

Fast Moment Estimation in Data Streams in Optimal Space

Daniel M. Kane[†] Jelani Nelson[‡] Ely Porat[§] David P. Woodruff[¶]

Abstract

We give a space-optimal algorithm with update time $O(\log^2(1/\varepsilon) \log \log(1/\varepsilon))$ for $(1 \pm \varepsilon)$ -approximating the p th frequency moment, $0 < p < 2$, of a length- n vector updated in a data stream. This provides a nearly exponential improvement in the update time complexity over the previous space-optimal algorithm of [Kane-Nelson-Woodruff, SODA 2010], which had update time $\Omega(1/\varepsilon^2)$.

1 Introduction

The problem of estimating frequency moments of a vector being updated in a data stream was first studied by Alon, Matias, and Szegedy [3] and has since received much attention [5, 6, 21, 22, 26, 28, 30, 40, 41]. Estimation of the second moment has applications to estimating join and self-join sizes [2] and to network anomaly detection [27, 37]. First moment estimation is useful in mining network traffic data [11], comparing empirical probability distributions, and several other applications (see [30] and the references therein). Estimating fractional moments between the 0th and 2nd has applications to entropy estimation for the purpose of network anomaly detection [20, 42], mining tabular data [9], image decomposition [17], and weighted sampling in turnstile streams [29]. It was also observed experimentally that the use of fractional moments in $(0, 1)$ can improve the effectiveness of standard clustering algorithms [1].

Formally in this problem, we are given up front a real number $p > 0$. There is also an underlying n -dimensional vector x which starts as $\vec{0}$. What follows is a sequence of m updates of the form $(i_1, v_1), \dots, (i_m, v_m) \in [n] \times \{-M, \dots, M\}$ for some $M > 0$. An update (i, v) causes the change $x_i \leftarrow x_i + v$. We would like to compute $F_p \stackrel{\text{def}}{=} \|x\|_p^p \stackrel{\text{def}}{=} \sum_{i=1}^n |x_i|^p$, also called the p th frequency moment of x . In many applications, it is required that the algorithm only use very limited space while processing the stream, e.g., in networking applications where x may be indexed by source-destination IP pairs and thus a router cannot afford to store the entire vector in memory, or in database applications where one wants a succinct “sketch” of some dataset, which can be compared with short sketches of other datasets for fast computation of various (dis)similarity measures.

Unfortunately, it is known that linear space ($\Omega(\min\{n, m\})$ bits) is required unless one allows for (a) *approximation*, so that we are only guaranteed to output a value in $[(1 - \varepsilon)F_p, (1 + \varepsilon)F_p]$ for some $0 < \varepsilon < 1/2$, and (b) *randomization*, so that the output is only guaranteed to be correct with some probability bounded away from 1, over the randomness used by the algorithm [3]. Furthermore, it is known that polynomial space is required for $p > 2$ [5, 7, 18, 23, 36], while it is known that

[†]Harvard University, Department of Mathematics. dankane@math.harvard.edu.

[‡]MIT Computer Science and Artificial Intelligence Laboratory. minilek@mit.edu.

[§]Department of Computer Science, Bar Ilan University. porately@macs.biu.ac.il.

[¶]IBM Almaden Research Center, 650 Harry Road, San Jose, CA, USA. dpwoodru@us.ibm.com.

Paper	Space	Update Time	Model	Which p
[3]	$O(\varepsilon^{-2} \log(mM))$	$O(\varepsilon^{-2})$	unrestricted updates	$p = 2$
[8, 37]	$O(\varepsilon^{-2} \log(mM))$	$O(1)$	unrestricted updates	$p = 2$
[13]	$O(\varepsilon^{-2} \log(mM))$	$O(\varepsilon^{-2})$	≤ 2 updates per coordinate	$p = 1$
[21, 28]	$O(\varepsilon^{-2} \log(n) \log(mM))$	$O(\varepsilon^{-2})$	unrestricted updates	$p \in (0, 2)$
[26]	$O(\varepsilon^{-2} \log(mM))$	$\tilde{O}(\varepsilon^{-2})$	unrestricted updates	$p \in (0, 2)$
[31]	$O(\varepsilon^{-2} \log(mM) \log(1/\varepsilon))$	$O(\log^2(mM))$	≤ 2 updates per coordinate	$p = 1$
[30]	$O(\varepsilon^{-2} \log(mM) \log(n))$	$O(\log(n/\varepsilon))$	unrestricted updates	$p = 1$
this work	$O(\varepsilon^{-2} \log(mM))$	$\tilde{O}(\log^2(1/\varepsilon))$	unrestricted updates	$p \in (0, 2)$

Figure 1: Comparison of our contribution to previous works on F_p -estimation in data streams. All space bounds hide an additive $O(\log \log n)$ term.

the space complexity for $0 < p \leq 2$ is only $\Theta(\varepsilon^{-2} \log(mM) + \log \log(n))$ bits to achieve success probability $2/3$ [3, 26], which can be amplified by outputting the median estimate of independent repetitions. In this work, we focus on this “feasible” regime for p , $0 < p \leq 2$, where logarithmic space is achievable.

While there has been much previous work on minimizing the space consumption in streaming algorithms, only recently have researchers begun to work toward minimizing *update time* [34, Question 1], i.e., the time taken to process a new update in the stream. We argue however that update time itself is an important parameter to optimize, and in some scenarios it may even be desirable to sacrifice space for speed. For example, in network traffic monitoring applications each packet is an update, and thus it is important that a streaming algorithm processing the packet stream be able to operate at network speeds (see for example the applications in [27, 37]). Note that if an algorithm has update time say, $\Omega(1/\varepsilon^2)$, then achieving a small error parameter such as $\varepsilon = .01$ could be intractable since this time is multiplied by the length of the stream. This is true even if the space required of the algorithm is small enough to fit in memory.

For $p = 2$, it is known that optimal space and $O(1)$ update time are simultaneously achievable [8, 37], improving upon the original F_2 algorithm of Alon, Matias, and Szegedy [3]. For $p = 1$ it is known that near-optimal, but not quite optimal, space and $O(\log(n/\varepsilon))$ update time are achievable [30]. Meanwhile, optimal (or even near-optimal) space for other $p \in (0, 2]$ is only known to be achievable with $\text{poly}(1/\varepsilon)$ update time [26].

Our Contribution: For all $0 < p < 2$ and $0 < \varepsilon < 1/2$ we give an algorithm for $(1 \pm \varepsilon)$ -approximating F_p with success probability at least $2/3$ which uses an optimal $O(\varepsilon^{-2} \log(mM) + \log \log n)$ bits of space with $O(\log^2(1/\varepsilon) \log \log(1/\varepsilon))$ update time.¹ This is a nearly exponential improvement in the time complexity of the previous space-optimal algorithm for every such p .

1.1 Previous Work

The complexity of streaming algorithms for moment estimation has a long history; see Figure 1 for a comparison of our result to that of previous work.

Alon, Matias, and Szegedy were the first to study moment estimation in data streams [3] and gave a space-optimal algorithm for $p = 2$. The update time was later brought down to an optimal

¹Throughout this document we say $g = \tilde{O}(f)$ if $g = O(f \cdot \text{polylog}(f))$. Similarly, $g = \tilde{\Omega}(f)$ if $g = \Omega(f/\text{polylog}(f))$.

$O(1)$ implicitly in [8] and explicitly in [37]. The work of [13] gave a space-optimal algorithm for $p = 1$, but under the restriction that each coordinate is updated at most twice, once positively and once negatively. Indyk [21] later removed this restriction, and also gave an algorithm handling all $0 < p < 2$, but at the expense of increasing the space by a $\log n$ factor. Li later [28] provided alternative estimators for all $0 < p < 2$, based on Indyk’s sketches. The extra $\log n$ factor in the space of these algorithms was later removed in [26], yielding optimal space. The algorithms of [13, 21, 26, 28] all required $\text{poly}(1/\varepsilon)$ update time. Nelson and Woodruff [31] gave an algorithm for $p = 1$ in the restricted setting where each coordinate is updated at most twice, as in [13], with space suboptimal by a $\log(1/\varepsilon)$ factor, and with update time $\log^2(mM)$. They also later gave an algorithm for $p = 1$ with unrestricted updates which was suboptimal by a $\log n$ factor, but had update time only $O(\log(n/\varepsilon))$ [30].

On the lower bound front, a lower bound of $\Omega(\min\{n, m, \varepsilon^{-2} \log(\varepsilon^2 mM)\} + \log \log(nmM))$ was shown in [26], together with an upper bound of $O(\varepsilon^{-2} \log(mM) + \log \log n)$ bits. For nearly the full range of parameters these are tight, since if $\varepsilon \leq 1/\sqrt{m}$ we can store the entire stream in memory in $O(m \log(nM)) = O(\varepsilon^{-2} \log(nM))$ bits of space (and we can ensure $n = O(m^2)$ via FKS hashing [14] with just an additive $O(\log \log n)$ bits increase in space), and if $\varepsilon \leq 1/\sqrt{n}$ we can store the entire vector in memory in $O(n \log(mM)) = O(\varepsilon^{-2} \log(mM))$ bits. Thus, a gap exists only when ε is very near $1/\sqrt{\min\{n, m\}}$. This lower bound followed many previous lower bounds for this problem, given in [3, 4, 24, 40, 41]. For the case $p > 2$ it was shown that $\Omega(n^{1-2/p})$ space is required [5, 7, 18, 23, 36], and this was shown to be tight up to $\text{poly}(\log(nmM)/\varepsilon)$ factors [6, 22].

1.2 Overview of our approach

At the top level, our algorithm follows the general approach set forth by [30] for F_1 -estimation. In that work, the coordinates $i \in \{1, \dots, n\}$ were split up into *heavy hitters*, and the remaining *light* coordinates. A ϕ -*heavy hitter* with respect to F_p is a coordinate i such that $|x_i|^p \geq \phi \|x\|_p^p$. A list L of ε^2 -heavy hitters with respect to F_1 were found by running the **CountMin** sketch of [10].

To estimate the contribution of the light elements to F_1 , [30] used $R = \Theta(1/\varepsilon^2)$ independent Cauchy sketches D_1, \dots, D_R (actually, D_j was a tuple of 3 independent Cauchy sketches). A *Cauchy sketch* of a vector x , introduced by Indyk [21], is the dot product of x with a random vector z with independent entries distributed according to the Cauchy distribution. This distribution has the property that $\langle z, x \rangle$ is itself a Cauchy random variable, scaled by $\|x\|_1$. Upon receiving an update to x_i in the stream, the update was fed to $D_{h(i)}$ for some hash function $h : [n] \rightarrow [R]$. At the end of the stream, the estimate of the contribution to F_1 from light elements was $(R/(R - |h(L)|)) \cdot \sum_{j \notin h(L)} \text{EstLi}_1(D_j)$, where EstLi_p is Li’s geometric mean estimator for F_p [28]. The analysis of [30] only used that Li’s geometric mean estimator is unbiased and has a good variance bound.

Our algorithm **LightEstimator** for estimating the contribution to F_p from light coordinates for $p \neq 1$ follows the same approach. Our main contribution here is to show that a variant of Li’s geometric mean estimator has bounded variance and is approximately unbiased (to within relative error ε) even when the associated p -stable random variables are only k -wise independent for $k = \Omega(1/\varepsilon^p)$. This variant allows us to avoid Nisan’s pseudorandom generator [32] and thus achieve optimal space. While the work of [26] also provided an estimator avoiding Nisan’s pseudorandom generator, their estimator is not known to be approximately unbiased, which makes it less useful in applications involving the average of many such estimators. We evaluate the necessary k -wise independent hash function quickly by a combination of buffering and fast multipoint evaluation of a collection

of pairwise independent polynomials. Our proof that bounded independence suffices uses the FT-mollification approach introduced in [26] and refined in [12], which is a method for showing that the expectation of some function is approximately preserved by bounded independence, via a smoothing operation (FT-mollification) and Taylor’s theorem. One novelty is that while [12, 26] only ever dealt with FT-mollifying indicator functions of regions in Euclidean space, here we must FT-mollify functions of the form $f(x) = |x|^{1/t}$. To achieve our results, we express $\mathbf{E}[f(x)] = \int_0^\infty f(x)\varphi_p(x)dx$ as $\int_0^\infty f'(x)(1 - \Phi_p(x))dx$ via integration by parts, where φ_p is the density function of the absolute value of the p -stable distribution, and Φ_p is the corresponding cumulative distribution function. We then note $1 - \Phi_p(x) = \mathbf{Pr}[|X| \geq x] = \mathbf{E}[I_{[x,\infty)\cup(-\infty,-x]}(X)]$ for X p -stable, where I_S is the indicator function of the set S . We then FT-mollify $I_{[x,\infty)\cup(-\infty,-x]}$, which is the indicator function of some set, to write $\mathbf{E}[f(x)]$ as a weighted integral of indicator functions, from which point we can apply the methods of [12, 26].

In order to estimate the contribution to F_p from coordinates in L , we develop a novel data structure we refer to as **HighEnd**. Suppose L contains all the α -heavy hitters, and every index in L is an $(\alpha/2)$ -heavy hitter. We would like to compute $\|x_L\|_p^p \pm O(\varepsilon) \cdot \|x\|_p^p$, where $\alpha = \Omega(\varepsilon^2)$. We maintain a matrix of counters $D_{j,k}$ for $(j,k) \in [t] \times [s]$ for $t = O(\log(1/\varepsilon))$ and $s = O(1/\alpha)$. For each $j \in [t]$ we have a hash function $h^j : [n] \rightarrow [s]$ and $g^j : [n] \rightarrow [r]$ for $r = O(\log(1/\varepsilon))$. The counter $D_{j,k}$ then stores $\sum_{h^j(v)=k} e^{2\pi i g^j(v)/r} x_v$ for $i = \sqrt{-1}$. That is, our data structure is similar to the **CountSketch** data structure of Charikar, Chen, and Farach-Colton [8], but rather than taking the dot product with a random sign vector in each counter, we take the dot product with a vector whose entries are random complex roots of unity. At the end of the stream, our estimate of the F_p -contribution from heavy hitters is

$$\mathbf{Re} \left[\sum_{w \in L} \left(\frac{3}{t} \sum_{k=1}^{t/3} e^{-2\pi i g^{j(w,k)}(w)/r} \cdot \text{sign}(x_w) \cdot D_{j(w,k), h^{j(w,k)}(w)} \right)^p \right].$$

The choice to use complex roots of unity is to ensure that our estimator is approximately unbiased, stemming from the fact that the real part of large powers of roots of unity is still 0 in expectation. Here $\mathbf{Re}[z]$ denotes the real part of z , and $j(w,k)$ denotes the k th smallest value $b \in [t]$ such that h^b isolates w from the other $w' \in L$ (if fewer than $t/3$ such b exist, we fail). The subroutine **Filter** for estimating the heavy hitter contribution for $p = 1$ in [30] did not use complex random variables, but rather just used the dot product with a random sign vector as in **CountSketch**. Furthermore, it required a $O(\log(1/\varepsilon))$ factor more space even for $p = 1$, since it did not average estimates across $\Omega(t)$ levels to reduce variance.

