Pseudorandom Hashing for Space-bounded Computation with Applications in Streaming

Praneeth Kacham[*]   Rasmus Pagh[†]   Mikkel Thorup[‡]   David P. Woodruff[§]

Abstract

We revisit Nisan's classical pseudorandom generator (PRG) for space-bounded computation (STOC 1990) and its applications in streaming algorithms. We describe a new generator, Hash-PRG, that can be thought of as a symmetric version of Nisan's generator over larger alphabets. Our generator allows a trade-off between seed length and the time needed to compute a given block of the generator's output. HashPRG can be used to obtain derandomizations with much better update time and *without sacrificing space* for a large number of data stream algorithms, for example:

- Andoni's F_p estimation algorithm for constant $p > 2$ (ICASSP, 2017) assumes a random oracle, but achieves optimal space and constant update time. Using HashPRG's time-space trade-off we eliminate the random oracle assumption while preserving the other properties. Previously no time-optimal derandomization was known. Using similar techniques, we give an algorithm for a relaxed version of ℓ_p sampling in a turnstile stream. Both of our algorithms use $\tilde{O}(d^{1-2/p})$ bits of space and have $O(1)$ update time.

- For $0 < p < 2$, the $1 \pm \varepsilon$ approximate F_p estimation algorithm of Kane et al., (STOC, 2011) uses an optimal $O(\varepsilon^{-2} \log d)$ bits of space but has an update time of $O(\log^2(1/\varepsilon) \log\log(1/\varepsilon))$. Using HashPRG, we show that if $1/\sqrt{d} \le \varepsilon \le 1/d^c$ for an arbitrarily small constant $c > 0$, then we can obtain a $1 \pm \varepsilon$ approximate F_p estimation algorithm that uses the optimal $O(\varepsilon^{-2} \log d)$ bits of space and has an update time of $O(\log d)$ in the Word RAM model, which is more than a quadratic improvement in the update time. We obtain similar improvements for entropy estimation.

- CountSketch, with the fine-grained error analysis of Minton and Price (SODA, 2014). For derandomization, they suggested a direct application of Nisan's generator, yielding a logarithmic multiplicative space overhead. With HashPRG we obtain an efficient derandomization yielding the same asymptotic space as when assuming a random oracle. Our ability to obtain a time-efficient derandomization makes crucial use of HashPRG's symmetry. We also give the first derandomization of a recent private version of CountSketch.

For a d-dimensional vector x being updated in a turnstile stream, we show that $\|x\|_\infty$ can be estimated up to an additive error of $\varepsilon\|x\|_2$ using $O(\varepsilon^{-2} \log(1/\varepsilon) \log d)$ bits of space. Additionally, the update time of this algorithm is $O(\log 1/\varepsilon)$ in the Word RAM model. We show that the space complexity of this algorithm is optimal up to constant factors. However, for vectors x with $\|x\|_\infty = \Theta(\|x\|_2)$, we show that the lower bound can be broken by giving an algorithm that uses $O(\varepsilon^{-2} \log d)$ bits of space which approximates $\|x\|_\infty$ up to an additive error of $\varepsilon\|x\|_2$. We use our aforementioned derandomization of the CountSketch data structure to obtain this algorithm, and using the time-space trade off of HashPRG, we show that the update time of this algorithm is also $O(\log 1/\varepsilon)$ in the Word RAM model.

[*]CMU. Email: pkacham@cs.cmu.edu
[†]BARC, University of Copenhagen. Email: pagh@di.ku.dk
[‡]BARC, University of Copenhagen. Email: mikkel2thorup@gmail.com
[§]CMU. Email: dwoodruf@cs.cmu.edu

Contents

1   Introduction

Space-efficient algorithms are a central theme in computer science. In many cases, the best such algorithms are *randomized*, which raises the question how to obtain the random bits. Nisan's classical pseudorandom generator [Nis92] shows that it is possible to expand a small random seed into a longer pseudorandom string that is essentially as good as full randomness if used in a space-bounded computation. In the context of *streaming algorithms* Nisan's generator has been used not only to reduce the need for random bits, but also because *storing* the seed allows us to re-create random values when they are needed, which is essentially a type of hashing [Cha02, Ind06, AGM12, DKS10, JST11, GKMS03, KNW10, FIS05]. Ideally we would like the space for the seed to be smaller than the space of the streaming algorithm, but a black-box application of Nisan's generator does not quite live up to this ideal: the seed length needed is larger than the space of the streaming algorithm by a multiplicative logarithmic factor. Furthermore, retrieving a block of the string output by the generator requires time proportional to the length of the seed, introducing a significant slowdown in many settings. Hence, obtaining streaming algorithms that are simultaneously space optimal and have a very fast update time—the time required to process each update in the stream—is challenging.

In a turnstile stream, a vector $x \in \mathbb{R}^d$ is initially set to 0^d and receives updates of the form $(i_1, v_1), (i_2, v_2), \ldots, (i_m, v_m) \in [d] \times \{-M, \ldots, M\}$. On receiving an update of the form (i_j, v_j), the vector x is updated as follows: $x_{i_j} \leftarrow x_{i_j} + v_j$. Given a function f with domain $\mathbb{R}^d$, at the end of the stream we want to output an approximation to $f(x)$ using space sublinear in d while processing the stream. Some examples of f are (i) the F_p moments $\sum_{i=1}^d |x_i|^p$ and (ii) the number of distinct elements in the stream, often denoted by $\|x\|_0$. Turnstile streaming algorithms typically apply a randomized linear map $S : \mathbb{R}^d \to \mathbb{R}^D$ to the vector x and show that Sx can be used to approximate $f(x)$ at the end of processing the stream. The advantage of S being a linear map (or linear sketch) is that on receiving an update (i_j, v_j) in the stream, the sketch Sx can be updated by simply adding $(S)_{*i_j} v_j$ to the current sketch. Here $(S)_{*i_j}$ denotes the i_j-th column of the matrix S. Note that to obtain sublinear space algorithms, we cannot store the full matrix S in memory. One of the techniques here is to describe the entries of the matrix S using hash functions that are k-wise independent for a small value of k. For example, the CountSketch matrix of [CCF04] can be described by using 4-wise independent hash functions and thus can be stored efficiently. Further, for any $j \in [d]$, the column $(S)_{*j}$ can be generated efficiently using the hash functions, thereby allowing for a fast update of the sketch in the stream.

Unfortunately, it is not always easy to show that a matrix S generated using hash functions sampled from hash families with limited independence is sufficient to approximate $f(x)$. Indyk [Ind06] showed that we can assume full independence when constructing S and later derandomize[1] the construction of S using pseudorandom generators for small-space computation. The crucial idea of Indyk is that the final state of a linear sketch depends only on the vector x at the end of the stream and not on the sequence of updates that result in the vector x. Thus, any algorithm that assumes the columns of S are sampled independently can be derandomized using a pseudorandom generator that fools small space algorithms as follows: suppose an algorithm needs r uniform random bits to sample a column of S. Fixing a vector $x \in \mathbb{R}^d$, we construct a small space algorithm that makes a single pass over an $r \cdot d$ length uniform random string reading r uniform random bits at a time, sampling the column S_{*j}, and updating the stored sketch by adding $S_{*j} x_j$. If each coordinate of Sx can be stored using t bits, such an algorithm only uses $D \cdot t$ bits of space. As the columns of S

---

[1] We use the term "derandomize" to denote any procedure that *lowers* the randomness required by the algorithm for example by replacing a uniform random string with a string sampled from a Pseudorandom generator that uses a smaller uniform random seed.

were sampled independently, we use the analysis assuming full independence to conclude that the sketch computed by the algorithm has certain desired properties with a certain probability. Now, the small space algorithm can be *fooled*[2] using Nisan's PRG with a seed length of $\Theta(Dt \log(rd/Dt))$, which is $\Theta(Dt \log d)$ in many cases as Dt is often much smaller than d. The argument essentially shows that sampling columns of $\boldsymbol{S}$ using blocks of bits in the pseudorandom string is sufficient to ensure that the properties of $\boldsymbol{S}x$, which were proved assuming independent sampling of columns, still hold.

To implement the derandomized algorithm in a stream, given an update (i_j, v_j), we need to generate the column $\boldsymbol{S}_{*i_j}$ on the fly. Nisan's PRG allows us to generate a block of bits of the pseudorandom string by sequentially evaluating $O(\log d)$ hash functions $h : \{0, 1\}^{Dt} \to \{0, 1\}^{Dt}$ on a Dt length random seed. Thus, if the hash functions $h : \{0, 1\}^{Dt} \to \{0, 1\}^{Dt}$ used by the generator can be evaluated in time T on any input, then the block of bits necessary to generate the column $\boldsymbol{S}_{*i_j}$ can be computed in time $O(T \log d)$. We think of the pseudorandom string as a "hash function" mapping an index $i \in [d]$ to the block of pseudorandom bits needed to generate the column $\boldsymbol{S}_{*i}$.

The above argument of Indyk gives a *black box* way to derandomize a streaming algorithm albeit with a logarithmic space blowup (Dt bits to $Dt \log d$ bits) and an update time of $O(T \log d)$. We use the standard Word RAM model with a word size of $w = \Omega(\log d)$ bits to measure the time complexity of the algorithm. If $Dt \gg \log d$, then the hash functions $h : \{0, 1\}^{Dt} \to \{0, 1\}^{Dt}$ are slow to compute, making the update time $O(T \log d)$ very large. We will mostly express space usage of algorithms in *bits*, but with some Word RAM upper bounds having space measured in *words*.

