vector is worst case, but the order of their arrival is uniformly random. As usual, the algorithms we initially consider are only permitted a single pass over the stream, though we later relax this constraint to permit adversarial ordering of updates, as well as both positive and negative integer updates (with magnitude bounded by $\text{poly}(n)$) to the coordinates of the frequency vector in the *turnstile* model. Random-order streams have been shown to be a natural assumption for problems of sorting and selecting in limited space [32] and many other real-world applications [21, 22, 16, 13]. Interestingly, there has sometimes been a significant qualitative difference between random-order streams and adversarial (or arbitrary) order streams [28, 9, 11]. In particular for F_p -moment estimation, the best known lower bound is $\Omega\left(\frac{1}{\varepsilon^2}\right)$ [12]. While the best known upper bound for $p \in (0, 2]$ on arbitrary order insertion-only streams is $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log n\right)$ [1, 27, 26], Braverman *et. al.* [11] gave an algorithm for $p \in (0, 2)$ on random-order streams that only used $\tilde{\mathcal{O}}\left(\frac{1}{\varepsilon^2} + \log n\right)$ space, where for a function $g(n, \varepsilon)$, we use $\tilde{\mathcal{O}}(g(n, \varepsilon))$ to denote a function bounded by $g(n, \varepsilon) \cdot \text{polylog}(g(n, \varepsilon))$. They also gave an algorithm for F_2 -moment estimation that only uses $\tilde{\mathcal{O}}\left(\frac{1}{\varepsilon^2} + \log n\right)$ space, but requires the assumption that $F_2 \geq F_1 \cdot \log n$, i.e., the second moment must be a logarithmic factor larger than the length of the stream.

The above works raise a number of important questions. A tantalizing open question, studied extensively in [25], and dating back to the original work of Alon, Matias, and Szegedy [1], is the exact space complexity of F_p -moment estimation in insertion-only streams. For $p > 2$, the best known upper bound for F_p -moment estimation on arbitrary order insertion-only streams is the minimum of $\mathcal{O}\left(\frac{n^{1-2/p}}{\varepsilon^{20}}\right)$ [9] and $\tilde{\mathcal{O}}\left(\frac{1}{\varepsilon^2} \cdot n^{1-2/p}\right)$ [17, 19]. For insertion-only streams, the best known lower bound is $\tilde{\Omega}\left(\frac{n^{1-2/p}}{\varepsilon^2 \log n}\right)$ [18]. We summarize these results in Figure 1. There are also major gaps in our understanding for F_p -moment estimation

Reference	Space Complexity	Order
[24, 31, 2]	$\tilde{\mathcal{O}}\left(n^{1-2/p} \varepsilon^{-\mathcal{O}(1)}\right)$	Arbitrary
[9]	$\mathcal{O}\left(n^{1-2/p} \varepsilon^{-20}\right)$	Arbitrary
[3]	$\tilde{\mathcal{O}}\left(n^{1-2/p} \varepsilon^{-2-6/p}\right)$	Arbitrary
[5]	$\tilde{\mathcal{O}}\left(n^{1-2/p} \varepsilon^{-2-4/p}\right)$	Arbitrary
[17, 19]	$\tilde{\mathcal{O}}\left(n^{1-2/p} \varepsilon^{-2}\right)$	Arbitrary
This work	$\tilde{\mathcal{O}}\left(n^{1-2/p} \varepsilon^{-4/p}\right)$	Random or Two-Pass Arbitrary
[1, 33]	$\Omega\left(n^{1-5/p} + \varepsilon^{-2}\right)$	Arbitrary
[14]	$\Omega\left(n^{1-2/p} \varepsilon^{-2/p}\right)$	Arbitrary
[34]	$\Omega\left(n^{1-2/p} \varepsilon^{-4/p} / \log^{\mathcal{O}(1)} n\right)$	Arbitrary, $\mathcal{O}(1)$ -Passes
[18]	$\Omega\left(n^{1-2/p} \varepsilon^{-2} / \log n\right)$	Arbitrary
[12]	$\Omega\left(n^{1-2.5/p} + \varepsilon^{-2}\right)$	Random
This work	$\Omega\left(n^{1-2/p} \varepsilon^{-2}\right)$	Arbitrary

■ **Figure 1** Summary of recent work for large frequency moment estimation.

for $p > 2$ in random order streams. The best known¹ lower bound is $\Omega(n^{1-2.5/p} + \varepsilon^{-2})$ [12], while no upper bounds that do better in random order streams than in arbitrary insertion streams are known.

1.1 Our Results

In this paper, we show a separation not only between random-order and arbitrary insertion-only streams for the F_p moment estimation problem, but also one-pass and multi-pass streams. We first show improved bounds for the F_p moment estimation problem for every $p > 2$ in the random-order insertion-only streaming model.

► **Theorem 1.1.** *For $p > 2$, there exists an algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment of a random-order insertion-only stream with probability at least $\frac{2}{3}$, while using total space (in bits) $\tilde{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p}\right)$.*

Theorem 1.1 utilizes the random-order model to improve the algorithm on arbitrary-order insertion-only streams using $\tilde{O}\left(\frac{1}{\varepsilon^2} \cdot n^{1-2/p}\right)$ space [17, 19], in terms of the dependence on $\frac{1}{\varepsilon}$. We then give an algorithm that uses roughly the same bounds even for the two-pass streaming model, even if the order of the updates is adversarial.

► **Theorem 1.2.** *For $p > 2$, there exists a two-pass streaming algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment with probability at least $\frac{2}{3}$.*

■ *If the stream is insertion-only, the algorithm uses $\tilde{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p}\right)$ bits of space (see Theorem 3.6).*

■ *If the stream has turnstile updates, the algorithm uses $\tilde{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p}\right)$ bits of space (see Theorem 3.5).*

Theorem 1.2 is the first algorithm to match the lower bound of $\tilde{\Omega}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p}\right)$ for multi-pass frequency moment estimation [34] up to $\log n$ factors.

By contrast, we give a lower bound for F_p estimation in the one-pass insertion-only streaming model when the order of the updates is adversarial.

► **Theorem 1.3.** *For any constant $p > 2$ and parameter $\varepsilon = \Omega\left(\frac{1}{n^{1/p}}\right)$, any one-pass insertion-only streaming algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment of an underlying frequency vector with probability at least $\frac{9}{10}$ requires $\Omega\left(\frac{1}{\varepsilon^2} \cdot n^{1-2/p}\right)$ bits of space.*

Theorem 1.3 improves the lower bound of $\Omega\left(\frac{1}{\varepsilon^2 \log n} \cdot n^{1-2/p}\right)$ for general insertion-only streams for $p > 2$ by [17]. Together, Theorem 1.1 and Theorem 1.3 for small enough $\varepsilon > 0$ show a somewhat surprising result that random-order streams are strictly easier than arbitrary insertion-only streams. Similarly, Theorem 1.2 and Theorem 1.3 together show that multiple passes are strictly easier than a single pass on arbitrary insertion-only streams.

1.2 Our Techniques

We first describe our random-order insertion-only algorithm. At a high level, our algorithm interleaves the recent heavy-hitters algorithm of [7] with a subsampling procedure to estimate the contributions of various level sets toward the frequency moment.

¹ After discussion with the authors, there appears to be an error in [20], which claims a stronger lower bound.

Approximate frequencies of heavy-hitters. The heavy-hitters algorithm of [7] partitions the updates of a random-order stream into blocks of updates. It then randomly chooses a number of coordinates from the universe $[n]$ to test in each block. The heavy-hitters will pass the test and subsequently be tracked across a number of following blocks before estimates for their frequencies are output. Due to the uniformity properties of the random-order stream, the heavy-hitters are sufficiently “spread out”, so the algorithm crucially outputs a $(1 + \varepsilon)$ -approximation to the frequency of each heavy-hitter; algorithms with the same guarantee and space complexity for arbitrary-order streams do not exist [18].

Using level sets to estimate frequency moments. Given the subroutine for finding $(1 + \varepsilon)$ -approximations to the frequencies of heavy-hitters, we now build upon a standard subsampling approach to estimate the frequency moment. Informally, we conceptually define the level sets so that level set Λ_i contains the coordinates $k \in [n]$ such that $f_k^p \in \left[\frac{F_p}{2^i}, \frac{2F_p}{2^i}\right]$. Since the level sets partition the universe, it is easy to see that if we define the contribution C_i of a level set Λ_i by the sum of the contributions of all their coordinates, $C_i := \sum_{k \in \Lambda_i} f_k^p$, then F_p is just the sum of all the contributions of the level sets, $F_p = \sum_i C_i$.

[24] showed that the contributions of each “significant” level set can be estimated by subsampling at exponentially smaller rates and considering the approximate frequencies of the heavy-hitters in each of the subsamples. For example, a single item with contribution F_p will be detected at the top level, while n items with contribution 1 will be detected at a subsampling level where there are roughly $\Theta\left(\frac{1}{\varepsilon^p}\right)$ survivors in expectation. Crucially, $(1 + \varepsilon)$ -approximations to the contribution of the surviving heavy-hitters in each subsampling level can then be rescaled by the sampling rate to obtain “good” approximations to the contributions of each significant level set; these very good estimates are not available in standard subsampling schemes for arbitrary order streams.

We adapt this approach for $p > 2$ to obtain Theorem 1.1. The first observation is that an item i with $f_i^p \geq \varepsilon^2 F_p$ should be identified to control the variance but also satisfies $f_i^2 \geq \varepsilon^{4/p} / n^{1-2/p} \cdot F_2$, so we can identify these items using the heavy-hitter algorithm of [7] with the corresponding threshold; this induces the overall $n^{1-2/p} / \varepsilon^{4/p}$ dependency. Moreover as we subsample, the space of the universe decreases in expectation from n to $n/2$ to $n/4$ and so forth. Thus determining the L_p -heavy hitters at lower subsampling rates can be done using significantly less space. We can take advantage of this by requiring that our heavy-hitter algorithm aggressively seeks heavy-hitters with lower thresholds at lower subsampling rates. We can thus achieve a geometric series and avoid an additional $\mathcal{O}(\log n)$ factor in our space.

From random-order to two-pass arbitrary order. As stated, our algorithm necessitates the random-order model so that the heavy-hitter subroutine can output $(1 + \varepsilon)$ -approximations to the frequencies of heavy-hitters across different subsampling levels; these approximations are then used to obtain $(1 + \varepsilon)$ -approximations to the contributions of each level set. The state-of-the-art heavy-hitter algorithms in insertion-only [8] or turnstile [15] streams with arbitrary arrival order do not give a $(1 + \varepsilon)$ -approximation to the frequency of each heavy-hitter while still using space dependency $\frac{1}{\varepsilon^2}$. Fortunately, our approach can be remedied in two-passes over the data stream by first using a pass to identify each of the heavy-hitters across the different subsampling levels and then using the second pass to exactly count their frequencies; note that since $n^{1-2/p} / \varepsilon^{4/p} \geq \frac{1}{\varepsilon^2}$ for $n \geq \frac{1}{\varepsilon^2}$, we are still permitted space to track the frequencies of $\min\left(n, \frac{1}{\varepsilon^2}\right)$ items. We can then obtain $(1 + \varepsilon)$ -approximations to the contributions of each significant level set, thus obtaining a $(1 + \varepsilon)$ -approximation to the F_p frequency moment.