For related problems, e.g., estimating F_p for $p > 2$, using complex roots of unity leads to suboptimal bounds [15]. Moreover, it seems that “similar” algorithms using sign variables in place of roots of unity do not work, as they have a constant factor bias in their expectation for which it is unclear how to remove. Our initial intuition was that an algorithm using p -stable random variables would be necessary to estimate the contribution to F_p from the heavy hitters. However, such approaches we explored suffered from too large a variance.

In parallel we must run an algorithm we develop to *find* the heavy hitters. Unfortunately, this algorithm, as well as **HighEnd**, use suboptimal space. To overcome this, we actually use a list of ε^2 -heavy hitters for $\epsilon = \varepsilon \cdot \log(1/\varepsilon)$. This then improves the space, at the expense of increasing the variance of **LightEstimator**. We then run $O((\epsilon/\varepsilon)^2)$ pairwise independent instantiations of **LightEstimator** in parallel and take the average estimate, to bring the variance down. This increases

some part of the update time of `LightEstimator` by a $\log^2(1/\varepsilon)$ factor, but this term turns out to anyway be dominated by the time to evaluate various hash functions. Though, even in the extreme case of balancing with $\varepsilon = 1$, our algorithm for finding the heavy hitters algorithm requires $\Omega(\log(n)\log(mM))$ space, which is suboptimal. We remedy this by performing a dimensionality reduction down to dimension $\text{poly}(1/\varepsilon)$ via hashing and dot products with random sign vectors. We then apply `HighEnd` to estimate the contribution from heavy hitters in this new vector, and we show that with high probability the correctness of our overall algorithm is still maintained.

1.3 Notation

For a positive integer r , we use $[r]$ to denote the set $\{1, \dots, r\}$. All logarithms are base-2 unless otherwise noted. For a complex number z , $\mathbf{Re}[z]$ is the real part of z , $\mathbf{Im}[z]$ is the imaginary part of z , \bar{z} is the complex conjugate of z , and $|z| \stackrel{\text{def}}{=} \sqrt{\bar{z}z}$. At times we consider random variables X taking on *complex* values. For such random variables, we use $\mathbf{Var}[X]$ to denote $\mathbf{E}[|X - \mathbf{E}[X]|^2]$. Note that the usual statement of Chebyshev’s inequality still holds under this definition.

For $x \in \mathbb{R}^n$ and $S \subseteq [n]$, x_S denotes the n -dimensional vector whose i th coordinate is x_i for $i \in S$ and 0 otherwise. For a probabilistic event \mathcal{E} , we use $\mathbf{1}_{\mathcal{E}}$ to denote the indicator random variable for \mathcal{E} . We sometimes refer to a constant as *universal* if it does not depend on other parameters, such as n, m, ε , etc. All space bounds are measured in bits. When measuring time complexity, we assume a word RAM with machine word size $\Omega(\log(nmM))$ so that standard arithmetic and bitwise operations can be performed on words in constant time. We use *reporting time* to refer to the time taken for a streaming algorithm to answer some query (e.g., “output an estimate of F_p ”).

Also, we can assume $n = O(m^2)$ by FKS hashing [14] with an additive $O(\log \log n)$ term in our final space bound; see Section A.1.1 of [26] for details. Thus, henceforth any terms involving n appearing in space and time bounds may be assumed at most m^2 . We also often assume that n, m, M, ε , and δ are powers of 2 (or sometimes 4), and that $1/\sqrt{n} < \varepsilon < \varepsilon_0$ for some universal constant $\varepsilon_0 > 0$. These assumptions are without loss of generality. We can assume $\varepsilon > 1/\sqrt{n}$ since otherwise we could store x explicitly in memory using $O(n \log(mM)) = O(\varepsilon^{-2} \log(mM))$ bits with constant update and reporting times. Finally, we assume $\|x\|_p^p \geq 1$. This is because, since x has integer entries, either $\|x\|_p^p \geq 1$, or it is 0. The case that it is 0 only occurs when x is the 0 vector, which can be detected in $O(\log(nmM))$ space by the AMS sketch [3].

1.4 Organization

An “ F_p ϕ -heavy hitter” is an index j such that $|x_j| \geq \phi \|x\|_p^p$. Sometimes we drop the “ F_p ” if p is understood from context. In Section 2, we give an efficient subroutine `HighEnd` for estimating $\|x_L\|_p^p$ to within additive error $\varepsilon \|x\|_p^p$, where L is a list containing all α -heavy hitters for some $\alpha > 0$, with the promise that no $i \in L$ is not an $\alpha/2$ -heavy hitter. In Section 3 we give a subroutine `LightEstimator` for estimating $\|x_{[n] \setminus L}\|_p^p$. Finally, in Section 4, we put everything together in a way that achieves optimal space and fast update time. We discuss how to compute L in Section A.1.

2 Estimating the contribution from heavy hitters

Before giving our algorithm `HighEnd` for estimating $\|x_L\|_p^p$, we first give a few necessary lemmas and theorems.

The following theorem gives an algorithm for finding the ϕ -heavy hitters with respect to F_p . This algorithm uses the dyadic interval idea of [10] together with a black-box reduction of the problem of finding F_p heavy hitters to the problem of estimating F_p . Our proof is in Section A.1. We note that our data structure both improves and generalizes that of [16], which gave an algorithm with slightly worse bounds that only worked in the case $p = 1$.

Theorem 1. *There is an algorithm $F_p\text{HH}$ satisfying the following properties. Given $0 < \phi < 1$ and $0 < \delta < 1$, with probability at least $1 - \delta$, $F_p\text{HH}$ produces a list L such that L contains all ϕ -heavy hitters and does not contain indices which are not $\phi/2$ -heavy hitters. For each $i \in L$, the algorithm also outputs $\text{sign}(x_i)$, as well as an estimate \tilde{x}_i of x_i satisfying $\tilde{x}_i^p \in [(6/7)|x_i|^p, (9/7)|x_i|^p]$. Its space usage is $O(\phi^{-1} \log(\phi n) \log(nmM) \log(\log(\phi n)/(\delta\phi)))$. Its update time is $O(\log(\phi n) \cdot \log(\log(\phi n)/(\delta\phi)))$. Its reporting time is $O(\phi^{-1}(\log(\phi n) \cdot \log(\log(\phi n)/(\delta\phi))))$.*

The following moment bound can be derived from the Chernoff bound via integration, and is most likely standard though we do not know the earliest reference. A proof can be found in [25].

Lemma 2. *Let X_1, \dots, X_n be such that X_i has expectation μ_i and variance σ_i^2 , and $X_i \leq K$ almost surely. Then if the X_i are ℓ -wise independent for some even integer $\ell \geq 2$,*

$$\mathbf{E} \left[\left(\sum_{i=1}^n X_i - \mu \right)^\ell \right] \leq 2^{O(\ell)} \cdot \left((\sigma\sqrt{\ell})^\ell + (K\ell)^\ell \right),$$

where $\mu = \sum_i \mu_i$ and $\sigma^2 = \sum_i \sigma_i^2$. In particular,

$$\mathbf{Pr} \left[\left| \sum_{i=1}^n X_i - \mu \right| \geq \lambda \right] \leq 2^{O(\ell)} \cdot \left((\sigma\sqrt{\ell}/\lambda)^\ell + (K\ell/\lambda)^\ell \right),$$

by Markov's inequality on the random variable $(\sum_i X_i - \mu)^\ell$.

Lemma 3 (Khinchine inequality [19]). *For $x \in \mathbb{R}^n$, $t \geq 2$, and uniformly random $z \in \{-1, 1\}^n$, $\mathbf{E}_z[|\langle x, z \rangle|^t] \leq \|x\|_2^t \cdot \sqrt{t}$.*

In the following lemma, and henceforth in this section, i denotes $\sqrt{-1}$.

Lemma 4. *Let $x \in \mathbb{R}^n$ be arbitrary. Let $z \in \{e^{2\pi i/r}, e^{2\pi i \cdot 2/r}, e^{2\pi i \cdot 3/r}, \dots, e^{2\pi i \cdot r/r}\}^n$ be a random such vector for $r \geq 2$ an even integer. Then for $t \geq 2$ an even integer, $\mathbf{E}_z[|\langle x, z \rangle|^t] \leq \|x\|_2^t \cdot 2^{t/2} \sqrt{t}$.*

Proof. Since x is real, $|\langle x, z \rangle|^2 = \left(\sum_{j=1}^n \mathbf{Re}[z_j] \cdot x_j \right)^2 + \left(\sum_{j=1}^n \mathbf{Im}[z_j] \cdot x_j \right)^2$. Then by Minkowski's inequality,

$$\begin{aligned} \mathbf{E}[|\langle x, z \rangle|^t] &= \mathbf{E} \left[\left| \left(\sum_{j=1}^n \mathbf{Re}[z_j] \cdot x_j \right)^2 + \left(\sum_{j=1}^n \mathbf{Im}[z_j] \cdot x_j \right)^2 \right|^{t/2} \right] \\ &\leq \left(2 \cdot \max \left\{ \mathbf{E} \left[\left(\sum_{j=1}^n \mathbf{Re}[z_j] \cdot x_j \right)^{2t} \right]^{1/2}, \mathbf{E} \left[\left(\sum_{j=1}^n \mathbf{Im}[z_j] \cdot x_j \right)^{2t} \right]^{1/2} \right\} \right)^{t/2} \\ &\leq 2^{t/2} \cdot \left(\mathbf{E} \left[\left(\sum_{j=1}^n \mathbf{Re}[z_j] \cdot x_j \right)^t \right] + \mathbf{E} \left[\left(\sum_{j=1}^n \mathbf{Im}[z_j] \cdot x_j \right)^t \right] \right). \end{aligned} \quad (1)$$

Since r is even, we may write $\mathbf{Re}[z_j]$ as $(-1)^{y_j}|\mathbf{Re}[z_j]|$ and $\mathbf{Im}[z_j]$ as $(-1)^{y'_j}|\mathbf{Im}[z_j]|$, where $y, y' \in \{-1, 1\}^n$ are random sign vectors chosen independently of each other. Let us fix the values of $|\mathbf{Re}[z_j]|$ and $|\mathbf{Im}[z_j]|$ for each $j \in [n]$, considering just the randomness of y and y' . Applying Lemma 3 to bound each of the expectations in Eq. (1), we obtain the bound $2^{t/2} \cdot \sqrt{t} \cdot (\|b\|_2^t + \|b'\|_2^t) \leq 2^{t/2} \cdot \sqrt{t} \cdot (\|b\|_2^2 + \|b'\|_2^2)^{t/2}$ where $b_j = \mathbf{Re}[z_j] \cdot x_j$ and $b'_j = \mathbf{Im}[z_j] \cdot x_j$. But this is just $2^{t/2} \cdot \sqrt{t} \cdot \|x\|_2^t$ since $|z_j|^2 = 1$. \blacksquare

2.1 The HighEnd data structure

In this section, we assume we know a subset $L \subseteq [n]$ of indices j so that

1. for all j for which $|x_j|^p \geq \alpha \|x\|_p^p$, $j \in L$,
2. if $j \in L$, then $|x_j|^p \geq (\alpha/2) \|x\|_p^p$,
3. for each $j \in L$, we know $\text{sign}(x_j)$.

for some $0 < \alpha < 1/2$ which we know. We also are given some $0 < \varepsilon < 1/2$. We would like to output a value $\|x_L\|_p^p \pm O(\varepsilon) \|x\|_p^p$ with large constant probability. We assume $1/\alpha = O(1/\varepsilon^2)$.

We first define the **BasicHighEnd** data structure. Put $s = \lceil 4/\alpha \rceil$. We choose a hash function $h : [n] \rightarrow [s]$ at random from an r_h -wise independent family for $r_h = \Theta(\log(1/\alpha))$. Also, let $r = \Theta(\log 1/\varepsilon)$ be a sufficiently large even integer. For each $j \in [n]$, we associate a random complex root of unity $e^{2\pi i g(j)/r}$, where $g : [n] \rightarrow [r]$ is drawn at random from an r_g -wise independent family for $r_g = r$. We initialize s counters b_1, \dots, b_s to 0. Given an update of the form (j, v) , add $e^{2\pi i g(j)/r} \cdot v$ to $b_{h(j)}$.

We now define the **HighEnd** data structure as follows. Define $T = \tau \cdot \max\{\log(1/\varepsilon), \log(2/\alpha)\}$ for a sufficiently large constant τ to be determined later. Define $t = 3T$ and instantiate t independent copies of the **BasicHighEnd** data structure. Given an update (j, v) , perform the update described above to each of the copies of **BasicHighEnd**. We think of this data structure as a $t \times s$ matrix of counters $D_{j,k}$, $j \in [t]$ and $k \in [s]$. We let g^j be the hash function g in the j th independent instantiation of **BasicHighEnd**, and similarly define h^j . We sometimes use g to denote the tuple (g^1, \dots, g^t) , and similarly for h .

We now define our estimator, but first we give some notation. For $w \in L$, let $j(w, 1) < j(w, 2) < \dots < j(w, n_w)$ be the set of n_w indices $j \in [t]$ such that w is *isolated* by h^j from other indices in L ; that is, indices $j \in [t]$ where no other $w' \in L$ collides with w under h^j .

Event \mathcal{E} . Define \mathcal{E} to be the event that $n_w \geq T$ for all $w \in L$.