Thus a naïve application of Nisan's PRG results in a sub-optimal algorithm space-wise, as well as a large update time. As numerous data stream algorithms use Nisan's PRG for derandomization, it is important to improve upon this. We significantly improve on the space overhead and update time for derandomization with our construction HashPRG. First, we show that by a careful analysis, in many problems, the pseudorandom string only has to fool an $O(\log d)$ space algorithm instead of the full $O(Dt)$ space algorithm. We show that such a reduction can be performed for the F_p moment estimation algorithms both for $p \in (0, 2)$ and for $p > 2$. We further show that using a symmetry property of HashPRG, we can reduce the derandomization of CountSketch to fooling an $O(\log d)$ space algorithm. We note that CountSketch in our context cannot be derandomized by an application of Nisan's PRG to fooling an $O(\log d)$ space algorithm; see below for further discussion. We, however, are able to give a reduction to fooling an $O(\log d)$ space algorithm which enjoys two properties (i) the space complexity of the derandomized algorithm will be $O(Dt + \log^2 d)$, which is $O(Dt)$ when $Dt = \Omega(\log^2(d))$ and (ii) the hash functions that are to be evaluated to generate a block of pseudorandom bits that correspond to a column of $\boldsymbol{S}$ will now map the domain $\{0, 1\}^{O(\log d)}$ to a range $\{0, 1\}^{O(\log d)}$. Such hash functions can be evaluated in $O(1)$ time in the Word RAM model.

Even with our reduction to a PRG needing to fool only an $O(\log d)$ space algorithm, the need to evaluate $O(\log d)$ hash functions one after another to compute a block of pseudorandom bits presents a barrier to obtaining fast update time. We show that in the case of $Dt = d^{\Omega(1)}$, the space-vs-time tradeoff of HashPRG lets us trade seed length for fast update time, by varying the number of hash functions needed in order to compute a block of pseudorandom bits.

---

[2]Formally, we say an algorithm using fully-random bits is fooled by a Pseudo Random Generator if the total variation distance between the distribution of outputs of the algorithm when using a string of fully random bits and when using a string sampled from the PRG is small.

1.1 Our Results

We construct a new pseudorandom generator, which we call HashPRG, that satisfies the same guarantees as Nisan's PRG but with an additional symmetry property. Our construction also allows a space-vs-time tradeoff and lets us compute any block of pseudorandom bits quickly if we increase the seed length.

Theorem 1.1 (Informal). *There is a constant $c > 0$ such that for any positive integers n, b and k satisfying $b^k \leq 2^{cn}$, there exists a pseudorandom generator parameterized by n, b and k that converts a random seed of length $O(bkn)$ bits to a bitstring of length $b^k \cdot n$ that cannot be distinguished from truly random bits by any Space(cn) algorithm making a single pass over the length $b^k \cdot n$ bitstring. A given n-bit block of this generator can be computed by evaluating k 2-wise independent hash functions mapping $\{0,1\}^n$ to $\{0,1\}^n$.*

Our generator uses the random seed of length $O(bkn)$ to sample $b \cdot k$ hash functions from a 2-wise independent hash family $\mathcal{H} = \{h : \{0,1\}^n \to \{0,1\}^n\}$ such as [Die96, Theorem 3(b)] and uses n bits as an additional random seed.

In the Word RAM model with a word size $\Omega(n)$, an arbitrary block of n bits in the pseudorandom string generated by our generator can be computed in time $O(k)$. Fixing a value of $b^k = t$, we obtain that HashPRG needs a seed of size $O(t^{1/k}kn)$ to be able to generate a pseudorandom string of length $t \cdot n$ supporting the computation of an arbitrary block of pseudorandom bits in time $O(k)$. By varying k, we get a space versus time trade off. Even more informally, our result can be stated as follows:

Theorem 1.2 (Informal, Compare with [Nis92, Theorem 1]). *A space S algorithm making a single pass over a length $R \leq \exp(S)$ random string can be fooled by a pseudorandom generator with a seed length $O((R/S)^{1/k} \cdot k \cdot S)$ where k is an integer parameter of HashPRG. A block of S random bits in the pseudorandom string can be computed by sequentially evaluating k hash functions mapping $\{0,1\}^S$ to $\{0,1\}^S$.*

Setting $k = \log(R/S)$ in the above theorem, we recover Nisan's result.

Another nice property of HashPRG is that the distribution of the pseudorandom string is symmetric in the following specific way. Let γ be sampled from HashPRG. Let γ be written as $\gamma_0 \circ \gamma_1 \circ \cdots \circ \gamma_{b^k-1}$ where $\circ$ denotes concatenation and each γ_i is a length-n block. For arbitrary $\ell \in \{0, \ldots, b^k - 1\}$, define $\gamma^{\oplus \ell} := \gamma_{0 \oplus \ell} \circ \gamma_{1 \oplus \ell} \circ \cdots \circ \gamma_{(b^k-1) \oplus \ell}$ where $\oplus$ denotes the bitwise xor operation. The construction of HashPRG ensures that $\gamma^{\oplus \ell}$ has the same distribution as γ. In all of our algorithms, we use a block of γ to generate appropriate random variables for the corresponding coordinate of the vector x that is being streamed. To analyze the properties of our streaming algorithms, we create *abstract algorithms*[3] that make a single pass over the pseudorandom string such that the behavior of the streaming algorithm is the same as the abstract algorithm. Since the distribution of the string γ is the same as the distribution of $\gamma^{\oplus \ell}$, the distribution of outputs of the abstract algorithm when run on $\gamma^{\oplus \ell}$ is the same as the distribution of the outputs when run directly on γ. So given a fixed ℓ, the abstract algorithm can "know", throughout its execution, the block of pseudorandom bits corresponding to the coordinate ℓ that our original streaming algorithm uses. "Knowing" these bits is important in our analysis of the CountSketch data structure.

Applications. We use HashPRG to obtain space-optimal algorithms for constant factor F_p ($p > 2$) estimation algorithms with an $O(1)$ update time in turnstile streams. Recall the turnstile stream

---

[3]Algorithms that are created only for analysis and are not implemented.

setting: a vector $x \in \mathbb{R}^d$ is being maintained in the stream. The vector x is initialized to 0^d and receives a stream of updates $(i_1, v_1), (i_2, v_2), \ldots, (i_m, v_m) \in [d] \times \{-M, \ldots, M\}$. Unless otherwise specified, we assume in all of our results that $m, M \leq \text{poly}(d)$. For $0 < p < \infty$, we define $F_p(x) := \sum_{i=1}^{d} |x_i|^p$ and $\ell_p(x) := (\sum_{i=1}^{d} |x_i|^p)^{1/p}$.

Theorem 1.3. *Given $p > 2$, there is a turnstile streaming algorithm that uses $O(d^{1-2/p} \log d)$ words of space and outputs a constant factor approximation to $\|x\|_p$. Further, the streaming algorithm processes each update to the stream in $O(1)$ time in the Word RAM model on a machine with $\Omega(\log d)$ word size.*

We also show that for $p > 2$, similar techniques used in obtaining the above theorem can be used to obtain an algorithm that performs a relaxed version of the approximate ℓ_p sampling in the stream. See Section 4.4.

We next give an algorithm to estimate F_p moments for $0 < p < 2$ in the high accuracy regime. We show that the algorithm of [KNPW11] can be implemented using HashPRG *without a space blowup* and also ensure that the algorithm has a faster update time.

Theorem 1.4. *Let $0 < p < 2$ be a parameter and $1/\sqrt{d} \leq \varepsilon \leq 1/d^c$ be the desired accuracy for a constant $0 < c \leq 1/2$. There is a streaming algorithm that uses $O(\varepsilon^{-2})$ words of space and outputs a $1 \pm \varepsilon$ approximation for $\|x\|_p^p$. The streaming algorithm processes each update to the stream in $O(\log d)$ time in the Word RAM model on a machine with $\Omega(\log d)$ word size.*

Using the reduction in [KNPW11] (Appendix A of the conference version), we can obtain a similarly improved update time for both additive and multiplicative entropy approximation in a stream.

We then derandomize the tighter analysis of CountSketch that [MP14] show assuming fully random hash functions. We prove that such tighter guarantees can be obtained even if the hash functions and sign functions in the CountSketch are generated from the pseudorandom string sampled using HashPRG, improving on the black-box derandomization using Nisan's generator.

Theorem 1.5 (Informal). *Given a table size t, number r of repetitions, and word size $w = \Omega(\log d)$, there exists a derandomization of CountSketch, $\text{CS}_{HashPRG} : \{-M, \ldots, M\}^d \rightarrow \{-2^w, \ldots, 2^w\}^{tr}$, with the following properties:*

1. *The space complexity of the data structure is $O(tr + \log d)$ words.*

2. *Given an update (i, v), the data structure can be updated in time $O(r \log d)$ in the Word RAM model.*

3. *For any $\alpha \in [0, 1]$, given any index ℓ, an estimate of x_l given by $\hat{x}_\ell$ can be constructed from $\text{CS}_{HashPRG}(x)$ such that*

$$\mathbf{Pr}[|x_\ell - \hat{x}_\ell| \geq \alpha \Delta] < 2 \exp(-\alpha^2 r) + 2^{-Cw}$$

for $\Delta = \|tail_t(x)\|_2 / \sqrt{t}$, where $tail_t(x) \in \mathbb{R}^d$ denotes the vector obtained after zeroing out the t entries with the highest absolute value in x.