Source of the separation. In summary, our improved upper bounds in both the random-order and multi-pass models exploit each model to obtain $(1 + \varepsilon)$ -approximate frequencies of the heavy-hitters in each subsampling level. Due to the uniformity of heavy-hitters across the stream in the random-order model, we are able to obtain $(1 + \varepsilon)$ -approximate frequencies by simply tracking the frequency of the heavy-hitters within a small block of the stream and then scaling by the entire length of the stream. For multi-pass models, we are able to identify heavy-hitters in the first pass and exactly track their frequencies in the second pass. By contrast, obtaining $(1 + \varepsilon)$ -approximate frequencies in adversarially ordered insertion-only streams with the same space guarantees in a single pass cannot be done [18].

Lower bound. To prove our improved lower bound, we first recall the standard approach for showing F_p moment estimation lower bounds in insertion-only streams for $p > 2$ that uses the multiplayer set disjointness problem, e.g., [4, 18]. In this problem, t players have binary vectors of length n and the promise is that the largest coordinate in the sum of all vectors is either (1) at most 1 or (2) exactly t . For $t = \varepsilon^{1/p} n^{1/p}$, the F_p frequency moment of the sum of the binary vectors differs by a factor of $(1 + \varepsilon)$ between the two cases, so that a $(1 + \varepsilon)$ -approximation to F_p solves the multiplayer set disjointness problem. The total communication complexity of the multiplayer set disjointness problem is $\Omega\left(\frac{n}{t}\right)$ so one of the t players communicates $\Omega\left(\frac{n}{t^2}\right)$ bits, and thus a lower bound of $\Omega\left(\frac{1}{\varepsilon^{2/p}} \cdot n^{1-2/p}\right)$ follows.

To improve the ε dependency in the lower bound to $\frac{1}{\varepsilon^2}$, we define the (t, ε, n) -player set disjointness estimation problem so that there are $t + 1$ players P_1, \dots, P_{t+1} with private coins in the standard blackboard model. The first t players each receive a vector $\mathbf{v}_s \in \{0, 1\}^n$ for $s \in [t]$, while player P_{t+1} receives an index $j \in [n]$ and a bit $c \in \{0, 1\}$. For $\mathbf{u} = \sum_{s \in [t]} \mathbf{v}_s$, the inputs are promised to satisfy $u_i \leq 1$ for each $i \neq j$ and either $u_j = 1$ or $u_j = t$, similar to the multiparty set disjointness problem. With probability at least $\frac{9}{10}$, P_{t+1} must differentiate between the three possible input cases (1) $u_j + \frac{ct}{\varepsilon} \leq t$, (2) $u_j + \frac{ct}{\varepsilon} \in \left\{\frac{t}{\varepsilon}, \frac{t}{\varepsilon} + 1\right\}$, or (3) $u_j + \frac{ct}{\varepsilon} = (1 + \varepsilon)\frac{t}{\varepsilon}$, where $\varepsilon \in (0, 1)$. We call coordinate $j \in [n]$ the *spike* location and show that the (t, ε, n) -player set disjointness estimation problem requires $\Omega\left(\frac{n}{t}\right)$ total communication by using a direct sum embedding to decompose the conditional information complexity into a sum of n single coordinate problems. The intuition is that the first t players do not know the spike location, so they must effectively solve the problem on each coordinate. We then bound the conditional information complexity of each single coordinate problem by the Hellinger distances between inputs for which the outputs differ. We thus apply the same reduction and set $t = \Theta\left(\frac{1}{\varepsilon} \cdot n^{1/p}\right)$ to obtain the desired lower bound.

We remark that the (t, ε, n) -player set disjointness estimation problem can be seen as a generalization of the augmented L_∞ promise problem introduced by [30]. In the augmented L_∞ promise problem, there are only three players, but each coordinate in the first two players' input vectors can be as large as εk for some fixed parameter k . [30] used the augmented L_∞ promise problem to give a lower bound of $\Omega\left(\frac{\log n}{\varepsilon^2} \cdot n^{1-2/p}\right)$ for F_p moment estimation on turnstile streams. However, since their reduction crucially allows players to use turnstile updates, we cannot adapt their techniques to obtain a reduction for insertion-only streams.

1.3 Preliminaries

For a positive integer n , we use the notation $[n]$ to denote the set of integers $\{1, 2, \dots, n\}$. For a frequency vector f with dimension n and $p \geq 2$, we define the F_p moment function by $F_p(f) := \sum_{k \in [n]} |f_k|^p$ and the L_p norm of f by $L_p(f) = (F_p(f))^{1/p}$. When f is defined through the updates of an insertion-only data stream, we simply use F_p to denote $F_p(f)$.

Thus for a subset $I \subseteq [n]$, the notation $F_p(I)$ is understood to mean $\sum_{k \in I} f_k^p$. We also use the notation $\|f\|_p$ to denote the L_p norm of f . The F_p moment estimation problem and the norm estimation problem are often used interchangeably, as a $(1 + \varepsilon)$ -approximation algorithm to one of these problems can be modified to output a $(1 + \varepsilon)$ -approximation to the other using a rescaling of ε , for constant $p > 0$.

We use \log and \ln to denote the base two logarithm and natural logarithms, respectively. We use $\text{poly}(n)$ to denote a fixed constant degree polynomial in n and $\frac{1}{\text{poly}(n)}$ to denote some arbitrary degree polynomial in n corresponding to the choice of constants in the algorithms. We use $\text{polylog}(n)$ to denote polylogarithmic factors of n .

The L_p -heavy hitters problem is to output all coordinates i such that $f_i \geq \varepsilon \cdot L_p$; such a coordinate is called a *heavy-hitter*. The problem allows coordinates j with $f_j \leq \varepsilon \cdot L_p$ to be output, provided that $f_j \geq \frac{\varepsilon}{2} \cdot L_p$. Moreover, each coordinate output by the algorithm must also have a frequency estimation with additive error at most $\frac{\varepsilon}{2} \cdot L_p$.

2 F_p Estimation for $p > 2$ in Random-Order Streams

In this section, we give our F_p estimation algorithm for $p > 2$. We first introduce an algorithm COUNTHH on random-order insertion-only streams that outputs an approximate frequency for each L_2 heavy-hitter.

► **Theorem 2.1** (Theorem 28 in [7]). *There exists a one-pass algorithm COUNTHH on random-order insertion-only streams that outputs a list H of ordered pairs (j, \hat{f}_j) such that $\hat{f}_j = (1 \pm 2\varepsilon)f_j$, for each $j \in H$, where f is the underlying frequency vector. Moreover, we have that $j \in H$ for each $j \in [n]$ with $f_j^2 \geq \varepsilon^2 \cdot F_2$ and $j \notin H$ for each j with $f_j^2 \leq \frac{\varepsilon^2}{2} \cdot F_2$. The algorithm uses $\mathcal{O}\left(\frac{1}{\varepsilon^2} (\log^2 \frac{1}{\varepsilon} + \log^2 \log m + \log n) + \log m\right)$ bits of space and succeeds with probability at least $38/39$.*

A standard heavy-hitter algorithm such as COUNTSKETCH [15] or BPTREE [8] outputs each $j \in [n]$ with $f_j^2 \geq \frac{\varepsilon^2}{2} \cdot F_2$ but only finds a constant-factor approximation to the frequency of each heavy-hitter. By comparison, COUNTHH crucially uses the random-order stream to output a $(1 + \varepsilon)$ -approximation to the frequency of each heavy-hitter; we defer the full details of the algorithm to the full version of the paper [36]. We use COUNTHH in a subsampling approach to approximate the contributions of each of the level sets, defined as follows:

► **Definition 2.2** (Level sets and contribution). *Given \widehat{F}_p such that $\widehat{F}_p \leq F_p \leq 1.01 \cdot \widehat{F}_p$ and a uniformly random $\zeta \in [1, 2]$, we define the level set Λ_i for each $i \in [\log 4n]$ so that*

$$\Lambda_i := \left\{ k \mid f_k^p \in \left[\frac{\zeta \cdot \widehat{F}_p}{2^i}, \frac{2\zeta \cdot \widehat{F}_p}{2^i} \right] \right\}.$$

Then we define the contribution C_i of level set Λ_i to be $C_i := \sum_{k \in \Lambda_i} f_k^p$. We define the fractional contribution ϕ_i of level set Λ_i to be the ratio $\phi_i := \frac{C_i}{\zeta \widehat{F}_p}$, so that $\phi_i \in [0, 1.01]$.

For a stream of length $m = \text{poly}(n)$, let α be an integer such that $2^\alpha > m^p$. We say a level set Λ_i is significant if its fractional contribution ϕ_i is at least $\frac{\varepsilon}{2^\alpha \log n}$. Otherwise, we say the level set is insignificant.

Observe that it suffices to obtain a multiplicative $(1 + \frac{\varepsilon}{2})$ -approximation to the contribution of each significant level set and estimate the contributions of the insignificant level sets to be zero, since there are at most $\alpha \log n$ level sets and thus the total additive error from the insignificant level sets is at most $\frac{\varepsilon}{2^\alpha \log n} \cdot F_p \cdot \alpha \log n = \frac{\varepsilon}{2} \cdot F_p$.

To estimate the contribution of each level set, we use a combination of COUNTHH and subsampling to approximate the frequencies of a number of heavy-hitters. The main idea is that subsampling induces a separate substream with a frequency vector with a smaller F_p . Thus the items in the level sets Λ_i with small i will be heavy-hitters of the frequency vector while the items in the level sets Λ_i with larger i will be heavy-hitters of the substreams with smaller sampling rates (if the items are subsampled), due to the lower F_p of the substreams. We can then use a $(1 + \varepsilon)$ -approximation to the frequency of each sampled heavy-hitter to estimate a $(1 + \varepsilon)$ -approximation to the contribution of the level set, due to uniformity properties of random-order streams. Note that this is where we crucially use COUNTHH over other possible heavy-hitter algorithms, due to its advantage of providing $(1 + \varepsilon)$ -approximate frequencies in the random-order model. We describe our algorithm for F_p estimation with $p > 2$ for random-order insertion-only streams in Algorithm 1.

■ **Algorithm 1** F_p Estimation in the Random-Order Insertion-Only Model, $p > 2$.

Input: Accuracy parameter $\varepsilon \in (0, 1)$, $\widehat{F}_p \leq F_p \leq 1.01 \cdot \widehat{F}_p$

Output: $(1 + \varepsilon)$ -approximation to F_p .