If \mathcal{E} does not hold, our estimator simply fails. Otherwise, define

$$x_w^* = \frac{1}{T} \cdot \sum_{k=1}^T e^{-2\pi i g^{j(w,k)}(w)/r} \cdot \text{sign}(x_w) \cdot D_{j(w,k), h^{j(w,k)}(w)}.$$

If $\mathbf{Re}[x_w^*] < 0$ for any $w \in L$, then we output fail. Otherwise, define

$$\Psi' = \sum_{w \in L} (x_w^*)^p.$$

Our estimator is then $\Psi = \mathbf{Re}[\Psi']$. Note x^* is a complex number. By z^p for complex z , we mean $|z|^p \cdot e^{ip \cdot \arg(z)}$, where $\arg(z) \in (-\pi, \pi]$ is the angle formed by the vector from the origin to z in the complex plane.

2.2 A useful random variable

For $w \in L$, we make the definitions

$$y_w \stackrel{\text{def}}{=} \frac{x_w^* - |x_w|}{|x_w|}, \quad \Phi_w \stackrel{\text{def}}{=} |x_w|^p \cdot \left(\sum_{k=0}^{r/3} \binom{p}{k} \cdot y_w^k \right)$$

as well as $\Phi \stackrel{\text{def}}{=} \sum_{w \in L} \Phi_w$. We assume \mathcal{E} occurs so that the y_w and Φ_w (and hence Φ) are defined. Also, we use the definition $\binom{p}{k} = (\prod_{j=0}^{k-1} (p-j))/k!$ (note p may not be an integer).

Our overall goal is to show that $\Psi = \|x_L\|_p^p \pm O(\varepsilon) \cdot \|x\|_p^p$ with large constant probability. Our proof plan is to first show that $|\Phi - \|x_L\|_p^p| = O(\varepsilon) \cdot \|x\|_p^p$ with large constant probability, then to show that $|\Psi' - \Phi| = O(\varepsilon) \cdot \|x\|_p^p$ with large constant probability, at which point our claim follows by a union bound and the triangle inequality since $|\Psi - \|x_L\|_p^p| \leq |\Psi' - \|x_L\|_p^p|$ since $\|x_L\|_p^p$ is real.

Before analyzing Φ , we define the following event.

Event \mathcal{D} . Let \mathcal{D} be the event that for all $w \in L$ we have

$$\frac{1}{T^2} \sum_{k=1}^T \sum_{\substack{v \notin L \\ h^j(w,k)(v) = h^j(w,k)(w)}} x_v^2 < \frac{(\alpha \cdot \|x\|_p^p)^{2/p}}{r}.$$

We also define

$$V = \frac{1}{T^2} \sum_{w \in L} \sum_{j=1}^t \sum_{\substack{v \notin L \\ h^j(w) = h^j(v)}} |x_w|^{2p-2} \cdot |x_v|^2.$$

Theorem 5. *Conditioned on h , $\mathbf{E}_g[\Phi] = \|x_L\|_p^p$ and $\mathbf{Var}_g[\Phi | \mathcal{D}] = O(V)$.*

Proof. By linearity of expectation,

$$\mathbf{E}_g[\Phi] = \sum_{w \in L} |x_w|^p \cdot \left[\sum_{k=0}^{r/3} \binom{p}{k} \mathbf{E}_g[y_w^k] \right] = \sum_{w \in L} |x_w|^p + \sum_{w \in L} |x_w|^p \cdot \sum_{k=1}^{r/3} \binom{p}{k} \mathbf{E}_g[y_w^k],$$

where we use that $\binom{p}{0} = 1$. Then $\mathbf{E}_g[y_w^k] = 0$ for $k > 0$ by using linearity of expectation and r_g -wise independence, since each summand involves at most $k < r$ r th roots of unity. Hence,

$$\mathbf{E}_g[\Phi] = \sum_{w \in L} |x_w|^p.$$

We now compute the variance. Note that if the g^j were each fully independent, then we would have $\mathbf{Var}_g[\Phi | \mathcal{D}] = \sum_{w \in L} \mathbf{Var}_g[\Phi_w | \mathcal{D}]$ since different Φ_w depend on evaluations of the g^j on disjoint $v \in [n]$. However, since $r_g > 2r/3$, $\mathbf{E}_g[|\Phi|^2]$ is identical as in the case of full independence of the g^j .

We thus have $\mathbf{Var}_g[\Phi \mid \mathcal{D}] = \sum_{w \in L} \mathbf{Var}_g[\Phi_w \mid \mathcal{D}]$ and have reduced to computing $\mathbf{Var}_g[\Phi_w \mid \mathcal{D}]$.

$$\begin{aligned} \mathbf{Var}_g[\Phi_w \mid \mathcal{D}] &= \mathbf{E}_g[|\Phi_w - \mathbf{E}_g[\Phi_w]|^2 \mid \mathcal{D}] \\ &= |x_w|^{2p} \cdot \mathbf{E}_g \left[\left| \sum_{k=1}^{r/3} \binom{p}{k} y_w^k \right|^2 \mid \mathcal{D} \right] \\ &= |x_w|^{2p} \cdot \left(p^2 \cdot \mathbf{E}_g[|y_w|^2 \mid \mathcal{D}] + \sum_{k=2}^{r/3} O(\mathbf{E}_g[|y_w|^{2k} \mid \mathcal{D}]) \right) \end{aligned}$$

We have

$$\mathbf{E}_g[|y_w|^2 \mid \mathcal{D}] \stackrel{\text{def}}{=} u_w^2 = \frac{1}{T^2} \sum_{k=1}^T \sum_{\substack{v \notin L \\ h^{j(w,k)}(v) = h^{j(w,k)}(w)}} \frac{x_v^2}{x_w^2}, \quad (2)$$

so that

$$\sum_{w \in L} p^2 \cdot \mathbf{E}_g[|y_w|^2 \mid \mathcal{D}] \leq p^2 V.$$

Eq. (2) follows since, conditioned on \mathcal{E} so that y_w is well-defined,

$$\mathbf{E}_g[|y_w|^2] = \frac{1}{T^2 x_w^2} \sum_{k=1}^T \sum_{k'=1}^T \sum_{\substack{v \notin L \\ h^{j(w,k)}(v) = h^{j(w,k)}(w)}} \sum_{\substack{v' \notin L \\ h^{j(w,k')}(v') = h^{j(w,k')}(w)}} \mathbf{E}[e^{-2\pi i(g^{j(w,k)}(v) - g^{j(w,k')}(v'))/r}] x_v x_{v'}.$$

When $j(w, k) \neq j(w, k')$ the above expectation is 0 since the g^j are independent across different j . When $j(w, k) = j(w, k')$ the above expectation is only non-zero for $v = v'$ since $r_g \geq 2$.

We also have for $k \geq 2$ that

$$\mathbf{E}_g[|y_w|^{2k} \mid \mathcal{D}] \leq 2^{O(k)} \cdot u_w^{2k} \cdot (2k)^k$$

by Lemma 4, so that

$$\sum_{k=2}^{r/3} \mathbf{E}_g[|y_w|^{2k} \mid \mathcal{D}] = O(u_w^2)$$

since \mathcal{D} holds and so the sum is dominated by its first term. Thus, $\mathbf{Var}_g[\Phi \mid \mathcal{D}] = O(V)$. \blacksquare

Lemma 6. $\mathbf{E}_h[V] \leq 3\alpha \cdot \|x\|_p^{2p}/(4T)$.

Proof. For any $w \in L$, $v \notin L$, and $j \in [t]$, we have $\Pr_h[h^j(w) = h^j(v)] = 1/s \leq \alpha/4$ since $r_h \geq 2$.

Thus,

$$\begin{aligned}
\mathbf{E}_h[V] &\leq \frac{\alpha}{4T^2} \sum_{\substack{w \in L \\ v \notin L \\ j \in [t]}} |x_w|^{2p-2} |x_v|^2 \\
&= \frac{3\alpha}{4T} \left(\sum_{w \in L} |x_w|^p |x_w|^{p-2} \right) \left(\sum_{v \notin L} |x_v|^2 \right) \\
&\leq \frac{3\alpha}{4T} \left(\sum_{w \in L} \|x\|_p^p (\alpha \cdot \|x\|_p^p)^{(p-2)/p} \right) \left(\frac{1}{\alpha} (\alpha \cdot \|x\|_p^p)^{2/p} \right) \\
&= \frac{3}{4} \cdot \alpha \cdot \|x\|_p^{2p} / T.
\end{aligned} \tag{3}$$

where Eq. (3) used that $\|x_{[n] \setminus L}\|_2^2$ is maximized when $[n] \setminus L$ contains exactly $1/\alpha$ coordinates v each with $|x_v|^p = \alpha \|x\|_p^p$, and that $|x_w|^{p-2} \leq (\alpha \cdot \|x\|_p^p)^{(p-2)/p}$ since $p \leq 2$. ■

Lemma 7. $\Pr_h[\mathcal{E}] \geq 1 - \varepsilon$.

Proof. For any $j \in [t]$, the probability that w is isolated by h^j is at least $1/2$, since the expected number of collisions with w is at most $1/2$ by pairwise independence of the h^j and the fact that $|L| \leq 2/\alpha$ so that $s \geq 2|L|$. If X is the expected number of buckets where w is isolated, the Chernoff bound gives $\Pr_h[X < (1 - \epsilon)\mathbf{E}_h[X]] < \exp(-\epsilon^2 \mathbf{E}_h[X]/2)$ for $0 < \epsilon < 1$. The claim follows for $\tau \geq 24$ by setting $\epsilon = 1/3$ then applying a union bound over $w \in L$. ■

Lemma 8. $\Pr_h[\mathcal{D}] \geq 63/64$.

Proof. We apply the bound of Lemma 2 for a single $w \in L$. Define $X_{j,v} = (x_v^2/T^2) \cdot \mathbf{1}_{h^j(v)=h^j(w)}$ and $X = \sum_{j=1}^t \sum_{v \notin L} X_{j,v}$. Note that X is an upper bound for the left hand side of the inequality defining \mathcal{D} , and thus it suffices to show a tail bound for X . In the notation of Lemma 2, we have $\sigma^2 \leq (3/(sT^3)) \cdot \|x_{[n] \setminus L}\|_4^4$, $K = (\alpha \cdot \|x\|_p^p)^{2/p}/T^2$, and $\mu = (3/(sT)) \cdot \|x_{[n] \setminus L}\|_2^2$. Since $\|x_{[n] \setminus L}\|_2^2$ and $\|x_{[n] \setminus L}\|_4^4$ are each maximized when there are exactly $1/\alpha$ coordinates $v \notin L$ with $|x_v|^p = \alpha \cdot \|x\|_p^p$,

$$\sigma^2 \leq \frac{3}{4T^3} \cdot (\alpha \cdot \|x\|_p^p)^{4/p}, \quad \mu \leq \frac{3}{4T} \cdot (\alpha \cdot \|x\|_p^p)^{2/p}.$$

Setting $\lambda = (\alpha \cdot \|x\|_p^p)^{2/p}/(2r)$, noting that $\mu < \lambda$ for τ sufficiently large, and assuming $\ell \leq r_h$ is even, we apply Lemma 2 to obtain

$$\Pr[X \geq 2\lambda] \leq 2^{O(\ell)} \cdot \left(\left(\frac{\sqrt{3}r \cdot \sqrt{\ell}}{T^{3/2}} \right)^\ell + \left(\frac{2r \cdot \ell}{T^2} \right)^\ell \right).$$

By setting τ sufficiently large and $\ell = \log(2/\alpha) + 6$, the above probability is at most $(1/64) \cdot (\alpha/2)$. The lemma follows by a union bound over all $w \in L$, since $|L| \leq 2/\alpha$. ■

We now define another event.

Event \mathcal{F} . Let \mathcal{F} be the event that for all $w \in L$ we have $|y_w| < 1/2$.

Lemma 9. $\Pr_g[\mathcal{F} \mid \mathcal{D}] \geq 63/64$.

Proof. \mathcal{D} occurring implies that $u_w \leq \sqrt{1/r} \leq \sqrt{1/(64(\log(2/\alpha) + 6))}$ (recall we assume $1/\alpha = O(1/\varepsilon^2)$ and pick $r = \Theta(\log(1/\varepsilon))$ sufficiently large, and u_w is as is defined in Eq. (2)), and we also have $\mathbf{E}_g[|y_w|^\ell \mid \mathcal{D}] < u_w^\ell \sqrt{\ell} 2^\ell$ by Lemma 4. Applying Markov's bound on the random variable $|y_w|^\ell$ for even $\ell \leq r_g$, we have $|y_w|^\ell$ is determined by r_g -wise independence of the g^j , and thus

$$\Pr_g[|y_w| \geq 1/2 \mid \mathcal{D}] < \left(\sqrt{\frac{16\ell}{64(\log(2/\alpha) + 6)}} \right)^\ell,$$

which equals $(1/64) \cdot (\alpha/2)$ for $\ell = \log(2/\alpha) + 6$. We then apply a union bound over all $w \in L$. ■

Lemma 10. Given \mathcal{F} , $|\Psi' - \Phi| < \varepsilon \|x_L\|_p^p$.

Proof. Observe

$$\Psi' = \sum_{w \in L} |x_w|^p \cdot (1 + y_w)^p.$$

We have that $\ln(1 + z)$, as a function of z , is holomorphic on the open disk of radius 1 about 0 in the complex plane, and thus $f(z) = (1 + z)^p$ is holomorphic in this region since it is the composition $\exp(p \cdot \ln(1 + z))$ of holomorphic functions. Therefore, $f(z)$ equals its Taylor expansion about 0 for all $z \in \mathbb{C}$ with $|z| < 1$ (see for example [39, Theorem 11.2]). Then since \mathcal{F} occurs, we can Taylor-expand f about 0 for $z = y_w$ and apply Taylor's theorem to obtain

$$\begin{aligned} \Psi' &= \sum_{w \in L} |x_w|^p \left(\sum_{k=0}^{r/3} \binom{p}{k} y_w^k \pm O\left(\binom{p}{r/3+1} \cdot |y_w|^{-r/3-1} \right) \right) \\ &= \Phi + O\left(\|x_L\|_p^p \cdot \left(\binom{p}{r/3+1} \cdot |y_w|^{-r/3-1} \right) \right) \end{aligned}$$

The lemma follows since $\binom{p}{r/3+1} < 1$ and $|y_w|^{-r/3-1} < \varepsilon$ for $|y_w| < 1/2$. ■

Theorem 11. The space used by HighEnd is $O(\alpha^{-1} \log(1/\varepsilon) \log(mM/\varepsilon) + O(\log^2(1/\varepsilon) \log n))$. The update time is $O(\log^2(1/\varepsilon))$. The reporting time is $O(\alpha^{-1} \log(1/\varepsilon) \log(1/\alpha))$. Also, $\Pr_{h,g}[|\Psi - \|x_L\|_p^p| < O(\varepsilon) \cdot \|x\|_p^p] > 7/8$.