We crucially use the symmetry property of HashPRG to obtain the above result. Note that when $t \cdot r \geq \log d$, the derandomization presents no asymptotic space blowup and retains the tighter analysis of estimation errors from [MP14]. Further, depending on the parameters t, r it is possible to obtain faster update time using the time-vs-space tradeoff offered by HashPRG. We similarly derandomize the utility analysis of Private CountSketch from [PT22] to obtain:

Theorem 1.6 (Informal). *Consider private CountSketch,* $\mathrm{PCS}(x) = \mathrm{CS}_{HashPRG}(x) + \nu$ *where* $\nu \sim N(0,\sigma^2)^D$, *with* r *repetitions and table size* t, *where HashPRG uses block size* w. *Given* $\ell \in [d]$ *we can compute an estimate* $\hat{x}_\ell$ *from* $\mathrm{PCS}(x)$ *such that for every* $\alpha \in [0,1]$ *and* $\Delta = \|tail_t(x)\|_2/\sqrt{t}$,

$$\mathbf{Pr}[|\hat{x}_\ell - x_\ell| \geq \alpha \max(\Delta, \sigma)] \leq 2\exp(-\Omega(\alpha^2 r)) + O(2^{-cw}) \ .$$

We show the following tight result for estimating $\|x\|_\infty$ in a turnstile stream.

Theorem 1.7 (Informal). *Let* x *be a an arbitrary* d *dimensional vector being maintained in a turnstile stream. Assuming that the coordinates of* x *are integers bounded in absolute value by* $\mathrm{poly}(d)$, *any streaming algorithm that estimates* $\|x\|_\infty$ *up to an additive error of* $\varepsilon\|x\|_2$ *for* $\varepsilon > ((\log d)/d)^{1/4}$ *must use* $\Omega(\varepsilon^{-2}\log d \log 1/\varepsilon)$ *bits. Matching this lower bound, there is an algorithm that uses* $O(\varepsilon^{-2}\log d \log 1/\varepsilon)$ *bits and outputs an approximation to* $\|x\|_\infty$ *up to an additive error of* $\varepsilon\|x\|_2$. *The update time of this algorithm is* $O(\log 1/\varepsilon)$ *in the Word RAM model with a word size* $\Omega(\log d)$.

Our upper bound beats the previous best result of [BGW20, Theorem 10]. Their algorithm uses $O(\varepsilon^{-2}\log d(\log\log d + \log 1/\varepsilon))$ bits of space. Our matching lower bound shows that our result cannot be improved without making additional assumptions on the vector x.

When $\|x\|_\infty = \Theta(\|x\|_2)$, we show that it is possible to break the lower bound in the above result by giving an algorithm that uses $O(\varepsilon^{-2}\log d)$ bits of space and estimates $\|x\|_\infty$ up to an additive error of $\varepsilon\|x\|_2$. The algorithm uses our derandomization of CountSketch with tight guarantees given in [MP14].

Theorem 1.8 (Informal). *Given a* d *dimensional vector* x *being updated in a turnstile stream, if* $\|x\|_\infty = \Theta(\|x\|_2)$ *there is a streaming algorithm that uses* $O(\varepsilon^{-2}\log d)$ *bits and approximates* $\|x\|_\infty$ *up to an additive error of* $\varepsilon\|x\|_2$ *with probability* $\geq 9/10$. *The update time of this algorithm is* $O(\log 1/\varepsilon)$ *in the Word RAM model with a word size* $\Omega(\log d)$.

1.2 Previous Work

F_p **estimation for** $p > 2$ The problem of estimating moments in a stream has been heavily studied in the streaming literature and has been a source of lot of techniques both from the algorithms and lower bounds perspective. A lower bound of $\Omega(d^{1-2/p})$ bits on the space complexity of a constant factor approximation algorithm was shown in [CKS03, BYJKS04, Jay09]. On the algorithms side, [AKO11] gives a sketching algorithm with $m = O(d^{1-2/p}\log d)$ rows using a technique called *precision sampling*; this improves additional polylogarithmic factors of earlier work [IW05, BGKS06]. For linear sketches, [ANPW13] shows a lower bound of $m = \Omega(d^{1-2/p}\log d)$ on the number of rows in the sketch, thereby proving that the algorithm of [AKO11] is tight up to constant factors for linear sketching algorithms. All upper and lower bounds mentioned here are for constant factor approximation algorithms. See [And17, AKO11] and references therein for the upper and lower bounds for $(1 \pm \varepsilon)$-approximate algorithms for F_p estimation.

Later, Andoni [And17] gave a simpler linear sketch using $m = O(d^{1-2/p}\log d)$ rows for F_p moment estimation. Andoni uses the min-stability property of exponential random variables to embed ℓ_p into ℓ_∞ and then uses a CountSketch data structure to estimate the maximum absolute value in the vector obtained by sketching x with scaled exponential random variables. The analysis assumes that the exponential random variables are sampled independently. To derandomize the algorithm, Andoni uses the pseudorandom generator of Nisan and Zuckerman [NZ96] which shows that any randomized algorithm that uses space s and $\mathrm{poly}(s)$ random bits can be simulated using $O(s)$ random bits. Since the space complexity of Andoni's algorithm is $d^{\Omega(1)}$, it can be derandomized with at

most a constant factor blowup in space complexity. A major drawback of using the pseudorandom generator of Nisan and Zuckerman is that the update time of the sketch in the stream is $d^{\Omega(1)}$, which is prohibitively large. In this work, we show that we can derandomize Andoni's algorithm while having an update time of $O(1)$ in the Word RAM model.

We note that there are many other algorithms for F_p-moment estimation, such as [IW05, BGKS06, BO10, MW10, Gan15, GW18], which are based on subsampling the input vector in $O(\log d)$ scales and running an ℓ_2-heavy hitters algorithm at each scale. Although it may be possible to amortize the update time of the $O(\log d)$ levels of subsampling, such algorithms cannot achieve an optimal $O(1)$ update time since all known algorithms for ℓ_2-heavy hitters in the turnstile streaming model require $\Omega(\log d)$ update time.

F_p **estimation for** $p < 2$ Indyk [Ind06] showed how to estimate F_p moments for $p \in (0, 2]$ up to a factor $1 \pm \varepsilon$ in turnstile streams using a space of $O(\varepsilon^{-2} \log d)$ words, which translates to $O(\varepsilon^{-2} \log^2(d))$ bits with our assumption on the values of m, M. This work used *p-stable distributions* and as discussed earlier, introduced the influential technique to derandomize streaming algorithms using pseudorandom generators. Li [Li08] used p-stable distributions to define the geometric mean estimator, which can be used to give an unbiased estimator for $\|x\|_p^p$ with low variance and such that the mean of $\Theta(\varepsilon^{-2})$ independent copies of the estimator gives a $1 \pm \varepsilon$ estimate for $\|x\|_p^p$. Li's algorithm can also be derandomized to use $O(\varepsilon^{-2} \log^2 d)$ bits.

The upper bound was then improved to $O(\varepsilon^{-2} \log d)$ bits by Kane, Nelson and Woodruff [KNW10]. They avoid the $O(\log d)$ factor blowup which is caused when derandomizing using Nisan's PRG by showing that Indyk's algorithm can be derandomized using k-wise independent random variables for a small value of k. They also mention that a "more prudent analysis" of the seed length required in Nisan's generator to derandomize Indyk's algorithm makes the space complexity $O(\varepsilon^{-2} \log d + (\log d)^2)$ bits. The idea there was to use Nisan to fool a median of dot products, but it was not able to exploit fast update time, as we do for $0 < p < 2$, without our large alphabet improvement to Nisan's PRG.

They also show that any algorithm that $1 \pm \varepsilon$ approximates $\|x\|_p^p$ in a turnstile stream must use $\Omega(\varepsilon^{-2} \log d)$ bits of space, hence resolving the space complexity of F_p moment estimation in turnstile streams. Although their algorithm uses an optimal amount of space, the update time of their algorithm is $\tilde{O}(\varepsilon^{-2})$ which is non-ideal when ε is small. Concurrent to their work, [AKO11] gave a streaming algorithm that has a fast update time of $O(\log d)$ per stream element but uses a sub-optimal space of $O(\varepsilon^{-2-p} \log^2 d)$ bits for $p \in [1, 2]$.

[KNPW11] then made progress by giving algorithms that are space-optimal while having a fast update time. They give an algorithm that uses $O(\varepsilon^{-2} \log d)$ bits of space and with an update time of $O(\log^2(1/\varepsilon) \log \log(1/\varepsilon))$ per stream element. They use multiple techniques such as estimating the contribution of heavy and light elements to F_p separately by using new data structures and hash functions drawn from limited independent hash families; they buffer updates and use fast multi-point evaluation of polynomials and amortize the time over multiple updates to obtain fast update times.

When $\varepsilon < 1/d^c$ for a small enough constant c, which is when the $\tilde{O}(\varepsilon^{-2})$ update time of earlier algorithms becomes prohibitive, we show that we can derandomize the algorithm of [KNPW11] using HashPRG. We show that there is a streaming algorithm using an optimal $O(\varepsilon^{-2} \log d)$ bits of space with an update time of $O(\log d)$ per stream element. Our algorithm updates the sketch immediately, removing the need to buffer updates and the use of fast multi-point polynomial evaluation from their algorithm. Our update time thus improves the previous update time of [KNPW11] from $O(\log^2 d \log \log d)$ to $O(\log d)$ for any polynomially small ε, making the first progress on this

problem in over 10 years.