- 1: Let $\zeta \in [1, 2]$ be chosen uniformly at random and α be a positive integer such that $2^\alpha > m^p$.
- 2: Let $\eta > \sum_{j \geq 0} 2^{-j(p/16-1/8)}$ be a sufficiently large constant.
- 3: $\gamma \leftarrow 2^{11}$, $n_i \leftarrow \min \left(\left(\frac{16\alpha p \log n}{\varepsilon^{1-2/p}} \right)^{(2p)/(p-2)}, \frac{10\gamma n}{2^i} \right)$, $\varepsilon_i = \frac{\varepsilon}{16\eta \cdot 2^{i(p/16-1/8)} \log \frac{1}{\varepsilon^2}}$
- 4: Let I_i^r be a (nested) subset of $[n]$ subsampled at rate $p_i := \min(1, 2^{-i\gamma})$.
- 5: Let H_i^r be the output of COUNTHH with threshold parameter $\frac{(\varepsilon_i)^{2/p}}{80\gamma} \cdot \left(\frac{1}{n_i}\right)^{1/2-1/p}$ on the substream induced by I_i^r .
- 6: **for** $i \in [\alpha \log n]$, $r \leq \mathcal{O}(\log \log n)$ **do**
- 7: $\ell_i := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$
- 8: Let S_i^r be the set of ordered pairs (j, \widehat{f}_j) in $H_{\ell_i}^r$ with $(\widehat{f}_j)^p \in \left[\frac{\zeta \widehat{F}_p}{2^i}, \frac{2\zeta \widehat{F}_p}{2^i} \right]$.
- 9: $\widehat{C}_i \leftarrow \text{median}_r \frac{1}{p\ell_i} \cdot \left(\sum_{(j, \widehat{f}_j) \in S_i^r} (\widehat{f}_j)^p \right)$
- 10: **return** $\widetilde{F}_p := \sum_i \widehat{C}_i$

► **Remark 2.3.** We remark that Algorithm 1 is written requiring an input \widehat{F}_p such that $\widehat{F}_p \leq F_p \leq 1.01 \cdot \widehat{F}_p$ in order to correctly index the level sets defined by Definition 2.2. Algorithm 1 can be rewritten by setting $X = (F_1)^p$ to be an upper bound on F_p and redefining the level sets in Definition 2.2 so that $\Lambda_i := \left\{ k \mid f_k^p \in \left[\frac{\zeta \cdot X}{2^i}, \frac{2\zeta \cdot X}{2^i} \right] \right\}$. Then we again have $F_p = \sum_i C_i$ across the contributions of all the level sets.

We first show that there exists a subsampling rate such that items in each level set will be reported as a heavy-hitter if they are successfully subsampled. Moreover, each item reported as a heavy-hitter will also be reported with an “accurate” estimate of its frequency.

► **Lemma 2.4.** *Let $\varepsilon \in (0, 1)$, Λ_i be a fixed level set, and let $\ell := \ell_i := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$. For a fixed r , let \mathcal{E}_1 be the event that $|I_\ell^r| \leq \frac{10\gamma n}{2^\ell}$ and let \mathcal{E}_2 be the event that $F_p(I_\ell^r) \leq \frac{10\gamma \widehat{F}_p}{2^\ell}$. Conditioned on \mathcal{E}_1 and \mathcal{E}_2 , there exists a (j, \widehat{f}_j) in S_i^r for each $j \in \Lambda_i \cap I_\ell^r$ such that with probability at least $\frac{7}{8}$,*

$$\left(1 - \frac{\varepsilon}{8\alpha \log n}\right) \cdot f_j^p \leq (\widehat{f}_j)^p \leq \left(1 + \frac{\varepsilon}{8\alpha \log n}\right) \cdot f_j^p.$$

Proof. Recall that Algorithm 1 considers H_ℓ^r across $r \leq \mathcal{O}(\log \log n)$ independent subsamples of $[n]$, to construct \widehat{C}_i by subsampling at rate $p_\ell := \min(1, 2^{-\ell}\gamma)$.

Suppose $2^i \cdot \varepsilon^2 < \gamma$, so that $\ell := \ell_i = \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\} \leq \log \gamma$ since $\varepsilon_k \leq \varepsilon$ for all k . Then $p_\ell = \min(1, 2^{-\ell}\gamma) = 1$ and all items are subsampled, i.e., $H_\ell^r = [n]$. By Definition 2.2 of the level sets, each item $j \in \Lambda_i$ satisfies $f_j^p \in \left[\frac{\zeta \cdot \widehat{F}_p}{2^i}, \frac{2\zeta \cdot \widehat{F}_p}{2^i} \right]$. Since $\varepsilon^2 \geq \frac{1}{2^i}$, then we have $f_j^p \geq \varepsilon^2 \cdot \widehat{F}_p$. We also have $\widehat{F}_p \leq F_p \leq 1.01 \cdot \widehat{F}_p$ and $n^{1/2-1/p} \cdot F_p^{1/p} \geq F_2^{1/2}$. Thus, each item in $j \in \Lambda_i$ satisfies

$$f_j \geq \frac{\varepsilon^{2/p}}{1.01^{1/p}} \cdot F_p^{1/p} \geq \frac{\varepsilon^{2/p}}{1.01^{1/p} n^{1/2-1/p}} \cdot F_2^{1/2},$$

so that $f_j^2 \geq \frac{\varepsilon^{4/p}}{1.01^{2/p} n^{1-2/p}} F_2$. Then by Theorem 2.1, each item $j \in \Lambda_i$ corresponds to some estimate $(\widehat{f}_j)^2$ reported by H_ℓ^r output by COUNTHH with threshold $\frac{(\varepsilon_\ell)^{2/p}}{80\gamma}$, as $\varepsilon_\ell \leq \varepsilon$ and $n_\ell = n$. Hence, \widehat{f}_j^p is a $\left(1 + \frac{\varepsilon}{8\alpha \log n}\right)$ -approximation to f_j^p . Furthermore, COUNTHH rounds the estimate of the frequency of each heavy-hitter to the nearest power of $\left(1 + \frac{\varepsilon}{16p\alpha \log n}\right)$. Hence, \widehat{f}_j^p is a $\left(1 + \frac{\varepsilon}{8\alpha \log n}\right)$ -approximation to f_j^p .

Now suppose $2^i \cdot \varepsilon^2 \geq \gamma$, so that $\ell := \ell_i = \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\} \geq \log \gamma$ and $p_\ell := \min(1, 2^{-\ell}\gamma) \geq \frac{\gamma}{2 \cdot 2^i \varepsilon^2}$. Again by Definition 2.2 of the level sets, each item $j \in \Lambda_i$ satisfies $f_j^p \in \left[\frac{\zeta \cdot \widehat{F}_p}{2^i}, \frac{2\zeta \cdot \widehat{F}_p}{2^i} \right]$. Now if $j \notin I_\ell^r$, then it will not be reported by COUNTHH. Thus we assume that $j \in I_\ell^r$.

Conditioning on the event \mathcal{E}_2 , we have that

$$F_p(I_\ell^r) \leq \frac{10\gamma F_p}{2^\ell} \leq \frac{20\gamma F_p}{2^i \varepsilon^2}.$$

Since $\widehat{F}_p \leq F_p \leq 1.01 \cdot \widehat{F}_p$, then we have that $f_j^p \geq \frac{\varepsilon^2}{80\gamma} \cdot F_p(I_\ell^r)$ and thus

$$f_j \geq \frac{\varepsilon^{2/p}}{(80\gamma)^{1/p}} \cdot F_p^{1/p}(I_\ell^r).$$

Instead of applying the inequality $n^{1/2-1/p} \cdot F_p^{1/p} \geq F_2^{1/2}$, we note that conditioning on the event \mathcal{E}_1 , we have that $|I_\ell^r| \leq n_\ell = \frac{10\gamma n}{2^\ell}$. Hence, the frequency vector defined by the substream I_ℓ^r potentially has much smaller support size than n . Thus we have

$$f_j \geq \frac{\varepsilon^{2/p}}{(80\gamma)^{1/p}} \cdot \left(\frac{1}{n_\ell}\right)^{1/2-1/p} \cdot F_2^{1/2}(I_\ell^r).$$

By Theorem 2.1, each item $j \in \Lambda_i \cap I_\ell^r$ corresponds to some estimate \widehat{f}_j reported by H_ℓ^r output by COUNTHH with threshold $\frac{(\varepsilon_\ell)^{2/p}}{80\gamma} \cdot \left(\frac{1}{n_\ell}\right)^{1/2-1/p}$, since $\varepsilon_\ell \leq \varepsilon$. Moreover, the estimate of the frequency of each heavy-hitter reported by COUNTHH is within a factor of $\left(1 + \frac{\varepsilon}{16p\alpha \log n}\right)$ of the true frequency, since $n_i \geq \left(\frac{16\alpha p \log n}{\varepsilon^{1-2/p}}\right)^{(2p)/(p-2)}$ implies that the threshold is at most $\frac{\varepsilon}{16p\alpha \log n}$. Hence for sufficiently small ε , $(\widehat{f}_j)^p$ is a $\left(1 + \frac{\varepsilon}{8\alpha \log n}\right)$ -approximation to f_j^p . ◀

In summary, an item $k \in \Lambda_i$ may not always be sampled by I_ℓ^r , but COUNTHH outputs a quantity \widehat{f}_k such that $(\widehat{f}_k)^p$ is a $\left(1 + \frac{\varepsilon}{8\alpha \log n}\right)$ -approximation to f_k^p if $k \in I_\ell^r$.

These approximations allow us to recover a $(1 + \varepsilon)$ -approximation to the F_p moment through Lemma 2.5, which is our main correctness statement and which we now show. Lemma 2.5 claims that the output \tilde{F}_p of our algorithm serves as a $(1 + \varepsilon)$ -approximation to the F_p moment of the underlying frequency vector. The proof of Lemma 2.5 first considers an idealized process and shows that we obtain “good” approximations to the contributions of each significant level set. Since the contributions of the insignificant level sets can be ignored, we thereby obtain a $(1 + \varepsilon)$ -approximation to the F_p moment. We then show that when the process is not idealized, the estimate \tilde{F}_p only occurs a small error and thus still guarantees a $(1 + \varepsilon)$ -approximation to the F_p moment.