Proof. We first argue correctness. By a union bound, \mathcal{E} and \mathcal{D} hold simultaneously with probability $31/32$. By Markov's inequality and Lemma 6, $V = O(\alpha \cdot \|x\|_p^{2p}/T)$ with probability $63/64$. We then have by Chebyshev's inequality and Theorem 5 that $|\Phi - \|x_L\|_p^p| = O(\varepsilon) \cdot \|x\|_p^p$ with probability $15/16$. Lemma 10 then implies $|\Psi' - \|x_L\|_p^p| = O(\varepsilon) \cdot \|x\|_p^p$ with probability $15/16 - \Pr[\neg \mathcal{F}] > 7/8$ by Lemma 9. In this case, the same must hold true for Ψ since $\Psi = \mathbf{Re}[\Psi']$ and $\|x_L\|_p^p$ is real.

Next we discuss space complexity. We start with analyzing the precision required to store the counters $D_{j,k}$. Since our correctness analysis conditions on \mathcal{F} , we can assume \mathcal{F} holds. We store the real and imaginary parts of each counter $D_{j,k}$ separately. If we store each such part to within precision $\gamma/(2mT)$ for some $0 < \gamma < 1$ to be determined later, then each of the real and imaginary parts, which are the sums of at most m summands from the m updates in the stream, is stored to within additive error $\gamma/(2T)$ at the end of the stream. Let \tilde{x}_w^* be our calculation of x_w^* with such

limited precision. Then, each of the real and imaginary parts of \tilde{x}_w^* is within additive error $\gamma/2$ of those for x_w^* . Since \mathcal{F} occurs, $|x_w^*| > 1/2$, and thus $\gamma/2 < \gamma|x_w^*|$, implying $|\tilde{x}_w^*| = (1 \pm O(\gamma))|x_w^*|$. Now we argue $\arg(\tilde{x}_w^*) = \arg(x_w^*) \pm O(\sqrt{\gamma})$. Write $x_w^* = a + ib$ and $\tilde{x}_w^* = \tilde{a} + i\tilde{b}$ with $\tilde{a} = a \pm \gamma/2$ and $\tilde{b} = b \pm \gamma/2$. We have $\cos(\arg(x_w^*)) = a/\sqrt{a^2 + b^2}$. Also, $\cos(\arg(\tilde{x}_w^*)) = (\tilde{a} \pm \gamma/2)/((1 \pm O(\gamma))\sqrt{a^2 + b^2}) = (1 \pm O(\gamma))\cos(\arg(x_w^*)) \pm O(\gamma) = \cos(\arg(x_w^*)) \pm O(\gamma)$, implying $\arg(\tilde{x}_w^*) = \arg(x_w^*) \pm O(\sqrt{\gamma})$. Our final output is $\sum_{w \in L} |\tilde{x}_w^*|^p \cdot \cos(p \cdot \arg(\tilde{x}_w^*))$. Since \cos never has derivative larger than 1 in magnitude, this is $\sum_{w \in L} [(1 \pm O(\gamma))|x_w^*|^p \cos(p \cdot \arg(x_w^*)) \pm O(\sqrt{\gamma}) \cdot (1 \pm O(\gamma))|x_w^*|^p]$. Since \mathcal{F} occurs, $|x_w^*|^p < (3/2)^p \cdot |x_w|^p$, and thus our overall error introduced from limited precision is $O(\sqrt{\gamma} \cdot \|x_L\|_p^p)$, and it thus suffices to set $\gamma = O(\varepsilon^2)$, implying each $D_{j,k}$ requires $O(\log(mM/\varepsilon))$ bits of precision. For the remaining part of the space analysis, we discuss storing the hash functions. The hash functions h^j, g^j each require $O(\log(1/\varepsilon) \log n)$ bits of seed, and thus in total consume $O(\log^2(1/\varepsilon) \log n)$ bits.

Finally we discuss time complexity. To perform an update, for each $j \in [t]$ we must evaluate g^j and h^j then update a counter. Each of g^j, h^j require $O(\log(1/\varepsilon))$ time to evaluate. For the reporting time, we can mark all counters with the unique $w \in L$ which hashes to it under the corresponding h^j (if a unique such w exists) in $|L| \cdot t \cdot r_h = O(\alpha^{-1} \log(1/\varepsilon) \log(1/\alpha))$ time. Then, we sum up the appropriate counters for each $w \in L$, using the Taylor expansion of $\cos(p \cdot \arg(z))$ up to the $\Theta(\log(1/\varepsilon))$ th degree to achieve additive error ε . Note that conditioned on \mathcal{F} , $\arg(x_w^*) \in (-\pi/4, \pi/4)$, so that $|p \cdot \arg(x_w^*)|$ is bounded away from $\pi/2$ for p bounded away from 2; in fact, one can even show via some calculus that $\arg(x_w^*) \in (-\pi/6, \pi/6)$ when \mathcal{F} occurs by showing that $\cos(\arg(x_w^*)) = \cos(\arg(1 - y_w))$ is minimized for $|y_w| \leq 1/2$ when $y_w = 1/4 + i\sqrt{3}/4$. Regardless, additive error ε is relative error $O(\varepsilon)$, since if $|p \cdot \arg(z)|$ is bounded away from $\pi/2$, then $|\cos(p \cdot \arg(z))| = \Omega(1)$. \blacksquare

3 Estimating the contribution from light elements

In this section, we show how to estimate the contribution to F_p from coordinates of x which are not heavy hitters. More precisely, given a list $L \subseteq [n]$ such that $|L| \leq 2/\varepsilon^2$ and $|x_i|^p \leq \varepsilon^2 \|x\|_p^p$ for all $i \notin L$, we describe a subroutine **LightEstimator** that outputs a value that is $\|x_{[n] \setminus L}\|_p^p \pm O(\varepsilon) \cdot \|x\|_p^p$ with probability at least $7/8$. This estimator is essentially the same as that given for $p = 1$ in [30], though in this work we show that (some variant of) the geometric mean estimator of [28] requires only bounded independence, in order that we may obtain optimal space.

Our description follows. We first need the following theorem, which comes from a derandomized variant of the geometric mean estimator. Our proof is in Section A.2.

Theorem 12. *For any $0 < p < 2$, there is a randomized data structure D_p , and a deterministic algorithm Est_p mapping the state space of the data structure to reals, such that*

1. $\mathbf{E}[\text{Est}_p(D_p(x))] = (1 \pm \varepsilon) \|x\|_p^p$
2. $\mathbf{E}[\text{Est}_p(D_p(x))^2] \leq C_p \cdot \|x\|_p^{2p}$

for some constant $C_p > 0$ depending only on p , and where the expectation is taken over the randomness used by D_p . Aside from storing a length- $O(\varepsilon^{-p} \log(nmM))$ random string, the space complexity is $O(\log(nmM))$. The update time is the time to evaluate a $\Theta(1/\varepsilon^p)$ -wise independent hash function over a field of size $\text{poly}(nmM)$, and the reporting time is $O(1)$.

We also need the following algorithm for fast multipoint evaluation of polynomials.

Theorem 13 ([38, Ch. 10]). *Let \mathbf{R} be a ring, and let $q \in \mathbf{R}[x]$ be a degree- d polynomial. Then, given distinct $x_1, \dots, x_d \in \mathbf{R}$, all the values $q(x_1), \dots, q(x_d)$ can be computed using $O(d \log^2 d \log \log d)$ operations over \mathbf{R} .*

The guarantees of the final `LightEstimator` are then given in Theorem 15, which is a modified form of an algorithm designed in [30] for the case $p = 1$. A description of the modifications of the algorithm in [30] needed to work for $p \neq 2$ is given in Remark 16, which in part uses the following uniform hash family of Pagh and Pagh [35].

Theorem 14 (Pagh and Pagh [35, Theorem 1.1]). *Let $S \subseteq U = [u]$ be a set of $z > 1$ elements, and let $V = [v]$, with $1 < v \leq u$. Suppose the machine word size is $\Omega(\log(u))$. For any constant $c > 0$ there is a word RAM algorithm that, using time $\log(z) \log^{O(1)}(v)$ and $O(\log(z) + \log \log(u))$ bits of space, selects a family \mathcal{H} of functions from U to V (independent of S) such that:*

1. *With probability $1 - O(1/z^c)$, \mathcal{H} is z -wise independent when restricted to S .*
2. *Any $h \in \mathcal{H}$ can be represented by a RAM data structure using $O(z \log(v))$ bits of space, and h can be evaluated in constant time after an initialization step taking $O(z)$ time.*

Theorem 15 ([30]). *Suppose we are given $0 < \varepsilon < 1$, and given a list $L \subseteq [n]$ at the end of the data stream such that $|L| \leq 2/\varepsilon^2$ and $|x_i|^p < \varepsilon^2 \|x\|_p^p$ for all $i \notin L$. Then, given access to a randomized data structure satisfying properties (1) and (2) of Theorem 12, there is an algorithm `LightEstimator` satisfying the following. The randomness used by `LightEstimator` can be broken up into a certain random hash function h , and another random string s . `LightEstimator` outputs a value Φ' satisfying $\mathbf{E}_{h,s}[\Phi'] = (1 \pm O(\varepsilon)) \|x_{[n] \setminus L}\|_p^p$, and $\mathbf{E}_h[\mathbf{Var}_s[\Phi']] = O(\varepsilon^2 \|x\|_p^{2p})$. The space usage is $O(\varepsilon^{-2} \log(nmM))$, the update time is $O(\log^2(1/\varepsilon) \log \log(1/\varepsilon))$, and the reporting time is $O(1/\varepsilon^2)$.*

Remark 16. The claim of Theorem 15 is not stated in the same form in [30], and thus we provide some explanation. The work of [30] only focused on the case $p = 1$. There, in Section 3.2, `LightEstimator` was defined² by creating $R = 4/\varepsilon^2$ independent instantiations of D_1 , which we label D_1^1, \dots, D_1^R (R chosen so that $R \geq 2|L|$), and picking a hash function $h : [n] \rightarrow [R]$ from a random hash family constructed as in Theorem 14 with $z = R$ and $c \geq 2$. Upon receiving an update to x_i in the stream, the update was fed to $D_1^{h(i)}$. The final estimate was defined as follows. Let $I = [R] \setminus h(L)$. Then, the estimate was $\Phi' = (R/|I|) \cdot \sum_{j \in I} \mathbf{Est}_1(D_1^j)$. In place of a generic D_1 , the presentation in [30] used Li's geometric mean estimator [28], though the analysis (Lemmas 7 and 8 of [30]) only made use of the generic properties of D_1 and \mathbf{Est}_1 given in Theorem 12. Let $s = (s_1, \dots, s_R)$ be the tuple of random strings used by the D_1^j , where the entries of s are pairwise independent. The analysis then showed that (a) $\mathbf{E}_{h,s}[\Phi'] = (1 \pm O(\varepsilon)) \|x_{[n] \setminus L}\|_1$, and (b) $\mathbf{E}_h[\mathbf{Var}_s[\Phi']] = O(\varepsilon^2 \|x\|_1^2)$. For (a), the same analysis applies for $p \neq 1$ when using \mathbf{Est}_p and D_p instead. For (b), it was shown that $\mathbf{E}_h[\mathbf{Var}_s[\Phi']] = O(\|x_{[n] \setminus L}\|_2^2 + \varepsilon^2 \|x_{[n] \setminus L}\|_1^2)$. The same analysis shows that $\mathbf{E}_h[\mathbf{Var}_s[\Phi']] = O(\|x_{[n] \setminus L}\|_{2p}^{2p} + \varepsilon^2 \|x_{[n] \setminus L}\|_p^p)$ for $p \neq 1$. Since L contains all the ε^2 -heavy hitters, $\|x_{[n] \setminus L}\|_{2p}^{2p}$ is maximized when there are $1/\varepsilon^2$ coordinates $i \in [n] \setminus L$ each with $|x_i|^p = \varepsilon^2 \|x\|_p^p$, in which case $\|x_{[n] \setminus L}\|_{2p}^{2p} = \varepsilon^2 \|x\|_p^{2p}$.

²The estimator given there was never actually named, so we name it `LightEstimator` here.

To achieve the desired update time, we buffer every $d = 1/\varepsilon^p$ updates then perform the fast multipoint evaluation of Theorem 13 in batch (note this does not affect our space bound since $p < 2$). That is, although the hash function h can be evaluated in constant time, updating any D_p^j requires evaluating a degree- $\Omega(1/\varepsilon^p)$ polynomial, which naïvely requires $\Omega(1/\varepsilon^p)$ time. Note that one issue is that the different data structures D_p^j use different polynomials, and thus we may need to evaluate $1/\varepsilon^p$ different polynomials on the $1/\varepsilon^p$ points, defeating the purpose of batching. To remedy this, note that these polynomials are themselves pairwise independent. That is, we can assume there are two coefficient vectors a, b of length $d+1$, and the polynomial corresponding to D_p^j is given by the coefficient vector $j \cdot a + b$. Thus, we only need to perform fast multipoint evaluation on the two polynomials defined by a and b . To achieve worst-case update time, this computation can be spread over the next d updates. If a query comes before d updates are batched, we need to perform $O(d \log d \log \log d)$ work at once, but this is already dominated by our $O(1/\varepsilon^2)$ reporting time since $p < 2$.