Estimation Error bounds with CountSketch The CountSketch data structure [CCF04] can be used to compute an estimate $\hat{x}_\ell$ of x_ℓ for each $\ell \in [d]$. In [CCF04], the authors show that with high probability the maximum estimation error $\|x - \hat{x}\|_\infty \leq \Delta$ where $\Delta = \|\text{tail}_k(x)\|_2 / \sqrt{k}$ if the table size $t = O(k)$ and the number of repetitions $r = O(\log d)$. Minton and Price [MP14] observed that even though the worst-case estimation error follows the above law, most coordinates in $\hat{x}$ have asymptotically smaller estimation error. Concretely, they show that for any $\alpha \in [0, 1]$ and any coordinate $\ell \in [d]$, $\mathbf{Pr}[|\hat{x}_\ell - x_\ell| \geq \alpha\Delta] \leq O(\exp(-\alpha^2 r))$.

They also show other applications of the above tighter analysis. While proving the above result, they assume that the sign functions used to construct the CountSketch data structure are fully random. They argue that their construction can be derandomized by incurring a factor $O(\log d)$ overhead in space using Nisan's PRG via the *black box* approach we described earlier. We derandomize CountSketch using HashPRG and show that even for the derandomized CountSketch construction, the above estimation error holds, albeit with an additional additive term from the failure of the pseudorandom generator. Our derandomized CountSketch data structure uses $O(r \cdot t + \log d)$ words of space. Note that $O(r \cdot t)$ words of space is anyway required to store the sketched vector and therefore when $\log d = O(r \cdot t)$, our derandomization increases the storage cost by at most a constant factor. Our derandomized CountSketch data structure has an update time of $O(r \log d)$ per stream element. We crucially use the symmetry property of HashPRG to reduce the problem to derandomizing a small space algorithm. While Minton and Price say that a derandomization of their algorithm with Nisan's PRG incurs an $O(\log d)$ factor space blow-up in the size of the CountSketch data structure, we observe that we can avoid the blow-up by a more careful analysis of derandomization using Nisan's PRG. We use the fact that Nisan's PRG is resilient to multiple passes [DPS11] as well to obtain the derandomization. See Section 6.2 for a discussion on how to derandomize CountSketch using Nisan's PRG and how the derandomization with Nisan's PRG compares against derandomization with HashPRG.

In the case of $rt = o(\log d)$, our derandomization is not ideal as the asymptotic space complexity of derandomized CountSketch is larger than a constant factor as compared to $O(r \cdot t)$ words of space. Jayaram and Woodruff [JW18] give an alternate derandomization of CountSketch with strong estimation error guarantees and show that their derandomized CountSketch needs $O(r \cdot t(\log \log d)^2)$ words of space. When $r \cdot t = o(\log d/(\log \log d)^2)$, their derandomization has a smaller space complexity than ours. They use a half-space fooling pseudorandom generator of Gopalan, Kane and Meka [GKM18] to derandomize CountSketch. However, the CountSketch data structure derandomized in [JW18] is a slight modification of the standard CountSketch data structure, and combined with the pseudorandom generator of [GKM18], leads to worse update times as compared to our derandomization of the standard CountSketch data structure. We also stress that in a number of applications of CountSketch, such as to the ℓ_2-heavy hitters problem, one has $r \cdot t = \Omega(\log d)$, and in this regime our derandomization is space-optimal, and not only significantly improves the update time, but also removes the additional $(\log \log d)^2$ factors in the space of [JW18].

Estimating $\|x\|_\infty$ Given a d dimensional vector x being maintained in a turnstile stream, it can be shown that approximating $\|x\|_\infty$ up to a multiplicative factor $C < \sqrt{2}$ requires $\Omega(d)$ bits of space by reducing from the INDEX problem. Notably, the lower bound for multiplicative approximation holds even in the stronger addition only model. Hence the problem of approximating $\|x\|_\infty$ up to an additive error of $\varepsilon\|x\|_2$ has gained more interest. Cormode, in the 2006 IITK workshop on

7

data streams[4], asked if it was possible to approximate $\|x\|_\infty$ to an additive error of $\varepsilon\|x\|_2$ using fewer than $O(\varepsilon^{-2}\log^2 d)$ bits of space. Later [BCIW16] for insertion only streams and [BGW20] for general turnstile streams answered the question in affirmative by giving an algorithm that uses only $O(\varepsilon^{-2}\log d\log\log d)$ bits of space.

1.3 Technical Overview

HashPRG We will briefly describe our construction and show how the construction and analysis differs from Nisan's. Given parameters n, b, k, our construction samples $b \cdot k$ independent hash functions $\boldsymbol{h}_i^{(j)} : \{0,1\}^n \to \{0,1\}^n$ for $(i,j) \in \{0,\ldots,k-1\}\times\{0,1,\ldots,b-1\}$ from a 2-wise independent hash family. We then sample a uniform random string $\boldsymbol{r} \sim \{0,1\}^n$, then the $b^k \cdot n$ pseudorandom string output by the generator is defined by $G_k(\boldsymbol{r},\boldsymbol{h}_1,\ldots,\boldsymbol{h}_k)$, where for any x,

$$G_0(x) := x$$
$$G_k(x,\boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-1}) := G_{k-1}(\boldsymbol{h}_{k-1}^{(0)}(x),\boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-2}) \circ \cdots \circ G_{k-1}(\boldsymbol{h}_{k-1}^{(b-1)}(x),\boldsymbol{h}_1,\ldots,\boldsymbol{h}_{k-1}).$$

Thus for $i \in \{0,\ldots,b^k-1\}$ if i is written as $(i_{k-1}i_{k-2}\cdots i_0)$ in base b, then the i-th block of n bits in the pseudorandom string $G_k(\boldsymbol{r},\boldsymbol{h}_1,\ldots,\boldsymbol{h}_k)$ is given by

$$\boldsymbol{h}_0^{(i_0)}(\cdots(\boldsymbol{h}_{k-1}^{(i_{k-1})}(\boldsymbol{r}))).$$

Our analysis of HashPRG is based on a new, simpler, and more precise analysis of Nisan's generator [Nis92] (see Appendix A for the definition of Nisan's PRG). Note that Nisan's generator corresponds to the special case of our generator where $b = 2$ and where for all i, x, we deterministically fix $\boldsymbol{h}_i^{(0)}(x) := x$. With $\boldsymbol{h}_i^{(0)}$ the identity function, Nisan only defines a single hash function $h_i = \boldsymbol{h}_i^{(1)}$ for each i.

We will now pinpoint where our analysis diverges from a more natural generalization of Nisan's [Nis92]. The first main difference is in Definition 3.4. It plays a role similar to Nisan's

Definition 3 Let $A \subset \{0,1\}^n$, $B \subset \{0,1\}^m$, $h : \{0,1\}^n \to \{0,1\}^m$, and $\varepsilon > 0$. We say that h is (ε, A, B)-independent if $|\mathbf{Pr}_{x\in\{0,1\}^n}[x \in A \text{ and } h(x) \in B] - \varrho(A)\varrho(B)| \le \varepsilon$, where $\varrho(A) = |A|/2^n$ and $\varrho(B) = |B|/2^m$.

However, our Definition 3.4 instead generalizes the following definition:

Definition 3' Let $A \subset \{0,1\}^{n+m}$, $h : \{0,1\}^n \to \{0,1\}^m$, and $\varepsilon > 0$. We say that h is (ϵ, A)-independent if $|\mathbf{Pr}_{x\in\{0,1\}^n}[(x, h(x)) \in A] - \varrho(A)| \le \varepsilon$, where $\varrho(A) = |A|/2^{n+m}$.

It turns out that replacing Definition 3 with Definition 3' in Nisan's proof yields a slightly tighter result. In the proof of [Nis92, Lemma 2], in step 2, all triplets of state nodes i, l, j are considered, where i is a start state, l is an intermediate state and j is an end state. Referring to Definition 3, Nisan uses $A = B_{i,l}^{h_1,\ldots,h_{k-1}} \subset \{0,1\}^n$ and $B = B_{l,j}^{h_1,\ldots,h_{k-1}} \subset \{0,1\}^n$. If instead we base the analysis on Definition 3', then we only need to consider pairs of state nodes i, j and use $A = B_{i,j}^{h_1,\ldots,h_{k-1}} \subset \{0,1\}^{2n}$. This affects the whole analysis of Nisan, but it turns out that it only gets simpler, and this was the starting point for our generalized analysis. Avoiding the intermediate state l was crucial for us because we would need $b-1$ intermediate states $l_1 \ldots l_{b-1}$, and this would lead to much worse bounds with larger b.

---

[4]Follow this link for the list of open problems.

Our construction also implies that the pseudorandom generator has a symmetry property. To see the symmetry, suppose we define $\boldsymbol{h}'_0, \ldots, \boldsymbol{h}'_{k-1}$ with $(\boldsymbol{h}'_0)^{(0)} = \boldsymbol{h}_0^{(1)}$, $(\boldsymbol{h}'_0)^{(1)} = \boldsymbol{h}_0^{(0)}$ and $(\boldsymbol{h}'_i)^{(j)} = \boldsymbol{h}_i^{(j)}$ in all other cases. Then the string $G_k(\boldsymbol{r}, \boldsymbol{h}'_0, \ldots, \boldsymbol{h}'_{k-1})$ is obtained by an appropriate permutation of blocks in the string $G_k(\boldsymbol{r}, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$. As both the strings $G_k(\boldsymbol{r}, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$ and $G_k(\boldsymbol{r}, \boldsymbol{h}'_0, \ldots, \boldsymbol{h}'_{k-1})$ are just as likely when sampling from HashPRG, we obtain the symmetry property.