► **Lemma 2.5.** *With probability at least $\frac{2}{3}$, we have that $|\tilde{F}_p - F_p| \leq \varepsilon \cdot F_p$.*

Proof. Let Λ_i be a fixed level set and let $\ell := \ell_i := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$. Let $k \in \Lambda_i \cap S_\ell^r$, so that $f_k^p \in \left[\frac{\zeta \cdot \widehat{F}_p}{2^i}, \frac{2\zeta \cdot \widehat{F}_p}{2^i} \right]$ and k is subsampled at the level in which its estimated frequency $(\widehat{f}_k)^p$ is used to estimate the contribution C_i of level set Λ_i . Then Lemma 2.4 shows that we obtain an estimate $(\widehat{f}_k)^p$ such that

$$\left(1 - \frac{\varepsilon}{8\alpha \log n}\right) \cdot f_k^p \leq (\widehat{f}_k)^p \leq \left(1 + \frac{\varepsilon}{8\alpha \log n}\right) \cdot f_k^p,$$

with constant probability. In an idealized process, we would have

$$\frac{\zeta \cdot \widehat{F}_p}{2^i} \leq (\widehat{f}_k)^p \leq \frac{2\zeta \cdot \widehat{F}_p}{2^i},$$

so that the estimate $(\widehat{f}_k)^p$ is used toward the estimation \widehat{C}_i of contribution C_i of level set Λ_i . However, this may not always be the case because the value of f_k^p may be near the boundary of the interval $\left[\frac{\zeta \cdot \widehat{F}_p}{2^i}, \frac{2\zeta \cdot \widehat{F}_p}{2^i} \right]$ and the value of the estimate $(\widehat{f}_k)^p$ may lie outside of the interval, so that the estimate $(\widehat{f}_k)^p$ is used toward the estimation of some other level set $\Lambda_{i'}$. We first analyze an idealized process, so that $(\widehat{f}_k)^p$ is correctly classified for all k across all level sets, and show that the output is a $(1 + \mathcal{O}(\varepsilon))$ -approximation to F_p . We then argue that because we randomize the boundaries of each interval due to the selection of ζ , the overall guarantee is only slightly worsened but remains a $(1 + \varepsilon)$ -approximation to F_p .

Idealized process. We first show that for an idealized process where $(\widehat{f}_k)^p$ is correctly classified for all k across all level sets, then for a fixed level set i , we have $|\widehat{C}_i - C_i| \leq \frac{\varepsilon_\ell}{2} \cdot F_p$ with probability at least $1 - \frac{1}{\text{polylog}(n)}$. Let \mathcal{E}_1 be the event that $|I_\ell^r| \leq \frac{10\gamma n}{2^\ell}$ and let \mathcal{E}_2 be the event that $F_p(I_\ell^r) \leq \frac{10\gamma F_p}{2^\ell}$. Lemma 2.4 shows that conditioned on \mathcal{E}_1 and \mathcal{E}_2 , then COUNTHH outputs a $\left(1 + \frac{\varepsilon}{8\alpha \log n}\right)$ -approximation to f_k^p if $k \in I_\ell^r$. We thus analyze the approximation \widehat{C}_i^r to C_i for a given set of subsamples, where we define

$$\widehat{C}_i^r := \frac{1}{p_{\ell_i}} \cdot \sum_{(k, f_k) \in S_i^r} (\widehat{f}_k)^p,$$

112:10 Separations for Estimating Large Frequency Moments on Data Streams

so that $\widehat{C}_i = \text{median}_r \widehat{C}_i^r$. Conditioned on \mathcal{E}_1 and \mathcal{E}_2 , we note that \widehat{C}_i^r is a $\left(1 + \frac{\varepsilon}{8\alpha \log n}\right)$ -approximation to

$$D_i^r := \frac{1}{p_{\ell_i}} \cdot \sum_{k \in S_i^r \cap \Lambda_i} f_k^p.$$

We thus analyze the expectation and variance of D_i^r .

We first analyze the expectation of D_i^r . Note that

$$\mathbb{E}[D_i^r] = \frac{1}{p_{\ell}} \cdot \sum_{k \in \Lambda_i} p_{\ell} \cdot f_k^p = C_i.$$

We next analyze the variance of D_i^r , which results from whether items $k \in \Lambda_i$ are sampled by I_{ℓ}^r . Since $p_{\ell} \geq \frac{\gamma}{2 \cdot 2^i \varepsilon_{\ell}^2}$, we have

$$\text{Var}(D_i^r) = \frac{1}{p_{\ell}^2} \cdot \sum_{k \in \Lambda_i} p_{\ell} f_k^{2p} \leq \sum_{k \in \Lambda_i} \frac{2 \cdot 2^i \varepsilon_{\ell}^2}{\gamma} \cdot f_k^{2p}.$$

Observe that for each $k \in \Lambda_i$, we have $f_k^{2p} \leq \frac{16(F_p)^2}{2^{2i}}$ and $\frac{|\Lambda_i|}{2^i} \leq \phi_i \leq 1$. Thus for $\gamma = 2^{11}$,

$$\text{Var}(D_i^r) \leq |\Lambda_i| \cdot \frac{2 \cdot 2^i \varepsilon_{\ell}^2}{\gamma} \cdot \frac{16(F_p)^2}{2^{2i}} \leq \phi_i \varepsilon_{\ell}^2 (F_p)^2 \leq \frac{\varepsilon_{\ell}^2}{64} (F_p)^2.$$

Hence, by Chebyshev's inequality, we have that

$$\Pr \left[|D_i^r - C_i| \geq \frac{\varepsilon_{\ell}}{2} \cdot F_p \right] \leq \frac{1}{16}.$$

Since $C_i \leq F_p$ and \widehat{C}_i^r is a $\left(1 + \frac{\varepsilon}{8\alpha \log n}\right)$ -approximation to D_i^r , then \widehat{C}_i^r gives an approximation to C_i with additive error at most $\left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_{\ell}\right) \cdot F_p$ with probability at least $\frac{15}{16}$, conditioned on the events \mathcal{E}_1 and \mathcal{E}_2 , and the correctness of the subroutine COUNTTHH. By Theorem 2.1, the correctness of COUNTTHH occurs with probability at least $\frac{38}{39}$. By a union bound, we have that conditioned on \mathcal{E}_1 and \mathcal{E}_2 , then

$$\Pr \left[|\widehat{C}_i^r - C_i| \leq \left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_{\ell}\right) \cdot F_p \right] \geq \frac{7}{8}.$$

For a fixed r , let \mathcal{E}_1 be the event that $|I_{\ell}^r| \leq \frac{10\gamma n}{2^{\ell}}$ and let \mathcal{E}_2 be the event that $F_p(I_{\ell}^r) \leq \frac{10\gamma F_p}{2^{\ell}}$. Recall that I_{ℓ}^r is a nested subset of $[n]$ subsampled at rate $p_{\ell} := \min(1, 2^{-\ell} \gamma)$. Thus we have $\mathbb{E}[|I_{\ell}^r|] \leq \frac{\gamma n}{2^{\ell}}$, so that by Markov's inequality, we have that $\Pr[\mathcal{E}_1] \geq \frac{9}{10}$. Similarly, we have $\mathbb{E}[F_p(I_{\ell}^r)] \leq \frac{\gamma F_p}{2^{\ell}}$, so that by Markov's inequality, we have that $\Pr[\mathcal{E}_2] \geq \frac{9}{10}$. Hence by a union bound, $\Pr[\mathcal{E}_1 \wedge \mathcal{E}_2] \geq \frac{8}{10}$.

By Lemma 2.4, conditioned on the events \mathcal{E}_1 and \mathcal{E}_2 , we have that $|\widehat{C}_i^r - C_i| \leq \left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_{\ell}\right) \cdot F_p$ for a fixed r , with probability at least $\frac{7}{8}$. Thus by a union bound, we have that $|\widehat{C}_i^r - C_i| \leq \left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_{\ell}\right) \cdot F_p$ for a fixed r , with probability at least $\frac{5}{8}$. Since $\widehat{C}_i = \text{median}_r \widehat{C}_i^r$ across $r \leq \mathcal{O}(\log \log n)$ iterations, then we have that $|\widehat{C}_i - C_i| \leq \left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_{\ell}\right) \cdot F_p$ with probability at least $1 - \frac{1}{\text{polylog}(n)}$.

By a union bound over the $\alpha \log n$ level sets, then with probability at least $1 - \frac{1}{\text{polylog}(n)}$, we have that $|\widehat{C}_i - C_i| \leq \left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_{\ell_i}\right) \cdot F_p$ simultaneously for all $i \in [\alpha \log n]$, where $\ell_i := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$. We form our estimate \tilde{F}_p to F_p by setting $\tilde{F}_p := \sum_i \widehat{C}_i$ and we have $F_p = \sum_i C_i$. Since $i \in [\alpha \log n]$, $\ell_i := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$, and $\varepsilon_{\ell_i} = \frac{\varepsilon}{16\eta \cdot 2^{\ell_i(p/16-1/8)} \log \frac{1}{\varepsilon^2}}$, then

$$\begin{aligned} |\tilde{F}_p - F_p| &= \left| \sum_i \widehat{C}_i - \sum_i C_i \right| \leq \sum_i |\widehat{C}_i - C_i| \leq \sum_i \left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_{\ell_i} \right) \cdot F_p \\ &\leq \sum_{i \in [\alpha \log n]} \frac{\varepsilon}{4\alpha \log n} \cdot F_p + \sum_{i: \ell_i \leq \log \frac{1}{\varepsilon^2}} \varepsilon_{\ell_i} \cdot F_p + \sum_{i: \ell_i > \log \frac{1}{\varepsilon^2}} \varepsilon_{\ell_i} \cdot F_p \\ &\leq \frac{\varepsilon}{4} \cdot F_p + \log \frac{1}{\varepsilon^2} \cdot \frac{\varepsilon}{16 \log \frac{1}{\varepsilon^2}} \cdot F_p + \frac{\varepsilon}{16} \cdot F_p, \end{aligned}$$

where the last bound on the last term follows from $\eta > \sum_{j \geq 0} 2^{-j(p/16-1/8)}$. Thus we have that $|\tilde{F}_p - F_p| \leq \frac{\varepsilon}{2} \cdot F_p$ with probability at least $1 - \frac{1}{\text{polylog}(n)}$ in an idealized process.

Effects of randomized boundaries. We say that for a fixed r , an item $k \in [n]$ is *misclassified* if there exists a level set Λ_i such that

$$\frac{\zeta \cdot F_p}{2^i} \leq f_k^p \leq \frac{2\zeta \cdot F_p}{2^i},$$

but for an estimate $(\widehat{f}_k)^p$ output by COUNTHH on the set $S_{\ell_i}^r$, we have

$$\frac{\zeta \cdot F_p}{2^i} \leq (\widehat{f}_k)^p \leq \frac{2\zeta \cdot F_p}{2^i}.$$

Recall that by Lemma 2.4, we have for any fixed value of ζ that

$$\left(1 - \frac{\varepsilon}{8\alpha \log n}\right) \cdot f_k^p \leq (\widehat{f}_k)^p \leq \left(1 + \frac{\varepsilon}{8\alpha \log n}\right) \cdot f_k^p.$$

Since $\zeta \in [1, 2]$, then the probability that item $k \in [n]$ is misclassified is at most $\frac{\varepsilon}{2\alpha \log n}$. Moreover, in the event that item $k \in \Lambda_i$ is misclassified, it can only be misclassified into either level set Λ_{i+1} or level set Λ_{i-1} , since $(\widehat{f}_k)^p$ is a $\left(1 \pm \frac{\varepsilon}{8\alpha \log n}\right)$ multiplicative approximation to f_k^p .