4 The final algorithm: putting it all together

To obtain our final algorithm, one option is to run `HighEnd` and `LightEstimator` in parallel after finding L , then output the sum of their estimates. Note that by the variance bound in Theorem 15, the output of a single instantiation of `LightEstimator` is $\|x_{[n] \setminus L}\|_p^p \pm O(\varepsilon)\|x\|_p^p$ with large constant probability. The downside to this option is that Theorem 1 uses space that would make our overall F_p -estimation algorithm suboptimal by $\text{polylog}(n/\varepsilon)$ factors, and `HighEnd` by an $O(\log(1/\varepsilon))$ factor for $\alpha = \varepsilon^2$ (Theorem 11). We can overcome this by a combination of balancing and universe reduction. Specifically, for balancing, notice that if instead of having L be a list of ε^2 -heavy hitters, we instead defined it as a list of ε^2 -heavy hitters for some $\epsilon > \varepsilon$, we could improve the space of both Theorem 1 and Theorem 11. To then make the variance in `LightEstimator` sufficiently small, i.e. $O(\varepsilon^2\|x\|_p^2)$, we could run $O((\epsilon/\varepsilon)^2)$ instantiations of `LightEstimator` in parallel and output the average estimate, keeping the space optimal but increasing the update time to $\Omega((\epsilon/\varepsilon)^2)$. This balancing gives a smooth tradeoff between space and update time; in fact note that for $\epsilon = 1$, our overall algorithm simply becomes a derandomized variant of Li's geometric mean estimator. We would like though to have $\epsilon \ll 1$ to have small update time.

Doing this balancing does not resolve all our issues though, since Theorem 1 is suboptimal by a $\log n$ factor. That is, even if we picked $\epsilon = 1$, Theorem 1 would cause our overall space to be $\Omega(\log(n) \log(mM))$, which is suboptimal. To overcome this issue we use universe reduction. Specifically, we set $N = 1/\varepsilon^{18}$ and pick hash functions $h_1 : [n] \rightarrow [N]$ and $\sigma : [n] \rightarrow \{-1, 1\}$. We define a new N -dimensional vector y by $y_i = \sum_{h_1(j)=i} \sigma(j)x_j$. Henceforth in this section, y , h_1 , and σ are as discussed here. Rather than computing a list L of heavy hitters of x , we instead compute a list L' of heavy hitters of y . Then, since y has length only $\text{poly}(1/\varepsilon)$, Theorem 1 is only suboptimal by $\text{polylog}(1/\varepsilon)$ factors and our balancing trick applies. The list L' is also used in place of L for both `HighEnd` and `LightEstimator`. Though, since we never learn L , we must modify the algorithm `LightEstimator` described in Remark 16. Namely, the hash function $h : [n] \rightarrow [R]$ in Remark 16 should be implemented as the composition of h_1 , and a hash function $h_2 : [N] \rightarrow [R]$ chosen as Theorem 14 (again with $z = R$ and $c = 2$). Then, we let $I = [R] \setminus h_2(L')$. The remaining parts of the algorithm remain the same.

There are several issues we must address to show that our universe reduction step still maintains correctness. Informally, we need that (a) any i which is a heavy hitter for y should have exactly

one $j \in [n]$ with $h_1(j) = i$ such that j was a heavy hitter for x , (b) if i is a heavy hitter for x , then $h_1(i)$ is a heavy hitter for y , and $|y_{h_1(i)}|^p = (1 \pm O(\varepsilon))|x_i|^p$ so that x_i 's contribution to $\|x\|_p^p$ is properly approximated by **HighEnd**, (c) $\|y\|_p^p = O(\|x\|_p^p)$ with large probability, since the error term in **HighEnd** is $O(\varepsilon \cdot \|y\|_p^p)$, and (d) the amount of F_p mass not output by **LightEstimator** because it collided with a heavy hitter for x under h_1 is negligible. Also, the composition $h = h_1 \circ h_2$ for **LightEstimator** does not satisfy the conditions of Theorem 14 even though h_1 and h_2 might do so individually. To see why, as a simple analogy, consider that the composition of two purely random functions is no longer random. For example, as the number of compositions increases, the probability of two items colliding increases as well. Nevertheless, the analysis of **LightEstimator** carries over essentially unchanged in this setting, since whenever considering the distribution of where two items land under h , we can first condition on them not colliding under h_1 . Not colliding under h_1 happens with $1 - O(\varepsilon^{18})$ probability, and thus the probability that two items land in two particular buckets $j, j' \in [R]$ under h is still $(1 \pm o(\varepsilon))/R^2$.

We now give our full description and analysis. We pick h_1 as in Theorem 14 with $z = R$ and $c = c_h$ a sufficiently large constant. We also pick σ from an $\Omega(\log N)$ -wise independent family. We run an instantiation of F_p HH for the vector y with $\phi = \varepsilon^2/(34C)$ for a sufficiently large constant $C > 0$. We also obtain a value $\tilde{F}_p \in [F_p/2, 3F_p/2]$ using the algorithm of [26]. We define L' to be the sublist of those w output by our F_p HH instantiation such that $|\tilde{y}_w|^p \geq (2\varepsilon^2/7)\tilde{F}_p$.

For ease of presentation in what follows, define L_ϕ to be the list of ϕ -heavy hitters of x with respect to F_p (“ L ”, without a subscript, always denotes the ε^2 -heavy hitters with respect to x), and define $z_i = \sum_{w \in h_1^{-1}(i) \setminus L_{\varepsilon^8}} \sigma(w)x_w$, i.e. z_i is the contribution to y_i from the significantly light elements of x .

Lemma 17. *For $x \in \mathbb{R}^n$, $\lambda > 0$ with λ^2 a multiple of 8, and random $z \in \{-1, 1\}^n$ drawn from a $(\lambda^2/4)$ -wise independent family, $\Pr[|\langle x, z \rangle| > \lambda \|x\|_2] < 2^{-\lambda^2/4}$.*

Proof. By Markov's inequality on the random variable $\langle x, z \rangle^{\lambda^2/4}$, $\Pr[|\langle x, z \rangle| > \lambda] < \lambda^{-\lambda^2/4} \cdot \mathbf{E}[\langle x, z \rangle^{\lambda^2/4}]$. The claim follows by applying Lemma 3. \blacksquare

Lemma 18. *For any $C > 0$, there exists ε_0 such that for $0 < \varepsilon < \varepsilon_0$, $\Pr[\|y\|_p^p > 17C\|x\|_p^p] < 2/C$.*

Proof. Condition on h_1 . Define $Y(i)$ to be the vector $x_{h_1^{-1}(i)}$. For any vector v we have $\|v\|_2 \leq \|v\|_p$ since $p < 2$. Letting \mathcal{E} be the event that no $i \in [N]$ has $|y_i| > 4\sqrt{\log N}\|Y(i)\|_p$, we have $\Pr[\mathcal{E}] \geq 1 - 1/N^4$ by Lemma 17. For $i \in [N]$, again by Lemma 17 any $i \in [N]$ has $|y_i| \leq 2t \cdot \|Y(i)\|_2 \leq 2t \cdot \|Y(i)\|_p$ with probability at least $1 - \max\{1/(2N), 2^{-t^2}\}$. Then for fixed $i \in [N]$,

$$\begin{aligned} \mathbf{E}[|y_i|^p \mid \mathcal{E}] &\leq 2^p \|Y(i)\|_p^p + \sum_{t=0}^{\infty} \Pr[(2 \cdot 2^t)^p \|Y(i)\|_p^p < |y_i|^p \leq (2 \cdot 2^{t+1})^p \|Y(i)\|_p^p \mid \mathcal{E}] \cdot (2 \cdot 2^{t+1})^p \|Y(i)\|_p^p \\ &\leq 2^p \|Y(i)\|_p^p + (1/\Pr[\mathcal{E}]) \cdot \sum_{t=0}^{\log(2\sqrt{\log N})} 2^{-2^{2t}} \cdot (2 \cdot 2^{t+1})^p \|Y(i)\|_p^p \\ &< 4 \|Y(i)\|_p^p + (1/\Pr[\mathcal{E}]) \cdot \sum_{t=0}^{\log(2\sqrt{\log N})} 2^{-2^{2t}} \cdot (2 \cdot 2^{t+1})^2 \|Y(i)\|_p^p \\ &< 17 \|Y(i)\|_p^p \end{aligned}$$

since $\Pr[\mathcal{E}] \geq 1 - 1/N^4$ and ε_0 is sufficiently small. Thus by linearity of expectation, $\mathbf{E}[\|y\|_p^p \mid \mathcal{E}] \leq 17\|x\|_p^p$, which implies $\|y\|_p^p \leq 17C\|x\|_p^p$ with probability $1 - 1/C$, conditioned on \mathcal{E} holding. We conclude by again using $\Pr[\mathcal{E}] \geq 1 - 1/N^4$. \blacksquare

Lemma 19. *With probability at least $1 - \text{poly}(\varepsilon)$ over σ , simultaneously for all $i \in [N]$ we have that $|z_i| = O(\sqrt{\log(1/\varepsilon)} \cdot \varepsilon^{6/p} \|x\|_p)$.*

Proof. By Lemma 17, any individual $i \in [N]$ has $|z_i| \leq 4\sqrt{\log(1/\varepsilon)} \cdot (\sum_{w \in h_1^{-1}(i) \setminus L_{\varepsilon^8}} |x_w|^2)^{1/2}$ with probability at least $1 - 1/N^4$. We then apply a union bound and use the fact that $\ell_p \leq \ell_2$ for $p < 2$, so that $|z_i| \leq 4\sqrt{\log(1/\varepsilon)} \cdot (\sum_{w \in h_1^{-1}(i) \setminus L_{\varepsilon^8}} |x_w|^p)^{1/p}$ (call this event \mathcal{E}) with probability $1 - \text{poly}(\varepsilon)$.

We now prove our lemma, i.e. we show that with probability $1 - \text{poly}(\varepsilon)$, $|z_i|^p = O(\log^{p/2} \varepsilon^6 \|x\|_p^p)$ simultaneously for all $i \in [N]$. We apply Lemma 2. Specifically, fix an $i \in [N]$. For all j with $|x_j|^p \leq \varepsilon^8 \|x\|_p^p$, let $X_j = |x_j|^p \cdot \mathbf{1}_{h_1(j)=i}$. Then, in the notation of Lemma 2, $\mu_j = |x_j|^p/N$, and $\sigma_j^2 \leq |x_j|^{2p}/N$, and thus $\mu = \|x\|_p^p/N$ and $\sigma^2 \leq \|x\|_{2p}^{2p}/N \leq \varepsilon^8 \|x\|_p^p/N$. Also, $K = \varepsilon^8 \|x\|_p^p$. Then if h_1 were ℓ -wise independent for $\ell = 10$, Lemma 2 would give

$$\Pr \left[\left| \sum_i X_i - \|x\|_p^p/N \right| > \varepsilon^6 \|x\|_p^p \right] < 2^{O(\ell)} \cdot (\varepsilon^{7\ell} + \varepsilon^{2\ell}) = O(\varepsilon/N).$$

A union bound would then give that with probability $1 - \varepsilon$, the F_p mass in any bucket from items i with $|x_i|^p \leq \varepsilon^8 \|x\|_p^p$ is at most $\varepsilon^6 \|x\|_p^p$. Thus by a union bound with event \mathcal{E} , $|z_i|^p = O(\log^{p/2} \varepsilon^6 \|x\|_p^p)$ for all $i \in [N]$ with probability $1 - \text{poly}(\varepsilon)$.

Though, h_1 is not 10-wise independent. Instead, it is selected as in Theorem 14. However, for any constant ℓ , by increasing the constant c_h in our definition of h_1 we can ensure that our ℓ th moment bound for $(\sum_i X_i - \mu)$ is preserved to within a constant factor, which is sufficient to apply Lemma 2. \blacksquare

Lemma 20. *With probability $1 - \text{poly}(\varepsilon)$, for all $w \in L$ we have $|y_{h_1(w)}|^p = (1 \pm O(\varepsilon))|x_w|^p$, and thus with probability $1 - \text{poly}(\varepsilon)$ when conditioned on $\|y\|_p^p \leq 17C\|x\|_p^p$, we have that if w is an α -heavy hitter for x , then $h_1(w)$ is an $\alpha/(34C)$ -heavy hitter for y .*

Proof. Let w be in L . We know from Lemma 19 that $|z_{h_1(w)}| \leq 2\sqrt{\log(1/\varepsilon)}\varepsilon^{6/p}\|x\|_p$ with probability $1 - \text{poly}(\varepsilon)$, and that the elements of L are perfectly hashed under h_1 with probability $1 - \text{poly}(\varepsilon)$. Conditioned on this perfect hashing, we have that $|y_{h_1(w)}| \geq |x_w| - 2\varepsilon^{6/p}\sqrt{\log(1/\varepsilon)}\|x\|_p$. Since for $w \in L$ we have $|x_w| \geq \varepsilon^{2/p}\|x\|_p$, and since $p < 2$, we have $|y_{h_1(w)}| \geq (1 - O(\varepsilon))|x_w|$.

For the second part of the lemma, $(1 - O(\varepsilon))|x_w| > |x_w|/2^{1/p}$ for ε_0 sufficiently small. Thus if w is an α -heavy hitter for x , then $h_1(w)$ is an $\alpha/(34C)$ -heavy hitter for y . \blacksquare

Finally, the following lemma follows from a Markov bound followed by a union bound.

Lemma 21. *For $w \in [n]$ consider the quantity $s_w = \sum_{\substack{v \neq w \\ h(v)=h(w)}} |x_v|^p$. Then, with probability at least $1 - O(\varepsilon)$, $s_w \leq \varepsilon^{15}\|x\|_p^p$ simultaneously for all $w \in L$.*

We now put everything together. We set $\epsilon = \varepsilon \log(1/\varepsilon)$. As stated earlier, we define L' to be the sublist of those w output by our $F_p\text{HH}$ instantiation with $\phi = \epsilon^2$ such that $|\tilde{y}_w|^p \geq (2\epsilon^2/7)\tilde{F}_p$. We interpret updates to x as updates to y to then be fed into HighEnd , with $\alpha = \epsilon^2/(34C)$. Thus both HighEnd and $F_p\text{HH}$ require $O(\varepsilon^{-2} \log(nmM/\varepsilon))$ space. We now define some events.