Extending the above switching argument, we get the following family of transformations that preserve the distribution of the pseudorandom string. Consider an arbitrary $\ell \in \{0, \ldots, b^k - 1\}$. Let $\ell_{k-1} \cdots \ell_0$ be the base b representation of ℓ. Suppose $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}$ be the hash functions sampled in the construction of HashPRG. Define $\boldsymbol{h}'_0, \ldots, \boldsymbol{h}'_{k-1}$ as

$$(\boldsymbol{h}'_i)^{(j)} := \boldsymbol{h}_i^{(j \oplus \ell_i)}$$

for all $i \in \{0, \ldots, k-1\}$ and $j \in \{0, \ldots, b-1\}$. Clearly, the joint distribution of $(\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$ is the same as $(\boldsymbol{h}'_0, \ldots, \boldsymbol{h}'_{k-1})$. Now, for any fixed x, we have

$$G_k(x, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})^{\oplus \ell} = G_k(x, \boldsymbol{h}'_0, \ldots, \boldsymbol{h}'_{k-1}).$$

Hence if $\boldsymbol{\gamma} \sim \text{HashPRG}$, then $\boldsymbol{\gamma}^{\oplus \ell}$ has the same distribution as $\boldsymbol{\gamma}$.

F_p **estimation for** $p > 2$ We derandomize Andoni's algorithm for F_p moment estimation [And17]. Andoni's algorithm can be seen as sketching in two stages: let $\boldsymbol{z} \in \mathbb{R}^d$ be a random vector such that $\boldsymbol{z}_i = \boldsymbol{E}_i^{-1/p} x_i$ where $\boldsymbol{E}_1, \ldots, \boldsymbol{E}_d$ are independent standard exponential random variables. By the min-stability property of exponential random variables, we obtain that $\|\boldsymbol{z}\|_\infty^p \sim \|x\|_p^p / \boldsymbol{E}$, where $\boldsymbol{E}$ is also a standard exponential random variable. Thus, the coordinate of maximum absolute value in $\boldsymbol{z}$ can be used to estimate $\|x\|_p$. Notice that we have not yet performed any dimensionality reduction. Then a CountSketch matrix $\boldsymbol{S}$ with $O(d^{1-2/p} \log d)$ rows is constructed using $O(\log d)$-wise independent hash functions, and this is applied to the vector $\boldsymbol{z}$ to obtain $\boldsymbol{f} = \boldsymbol{S}\boldsymbol{z}$. Andoni argues using the properties of exponential random variables that all of the following hold true simultaneously with a large constant probability: (i) $\|\boldsymbol{z}\|_\infty$ is a constant factor approximation to $\|x\|_p$, (ii) there are only $O((\log d)^p)$ coordinates in $\boldsymbol{z}$ with absolute value greater than $\|x\|_p / c \log d$ and (iii) $\|\boldsymbol{z}\|_2^2 \leq O(d^{1-2/p} \|x\|_2^2)$. Conditioned on these properties of $\boldsymbol{z}$, Andoni argues that $\|\boldsymbol{f}\|_\infty \approx \|\boldsymbol{z}\|_\infty \approx \|x\|_p$ with a large probability. Hence, $\|\boldsymbol{f}\|_\infty$ is a constant factor approximation for $\|x\|_p$ with a large constant probability.

We only have to derandomize the exponential random variables as $\boldsymbol{S}$ is constructed using $O(\log d)$-wise independent hash functions which can be efficiently stored and evaluated [CPT15]. So, we want to show that each of the three properties of the vector $\boldsymbol{z}$ hold even when the exponential random variables are sampled using a pseudorandom string. Now fix a vector x. There is an $O(\log d)$ space algorithm that makes a single pass over the string used to generate exponential random variables and (1) computes $\|\boldsymbol{z}\|_\infty$, (2) computes the number of coordinates in $\boldsymbol{z}$ with absolute value at least $\|x\|_p / c \log d$, and (3) computes $\|\boldsymbol{z}\|_2^2$. The algorithm is simple: it goes over a block of the string to generate $\boldsymbol{E}_1$, sets $\|\boldsymbol{z}\|_\infty = |\boldsymbol{E}_1^{-1/p} x_1|$, increases a counter if $|\boldsymbol{E}_1^{-1/p} x_1| \geq \|x\|_p / (c \log d)$ and sets $\|\boldsymbol{z}\|_2^2 = (\boldsymbol{E}_1^{-1/p} x_1)^2$, and proceeds to read the next block of bits to generate $\boldsymbol{E}_2$ and update the variables using the value of $\boldsymbol{E}_2$ accordingly and continues so on. Now, using any PRG that fools an $O(\log d)$ space algorithm, we obtain that the random variables $\boldsymbol{E}_i$ constructed using the pseudorandom string also make the vector $\boldsymbol{z}$ have each of the three properties. Using the time-vs-space trade-off of HashPRG, we obtain that any block of the pseudorandom string can be obtained in $O(1)$ time if the seed length of HashPRG is d^ϵ for a small constant $\epsilon > 0$, thus

making the time to compute a block of pseudorandom bits $O(1)$ in the Word RAM model. Further, the $O(\log d)$-wise independent hash functions necessary to map the vector $\boldsymbol{z}$ to $\boldsymbol{f}$ when sampled from the constructions of [CPT15] allow for the hash functions to be evaluated in $O(1)$ time. The data structures that allow the hash value to be computed in $O(1)$ time can be stored in d^ϵ bits of space as well for any constant $\epsilon > 0$. Thus, the overall update time of the algorithm is $O(1)$ in the Word RAM model. If ϵ is chosen smaller than $1 - 2/p$, then the asymptotic space complexity of our derandomized F_p moment estimation algorithm remains $O(d^{1-2/p} \log d)$ words, while having a very fast $O(1)$ update time.

We note that the F_p estimation problem for $p > 2$ is ideally suited for HashPRG — it is precisely because the algorithm uses a large amount of memory already that we are able to use our generator over a large alphabet without a space overhead, which then allows us to remove the $O(\log d)$ factor in the update time and achieve constant time. Moreover, it is critical that we can keep track of various quantities needed to fool the algorithm with only $O(1)$ words of memory, as this is also needed for fast update time in our derandomization.

F_p estimation for $p < 2$ We give an alternate derandomization of the algorithm of [KNPW11]. Their algorithm is based on Li's geometric mean estimator [Li08] using p-stable random variables. They introduce two new data structures they call HighEnd and LightEstimator. They also concurrently run an ℓ_p heavy hitters algorithm and at the end of processing the stream, they use the heavy hitters data structure to find a set $L \subseteq [d]$ of coordinates such that $\{\, i \mid |x_i|^p \geq (\alpha/2)\|x\|_p^p \,\} \supseteq L \supseteq \{\, i \mid |x_i|^p \geq \alpha\|x\|_p^p \,\}$. They then use the HighEnd data structure to estimate $\|x_L\|_p^p$ up to a factor of $1 \pm \varepsilon$. By definition of L, all the coordinates in $x_{[d]\setminus L}$ have a small magnitude. Using this fact, they show that their LightEstimator data structure can be used to give a low variance estimator for $\|x_{[d]\setminus L}\|_p^p$. By running multiple independent copies of LightEstimator concurrently, they obtain an accurate estimate for $\|x_{[d]\setminus L}\|_p^p$ and then output the sum of estimates of $\|x_L\|_p^p$ and $\|x_{[d]\setminus L}\|_p^p$.

The HighEnd estimator can be maintained using $O(\varepsilon^{-2} \log d)$ bits and has an update time of $O(\log d)$. Hence, we focus on giving a more efficient derandomization of LightEstimator. At a high level, they hash coordinates of x into $O(1/\alpha)$ buckets and for each bucket they maintain Li's estimator for the F_p moment of coordinates hashed into that bucket. At the end of the stream, the set L is revealed and they output the sum of Li's estimators of the buckets into which none of the elements of L are hashed into. They scale the sum appropriately to obtain an unbiased estimator to $\|x_{[d]\setminus L}\|_p^p$. They show that hashing of the coordinates can be quickly performed using the hash family of [PP08]. The only thing that remains is to derandomize Li's estimator in each individual bucket. They prove that the p-stable random variables in Li's estimator can be derandomized by using $O(1/\varepsilon^p)$-wise independent random variables and show that this is sufficient to obtain algorithms with an optimal space complexity of $O(\varepsilon^{-2} \log d)$ bits and an update time of $O(\log^2(1/\varepsilon) \log \log(1/\varepsilon))$. While the other parts of their algorithm have fast update times, the LightEstimator derandomized using limited independent p-stable random variables leads to a slow update time. We give an alternate derandomization of LightEstimator using HashPRG.

Fix a particular bucket b and the hash function h that hashes the coordinates into one of the buckets. We give an $O(\log d)$ space algorithm that makes a single pass over the string used to generate p-stable random variables and computes Li's estimator for bucket b. The algorithm simply makes a pass over the string, and if a coordinate i gets hashed into bucket b, it uses the block of bits in the string corresponding to the i-th coordinate to generate p-stable random variables and updates Li's estimator for bucket b using the generated p-stable random variables. The existence of such an algorithm implies that the expectation of Li's estimator is fooled by HashPRG and therefore the expectation of the sum of Li's estimators over all the buckets is fooled as well by HashPRG. We

now need to bound the variance of the sum of estimators of the b buckets. Here to fool the variance, that is, we only have to fool $\mathbf{E}[\mathbf{Est}_b \cdot \mathbf{Est}_{b'}]$ for pairs of buckets b, b'. The idea of using a PRG to fool the variance for a streaming problem has also been used in [CIW23]. Here $\mathbf{Est}_b$ denotes the result of Li's estimator for bucket b. Now, for any fixed pair of buckets b, b' we again have that there is an $O(\log d)$ space algorithm that computes Li's estimators for buckets b and b' simultaneously while making a single pass over the string used to generate p-stable random variables. This shows that HashPRG fools $\mathbf{E}[\mathbf{Est}_b \cdot \mathbf{Est}_{b'}]$ and hence the overall variance of the estimator. As we only need to fool the mean and variance, we obtain a derandomization of p-stable random variables using HashPRG. Additionally, when $\varepsilon < 1/d^c$, we can use the time-vs-space tradeoff of HashPRG to obtain an update time of $O(\log d)$ without changing the asymptotic space complexity of the algorithm. Note that the case of ε being small is actually the setting for which we would most like to improve the update time of [KNPW11].