Thus in the event that item $k \in [n]$ is misclassified, then $(\widehat{f}_k)^p$ will be rescaled by an incorrect probability, but only by at most a factor of two. Hence the error in the computation of the contribution of f_k^p to some level set Λ_i is at most $2f_k^p$. Then in expectation across all $k \in [n]$, the error due to the misclassification is at most $2F_p \cdot \frac{\varepsilon}{2\alpha \log n} = \frac{\varepsilon}{\alpha \log n} \cdot F_p$. Hence by Markov's inequality for sufficiently large n , the misclassification error is at most an additive $\frac{\varepsilon}{2} \cdot F_p$ with probability at least $\frac{3}{4}$. Therefore in total, we have that $|\tilde{F}_p - F_p| \leq \varepsilon \cdot F_p$ with probability at least $\frac{2}{3}$. \blacktriangleleft

It remains to analyze the space complexity of the algorithm as well as remove some additional unnecessary assumptions.

► Theorem 1.1. *For $p > 2$, there exists an algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment of a random-order insertion-only stream with probability at least $\frac{2}{3}$, while using total space (in bits) $\tilde{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p}\right)$.*

Proof. We first observe that we only require a 1.01-approximation \widehat{F}_p to F_p as the input of Algorithm 1 to create the level sets and analyze accordingly. However, we can instead define the level sets using any upper bound X on \widehat{F}_p by setting level set Λ_i to be the indices $k \in [n]$ for which $f_k^p \in \left[\frac{\zeta X}{2^i}, \frac{2\zeta X}{2^i}\right]$, rather than $\left[\frac{\zeta \widehat{F}_p}{2^i}, \frac{2\zeta \widehat{F}_p}{2^i}\right]$, and the same analysis will follow (see Remark 2.3). Intuitively, the fluidity of the definition of the level sets can be seen from the fact that we randomize the boundaries by going through a multiplicative ζ chosen randomly from $[1, 2]$ anyway, and additional empty level sets will not change the approximation guarantee. Thus by Lemma 2.5, there exists an algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment.

It remains to analyze the space complexity. By Theorem 2.1, COUNTHH with threshold ε requires $\mathcal{O}\left(\frac{1}{\varepsilon^2} (\log^2 \frac{1}{\varepsilon} + \log^2 \log m + \log n) + \log m\right)$ bits of space. Algorithm 1 runs a separate instance of COUNTHH with threshold $\frac{(\varepsilon_i)^{2/p}}{80\gamma} \cdot \left(\frac{1}{n_i}\right)^{1/2-1/p}$ to output a set H_i^r for $r \leq \mathcal{O}(\log \log n)$, where γ is a sufficiently large constant. Thus the total space for the COUNTHH instances across all i for a fixed r is at most

$$C_1 \log n + (C_1^2 \log n) \cdot \sum_{i \in [\alpha \log n]} \frac{C_1 (n_i)^{1-2/p}}{(\varepsilon_i)^{4/p}},$$

for some positive constants $C_1, \alpha > 0$. In particular, we have $\varepsilon_i = \frac{\varepsilon}{16\eta \cdot 2^{i(p/16-1/8)} \log \frac{1}{\varepsilon^2}}$ and $n_i = \min\left(\left(\frac{16\alpha p \log n}{\varepsilon^{1-2/p}}\right)^{(2p)/(p-2)}, \frac{10\gamma n}{2^i}\right)$ for a sufficiently large constant η . Then the total space for the COUNTHH instances across all i for a fixed r is at most

$$C_1 \log n + (C_1 \log^2 n) \cdot \sum_{i=1}^{\alpha \log n} \left(\frac{C_2 n^{1-2/p}}{\varepsilon^{4/p}} \cdot \frac{2^{i(1/4-1/(2p))}}{2^{i(1-2/p)}} \log^2 \frac{1}{\varepsilon} + \frac{C_2 \log^2 n}{\varepsilon^2} \right),$$

for some positive constants $C_1, C_2 > 0$. Observe that $\sum_{i=1}^{\infty} \frac{2^{i(1/4-1/(2p))}}{2^{i(1-2/p)}} = \sum_{i=1}^{\infty} \frac{1}{2^{3i(1-2/p)/4}}$ is a geometric series that is upper bounded by an absolute constant. Hence, the total space across all i for a fixed r is at most $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log^2 n \log^2 \frac{1}{\varepsilon} + \frac{1}{\varepsilon^2} \log^5 n\right)$ and the total space across all $r \leq \mathcal{O}(\log \log n)$ is

$$\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log^2 n \log^2 \frac{1}{\varepsilon} \log \log n + \frac{1}{\varepsilon^2} \log^5 n \log \log n\right),$$

which is $\tilde{\mathcal{O}}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p}\right)$ space in total, since $n \geq \frac{1}{\varepsilon^2}$ implies $\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \geq \frac{1}{\varepsilon^2}$. \blacktriangleleft

3 F_p Estimation for $p > 2$ in Two-Pass Streams

In this section, we consider two-pass algorithms for F_p estimation, with $p > 2$. For turnstile streams, we require the guarantees of the well-known COUNTSKETCH algorithm for finding heavy-hitters.

► **Theorem 3.1** ([15]). *There exists an algorithm COUNTSKETCH that reports all items $i \in [n]$ such that $f_i \geq \varepsilon L_2$ and no items $j \in [n]$ such that $f_j \leq \frac{\varepsilon}{2} \cdot L_2$ in the turnstile streaming algorithm. The algorithm uses $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log^2 n\right)$ bits of space and succeeds with probability $1 - \frac{1}{\text{poly}(n)}$.*

For insertion-only streams, we use the more space-efficient BPTREE algorithm for finding heavy-hitters.

► **Theorem 3.2** ([8]). *There exists an algorithm BPTREE that reports all items $i \in [n]$ such that $f_i \geq \varepsilon L_2$ and no items $j \in [n]$ such that $f_j \leq \frac{\varepsilon}{2} \cdot L_2$ in the turnstile streaming algorithm. The algorithm uses $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log n\right)$ bits of space and succeeds with probability 0.99.*

Recall that the main idea of our one-pass algorithm in the random-order insertion-only streaming model was to subsample at different rates and find $(1 + \varepsilon)$ -approximations to the frequencies of the heavy-hitters in each subsampling rate, which are then used to form estimates of the contributions of each level set and ultimately estimate of the frequency moment. The random-order setting crucially allowed us to obtain $(1 + \varepsilon)$ -approximations to the frequencies of the heavy-hitters. By contrast, in an adversarial-order setting, by the time a possible heavy-hitter is detected, a significant fraction of its frequency may have already appeared in the stream if we use a heavy-hitter algorithm with space dependency $\frac{1}{\varepsilon^2}$. Fortunately, in a two-pass setting, we can use the first pass to identify possible heavy-hitters and the second pass to find their frequencies. We give the full details in Algorithm 2, where the subroutine HEAVYHITTERS denotes the algorithm COUNTSKETCH of Theorem 3.1 for turnstile streams and BPTREE of Theorem 3.2 for insertion-only streams.

■ **Algorithm 2** F_p Estimation on Two-Pass Streams, $p > 2$.

Input: Accuracy parameter $\varepsilon \in (0, 1)$, $\widehat{F}_p \leq F_p \leq 1.01 \cdot \widehat{F}_p$

Output: $(1 + \varepsilon)$ -approximation to F_p

- 1: Let $\eta > \sum_{j \geq 0} 2^{-j(p/16-1/8)}$ be a sufficiently large constant.
- 2: $\gamma \leftarrow 2^{11}$, $n_i \leftarrow \frac{10\gamma n}{2^i}$, $\varepsilon_i = \frac{\varepsilon}{16\eta \cdot 2^{i(p/16-1/8)} \log \frac{1}{\varepsilon^2}}$
- 3: Let I_i^r be a (nested) subset of $[n]$ subsampled at rate $p_i := \min(1, 2^{-i}\gamma)$.
- 4: **for first pass** $i \in [\alpha \log n]$, $r \leq \mathcal{O}(\log \log n)$: **do**
- 5: Let H_i^r be the output of HEAVYHITTERS with threshold parameter $\frac{(\varepsilon_i)^{2/p}}{80\gamma}$.
 $\left(\frac{1}{n_i}\right)^{1/2-1/p}$ on the substream induced by I_i^r .
- 6: **for second pass do**
- 7: Track the frequency f_k for any coordinate $k \in \cup H_i^r$.
- 8: Let S_i^r be the items $k \in [n]$ with $f_k^p \in \left[\frac{\widehat{F}_p}{2^i}, \frac{2\widehat{F}_p}{2^i}\right]$.
- 9: $\ell_i \leftarrow \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$
- 10: $D_i^r \leftarrow \frac{1}{p_{\ell_i}} \cdot \left(\sum_{k \in S_i^r} f_k^p\right)$
- 11: $\widehat{C}_i \leftarrow \text{median}_r D_i^r$
- 12: **return** $\widetilde{F}_p := \sum_i \widehat{C}_i$

Using COUNTSKETCH as the subroutine for HEAVYHITTERS on two-pass turnstile streams and BPTREE as the subroutine for HEAVYHITTERS on two-pass insertion-only streams, we obtain the following guarantees for Algorithm 2. We first show that there exists a subsampling rate such that items in each level set will be reported as a heavy-hitter. The proof is similar to the proof of Lemma 2.4, but we no longer require each reported item to also be reported with a $(1 + \varepsilon)$ -approximation to their frequency. In fact, this cannot be done in a single pass; we will instead track their frequencies in the second pass.

► **Lemma 3.3.** *Let $\varepsilon \in (0, 1)$, Λ_i be a fixed level set, and let $\ell := \ell_i := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$. For a fixed r , let \mathcal{E}_1 be the event that $|I_\ell^r| \leq \frac{10\gamma n}{2^\ell}$ and let \mathcal{E}_2 be the event that $F_p(I_\ell^r) \leq \frac{10\gamma \widehat{F}_p}{2^\ell}$. Conditioned on \mathcal{E}_1 and \mathcal{E}_2 , then HEAVYHITTERS reports $j \in S_i^r$ for each $j \in \Lambda_i \cap I_\ell^r$ with probability at least $1 - \frac{1}{\text{poly}(n)}$.*

112:14 Separations for Estimating Large Frequency Moments on Data Streams

Proof. We consider casework on the value of i . First suppose $2^i \cdot \varepsilon^2 < \gamma$, so that $\ell := \ell_i = \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\} \leq \log \gamma$ since $\varepsilon_k \leq \varepsilon$ for all k . By Definition 2.2 of the level sets, each item $j \in \Lambda_i$ satisfies $f_j^p \in \left[\frac{\zeta \cdot \widehat{F}_p}{2^i}, \frac{2\zeta \cdot \widehat{F}_p}{2^i} \right]$. We also have $\widehat{F}_p \leq F_p \leq 1.01 \cdot \widehat{F}_p$ and $n^{1/2-1/p} \cdot F_p^{1/p} \geq F_2^{1/2}$. Thus, each item in $j \in \Lambda_i$ satisfies

$$f_j \geq \frac{\varepsilon^{2/p}}{1.01^{1/p}} \cdot F_p^{1/p} \geq \frac{\varepsilon^{2/p}}{1.01^{1/p} n^{1/2-1/p}} \cdot F_2^{1/2},$$

so that $f_j^2 \geq \frac{\varepsilon^{4/p}}{1.01^{2/p} n^{1/2-1/p}} F_2$. Then by Theorem 3.1 or Theorem 3.2, each item $j \in \Lambda_i$ is reported by COUNTSKETCH or BPTREE with threshold $\frac{(\varepsilon_\ell)^{2/p}}{80\gamma}$, as $\varepsilon_\ell \leq \varepsilon$ and $n_\ell = n$.