Event \mathcal{A} . L_{ϵ^8} is perfectly hashed under h_1 , and $\forall i \in [N], |z_i|^p = O(\log(1/\epsilon)^{p/2} \cdot \epsilon^6 \|x\|_p^p)$.

Event \mathcal{B} . $\forall w \in L_{\epsilon^2}$, $h_1(w)$ is output as an $\epsilon^2/(34C)$ -heavy hitter by $F_p\text{HH}$.

Event \mathcal{C} . $\forall w \in L_{\epsilon^2/18}$, $|y_{h_1(w)}| = (1 \pm O(\epsilon))|x_w|$.

Event \mathcal{D} . $\tilde{F}_p \in [(1/2) \cdot \|x\|_p^p, (3/2) \cdot \|x\|_p^p]$, and **HighEnd**, **LightEstimator**, and $F_p\text{HH}$ succeed.

Now, suppose \mathcal{A} , \mathcal{B} , \mathcal{C} , and \mathcal{D} all occur. Then for $w \in L_{\epsilon^2}$, w is output by $F_p\text{HH}$, and furthermore $|y_{h_1(w)}|^p \geq (1 - O(\epsilon))|x_w|^p \geq |x_w|^p/2 \geq \epsilon^2 \|x\|_p^p/2$. Also, $\tilde{y}_{h_1(w)}^p \geq (6/7) \cdot |y_{h_1(w)}|^p$. Since $\tilde{F}_p \leq 3\|x\|_p^p/2$, we have that $h_1(w) \in L'$. Furthermore, we also know that for i output by $F_p\text{HH}$, $\tilde{y}_i^p \leq (9/7) \cdot |y_i|^p$, and thus $i \in L'$ implies $|y_i|^p \geq (\epsilon^2/9) \cdot \|x\|_p^p$. Notice that by event \mathcal{A} , each y_i is z_i , plus potentially $x_{w(i)}$ for some $x_{w(i)} \in L_{\epsilon^8}$. If $|y_i|^p \geq (\epsilon^2/9) \cdot \|x\|_p^p$, then there must exist such a $w(i)$, and furthermore it must be that $|x_{w(i)}|^p \geq (\epsilon^2/18) \cdot \|x\|_p^p$. Thus, overall, L' contains $h_1(w)$ for all $w \in L_{\epsilon^2}$, and furthermore if $i \in L'$ then $w(i) \in L_{\epsilon^2/18}$.

Since L' contains $h_1(L_{\epsilon^2})$, **LightEstimator** outputs $\|x_{[n] \setminus h^{-1}(L')}\|_p^p \pm O(\epsilon \|x\|_p^p)$. Also, **HighEnd** outputs $\|y_{L'}\| \pm O(\epsilon) \cdot \|y\|_p^p$. Now we analyze correctness. We have $\Pr[\mathcal{A}] = 1 - \text{poly}(\epsilon)$, $\Pr[\mathcal{B} \mid \|y\|_p^p \leq 17C\|x\|_p^p] = 1 - \text{poly}(\epsilon)$, $\Pr[\mathcal{C}] = 1 - \text{poly}(\epsilon)$, and $\Pr[\mathcal{D}] \geq 5/8$. We also have $\Pr[\|y\|_p^p \leq 17C\|x\|_p^p] \geq 1 - 2/C$. Thus by a union bound and setting C sufficiently large, we have $\Pr[\mathcal{A} \wedge \mathcal{B} \wedge \mathcal{C} \wedge \mathcal{D} \wedge (\|y\|_p^p \leq 17C\|x\|_p^p)] \geq 9/16$. Define L_{inv} to be the set $\{w(i)\}_{i \in L'}$, i.e. the heavy hitters of x corresponding to the heavy hitters in L' for y . Now, if all these events occur, then $\|x_{[n] \setminus h^{-1}(L')}\|_p^p = \|x_{[n] \setminus L_{\text{inv}}}\|_p^p \pm O(\epsilon^{15})\|x\|_p^p$ with probability $1 - O(\epsilon)$ by Lemma 21. We also have, since \mathcal{C} occurs and conditioned on $\|y\|_p^p = O(\|x\|_p^p)$, that $\|y_{L'}\| \pm O(\epsilon) \cdot \|y\|_p^p = \|x_{L_{\text{inv}}}\|_p^p \pm O(\epsilon) \cdot \|x\|_p^p$. Thus, overall, our algorithm outputs $\|x\|_p^p \pm O(\epsilon) \cdot \|x\|_p^p$ with probability $17/32 > 1/2$ as desired. Notice this probability can be amplified to $1 - \delta$ by outputting the median of $O(\log(1/\delta))$ independent instantiations.

We further note that for a single instantiation of **LightEstimator**, we have $\mathbf{E}_h[\mathbf{Var}_s[\Phi']] = O(\epsilon^2 \|x\|_p^{2p})$. Once h is fixed, the variance of Φ' is simply the sum of variances across the D_j for $j \notin h_1(L')$. Thus, it suffices for the D_j to use pairwise independent randomness. Furthermore, in repeating $O((\epsilon/\epsilon)^2)$ parallel repetitions of **LightEstimator**, it suffices that all the D_j across all parallel repetitions use pairwise independent randomness, and the hash function h can remain the same. Thus, as discussed in Remark 16, the coefficients of the degree- $O(1/\epsilon^p)$ polynomials used in all D_j combined can be generated by just two coefficient vectors, and thus the update time of **LightEstimator** with $O((\epsilon/\epsilon)^2)$ parallel repetitions is just $O((\epsilon/\epsilon)^2 + O(\log^2(1/\epsilon) \log \log(1/\epsilon))) = O(\log^2(1/\epsilon) \log \log(1/\epsilon))$. Thus overall, we have the following theorem.

Theorem 22. *There exists an algorithm such that given $0 < p < 2$ and $0 < \epsilon < 1/2$, the algorithm outputs $(1 \pm \epsilon)\|x\|_p^p$ with probability $2/3$ using $O(\epsilon^{-2} \log(nmM/\epsilon))$ space. The update time is $O(\log^2(1/\epsilon) \log \log(1/\epsilon))$. The reporting time is $O(\epsilon^{-2} \log^2(1/\epsilon) \log \log(1/\epsilon))$.*

The space bound above can be assumed $O(\epsilon^{-2} \log(mM) + \log \log n)$ by comments in Section 1.3.

References

- [1] Charu C. Aggarwal, Alexander Hinneburg, and Daniel A. Keim. On the surprising behavior of distance metrics in high dimensional spaces. In *ICDT*, pages 420–434, 2001.
- [2] Noga Alon, Phillip B. Gibbons, Yossi Matias, and Mario Szegedy. Tracking join and self-join sizes in limited storage. *J. Comput. Syst. Sci.*, 64(3):719–747, 2002.

- [3] Noga Alon, Yossi Matias, and Mario Szegedy. The Space Complexity of Approximating the Frequency Moments. *J. Comput. Syst. Sci.*, 58(1):137–147, 1999.
- [4] Ziv Bar-Yossef. *The Complexity of Massive Data Set Computations*. PhD thesis, University of California at Berkeley, 2002.
- [5] 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.
- [6] Lakshminath Bhuvanagiri, Sumit Ganguly, Deepanjan Kesh, and Chandan Saha. Simpler algorithm for estimating frequency moments of data streams. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 708–713, 2006.
- [7] Amit Chakrabarti, Subhash Khot, and Xiaodong Sun. Near-optimal lower bounds on the multi-party communication complexity of set disjointness. In *Proceedings of the 18th Annual IEEE Conference on Computational Complexity (CCC)*, pages 107–117, 2003.
- [8] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. In *Proceedings of the 29th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 693–703, 2002.
- [9] Graham Cormode, Piotr Indyk, Nick Koudas, and S. Muthukrishnan. Fast mining of massive tabular data via approximate distance computations. In *ICDE*, pages 605–, 2002.
- [10] Graham Cormode and S. Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *J. Algorithms*, 55(1):58–75, 2005.
- [11] Graham Cormode, S. Muthukrishnan, and Irina Rozenbaum. Summarizing and mining inverse distributions on data streams via dynamic inverse sampling. In *VLDB*, pages 25–36, 2005.
- [12] Ilias Diakonikolas, Daniel M. Kane, and Jelani Nelson. Bounded independence fools degree-2 threshold functions. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, to appear (see also *CoRR abs/0911.3389*), 2010.
- [13] Joan Feigenbaum, Sampath Kannan, Martin Strauss, and Mahesh Viswanathan. An approximate L1-difference algorithm for massive data streams. *SIAM J. Comput.*, 32(1):131–151, 2002.
- [14] Michael L. Fredman, János Komlós, and Endre Szemerédi. Storing a sparse table with $O(1)$ worst case access time. *J. ACM*, 31(3):538–544, 1984.
- [15] Sumit Ganguly. Estimating frequency moments of data streams using random linear combinations. In *Proceedings of the 8th International Workshop on Randomization and Computation (RANDOM)*, pages 369–380, 2004.
- [16] Sumit Ganguly, Abhayendra N. Singh, and Satyam Shankar. Finding frequent items over general update streams. In *Proceedings of the 20th International Conference on Scientific and Statistical Database Management (SSDBM)*, pages 204–221, 2008.

- [17] Davi Geiger, Tyng-Luh Liu, and Michael J. Donahue. Sparse representations for image decompositions. *International Journal of Computer Vision*, 33(2):139–156, 1999.
- [18] Andre Gronemeier. Asymptotically optimal lower bounds on the NIH-multi-party information complexity of the AND-function and disjointness. In *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 505–516, 2009.
- [19] Uffe Haagerup. The best constants in the Khintchine inequality. *Studia Math.*, 70(3):231–283, 1982.
- [20] Nicholas J. A. Harvey, Jelani Nelson, and Krzysztof Onak. Sketching and streaming entropy via approximation theory. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 489–498, 2008.
- [21] Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *J. ACM*, 53(3):307–323, 2006.
- [22] 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.
- [23] T. S. Jayram. Hellinger strikes back: A note on the multi-party information complexity of AND. In *APPROX-RANDOM*, pages 562–573, 2009.
- [24] T. S. Jayram, Ravi Kumar, and D. Sivakumar. The one-way communication complexity of hamming distance. *Theory of Computing*, 4(1):129–135, 2008.
- [25] Daniel M. Kane and Jelani Nelson. A derandomized sparse Johnson-Lindenstrauss transform. *CoRR*, abs/1006.3585, 2010.
- [26] Daniel M. Kane, Jelani Nelson, and David P. Woodruff. On the exact space complexity of sketching and streaming small norms. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1161–1178, 2010.
- [27] Balachander Krishnamurthy, Subhabrata Sen, Yin Zhang, and Yan Chen. Sketch-based change detection: methods, evaluation, and applications. In *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement (IMC)*, pages 234–247, 2003.
- [28] Ping Li. Estimators and tail bounds for dimension reduction in l_p ($0 < p \leq 2$) using stable random projections. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 10–19, 2008.
- [29] Morteza Monemizadeh and David P. Woodruff. 1-pass relative-error l_p sampling with applications. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1143–1160, 2010.
- [30] Jelani Nelson and David Woodruff. Fast Manhattan sketches in data streams. In *Proceedings of the 29th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, pages 99–110, 2010.

- [31] Jelani Nelson and David P. Woodruff. A near-optimal algorithm for L1-difference. *CoRR*, abs/0904.2027, 2009.
- [32] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.
- [33] John P. Nolan. *Stable Distributions - Models for Heavy Tailed Data*. Birkhäuser, 2010. In progress, Chapter 1 online at <http://academic2.american.edu/~jpnolan>.
- [34] Open Problems in Data Streams and Related Topics. IITK Workshop on Algorithms for Data Streams, 2006. <http://www.cse.iitk.ac.in/users/sganguly/data-stream-probs.pdf>.
- [35] Anna Pagh and Rasmus Pagh. Uniform hashing in constant time and linear space. *SIAM J. Comput.*, 38(1):85–96, 2008.
- [36] Michael E. Saks and Xiaodong Sun. Space lower bounds for distance approximation in the data stream model. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 360–369, 2002.
- [37] Mikkel Thorup and Yin Zhang. Tabulation based 4-universal hashing with applications to second moment estimation. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 615–624, 2004.
- [38] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.
- [39] M.W. Wong. *Complex Analysis*, volume 2. World Scientific, 2008.
- [40] David P. Woodruff. Optimal space lower bounds for all frequency moments. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 167–175, 2004.
- [41] David P. Woodruff. *Efficient and Private Distance Approximation in the Communication and Streaming Models*. PhD thesis, Massachusetts Institute of Technology, 2007.
- [42] Haiquan (Chuck) Zhao, Ashwin Lall, Mitsunori Ogihara, Oliver Spatscheck, Jia Wang, and Jun Xu. A data streaming algorithm for estimating entropies of od flows. In *Internet Measurement Conference*, pages 279–290, 2007.
- [43] Vladimir Mikhailovich Zolotarev. *One-dimensional Stable Distributions*. Vol. 65 of Translations of Mathematical Monographs, American Mathematical Society, 1986.

A Appendix

A.1 A heavy hitter algorithm for F_p

Note that F_p Report, F_p Update, and F_p Space below can be as in the statement in Section 2 by using the algorithm of [26].