Entropy estimation   An improved update time for F_p moment estimation for $p \in (0, 2)$ also leads to improved update time for entropy estimation using the algorithm of [HNO08]. See Section A in the conference version of [KNPW11] for a discussion on how the update time of F_p moment estimation algorithms translate to update time of approximate entropy estimation algorithms.

CountSketch   CountSketch is a randomized linear map $\mathrm{CS} : \mathbb{R}^d \to \mathbb{R}^D$ defined by two parameters: (i) the table size t and (ii) the number of repetitions r. Here $D = r \cdot t$. For each $i \in [r]$, we have a hash function $\boldsymbol{g}_i : [d] \to [t]$ and a sign function $\boldsymbol{s}_i : [d] \to \{+1, -1\}$. Indexing the coordinates of $\mathrm{CS}(x)$ by $(i, j) \in [r] \times [t]$, we define $\mathrm{CS}(x)_{i,j} = \sum_{\ell \in [d]} [\boldsymbol{g}_i(\ell) = j] \boldsymbol{s}_i(\ell) x_\ell$. Thus, for each repetition i, the coordinate x_ℓ is multiplied with a sign $\boldsymbol{s}_i(\ell)$ and added to the $\boldsymbol{g}_i(\ell)$-th bucket. For each $\ell \in [d]$, we can define $\hat{x}_\ell = \mathrm{median}(\{ \boldsymbol{s}_i(\ell) \cdot \mathrm{CS}(x)_{i, \boldsymbol{g}_i(\ell)} \mid i \in [r] \})$. The randomness in a CountSketch data structure is from the hash functions $\boldsymbol{g}_i$ and the sign functions $\boldsymbol{s}_i$. Assuming that the hash functions $\boldsymbol{h}_i$ and $\boldsymbol{s}_i$ for $i \in [t]$ are drawn independently from 2-wise and 4-wise independent hash families respectively, [CCF04] showed that if $r = O(\log d)$, then with probability at least $1 - 1/\mathrm{poly}(d)$, $\|x - \hat{x}\|_\infty \le \Delta$ for $\Delta = \|\mathrm{tail}_k(x)\|_2/\sqrt{k}$ if $t = O(k)$. As discussed in Section 1.2, Minton and Price assume that the hash functions $\boldsymbol{g}_i$ and the sign functions $\boldsymbol{s}_i$ are fully random to give probability bounds on estimation error for any particular index $\ell \in [d]$. We derandomize their construction by using a pseudorandom generator to sample the hash functions $\boldsymbol{g}_i$ and the sign functions $\boldsymbol{s}_i$. Suppose we treat a bitstring γ as $t \cdot d$ blocks of equal length. We index the blocks of γ by $(i, \ell) \in [r] \times [d]$. We use the block $\gamma_{i,\ell}$ to define $\boldsymbol{g}_i(\ell)$ and $\boldsymbol{s}_i(\ell)$ in the natural way. If γ is sampled uniformly at random, then clearly we have that $\boldsymbol{g}_i$ and $\boldsymbol{s}_i$ constructed using the string γ are fully random and the CountSketch data structure constructed using such hash functions satisfies the guarantees given by Minton and Price. We need to define which block of the string corresponds to which (i, ℓ), since when we receive an update in the stream, we need to be able to extract the corresponding block from the pseudorandom string. Now, we want to show that even if γ is sampled from HashPRG, we obtain similar guarantees on the estimation error. Fix a vector x and coordinate ℓ. Our strategy to derandomize has been to give an algorithm using space $O(\log d)$ bits that makes a single pass over the randomness and computes the quantity of interest. Now, in a black box way, we can conclude that the distribution of the quantity of interest does not change much even if a string sampled from HashPRG is used instead of a fully random string.

We shall try to employ the same strategy here. Fixing an $x \in \mathbb{R}^d$, $\ell \in [d]$ and $\alpha \in [0, 1]$, we want to walk over the string γ and count the number of repetitions $i \in [r]$ for which the value $\boldsymbol{s}_i(\ell) \cdot \mathrm{CS}(x)_{i, \boldsymbol{g}_i(\ell)} > x_\ell + \alpha\Delta$ and the number of repetitions $i \in [r]$ for which the value $\boldsymbol{s}_i(\ell) \cdot \mathrm{CS}(x)_{i, \boldsymbol{g}_i(\ell)} < x_\ell - \alpha\Delta$. Clearly the estimator $\hat{x}_\ell \in [x_\ell - \alpha\Delta, x_\ell + \alpha\Delta]$ if and only if both

counts are smaller than $r/2$ (for odd r). Hence, if both the counts can be computed by a small space algorithm, we can derandomize CountSketch using HashPRG. We immediately hit a roadblock while trying to design such an algorithm. As we fixed an index $\ell \in [d]$, for repetition $i \in [r]$, the quantity of interest is $\mathrm{CS}(x)_{i,\boldsymbol{g}_i(\ell)} := \sum_{\ell' \in [d]} [\boldsymbol{g}_i(\ell') = \boldsymbol{g}_i(\ell)] \boldsymbol{s}_i(\ell') x_{\ell'}$. As the string γ is ordered in increasing order of ℓ' for each repetition i, the algorithm does not know the value of $\boldsymbol{g}_i(\ell)$ until it gets to the block $\gamma_{i,\ell}$. So, for indices $\ell' < \ell$, the algorithm is not aware if $\boldsymbol{g}_i(\ell')$ equals $\boldsymbol{g}_i(\ell)$ or not and hence cannot track the value of $\mathrm{CS}(x)_{i,\boldsymbol{g}_i(\ell)}$ with a single pass over γ. To solve the problem, we use the symmetric property unique to HashPRG. As we mentioned in Section 1.1, if $\boldsymbol{\gamma} \sim$ HashPRG, then for any integer m, the string $\boldsymbol{\gamma}^{\boxplus m}$ has the same distribution as $\boldsymbol{\gamma}$. Making use of this symmetry property, we can now assume that the block $\boldsymbol{\gamma}_{i,j}$ actually corresponds to the hash values of the index $j \boxplus \ell$ for iteration i as the joint distribution of hash and sign values defined by the earlier ordering of the blocks will be the same as the new ordering of the blocks by the symmetry property.

Thus, reading the block $\boldsymbol{\gamma}_{i,1}$, the algorithm immediately knows which bucket the index ℓ gets hashed into in the i-th repetition of CountSketch. Now in a single pass over the blocks $\boldsymbol{\gamma}_{i,1}, \ldots, \boldsymbol{\gamma}_{i,d}$, an $O(\log d)$ space algorithm can compute $\boldsymbol{s}_i(\ell) \mathrm{CS}(x)_{i,\boldsymbol{g}_i(\ell)}$ and at the end of traversing the blocks corresponding to the i-th iteration, the algorithm checks if $\boldsymbol{s}_i(\ell) \mathrm{CS}(x)_{i,\boldsymbol{g}_i(\ell)}$ is $> x_\ell + \alpha \Delta$ or $< x_\ell - \alpha \Delta$ and updates the corresponding counters accordingly. Thus, we have an $O(\log d)$ space algorithm and HashPRG fools an $O(\log d)$ space algorithm and can be used to derandomize CountSketch with the stronger guarantees as given by Minton and Price without incurring a space blowup when $\log d = O(r \cdot t)$.

For each update in the stream, and for each of the r repetitions, we need to compute a block of the pseudorandom string which takes $O(\log d)$ time in the Word RAM model with a word size $w = \Omega(\log d)$. Hence, our derandomized CountSketch with strong guarantees from [MP14] has an update time of $O(r \log d)$. When $r \cdot t = d^{\Omega(1)}$, using the space-vs-time trade-off of HashPRG, we can obtain an update time of $O(r)$.

Private CountSketch   In a recent work, Pagh and Thorup [PT22] gave an improved analysis of the estimation error of differentially private CountSketch. After computing $\mathrm{CS}(x)$ in the stream, they compute $\mathrm{PCS}(x) = \mathrm{CS}(x) + \boldsymbol{\nu}$ where $\boldsymbol{\nu} \sim N(0, \sigma^2)^D$. They show that for an appropriate value of σ, $\mathrm{PCS}(x)$ is (ε, δ)-differentially private. They also show that for any $\ell \in [d]$, the estimator $\hat{x}_\ell$ computed using $\mathrm{PCS}(x)$ also concentrates heavily around x_ℓ, and gave similar concentration bounds to that of Minton and Price [MP14]. The analysis of [PT22] assumes that the hash functions and sign functions are fully random – the same assumption as in [MP14]. We derandomize the Private CountSketch construction using HashPRG, which is similar to our derandomization of [MP14].