Otherwise, suppose $2^i \cdot \varepsilon^2 \geq \gamma$, so that $\ell := \ell_i = \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\} \geq \log \gamma$ and $p_\ell := \min(1, 2^{-\ell} \gamma) \geq \frac{\gamma}{2 \cdot 2^{\ell} \varepsilon_\ell^2}$. Again by Definition 2.2 of the level sets, each item $j \in \Lambda_i$ satisfies $f_j^p \in \left[\frac{\zeta \cdot \widehat{F}_p}{2^i}, \frac{2\zeta \cdot \widehat{F}_p}{2^i} \right]$. If $j \notin I_\ell^r$, then j certainly will not be reported by HEAVYHITTERS (regardless of whether HEAVYHITTERS is COUNTSKETCH or BPTREE). Hence, we assume that $j \in I_\ell^r$.

Conditioning on the event \mathcal{E}_2 ,

$$F_p(I_\ell^r) \leq \frac{10\gamma F_p}{2^\ell} \leq \frac{20\gamma F_p}{2^i \varepsilon^2}.$$

Given $\widehat{F}_p \leq F_p \leq 1.01 \cdot \widehat{F}_p$, then $f_j^p \geq \frac{\varepsilon^2}{80\gamma} \cdot F_p(I_\ell^r)$. Therefore,

$$f_j \geq \frac{\varepsilon^{2/p}}{(80\gamma)^{1/p}} \cdot F_p^{1/p}(I_\ell^r).$$

Rather than applying the inequality $n^{1/2-1/p} \cdot F_p^{1/p} \geq F_2^{1/2}$, we observe that conditioning on the event \mathcal{E}_1 , it follows that $|I_\ell^r| \leq n_\ell = \frac{10\gamma n}{2^\ell}$. Thus, the frequency vector defined by the substream I_ℓ^r has significantly smaller support size than n , which we can leverage to use a heavy-hitter algorithm with a lower threshold. We have

$$f_j \geq \frac{\varepsilon^{2/p}}{(80\gamma)^{1/p}} \cdot \left(\frac{1}{n_\ell} \right)^{1/2-1/p} \cdot F_2^{1/2}(I_\ell^r).$$

Therefore by Theorem 3.1 or Theorem 3.2, each item $j \in \Lambda_i \cap I_\ell^r$ will be reported by H_ℓ^r output by HEAVYHITTERS with threshold $\frac{(\varepsilon_\ell)^{2/p}}{80\gamma} \cdot \left(\frac{1}{n_\ell} \right)^{1/2-1/p}$, since $\varepsilon_\ell \leq \varepsilon$. ◀

We now show our main correctness statement for our two-pass algorithms. Lemma 3.4 proves that the output \tilde{F}_p of our algorithm gives a $(1 + \varepsilon)$ -approximation to the F_p moment of the underlying frequency vector. Although the guarantees of Lemma 3.4 are similar to the guarantees of Lemma 2.5 are similar, the proof of Lemma 3.4 is much simpler. We show that we obtain $(1 + \varepsilon)$ -approximations to the contributions of each significant level set, thus obtaining a $(1 + \varepsilon)$ -approximation to the F_p moment, since the contributions of the insignificant level sets can be ignored. Unlike Lemma 2.5, we need not concern about an idealized process since the frequency of each heavy-hitter is exactly tracked in the second pass of the algorithm over the data stream, so no heavy-hitters can be accidentally misclassified into an incorrect level set.

► **Lemma 3.4.** *With probability at least $\frac{2}{3}$, we have that $|\tilde{F}_p - F_p| \leq \varepsilon \cdot F_p$.*

Proof. Let Λ_i be a fixed level set and $\ell := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$. Let \mathcal{E}_1 be the event that $|I_\ell^r| \leq \frac{10\gamma n}{2^\ell}$ and let \mathcal{E}_2 be the event that $F_p(I_\ell^r) \leq \frac{10\gamma F_p}{2^\ell}$. By Lemma 3.3, COUNTSKETCH returns each $k \in \Lambda_i \cap I_\ell^r$ in S_i^r , conditioned on \mathcal{E}_1 and \mathcal{E}_2 . Thus in the second pass, Algorithm 2 tracks the contribution f_k^p explicitly, by tracking f_k . We first analyze the approximation \widehat{C}_i^r to C_i for a given set of subsamples, where we define

$$D_i^r := \frac{1}{p_\ell} \cdot \sum_{k \in \Lambda_i \cap I_\ell^r} f_k^p.$$

so that $\widehat{C}_i = \text{median}_r D_i^r$. Conditioned on \mathcal{E}_1 and \mathcal{E}_2 , we note that

$$\mathbb{E}[D_i^r] = \frac{1}{p_\ell} \cdot \sum_{k \in \Lambda_i} p_\ell \cdot f_k^p = C_i.$$

We also have

$$\text{Var}(D_i^r) = \frac{1}{p_\ell^2} \cdot \sum_{k \in \Lambda_i} p_\ell f_k^{2p} \leq \sum_{k \in \Lambda_i} \frac{2 \cdot 2^i \varepsilon_\ell^2}{\gamma} \cdot f_k^{2p},$$

since $p_\ell \geq \frac{\gamma}{2 \cdot 2^i \varepsilon_\ell^2}$. Observe that for each $k \in \Lambda_i$, we have $f_k^{2p} \leq \frac{16(F_p)^2}{2^{2i}}$ and $\frac{|\Lambda_i|}{2^i} \leq \phi_i \leq 1$. Thus for $\gamma = 2^{11}$,

$$\text{Var}(D_i^r) \leq |\Lambda_i| \cdot \frac{2 \cdot 2^i \varepsilon_\ell^2}{\gamma} \cdot \frac{16(F_p)^2}{2^{2i}} \leq \phi_i \varepsilon_\ell^2 (F_p)^2 \leq \frac{\varepsilon_\ell^2}{64} (F_p)^2.$$

Hence, by Chebyshev's inequality, we have that

$$\Pr \left[|D_i^r - C_i| \geq \frac{\varepsilon_\ell}{2} \cdot F_p \right] \leq \frac{1}{16}.$$

In summary, D_i^r gives an approximation to C_i with additive error at most $\frac{\varepsilon_\ell}{2} \cdot F_p$ with probability at least $\frac{15}{16}$, conditioned on the events \mathcal{E}_1 and \mathcal{E}_2 , and the correctness of the subroutine COUNTSKETCH. Since COUNTSKETCH fails with probability at most $1 - \frac{1}{\text{poly}(n)}$ by Theorem 3.1, then by a union bound, we have that conditioned on \mathcal{E}_1 and \mathcal{E}_2 ,

$$\Pr \left[|D_i^r - C_i| \leq \left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_\ell \right) \cdot F_p \right] \geq \frac{7}{8}.$$

For a fixed r , let \mathcal{E}_1 be the event that $|I_\ell^r| \leq \frac{10\gamma n}{2^\ell}$ and let \mathcal{E}_2 be the event that $F_p(I_\ell^r) \leq \frac{10\gamma F_p}{2^\ell}$. Recall that I_ℓ^r is a nested subset of $[n]$ subsampled at rate $p_\ell := \min(1, 2^{-\ell}\gamma)$, so that $\mathbb{E}[|I_\ell^r|] \leq \frac{\gamma n}{2^\ell}$. Thus by Markov's inequality, $\Pr[\mathcal{E}_1] \geq \frac{9}{10}$. Similarly, $\mathbb{E}[F_p(I_\ell^r)] \leq \frac{\gamma F_p}{2^\ell}$. Thus by Markov's inequality, $\Pr[\mathcal{E}_2] \geq \frac{9}{10}$. By a union bound, we first have $\Pr[\mathcal{E}_1 \wedge \mathcal{E}_2] \geq \frac{8}{10}$. Applying another union bound, we have that $|D_i^r - C_i| \leq \frac{\varepsilon_\ell}{2} \cdot F_p$ for a fixed r , with probability at least $\frac{5}{8}$. Since $\widehat{C}_i = \text{median}_r D_i^r$ across $r \leq \mathcal{O}(\log \log n)$ iterations, then we have that $|\widehat{C}_i - C_i| \leq \frac{\varepsilon_\ell}{2} \cdot F_p$ with probability at least $1 - \frac{1}{\text{polylog}(n)}$.

By a union bound over the indices $i \in [\alpha \log n]$, corresponding to the $\alpha \log n$ level sets, then with probability at least $1 - \frac{1}{\text{polylog}(n)}$, we have that $|\widehat{C}_i - C_i| \leq \frac{\varepsilon_\ell}{2} \cdot F_p$ simultaneously for all $i \in [\alpha \log n]$, where $\ell_i := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$. We form our estimate \tilde{F}_p to F_p by setting $\tilde{F}_p := \sum_i \widehat{C}_i$ and we have $F_p = \sum_i C_i$. Since $i \in [\alpha \log n]$, $\ell_i := \min\{k : 2^k > 2^i \cdot \varepsilon_k^2\}$, and $\varepsilon_{\ell_i} = \frac{\varepsilon}{16\eta \cdot 2^{\ell_i(p/16-1/8)} \log \frac{1}{\varepsilon^2}}$, then

$$\begin{aligned}
 |\tilde{F}_p - F_p| &= \left| \sum_i \widehat{C}_i - \sum_i C_i \right| \leq \sum_i |\widehat{C}_i - C_i| \leq \sum_i \left(\frac{\varepsilon}{4\alpha \log n} + \varepsilon_{\ell_i} \right) \cdot F_p \\
 &\leq \sum_{i \in [\alpha \log n]} \frac{\varepsilon}{4\alpha \log n} \cdot F_p + \sum_{i: \ell_i \leq \log \frac{1}{\varepsilon^2}} \varepsilon_{\ell_i} \cdot F_p + \sum_{i: \ell_i > \log \frac{1}{\varepsilon^2}} \varepsilon_{\ell_i} \cdot F_p \\
 &\leq \frac{\varepsilon}{4} \cdot F_p + \log \frac{1}{\varepsilon^2} \cdot \frac{\varepsilon}{16 \log \frac{1}{\varepsilon^2}} \cdot F_p + \frac{\varepsilon}{16} \cdot F_p,
 \end{aligned}$$

where the last bound on the last term follows from $\eta > \sum_{j \geq 0} 2^{-j(p/16-1/8)}$. Therefore with probability at least $1 - \frac{1}{\text{polylog}(n)}$, we have that $|\tilde{F}_p - F_p| \leq \frac{\varepsilon}{2} \cdot F_p$. \blacktriangleleft

We now justify the full guarantees claimed by Theorem 1.2. We first handle two passes over a turnstile stream.