Theorem 1 (restatement). *There is an algorithm F_p HH satisfying the following properties. Given $0 < \phi, \delta < 1/2$ and black-box access to an F_p -estimation algorithm F_p Est(ε', δ') with $\varepsilon' = 1/7$ and $\delta' = \phi\delta/(12(\log(\phi n) + 1))$, F_p HH produces a list L such that L contains all ϕ -heavy hitters and does not contain indices which are not $\phi/2$ -heavy hitters with probability at least $1 - \delta$. For each $i \in L$, the algorithm also outputs $\text{sign}(x_i)$, as well as an estimate \tilde{x}_i of x_i satisfying $\tilde{x}_i^p \in [(6/7)|x_i|^p, (9/7)|x_i|^p]$. Its space usage is $O(\phi^{-1} \log(\phi n) \cdot F_p\text{Space}(\varepsilon', \delta') + \phi^{-1} \log(1/(\delta\phi)) \log(nmM))$. Its update time is $O(\log(\phi n) \cdot F_p\text{Update}(\varepsilon', \delta') + \log(1/(\delta\phi)))$. Its reporting time is $O(\phi^{-1}(\log(\phi n) \cdot F_p\text{Report}(\varepsilon', \delta') + \log(1/(\delta\phi))))$. Here, $F_p\text{Report}(\varepsilon', \delta')$, $F_p\text{Update}(\varepsilon', \delta')$, and $F_p\text{Space}(\varepsilon', \delta')$ are the reporting time, update time, and space consumption of F_p Est when a $(1 \pm \varepsilon')$ -approximation to F_p is desired with probability at least $1 - \delta'$.*

Proof. First we argue with $\delta' = \phi\delta/(12(\log n + 1))$. We assume without loss of generality that n is a power of 2. Consider the following data structure $\text{Basic}F_p\text{HH}(\phi', \delta, \varepsilon', k)$, where $k \in \{0, \dots, \log n\}$. We set $R = \lceil 1/\phi' \rceil$ and pick a function $h : \{0, \dots, 2^k - 1\} \rightarrow [R]$ at random from a pairwise independent hash family. We also create instantiations D_1, \dots, D_R of $F_p\text{Est}(\varepsilon', 1/5)$. This entire structure is then repeated independently in parallel $T = \Theta(\log(1/\delta))$ times, so that we have hash functions h_1, \dots, h_T , and instantiations D_i^j of $F_p\text{Est}$ for $i, j \in [R] \times [T]$. For an integer x in $[n]$, let $\text{prefix}(x, k)$ denote the length- k prefix of $x - 1$ when written in binary, treated as an integer in $\{0, \dots, 2^k - 1\}$. Upon receiving an update (i, v) in the stream, we feed this update to $D_{h_j(\text{prefix}(i, k))}^j$ for each $j \in [T]$.

For $t \in \{0, \dots, 2^k - 1\}$, let $F_p(t)$ denote the F_p value of the vector x restricted to indices $i \in [n]$ with $\text{prefix}(i) = t$. Consider the procedure $\text{Query}(t)$ which outputs the median of F_p -estimates given by $D_{h_j(t)}^j$ over all $j \in [T]$. We now argue that the output of $\text{Query}(t)$ is in the interval $[(1 - \varepsilon') \cdot F_p(t), (1 + \varepsilon') \cdot (F_p(t) + 5\phi' \|x\|_p^p)]$, i.e. $\text{Query}(t)$ “succeeds”, with probability at least $1 - \delta$.

For any $j \in [T]$, consider the actual F_p value $F_p(t)^j$ of the vector x restricted to coordinates i such that $h_j(\text{prefix}(i, k)) = h_j(t)$. Then $F_p(t)^j = F_p(t) + R(t)^j$, where $R(t)^j$ is the F_p contribution of the i with $\text{prefix}(i, k) \neq t$, yet $h_j(\text{prefix}(i, k)) = h_j(t)$. We have $R(t)^j \geq 0$ always, and furthermore $\mathbf{E}[R(t)^j] \leq \|x\|_p^p / R$ by pairwise independence of h_j . Thus by Markov’s inequality, $\mathbf{Pr}[R(t)^j > 5\phi' \|x\|_p^p] < 1/5$. Note for any fixed $j \in [T]$, the F_p -estimate output by $D_{h_j(t)}^j$ is in $[(1 - \varepsilon') \cdot F_p(t), (1 + \varepsilon') \cdot (F_p(t) + 5\phi' \|x\|_p^p)]$ as long as both the events “ $D_{h_j(t)}^j$ successfully gives a $(1 \pm \varepsilon')$ -approximation” and “ $R(t)^j \leq 5\phi' \|x\|_p^p$ ” occur. This happens with probability at least $3/5$. Thus, by a Chernoff bound, the output of $\text{Query}(t)$ is in the desired interval with probability at least $1 - \delta$.

We now define the final F_p HH data structure. We maintain one global instantiation D of $F_p\text{Est}(1/7, \delta/2)$. We also use the dyadic interval idea for L_1 -heavy hitters given in [10]. Specifically, we imagine building a binary tree \mathcal{T} over the universe $[n]$ (without loss of generality assume n is a power of 2). The number of levels in the tree is $\ell = 1 + \log n$, where the root is at level 0 and the leaves are at level $\log n$. For each level $j \in \{0, \dots, \ell\}$, we maintain an instantiation B_j of $\text{Basic}F_p\text{HH}(\phi/80, \delta', 1/7, j)$ for δ' as in the theorem statement. When we receive an update (i, v) in the stream, we feed the update to D and also to each B_j .

We now describe how to answer a query to output the desired list L . We first query D to obtain \tilde{F}_p , an approximation to F_p . We next initiate an iterative procedure on our binary tree, beginning at the root, which proceeds level by level. The procedure is as follows. Initially, we set $L = \{0\}$, $L' = \emptyset$, and $j = 0$. For each $i \in L$, we perform $\text{Query}(i)$ on B_j then add $2i$ and $2i + 1$ to L' if the output of $\text{Query}(i)$ is at least $3\phi\tilde{F}_p/4$. After processing every $i \in L$, we then set $L \leftarrow L'$ then $L' \leftarrow \emptyset$, and we increment j . This continues until $j = 1 + \log n$, at which point we halt and return L . We now show why the list L output by this procedure satisfies the claim in the theorem statement. We condition on the event \mathcal{E} that $\tilde{F}_p = (1 \pm 1/7)F_p$, and also on the event \mathcal{E}' that every query made throughout the recursive procedure is successful. Let i be such that $|x_i|^p \geq \phi F_p$. Then, since $F_p(\text{prefix}(i, j)) \geq |x_i|^p$ for any j , we always have that $\text{prefix}(i, j) \in L$ at the end of the j th round of our iterative procedure, since $(6/7)|x_i|^p \geq (3/4)\phi\tilde{F}_p$ given \mathcal{E} . Now, consider an i such that $|x_i|^p < (\phi/2)F_p$. Then, $(8/7) \cdot (|x_i|^p - 5 \cdot (\phi/80)) < 3\phi\tilde{F}_p/4$, implying i is not included in the final output list. Also, note that since the query at the leaf corresponding to $i \in L$ is successful, then by definition of a successful query, we are given an estimate \tilde{x}_i^p of $|x_i|^p$ by the corresponding BasicF_pHH structure satisfying $\tilde{x}_i^p \in [(6/7)|x_i|^p, (8/7)|x_i|^p + (\phi/16)F_p]$, which is $[(6/7)|x_i|^p, (9/7)|x_i|^p]$ since $|x_i|^p \geq (\phi/2)F_p$.

We now only need to argue that \mathcal{E} and \mathcal{E}' occur simultaneously with large probability. We have $\Pr[\mathcal{E}] \geq 1 - \delta/2$. For \mathcal{E}' , note there are at most 2ϕ $\phi/2$ -heavy hitters at any level of the tree, where at level j we are referring to heavy hitters of the 2^j -dimensional vector y_j satisfying $(y_j)_i^p = \sum \text{prefix}(t, j) = i|x_t|^p$. As long as the $\text{Query}(\cdot)$ calls made for all $\phi/2$ -heavy hitters and their two children throughout the tree succeed (including at the root), \mathcal{E}' holds. Thus, $\Pr[\mathcal{E}'] \geq 1 - \delta' \cdot 6(\log n + 1)\phi^{-1} = 1 - \delta/2$. Therefore, by a union bound $\Pr[\mathcal{E} \wedge \mathcal{E}'] \geq 1 - \delta$.

Finally, notice that the number of levels in F_pHH can be reduced from $\log n$ to $\log n - \log \lceil 1/\phi \rceil = O(\log(\phi n))$ by simply ignoring the top $\log \lceil 1/\phi \rceil$ levels of the tree. Then, in the topmost level of the tree which we maintain, the universe size is $O(1/\phi)$, so we can begin our reporting procedure by querying all these universe items to determine which subtrees to recurse upon.

To recover $\text{sign}(x_w)$ for each $w \in L$, we use the CountSketch data structure of [8] with $T = (21 \cdot 2^p)/\phi$ columns and $C = \Theta(\log(1/(\delta\phi)))$ rows; the space is $O(\phi^{-1} \log(1/(\delta\phi)) \log(nmM))$, and the update time is $O(\log(1/(\delta\phi)))$. CountSketch operates by, for each row i , having a pairwise independent hash function $h_i : [n] \rightarrow [T]$ and a 4-wise independent hash function $\sigma_i : [n] \rightarrow \{-1, 1\}$. There are $C \cdot T$ counters $A_{i,j}$ for $(i, j) \in [C] \times [T]$. Counter $A_{i,j}$ maintains $\sum_{h_i(v)=j} \sigma_i(v) \cdot x_v$. For $(i, j) \in [C] \times [T]$, let x^i be the vector x restricted to coordinates v with $h_i(v) = h_i(w)$, other than w itself. Then for fixed i , the expected contribution to $\|x^i\|_p^p$ is at most $\|x\|_p^p/T$, and thus is at most $10\|x\|_p^p/T$ with probability 9/10 by Markov's inequality. Conditioned on this event, $|x_w| > \|x^i\|_p/2 \geq \|x^i\|_2/2$. The analysis of CountSketch also guarantees $|A_{i,h_i(w)} - \sigma_i(w)x_w| \leq 2\|x^i\|_2$ with probability at least 2/3, and thus by a union bound, $|x_w| > |A_{i,h_i(w)} - \sigma_i(w)x_w|$ with probability at least 11/20, in which case $\sigma_i(w) \cdot \text{sign}(A_{i,h_i(w)}) = \text{sign}(x_w)$. Thus, by a Chernoff bound over all rows, together with a union bound over all $w \in L$, we can recover $\text{sign}(x_w)$ for all $w \in L$ with probability $1 - \delta$. \blacksquare

A.2 Proof of Theorem 12

In this section we prove Theorem 12. The data structure and estimator we give is a slightly modified version of the geometric mean estimator of Li [28]. Our modification allows us to show that only bounded independence is required amongst the p -stable random variables in our data structure.

Before giving our D_p and Est_p , we first define the p -stable distribution.

Definition 23 (Zolotarev [43]). *For $0 < p < 2$, there exists a probability distribution \mathcal{D}_p called the p -stable distribution satisfying the following property. For any positive integer n and vector $x \in \mathbb{R}^n$, if $Z_1, \dots, Z_n \sim \mathcal{D}_p$ are independent, then $\sum_{j=1}^n Z_j x_j \sim \|x\|_p Z$ for $Z \sim \mathcal{D}_p$.*

Li's *geometric mean estimator* is as follows. For some positive integer $t > 2$, select a matrix $A \in \mathbb{R}^{t \times n}$ with independent p -stable entries, and maintain $y = Ax$ in the stream. Given y , the estimate of $\|x\|_p^p$ is then $C_{t,p} \cdot (\prod_{j=1}^t |y_j|^{p/t})$ for some constant $C_{t,p}$. For Theorem 12, we make the following adjustments. First, we require $t > 4$. Next, for any fixed row of A we only require that the entries be $\Omega(1/\varepsilon^p)$ -wise independent, though the rows themselves we keep independent. Furthermore, in parallel we run the algorithm of [26] with constant error parameter to obtain a value \tilde{F}_p in $[\|x\|_p^p/2, 3\|x\|_p^p/2]$. The D_p data structure of Theorem 12 is then simply y , together with the state maintained by the algorithm of [26]. The estimator Est_p is $\min\{C_{t,p} \cdot (\prod_{j=1}^t |y_j|^{p/t}), \tilde{F}_p/\varepsilon\}$. To state the value $C_{t,p}$, we use the following theorem.

Theorem 24 ([43, Theorem 2.6.3]). *For $Q \sim \mathcal{D}_p$ and $-1 < \lambda < p$,*

$$\mathbf{E}[|Q|^\lambda] = \frac{2}{\pi} \Gamma\left(1 - \frac{\lambda}{p}\right) \Gamma(\lambda) \sin\left(\frac{\pi\lambda}{2}\right).$$

Theorem 24 implies that we should set

$$C_{t,p} = \left[\frac{2}{\pi} \cdot \Gamma\left(1 - \frac{1}{t}\right) \cdot \Gamma\left(\frac{p}{t}\right) \cdot \sin\left(\frac{\pi p}{2t}\right) \right]^{-t}.$$

To carry out our analysis, we will need the following theorem, which gives a way of producing a smooth approximation of the indicator function of an interval while maintaining good bounds on high order derivatives.

Theorem 25 ([12]). *For any interval $[a, b] \subseteq \mathbb{R}$ and integer $c > 0$, there exists a nonnegative function $\tilde{I}_{[a,b]}^c : \mathbb{R} \rightarrow \mathbb{R}$ satisfying the following properties:*

- i. $\|(\tilde{I}_{[a,b]}^c)^{(\ell)}\|_\infty \leq (2c)^\ell$ for all $\ell \geq 0$.
- ii. For any $x \in \mathbb{R}$, $|\tilde{I}_{[a,b]}^c(x) - I_{[a,b]}(x)| \leq \min\{1, 5/(2c^2 \cdot d(x, \{a, b\})^2)\}$.

We also need the following lemma of [26], which argues that smooth, bounded functions have their expectations approximately preserved when their input is a linear form evaluated at boundedly independent p -stable random variables, as opposed to completely independent p -stable random variables.

Lemma 26 ([26, Lemma 2.2]). *There exists an $\varepsilon_0 > 0$ such that the following holds. Let $0 < \varepsilon < \varepsilon_0$ and $0 < p < 2$ be given. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfy $\|f^{(\ell)}\|_\infty = O(\alpha^\ell)$ for all $\ell \geq 0$, for some α satisfying $\alpha^p \geq \log(1/\varepsilon)$. Let $k = \alpha^p$. Let $x \in \mathbb{R}^n$ satisfy $\|x\|_p = O(1)$. Let R_1, \dots, R_n be drawn from a $3Ck$ -wise independent family of p -stable random variables for $C > 0$ a sufficiently large constant, and let Q be the product of $\|x\|_p$ and a p -stable random variable. Then $|\mathbf{E}[f(R)] - \mathbf{E}[f(Q)]| = O(\varepsilon)$.*

We now prove a tail bound for linear forms over k -wise independent p -stable random variables. Note that for a random variable X whose moments are bounded, one has $\Pr[X - \mathbf{E}[X] > t] \leq \mathbf{E}[(X - \mathbf{E}[X])^k]/t^k$ by applying Markov's inequality to the random variable $(X - \mathbf{E}[X])^k$ for some even integer $k \geq 2$. Unfortunately, for $0 < p < 2$, it is known that even the second moment of \mathcal{D}_p is already infinite, so this method cannot be applied. We instead prove our tail bound via FT-mollification of $I_{[t, \infty)}$, since $\Pr[X \geq t] = \mathbf{E}[I_{[t, \infty)}(X)]$.