Approximating $\|x\|_\infty$ We consider the problem of approximating $\|x\|_\infty$ up to an additive error $\varepsilon \|x\|_2$. Directly using the CountSketch data structure, we can approximate each coordinate of the vector x up to an additive error of $\varepsilon \|x\|_2$ using $O(\varepsilon^{-2}(\log d)^2)$ bits. In the space complexity, a $\log d$ factor comes from the fact that we have to store integers with magnitudes $\mathrm{poly}(d)$ and other $\log d$ factor is because of $O(\log d)$ repetitions of CountSketch to be able to take a union bound over the reconstruction error of all $\mathrm{poly}(d)$ coordinates. We show that there is a simple linear sketching technique to reduce the dimension from d to $\mathrm{poly}(1/\varepsilon)$ while preserving the $\|x\|_\infty$ up to an additive error of $\varepsilon \|x\|_2$. Then the CountSketch data structure, using a space of $O(\varepsilon^{-2} \log 1/\varepsilon \log d)$ bits, can be used to approximate the $\|\cdot\|_\infty$ of the sketched vector thereby approximating $\|x\|_\infty$. We use the randomized map $\boldsymbol{L} : \mathbb{R}^d \to \mathbb{R}^t$ defined as follows to reduce the dimension: let $\boldsymbol{h} : [d] \to [t]$ be drawn at random from a 2-wise independent hash family and $\boldsymbol{s} : [d] \to \{+1, -1\}$ be drawn at random from a 4-wise independent hash family. Define $(\boldsymbol{L}x)_i = \sum_{j \in [d] : \boldsymbol{h}(i) = j} \boldsymbol{s}(j) x_j$. Using simple variance

computations, we show that if $t = \text{poly}(1/\varepsilon)$, with a high probability, $\|\boldsymbol{L}x\|_\infty = \|x\|_\infty \pm \varepsilon\|x\|_2$ and that $\|\boldsymbol{L}x\|_2 \le 2\|x\|_2$. Note that a turnstile update to one coordinate of x simply translates to a turnstile update to one coordinate of $\boldsymbol{L}x$ and hence the two stage sketching algorithm can be efficiently implemented in a stream. A similar universe reduction was employed in [KNPW11] but the technique was previously not explored in the context of ℓ_∞ estimation.

We further show that this simple algorithm is tight by showing an $\Omega(\varepsilon^{-2} \log 1/\varepsilon \log d)$ bits lower bound on any algorithm that approximates $\|x\|_\infty$ up to an additive $\varepsilon\|x\|_2$. We define a communication problem called "Augmented Sparse Set-Disjointness". In this one-way communication problem, Alice receives sets $A_1, \ldots, A_t$ with the property that $|A_i| = k$ and $A_i \subseteq [n]$ for all $i \in [t]$. Bob similarly receives the sets $B_1, \ldots, B_t$ with the same properties. Given an index $i \in [t]$ and the sets $A_1, \ldots, A_{i-1}$, using a single message M (possibly randomized) from Alice, Bob has to compute if $B_i \cap A_i = \emptyset$ or not. In the case of $t = 1$, [DKS10] show that if $n \ge k^2$, then the sparse set disjointness problem has a communication lower bound of $\Omega(k \log k)$. Note that for $t = 1$, simply sending the entire set A_1 to Bob requires a communication of $O(k \log n)$ bits. Surprisingly, they show that there is a protocol using $O(k \log k)$ bits to solve the problem. Extending their ideas, we show that the Augmented Sparse Set-Disjointness problem has a communication lower bound of $\Omega(tk \log k)$. Embedding an instance of the Sparse Set-Disjointness into approximating $\|x\|_\infty$ for an appropriate vector x, we show a lower bound of $\Omega(\varepsilon^{-2} \log 1/\varepsilon \log d)$ bits.

The hard distribution in the lower bound has the property that the vector x constructed satisfies $\|x\|_\infty = O(\varepsilon\|x\|_2)$. We show that it is possible to break the lower bound if we assume that $\|x\|_\infty \ge c\|x\|_2$ for a constant c. This is the case when the coordinates of x follow Zipf's law where the i-th largest coordinate has a value approximately $i^{-\alpha}$ for $\alpha > 0.5$. The algorithm is again the two stage sketch we described but in the second stage, instead of using CountSketch as described in [CCF04] using constant wise independent hash functions, we use the tighter CountSketch guarantees that we obtain by derandomizing the analysis of [MP14] using HashPRG. We show that after the first level sketching, only a few large coordinates need to be estimated to large accuracy while the rest of the coordinates can have larger estimation errors. Using this insight, we show that $O(\varepsilon^{-2} \log d)$ bits of space is sufficient to estimate $\|x\|_\infty$ up to an additive error of $\varepsilon\|x\|_2$.

2 Preliminaries

Notation. For an integer $n \ge 1$, let $[n]$ denote $\{1, 2, \ldots, n\}$. For a predicate P let $[P]$ have the value 1 when P is true and the value 0 when P is false. Let $\text{tail}_t(x)$ denote the vector derived from x by changing the t entries with largest absolute value to zero. We use bold symbols such as $\boldsymbol{h}, \boldsymbol{x}, \boldsymbol{N}, \ldots$ to denote that these objects are explicitly sampled from an appropriate distribution.

For a real-valued matrix M define $\|M\| = \sup_{x \ne 0} \|xM\|_1 / \|x\|_1$ where $\|x\|_1 = \sum_i |x_i|$. (This matrix norm is sometimes written $\|M\|_1$, but since we do not need other matrix norms we omit the subscript.)

Model of Computation. All of our running times are in the Word RAM model with a word size $O(\log d)$ unless otherwise mentioned. We assume that all elementary operations on words can be performed in $O(1)$ time.

Definition 2.1 (k-wise independence). A family of hash functions $\mathcal{H} = \{\, h : [u] \to [v] \,\}$ is said to be k-wise independent if for any k distinct keys $x_1, x_2, \ldots, x_k \in [u]$ and not necessarily distinct values $y_1, y_2, \ldots, y_k \in [v]$

$$\Pr_{\boldsymbol{h} \sim \mathcal{H}}[\boldsymbol{h}(x_1) = y_1 \wedge \cdots \wedge \boldsymbol{h}(x_k) = y_k] = \frac{1}{v^k}.$$

The definition states that if $\boldsymbol{h} \sim \mathcal{H}$, then for any $x \in [u]$, the random variable $\boldsymbol{h}(x)$ is uniformly distributed over $[v]$ and for any k distinct keys $x_1, \ldots, x_k \in [u]$, the random variables $\boldsymbol{h}(x_1), \ldots, \boldsymbol{h}(x_k)$ are independent.

We now state a randomized construction of a hash family from [CPT15] that lets us evaluate the sampled hash function quickly on any input.

Theorem 2.2 (Corollary 3 in [CPT15]). *There exists a randomized data structure that takes as input positive integers $u, v = u^{O(1)}, t, k = u^{O(1/t)}$ and selects a family of function $\mathcal{F}$ from $[u]$ to $[v]$. In the Word RAM model with word length $\Theta(\log u)$ the data structure satisfies the following:*

1. *The space used to represent the family $\mathcal{F}$ as well as a function $f \in \mathcal{F}$ is $O(ku^{1/t}t)$ bits.*

2. *The evaluation time of any function $f \in \mathcal{F}$ on any input is $O(t \log t)$.*

3. *With probability $\geq 1 - u^{-1/t}$, we have that $\mathcal{F}$ is a k-wise independent family.*

Throughout the paper, we use this construction with a constant t and k at most $O(\log u)$.

We now state the guarantees of the hash family construction from [PP08]. Whereas the above construction gives a randomized hash family $\mathcal{F}$ that is k-wise independent with some probability, the following construction gives a randomized hash family $\mathcal{H}$ that when restricted to any fixed subset S of a certain size is a uniform hash family with some probability. We use this construction when we want to ensure that all elements of a small underlying set (but unknown to the streaming algorithm) are hashed to uniformly random locations.

Theorem 2.3 (Theorem 1 of [PP08]). *Let $S \subseteq U = [u]$ be a set of $z > 1$ elements and let $V = [v]$ for any $1 \leq v \leq u$. Suppose the machine word size is $\Omega(\log u)$. For any constant $C > 0$, there is a Word RAM algorithm that, using $\log(z)(\log v)^{O(1)}$ time and $O(\log z + \log \log u)$ bits of space, selects a family of hash functions $\mathcal{H}$ from U to V (independent of S) such that*

- *$\mathcal{H}$ is z-wise independent (in other words, uniform) when restricted to S, with probability $1 - O(1/z^C)$.*

- *Any function $h \in \mathcal{H}$ can be represented using $O(z \log v)$ bits and h can can be evaluated on any $x \in U$ in $O(1)$ time in the Word RAM model. The data structure (or representation) of a random function from the family $\mathcal{H}$ can be constructed in $O(z)$ time.*

We now define ε pseudorandom generators for any given class of algorithms $\mathcal{C}$.

Definition 2.4. A generator $G : \{0, 1\}^n \to \{0, 1\}^m$ is called an ε pseudorandom generator for the class of algorithms $\mathcal{C}$ if for any $C \in \mathcal{C}$,

$$\left| \Pr_{\boldsymbol{x} \sim U_m} [C(\boldsymbol{x}) \text{ accepts}] - \Pr_{\boldsymbol{y} \sim U_n} [C(G(\boldsymbol{y})) \text{ accepts}] \right| < \varepsilon.$$

Here U_m denotes the uniform distribution over $\{0, 1\}^m$. For a generator G, the quantity n is called the *seed length*. In this work $\mathcal{C}$ is taken to be the set of space bounded algorithms with an appropriate space parameter s. The algorithms have a read/write space of s bits and a read only tape that contains an input. The algorithms are allowed to only stream over the m length random/pseudorandom string. Although the above definition is in terms of accept/reject, it can extended to general functions by instead considering the total variation distance between the distributions of the outputs.