► Theorem 3.5. *For $p > 2$, there exists a two-pass turnstile streaming algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment with probability at least $\frac{2}{3}$, while using $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log^2 n \log \log n\right)$ bits of space.*

Proof. Our analysis is similar to the proof of Theorem 1.1. We again observe that by redefining the level sets according to any upper bound on F_p , we do not require a 1.01-approximation \widehat{F}_p to F_p as the input of Algorithm 2. Thus by Lemma 3.4, there exists an algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment on two pass turnstile streams and it remains to analyze the space complexity. By Theorem 3.1, COUNTSKETCH with threshold ε requires $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log^2 n\right)$ bits of space to output the indices of the heavy-hitters. Algorithm 2 runs a separate instance of COUNTSKETCH with threshold $\frac{(\varepsilon_i)^{2/p}}{80\gamma} \cdot \left(\frac{1}{n_i}\right)^{1/2-1/p}$ to output a set H_i^r for $r \leq \mathcal{O}(\log \log n)$, where γ is a sufficiently large constant. Thus the total space in the first pass across all indices i for a fixed r is at most

$$(C_1 \log^2 n) \cdot \sum_{i \in [\alpha \log n]} \frac{(n_i)^{1-2/p}}{(\varepsilon_i)^{4/p}},$$

for some positive constants $C_1, \alpha > 0$. In particular, we have $n_i = \frac{10\gamma n}{2^i}$ and $\varepsilon_i = \frac{\varepsilon}{16\eta \cdot 2^{i(p/16-1/8)} \log \frac{1}{\varepsilon^2}}$ for a sufficiently large constant η . Then the total space in the first pass across all indices i for a fixed r is at most

$$(C_1 \log^2 n) \cdot \sum_{i=1}^{\infty} \frac{C_2 n^{1-2/p}}{\varepsilon^{4/p}} \cdot \frac{2^{i(1/4-1/(2p))}}{2^{i(1-2/p)}},$$

for some positive constants $C_1, C_2 > 0$. Since $\sum_{i=1}^{\infty} \frac{2^{i(1/4-1/(2p))}}{2^{i(1-2/p)}} = \sum_{i=1}^{\infty} \frac{1}{2^{3i(1-2/p)/4}}$ is a geometric series that is upper bounded by a fixed constant, the total space in the first pass across all i for a fixed r is $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log^2 n\right)$. Because $r \in \mathcal{O}(\log \log n)$, then the total space is $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log^2 n \log \log n\right)$.

In the second pass, we track the frequencies of each item reported by some instance of COUNTSKETCH. Since at most $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log \log n\right)$ indices can be reported across all instances of COUNTSKETCH and $\mathcal{O}(\log n)$ bits of space can be used to track the frequency of each reported index, then the total space for the second pass is at most $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log n \log \log n\right)$. Thus, the total space is $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log^2 n \log \log n\right)$. \blacktriangleleft

Finally, we note that to report at most $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log \log n\right)$ indices of possible heavy-hitters in turnstile streams, COUNTSKETCH uses at most $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log^2 n \log \log n\right)$ space. By the same reasoning, we use space $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log n \log \log n\right)$ in insertion-only streams by using the more space efficient BPTREE.

► **Theorem 3.6.** *For $p > 2$, there exists a two-pass insertion-only streaming algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment with probability at least $\frac{2}{3}$, while using $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log n \log \log n\right)$ bits of space.*

Proof. The proof of correctness is exactly the same as that of Theorem 3.5 since using BPTREE as the subroutine for HEAVYHITTERS rather than COUNTSKETCH offers the same guarantee for insertion-only streams. By Theorem 3.2, BPTREE with threshold ε requires $\mathcal{O}\left(\frac{1}{\varepsilon^2} \log n\right)$ bits of space to output the indices of the heavy-hitters. To analyze the space complexity, note that Algorithm 2 runs a separate instance of BPTREE with threshold $\frac{(\varepsilon_i)^{2/p}}{80^\gamma} \cdot \left(\frac{1}{n_i}\right)^{1/2-1/p}$ to output a set H_i^r for $r \leq \mathcal{O}(\log \log n)$, where γ is a sufficiently large constant. Hence, the first pass across all indices i for a fixed r uses space at most

$$(C_1 \log n) \cdot \sum_{i \in [\alpha \log n]} \frac{(n_i)^{1-2/p}}{(\varepsilon_i)^{4/p}},$$

for some absolute constants $C_1, \alpha > 0$. Since $n_i = \frac{10\gamma n}{2^i}$ and $\varepsilon_i = \frac{\varepsilon}{16\eta \cdot 2^{i(p/16-1/8)} \log \frac{1}{\varepsilon^2}}$ for a sufficiently large constant η , the first pass uses space at most

$$(C_1 \log n) \cdot \sum_{i=1}^{\infty} \frac{C_2 n^{1-2/p}}{\varepsilon^{4/p}} \cdot \frac{2^{i(1/4-1/(2p))}}{2^{i(1-2/p)}},$$

for some absolute constants $C_1, C_2 > 0$, across all indices i for a fixed r . As $\sum_{i=1}^{\infty} \frac{2^{i(1/4-1/(2p))}}{2^{i(1-2/p)}} = \sum_{i=1}^{\infty} \frac{1}{2^{3i(1-2/p)/4}}$ is a geometric series that is upper bounded by some constant, then the total space in the first pass is $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log n\right)$ across all indices i for a fixed r . Since $r \leq \mathcal{O}(\log \log n)$, then the total space in the first pass is $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log n \log \log n\right)$.

The second pass tracks the frequencies of each item reported by some instance of BPTREE. Since at most $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log \log n\right)$ indices can be reported across all instances of BPTREE and $\mathcal{O}(\log n)$ bits of space can be used to track the frequency of each reported index, then the total space for the second pass is at most $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log n \log \log n\right)$. Thus, the total space is $\mathcal{O}\left(\frac{1}{\varepsilon^{4/p}} \cdot n^{1-2/p} \log n \log \log n\right)$. ◀

4 Lower Bounds

In this section, we first consider the standard *blackboard* communication model, where a number of players each have a local input and the goal is to solve some predetermined communication problem by sending messages to a shared medium. Each player is assumed to have access to an unlimited amount of private randomness. The sequence of messages on the shared blackboard is called the *transcript* and the maximum length of the transcript over all inputs is the *communication cost* of a given protocol. The communication complexity of f , denoted by $R_\delta(f)$, is the minimal communication cost of all protocols that succeed with probability at least $1 - \delta$ for all legal inputs to f .

► **Definition 4.1.** In the (t, ε, n) -player set disjointness estimation problem (t, ε, n) – DISJINFTY, there are $t + 1$ players P_1, \dots, P_{t+1} with private coins in the standard blackboard model. For $s \in [t]$, each player P_s receives a vector $\mathbf{v}_s \in \{0, 1\}^n$ and player P_{t+1} receives both an index $j \in [n]$ and a bit $c \in \{0, 1\}$. For $\mathbf{u} = \sum_{s \in [t]} \mathbf{v}_s$, the inputs are promised to satisfy $u_i \leq 1$ for each $i \neq j$ and either $u_j = 1$ or $u_j = t$. With probability at least $\frac{9}{10}$, P_{t+1} must differentiate between the three possible input cases:

- (1) $u_j + \frac{ct}{\varepsilon} \leq t$
- (2) $u_j + \frac{ct}{\varepsilon} \in \left\{ \frac{t}{\varepsilon}, \frac{t}{\varepsilon} + 1 \right\}$
- (3) $u_j + \frac{ct}{\varepsilon} = (1 + \varepsilon) \frac{t}{\varepsilon}$,

where $\varepsilon \in (0, 1)$. We call coordinate $j \in [n]$ the spike location.

Direct sum embedding. We use the direct sum technique by showing that even in the case where $u_i \leq 1$ for all $i \in [n]$, the information cost of the players is sufficiently high. We describe the embedding performed by each player to sample coordinates independently conditioned on the auxiliary variable D . We define the auxiliary variable $D = (D_1, \dots, D_n)$ by first defining a distribution μ for D_i and $v_{s,i}$, for each fixed coordinate $i \in [n]$ and across all players $s \in [t]$:

- (1) $D_i \sim [t]$ uniformly at random, so that D_i determines a player whose input bit will be randomized while the remaining players have input bit zero.
- (2) Conditioned on $D_i = s$, then each player P_r with $r \neq s$ sets $v_{r,i} = 0$ while player P_s independently and uniformly sets $v_{s,i} \in \{0, 1\}$.

We define the distribution $\zeta := \mu^n$ so that the auxiliary random variable $D = (D_1, \dots, D_n)$ is the vector whose i -th coordinate determines the player P_{D_i} whose i -th bit $v_{D_i,i}$ in their input vector will vary, while the other players P_s have i -th bit $v_{s,i}$ zero in their vectors, for $s \neq D_i$. Observe that under the distribution ζ , we indeed have $u_i \leq 1$ for each coordinate $i \in [n]$ of $\mathbf{u} = \sum_{s \in [t]} \mathbf{v}_s$. Finally, we fix the input $c = 0$ to player P_{t+1} and pick $j \in [n]$ according to any arbitrary distribution.

► **Theorem 4.2.** The total communication complexity for the (t, ε, n) -player set disjointness estimation problem is $\Omega\left(\frac{n}{t}\right)$.

We remark that the communication complexity of Theorem 4.2 matches the communication complexity of t -player set disjointness. However, since the problem requires correctness on all inputs, then we can distinguish between the possible input cases by focusing on the specific coordinate j given to player P_{t+1} . By contrast, the reduction of [18] from t -player set disjointness requires an algorithm to “test” all coordinates $i \in [n]$ for the spike location. Thus the reduction requires that an F_p moment estimation algorithm succeeds with probability $1 - \frac{1}{\text{poly}(n)}$, thereby losing a multiplicative $\mathcal{O}(\log n)$ factor and achieving $\Omega\left(\frac{n^{1-2/p}}{\varepsilon^2 \log n}\right)$ in the space lower bound for F_p moment estimation. Since our communication problem gives the specific spike location as input, our reduction only requires an F_p moment estimation algorithm that succeeds with constant probability. We remark that a similar technique was used in [30] for the L_∞ promise problem.