We will need to refer to the following lemma.

Lemma 27 (Nolan [33, Theorem 1.12]). *For fixed $0 < p < 2$, the probability density function φ_p of the p -stable distribution satisfies $\varphi_p(x) = O(1/(1 + |x|^{p+1}))$ and is an even function. The cumulative distribution function satisfies $\Phi_p(x) = O(|x|^{-p})$.*

We now prove our tail bound.

Lemma 28. *Suppose $x \in \mathbb{R}^n$, $\|x\|_p = 1$, $0 < \varepsilon < 1$ is given, and R_1, \dots, R_n are k -wise independent p -stable random variables for $k \geq 2$. Let $Q \sim \mathcal{D}_p$. Then for all $t \geq 0$, $R = \sum_{i=1}^n R_i x_i$ satisfies*

$$|\Pr[|Q| \geq t] - \Pr[|R| \geq t]| = O(k^{-1/p}/(1 + t^{p+1}) + k^{-2/p}/(1 + t^2) + 2^{-\Omega(k)}).$$

Proof. We have $\Pr[|Z| \geq t] = \mathbf{E}[I_{[t, \infty)}(Z)] + \mathbf{E}[I_{(-\infty, -t]}(Z)]$ for any random variable Z , and thus we will argue $|\mathbf{E}[I_{[t, \infty)}(Q)] - \mathbf{E}[I_{[t, \infty)}(R)]| = O(k^{-1/p}/(1 + t^{p+1}) + k^{-2/p}/(1 + t^2) + 2^{-\Omega(k)})$; a similar argument shows the same bound for $|\mathbf{E}[I_{(-\infty, -t]}(Q)] - \mathbf{E}[I_{(-\infty, -t]}(R)]|$.

We argue the following chain of inequalities for $c = k^{1/p}/(3C)$, for C the constant in Lemma 26, and we define $\gamma = k^{-1/p}/(1 + t^{p+1}) + k^{-2/p}/(1 + t^2)$:

$$\mathbf{E}[I_{[t, \infty)}(Q)] \approx_\gamma \mathbf{E}[\tilde{I}_{[t, \infty)}^c(Q)] \approx_{2^{-c^p}} \mathbf{E}[\tilde{I}_{[t, \infty)}^c(R)] \approx_{\gamma+2^{-c^p}} \mathbf{E}[I_{[t, \infty)}(R)].$$

$\mathbf{E}[\tilde{I}_{[t, \infty)}(Q)] \approx_\gamma \mathbf{E}[\tilde{\mathbf{I}}_{[t, \infty)}^c(Q)$: Assume $t \geq 1$. We have

$$\begin{aligned} |\mathbf{E}[I_{[t, \infty)}(Q)] - \mathbf{E}[\tilde{I}_{[t, \infty)}^c(Q)]| &\leq \mathbf{E}[|I_{[t, \infty)}(Q) - \tilde{I}_{[t, \infty)}^c(Q)|] \\ &\leq \Pr[|Q - t| \leq 1/c] + \left(\sum_{s=1}^{\log(ct)-1} \Pr[|Q - t| \leq 2^s/c] \cdot O(2^{-2s}) \right) \quad (4) \\ &\quad + \Pr[|Q - t| > t/2] \cdot O(c^{-2}t^{-2}) \\ &= O(1/(c \cdot t^{p+1})) + O(c^{-2}t^{-2}) \end{aligned}$$

since $\Pr[|Q - t| \leq 2^s/c]$ is $O(2^s/(c \cdot t^{p+1}))$ as long as $2^s/c \leq t/2$.

In the case $0 < t < 1$, we repeat the same argument as above but replace Eq. (4) with a summation from $s = 1$ to ∞ , and also remove the additive $\Pr[|Q - t| > t/2] \cdot O(c^{-2}t^{-2})$ term. Doing so gives an overall upper bound of $O(1/c)$ in this case.

$\mathbf{E}[\tilde{\mathbf{I}}_{[t, \infty)}^c(Q)] \approx_{2^{-c^p}} \mathbf{E}[\tilde{\mathbf{I}}_{[t, \infty)}^c(\mathbf{R})$: This follows from Lemma 26 with $\varepsilon = 2^{-c^p}$ and $\alpha = c$.

$\mathbf{E}[\tilde{\mathbf{I}}_{[t, \infty)}^c(\mathbf{R})] \approx_{\gamma+2^{-c^p}} \mathbf{E}[\mathbf{I}_{[t, \infty)}(\mathbf{R})$: We would like to apply the same argument as when showing $\mathbf{E}[\tilde{I}_{[t, \infty)}^c(Q)] \approx_\gamma \mathbf{E}[I_{[t, \infty)}(Q)]$ above. The trouble is, we must bound $\Pr[|R - t| > t/2]$ and $\Pr[|R - t| \leq 2^s/c]$ given that the R_i are only k -wise independent. For the first probability, we above only used that $\Pr[|Q - t| > t/2] \leq 1$, which still holds with Q replaced by R .

For the second probability, observe $\Pr[|R - t| \leq 2^s/c] = \mathbf{E}[I_{[t-2^s/c, t+2^s/c]}(R)]$. Define $\delta = 2^s/c + b/c$ for a sufficiently large constant $b > 0$ to be determined later. Then, arguing as above, we have $\mathbf{E}[\tilde{I}_{[t-\delta, t+\delta]}^c(R)] \approx_{2^{-c^p}} \mathbf{E}[\tilde{I}_{[t-\delta, t+\delta]}^c(Q)] \approx_\gamma \mathbf{E}[I_{[t-\delta, t+\delta]}(Q)]$, and we also know $\mathbf{E}[I_{[t-\delta, t+\delta]}(Q)] = O(\mathbf{E}[I_{[t-2^s/c, t+2^s/c]}(Q)]) = O(\Pr[|Q - t| \leq 2^s/c]) = O(2^s/(c \cdot t^{p+1}))$. Now, for $x \notin [t - 2^s/c, t + 2^s/c]$, $I_{[t-2^s/c, t+2^s/c]}(x) = 0$ while $I_{[t-\delta, t+\delta]}(x) = 1$. For $x \in [t - 2^s/c, t + 2^s/c]$, the distance from x to $\{t - \delta, t + \delta\}$ is at least b/c , implying $\tilde{I}_{[t-\delta, t+\delta]}^c(x) \geq 1/2$ for b sufficiently large by item (ii) of Lemma 25. Thus, $2 \cdot \tilde{I}_{[t-\delta, t+\delta]}^c \geq I_{[t-2^s/c, t+2^s/c]}$ on \mathbb{R} , and thus in particular, $\mathbf{E}[I_{[t-2^s/c, t+2^s/c]}(R)] \leq 2 \cdot \mathbf{E}[\tilde{I}_{[t-\delta, t+\delta]}^c(R)]$. Thus, in summary, $\mathbf{E}[I_{[t-2^s/c, t+2^s/c]}(R)] = O(2^s/(c \cdot t^{p+1}) + \gamma + 2^{-c^p})$. ■

We now prove the main lemma of this section, which implies Theorem 12.

Lemma 29. *Let $x \in \mathbb{R}^n$ be such that $\|x\|_p = 1$, and suppose $0 < \varepsilon < 1/2$. Let $0 < p < 2$, and let t be any constant greater than $4/p$. Let R_1, \dots, R_n be k -wise independent p -stable random variables for $k = \Omega(1/\varepsilon^p)$, and let Q be a p -stable random variable. Define $f(x) = \min\{|x|^{1/t}, T\}$, for $T = 1/\varepsilon$. Then, $|\mathbf{E}[f(R)] - \mathbf{E}[|Q|^{1/t}]| = O(\varepsilon)$ and $\mathbf{E}[f^2(R)] = O(\mathbf{E}[|Q|^{2/t}])$.*

Proof. We first argue $|\mathbf{E}[f(R)] - \mathbf{E}[|Q|^{1/t}]| = O(\varepsilon)$. We argue through the chain of inequalities

$$\mathbf{E}[|Q|^{1/t}] \approx_\varepsilon \mathbf{E}[f(Q)] \approx_\varepsilon \mathbf{E}[f(R)].$$

$\mathbf{E}[|Q|^{1/t}] \approx_\varepsilon \mathbf{E}[f(Q)]$: We have

$$\begin{aligned} |\mathbf{E}[|Q|^{1/t}] - \mathbf{E}[f(Q)]| &= 2 \int_{T^t}^{\infty} (x^{1/t} - T) \cdot \varphi_p(x) dx \\ &= \int_{T^t}^{\infty} (x^{1/t} - T) \cdot O(1/x^{p+1}) dx \\ &= O\left(T^{1-tp} \cdot \left(\frac{t}{pt-1} + \frac{1}{p}\right)\right) \\ &= O(1/(Tp)) \\ &= O(\varepsilon) \end{aligned}$$

$\mathbf{E}[f(Q)] \approx_\varepsilon \mathbf{E}[f(\mathbf{R})]$: Let φ_p^+ be the probability density function corresponding to the distribution of $|Q|$, and let Φ_p^+ be its cumulative distribution function. Then, by integration by parts and

Lemma 28,

$$\begin{aligned}
\mathbf{E}[f(Q)] &= \int_0^{T^t} x^{1/t} \varphi_p^+(x) dx + T \cdot \int_{T^t}^{\infty} \varphi_p^+(x) dx \\
&= -[x^{1/t} \cdot (1 - \Phi_p^+(x))]_0^{T^t} - T \cdot [(1 - \Phi_p^+(x))]_{T^t}^{\infty} + \frac{1}{t} \int_0^{T^t} \frac{1}{x^{1-1/t}} (1 - \Phi_p^+(x)) dx \\
&= \frac{1}{t} \int_0^{T^t} \frac{1}{x^{1-1/t}} \cdot \mathbf{Pr}[|Q| \geq x] dx \\
&= \frac{1}{t} \int_0^{T^t} \frac{1}{x^{1-1/t}} \cdot (\mathbf{Pr}[|R| \geq x] + O(k^{-1/p} 1/(1+x^{p+1}) + k^{-2/p}/(1+x^2) + 2^{-\Omega(k)})) dx \\
&= \mathbf{E}[f(R)] + \int_0^1 x^{1/t-1} \cdot O(k^{-1/p} + k^{-2/p} + 2^{-\Omega(k)}) dx \\
&\quad + \int_1^{T^t} x^{1/t-1} \cdot O(k^{-1/p}/x^{p+1} + k^{-2/p}/x^2 + 2^{-\Omega(k)}) dx \tag{5} \\
&= \mathbf{E}[f(R)] + O(\varepsilon) + O\left(\frac{1}{k^{1/p}} \cdot \left(\frac{1}{T^{tp+t-1}} - 1\right) \cdot \frac{1}{\frac{1}{t} - p - 1}\right) \\
&\quad + O\left(\frac{1}{k^{2/p}} \cdot \left(\frac{1}{T^{2t-1}} - 1\right) \cdot \frac{1}{\frac{1}{t} - 2}\right) + O(2^{-\Omega(k)} \cdot (T-1)) \\
&= \mathbf{E}[f(R)] + O(\varepsilon)
\end{aligned}$$

We show $\mathbf{E}[f^2(R)] = O(|Q|^{2/t})$ similarly. Namely, we argue through the chain of inequalities

$$\mathbf{E}[|Q|^{2/t}] \approx_{\varepsilon} \mathbf{E}[f^2(Q)] \approx_{\varepsilon} \mathbf{E}[f^2(R)],$$

which proves our claim since $\mathbf{E}[|Q|^{2/t}] = \Omega(1)$ by Theorem 24.

$\mathbf{E}[|\mathbf{Q}|^{1/t}] \approx_{\varepsilon} \mathbf{E}[\mathbf{f}^2(\mathbf{Q})]$: We have

$$\begin{aligned}
|\mathbf{E}[|\mathbf{Q}|^{2/t}] - \mathbf{E}[f^2(Q)]| &= 2 \int_{T^t}^{\infty} (x^{2/t} - T^2) \cdot \varphi_p(x) dx \\
&= \int_{T^t}^{\infty} (x^{2/t} - T^2) \cdot O(1/x^{p+1}) dx \\
&= O\left(T^{2-tp} \cdot \left(\frac{t}{pt-2} - \frac{1}{p}\right)\right) \\
&= O(1/(Tp)) \\
&= O(\varepsilon)
\end{aligned}$$

$\mathbf{E}[\mathbf{f}^2(\mathbf{Q})] \approx_{\varepsilon} \mathbf{E}[\mathbf{f}^2(\mathbf{R})]$: This is argued nearly identically as in our proof that $\mathbf{E}[f(Q)] \approx_{\varepsilon} \mathbf{E}[f(R)]$

above. The difference is that our error term now corresponding to Eq. (5) is

$$\begin{aligned}
& \int_0^1 x^{2/t-1} \cdot O(k^{-1/p} + k^{-2/p} + 2^{-\Omega(k)}) dx + \int_1^{T^t} x^{2/t-1} \cdot O(k^{-1/p}/x^{p+1} + k^{-2/p}/x^2 + 2^{-\Omega(k)}) dx \\
&= O(\varepsilon) + O\left(\frac{1}{k^{1/p}} \cdot \left(\frac{1}{T^{tp+t-2}} - 1\right) \cdot \frac{1}{\frac{2}{t} - p - 1}\right) \\
&\quad + O\left(\frac{1}{k^{2/p}} \cdot \left(\frac{1}{T^{2t-2}} - 1\right) \cdot \frac{1}{\frac{2}{t} - 2}\right) + O(2^{-\Omega(k)} \cdot (T^2 - 1)) \\
&= O(\varepsilon)
\end{aligned}$$

■