Reduction. Let $t = \Theta\left(\frac{1}{\varepsilon} \cdot n^{1/p}\right)$. We reduce $(1 + \varepsilon)$ -approximate F_p moment estimation to an instance of (t, ε) – DISJINFTY as follows. Let \mathcal{A} be any fixed randomized one-pass insertion-only streaming algorithm that outputs a $\left(1 + \frac{\varepsilon}{3}\right)$ -approximation to the F_p moment of the underlying frequency vector with probability at least $\frac{2}{3}$. Recall that the first t players each receive a vector \mathbf{v}_s with $s \in [t]$ and player P_{t+1} receives both a special index $j \in [n]$ and a bit $c \in \{0, 1\}$.

- For each $s \in [t]$, player P_s takes the state of the algorithm \mathcal{A} , inserts the coordinates of vector \mathbf{v}_s , and passes the state of the algorithm to player P_{s+1} .
- Player P_{t+1} takes the state of \mathcal{A} , adds the vector $\frac{ct}{\varepsilon} \cdot \mathbf{e}_j$, where \mathbf{e}_j is the elementary vector with a one in position j and zeros elsewhere, and then queries the state of \mathcal{A} to obtain a $(1 + \frac{\varepsilon}{3})$ -approximation to the F_p moment of underlying frequency vector.

► **Theorem 4.3.** *For any constant $p > 2$ and parameter $\varepsilon = \Omega(\frac{1}{n^{1/p}})$, any one-pass insertion-only streaming algorithm that outputs a $(1 + \varepsilon)$ -approximation to the F_p moment of an underlying frequency vector with probability at least $\frac{9}{10}$ requires $\Omega(\frac{1}{\varepsilon^2} \cdot n^{1-2/p})$ bits of space.*

Proof. Player P_{t+1} receives the state of \mathcal{A} on the input $\mathbf{u} = \sum_{s \in [t]} \mathbf{v}_s$ and induces a frequency vector $\mathbf{x} := \mathbf{u} + \frac{ct}{\varepsilon} \cdot \mathbf{e}_j$. Recall that the inputs are promised to satisfy $u_i \leq 1$ for each $i \neq j$ and either $u_j = 1$ or $u_j = t$. If $c = 0$, then $\mathbf{x} = \mathbf{u}$ so that for a constant $C > 0$ and $t = \frac{C}{\varepsilon} \cdot n^{1/p}$,

$$\|\mathbf{x}\|_p^p \leq F_0 + t^p \leq n + \frac{C^p}{\varepsilon^p} \cdot n.$$

If $c = 1$ and $u_j \leq 1$, then for $t = \frac{C}{\varepsilon} \cdot n^{1/p}$, we have $\|\mathbf{x}\|_p^p \geq (\frac{t}{\varepsilon})^p = \frac{C^p}{\varepsilon^{2p}} \cdot n$ and thus

$$\|\mathbf{x}\|_p^p \leq F_0 + \left(1 + \frac{t}{\varepsilon}\right)^p \leq n + p + \frac{p \cdot C^p}{\varepsilon^{2p}} \cdot n.$$

Finally, if $c = 1$ and $u_j = t$, then

$$\|\mathbf{x}\|_p^p \geq \left(t + \frac{t}{\varepsilon}\right)^p = \left(1 + \frac{1}{\varepsilon}\right)^p \cdot \frac{C^p}{\varepsilon^{2p}} \cdot n.$$

For sufficiently small $\varepsilon \in (0, 1)$ with $\varepsilon = \Omega(\frac{1}{n^{1/p}})$ and constant $p > 2$, there exists a constant $C > 0$ such that these three cases are separated by a multiplicative $(1 + \varepsilon)$. Since \mathcal{A} outputs a $(1 + \frac{\varepsilon}{3})$ -approximation to the F_p moment of the underlying frequency vector, then player P_{t+1} obtains a $(1 + \frac{\varepsilon}{3})$ -approximation to $\|\mathbf{x}\|_p^p$ and can differentiate between the three cases, thus solving (t, ε) -DISJINFTY with probability at least $\frac{9}{10}$. By Theorem 4.2, \mathcal{A} uses $\Omega(\frac{n}{t})$ total communication across the $t + 1$ players. Therefore, the space complexity of \mathcal{A} is $\Omega(\frac{n}{t^2}) = \Omega(\frac{1}{\varepsilon^2} \cdot n^{1-2/p})$ for $t = \Theta(\frac{1}{\varepsilon} \cdot n^{1/p})$. The result then follows from rescaling ε . ◀

References

- 1 Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. *J. Comput. Syst. Sci.*, 58(1):137–147, 1999.
- 2 Alexandr Andoni. High frequency moments via max-stability. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, pages 6364–6368, 2017.
- 3 Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Streaming algorithms via precision sampling. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS*, pages 363–372, 2011.
- 4 Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004.
- 5 Lakshminath Bhuvanagiri, Sumit Ganguly, Deepanjan Kesh, and Chandan Saha. Simpler algorithm for estimating frequency moments of data streams. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 708–713, 2006.
- 6 Jaroslaw Blasiok, Jian Ding, and Jelani Nelson. Continuous monitoring of ℓ_p norms in data streams. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*, volume 81, pages 32:1–32:13, 2017.

- 7 Mark Braverman, Sumegha Garg, and David P. Woodruff. The coin problem with applications to data streams. *Electron. Colloquium Comput. Complex.*, 27:139, 2020.
- 8 Vladimir Braverman, Stephen R. Chestnut, Nikita Ivkin, Jelani Nelson, Zhengyu Wang, and David P. Woodruff. Bptree: An ℓ_2 heavy hitters algorithm using constant memory. In *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS*, pages 361–376, 2017.
- 9 Vladimir Braverman, Jonathan Katzman, Charles Seidell, and Gregory Vorsanger. An optimal algorithm for large frequency moments using $O(n^{1-2/k})$ bits. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*, pages 531–544, 2014.
- 10 Vladimir Braverman and Rafail Ostrovsky. Approximating large frequency moments with pick-and-drop sampling. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop, APPROX, and 17th International Workshop, RANDOM. Proceedings*, 2013.
- 11 Vladimir Braverman, Emanuele Viola, David P. Woodruff, and Lin F. Yang. Revisiting frequency moment estimation in random order streams. In *45th International Colloquium on Automata, Languages, and Programming, ICALP*, pages 25:1–25:14, 2018.
- 12 Amit Chakrabarti, Graham Cormode, and Andrew McGregor. Robust lower bounds for communication and stream computation. *Theory Comput.*, 12(1):1–35, 2016.
- 13 Amit Chakrabarti, T. S. Jayram, and Mihai Patrascu. Tight lower bounds for selection in randomly ordered streams. In Shang-Hua Teng, editor, *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 720–729, 2008.
- 14 Amit Chakrabarti, Subhash Khot, and Xiaodong Sun. Near-optimal lower bounds on the multi-party communication complexity of set disjointness. In *18th Annual IEEE Conference on Computational Complexity*, pages 107–117, 2003.
- 15 Moses Charikar, Kevin C. Chen, and Martin Farach-Colton. Finding frequent items in data streams. *Theor. Comput. Sci.*, 312(1):3–15, 2004.
- 16 Erik D. Demaine, Alejandro López-Ortiz, and J. Ian Munro. Frequency estimation of internet packet streams with limited space. In *Algorithms - ESA, 10th Annual European Symposium, Proceedings*, pages 348–360, 2002.
- 17 Sumit Ganguly. Polynomial estimators for high frequency moments. *CoRR*, abs/1104.4552, 2011. [arXiv:1104.4552](https://arxiv.org/abs/1104.4552).
- 18 Sumit Ganguly. A lower bound for estimating high moments of a data stream. *CoRR*, abs/1201.0253, 2012. [arXiv:1201.0253](https://arxiv.org/abs/1201.0253).
- 19 Sumit Ganguly and David P. Woodruff. High probability frequency moment sketches. In *45th International Colloquium on Automata, Languages, and Programming, ICALP*, 2018.
- 20 Sudipto Guha and Zhiyi Huang. Revisiting the direct sum theorem and space lower bounds in random order streams. In *International Colloquium on Automata, Languages, and Programming*, pages 513–524. Springer, 2009.
- 21 Sudipto Guha and Andrew McGregor. Approximate quantiles and the order of the stream. In *Proceedings of the Twenty-Fifth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 273–279, 2006.
- 22 Sudipto Guha and Andrew McGregor. Lower bounds for quantile estimation in random-order and multi-pass streaming. In *Automata, Languages and Programming, 34th International Colloquium, ICALP, Proceedings*, pages 704–715, 2007.
- 23 Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *J. ACM*, 53(3):307–323, 2006.
- 24 Piotr Indyk and David P. Woodruff. Optimal approximations of the frequency moments of data streams. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 202–208, 2005.

- 25 Rajesh Jayaram and David P. Woodruff. Towards optimal moment estimation in streaming and distributed models. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*, pages 29:1–29:21, 2019.
- 26 Daniel M. Kane, Jelani Nelson, Ely Porat, and David P. Woodruff. Fast moment estimation in data streams in optimal space. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC*, pages 745–754, 2011.
- 27 Daniel M. Kane, Jelani Nelson, and David P. Woodruff. On the exact space complexity of sketching and streaming small norms. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1161–1178, 2010.
- 28 Christian Konrad, Frédéric Magniez, and Claire Mathieu. Maximum matching in semi-streaming with few passes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 15th International Workshop, APPROX, and 16th International Workshop, RANDOM*, pages 231–242, 2012.
- 29 Ping Li. Estimators and tail bounds for dimension reduction in ℓ_α ($0 < \alpha \leq 2$) using stable random projections. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 10–19, 2008.
- 30 Yi Li and David P. Woodruff. A tight lower bound for high frequency moment estimation with small error. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop, APPROX, and 17th International Workshop, RANDOM. Proceedings*, pages 623–638, 2013.
- 31 Morteza Monemizadeh and David P. Woodruff. 1-pass relative-error L_p -sampling with applications. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1143–1160, 2010.
- 32 J. Ian Munro and Mike Paterson. Selection and sorting with limited storage. *Theor. Comput. Sci.*, 12:315–323, 1980.
- 33 David P. Woodruff. Optimal space lower bounds for all frequency moments. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 167–175, 2004.
- 34 David P. Woodruff and Qin Zhang. Tight bounds for distributed functional monitoring. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC*, pages 941–960, 2012.
- 35 David P. Woodruff and Samson Zhou. Tight bounds for adversarially robust streams and sliding windows via difference estimators. *CoRR*, abs/2011.07471, 2020. [arXiv:2011.07471](https://arxiv.org/abs/2011.07471).
- 36 David P. Woodruff and Samson Zhou. Separations for estimating large frequency moments on data streams. *CoRR*, abs/2105.03773, 2021. [arXiv:2105.03773](https://arxiv.org/abs/2105.03773).