14

3 HashPRG

In this section we present a new pseudorandom generator for space-bounded computations, which is going to be our main tool for derandomizing streaming algorithms. The starting point is the classical generator of Nisan [Nis92] (summarized in Appendix A), which we extend to provide a trade-off between seed length and the time to compute an arbitrary output block. To make our treatment easy to access for readers familiar with Nisan's generator, we will follow the same proof outline, but make crucial changes to avoid a union bound over all possible intermediate states. An advantage of our generator over Nisan's, even in the setting where the seed lengths are the same, is that it has a certain *symmetry* property that we need in our applications.

3.1 HashPRG Construction

Let $b \geq 2$ be an integer. Consider $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}$ where $\boldsymbol{h}_i := (\boldsymbol{h}_i^{(0)}, \ldots, \boldsymbol{h}_i^{(b-1)})$ is a vector of b hash functions with each $\boldsymbol{h}_i^{(j)} : \{0,1\}^n \to \{0,1\}^n$ being a hash function drawn independently from a 2-wise independent hash family $\mathcal{H}$. We slightly abuse terminology and also refer to the vectors $\boldsymbol{h}_i$ as hash functions.

Using hash functions $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}$, define a generator $G_k : \{0,1\}^n \to \{0,1\}^{n \cdot b^k}$ recursively:

$$G_0(x) = x$$
$$G_k(x, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}) = G_{k-1}(\boldsymbol{h}_{k-1}^{(0)}(x), \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}) \circ \cdots \circ G_{k-1}(\boldsymbol{h}_{k-1}^{(b-1)}(x), \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}).$$

Here $\circ$ denotes concatenation. Note that Nisan's generator can be obtained by setting $b = 2$ and deterministically setting $\boldsymbol{h}_i^{(0)}(x) := x$ for all i, x.

By construction we have, for $x \in \{0,1\}^n$, that $G_k(x, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$ is a bitstring of length $n \cdot b^k$. We look at the string $G_k(x, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$ as concatenation of b^k chunks each of length n, chunks indexed by $0, \ldots, b^k - 1$. Let $j \in \{0, \ldots, b^k - 1\}$ be written as $j_{k-1} \cdots j_0$ in base b. Then the jth chunk of the string $G_k(x, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$ is given by

$$\boldsymbol{h}_0^{(j_0)}(\boldsymbol{h}_1^{(j_1)}(\cdots \boldsymbol{h}_{k-1}^{(j_{k-1})}(x))).$$

To define the power of our pseudorandom generator we need the following notation. Let Q be an arbitrary finite state machine with 2^w states over the alphabet $\{0,1\}^n$. Let D be any distribution over the strings of length $b^k \cdot n$, encoding b^k steps of the FSM. Let $Q(D)$ be a $2^w \times 2^w$ matrix where $[Q(D)]_{ij}$ is the probability that the FSM starting in state i goes to state j after performing b^k steps based on an input drawn from D. Let U_n denote the uniform distribution over $\{0,1\}^n$. We will prove the following lemma.

Lemma 3.1. *There exists a constant $c > 0$, given integers n and $w \leq cn$ and parameters b, k with $b^k \leq 2^{cn}$, for any FSM Q with at most 2^w states, if $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1} : \{0, \ldots, b-1\} \to \{0,1\}^n \to \{0,1\}^n$ are drawn independently from $\mathcal{H}^b$, where $\mathcal{H}$ is any family of 2-wise independent hash functions, then with probability $\geq 1 - 2^{-cn}$ (over the draw of $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}$),*

$$\|Q(G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})) - Q((U_n)^{b^k})\| \leq 2^{-cn}.$$

Here $G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$ denotes uniform distribution over the set $\{G_k(x, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}) \mid x \in \{0,1\}^n\}$.

By definition of the matrix norm $\| \cdot \|$ (see section 2) this implies that that with probability at least $1 - 2^{-cn}$ over the hash functions $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}$, we have that the total variation distance

between the distribution of final state using a random string drawn from $(U_n)^{b^k}$ and a random string drawn from $G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$ is at most 2^{-cn}.

Using the above lemma, we will then prove the following theorem which we will use through out the paper.

Theorem 3.2. *There exists a constant $c > 0$ such that given any parameters n, b and k satisfying $b^k \leq 2^{cn}$, there exists a generator which we call $HashPRG : \{0,1\}^{O(bkn)} \rightarrow \{0,1\}^{b^k \cdot n}$ such that HashPRG is an $O(2^{-cn})$ pseudorandom generator for the class of Finite State Machines over alphabet $\{0,1\}^n$ with at most 2^{cn} states. For any seed $\boldsymbol{r}$, the i-th block of bits in $HashPRG(\boldsymbol{r})$ can be computed in time $O(k)$ on a machine with Word Size $\Omega(n)$.*

Proof. Let $p_{i,j}^{h_0,\ldots,h_{k-1}}$ be $[Q(G_k(*, h_0, \ldots, h_{k-1}))]_{i,j}$. Let $p_{i,j}$ be the probability the FSM starting in state i goes to a state j after b^k steps on an input $\boldsymbol{r}$ where $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1} \sim \mathcal{H}^b$ and $\boldsymbol{r} \sim G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$. Clearly, for any pair of states i, j,

$$p_{i,j} = \frac{1}{|\mathcal{H}^b|^k} \sum_{(h_0,\ldots,h_{k-1}) \in (\mathcal{H}^b)^k} p_{i,j}^{h_0,\ldots,h_{k-1}}.$$

Let $q_{i,j} = [Q((U_n)^{b^k})]_{i,j}$. Now, for any i,

$$\sum_j |p_{i,j} - q_{i,j}| \leq \sum_j \frac{1}{|\mathcal{H}^b|^k} \sum_{(h_0,\ldots,h_{k-1}) \in (\mathcal{H}^b)^k} |p_{i,j}^{h_0,\ldots,h_{k-1}} - q_{i,j}|$$

$$= \frac{1}{|\mathcal{H}^b|^k} \sum_{(h_0,\ldots,h_{k-1}) \in (\mathcal{H}^b)^k} (\sum_j |p_{i,j}^{h_0,\ldots,h_{k-1}} - q_{i,j}|).$$

By Lemma 3.1, $|\{(h_0, \ldots, h_{k-1}) \mid \sum_j |p_{i,j}^{h_0,\ldots,h_{k-1}} - q_{i,j}| \leq 2^{-cn}\}| \geq |\mathcal{H}^b|^k(1 - 2^{-cn})$ and using the fact that for any $(h_0, \ldots, h_{k-1})$, $|p_{i,j}^{h_0,\ldots,h_{k-1}} - q_{i,j}| \leq 1$, we obtain that

$$\sum_j |p_{i,j} - q_{i,j}| \leq 2^{-cn} + 1 \cdot 2^{-cn} \leq 2 \cdot 2^{-cn}.$$

As the above holds for any i, we obtain that $\|Q(HashPRG) - Q((U_n)^{b^k})\| \leq 2 \cdot 2^{-cn}$. Where we overload the notation HashPRG to also denote the distribution of string $\boldsymbol{\gamma}$ obtained by first sampling $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}$ and then sampling $\boldsymbol{\gamma} \sim G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$. The $O(bkn)$ bits of random seed is used to sample bk independent hash functions from the 2-wise independent hash family construction of Dietzfelbinger [Die96, Theorem 3(b)]. The bk independent hash functions are used to construct the hash functions $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}$ and n bits from the random string are used to get an $\boldsymbol{x} \sim \{0,1\}^n$. The output of HashPRG is then $G_k(\boldsymbol{x}, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$. Clearly, by the definition of G_k, any block of bits in the output string can be computed in $O(k)$ time. $\qquad \square$

A special case of the above theorem we use is when $n = O(\log d)$. By choosing $b = d^\epsilon$ and $k = O(1/\epsilon)$ for a small enough constant ϵ, we obtain that HashPRG with these parameters fools any FSM with $\text{poly}(d)$ states over an alphabet $\{0,1\}^n$. And that on a machine with word size $\Omega(\log d)$, any block of bits in the output of HashPRG can be computed in time $O(1)$ since ϵ is a constant.

3.2 HashPRG Analysis

We will be reasoning about matrices that represent transition probabilities of Q under different input distributions. We start with some simple facts about norms of matrices. Recall that $\|M\| := \sup_{x:\|x\|_\infty=1} \|Mx\|_\infty$. First, for any two matrices A and B we have $\|A + B\| \leq \|A\| + \|B\|$ and $\|AB\| \leq \|A\|\|B\|$. We will also need the following lemma about the norm of differences of matrix powers:

Lemma 3.3. *For integer $b \geq 1$ and square real matrices M and N with $\|M\| \leq 1$ and $\|N\| \leq 1$, $\|M^b - N^b\| \leq b\|M - N\|$.*

Proof. The proof is by induction on b. The statement clearly holds for $b = 1$. For the induction step we have:

$$
\begin{aligned}
\|M^b - N^b\| &\leq \|M^b - N^{b-1}M\| + \|N^{b-1}M - N^b\| \\
&= \|(M^{b-1} - N^{b-1})M\| + \|N^{b-1}(M - N)\| \\
&\leq \|M^{b-1} - N^{b-1}\|\|M\| + \|N\|^{b-1}\|M - N\| \\
&\leq (b-1)\|M - N\| + \|M - N\| = b\|M - N\| \ .
\end{aligned}
$$

This first uses triangle inequality, the second inequality uses that the norm of a product is bounded by the product of the norms, and the third inequality uses the induction hypothesis as well as the assumption $\|M\| \leq 1$, $\|N\| \leq 1$. $\square$

For fixed hash functions $h_0, \ldots, h_{k-1}$, let $G_k(*, h_0, \ldots, h_{k-1})$ be the distribution of $G_k(\boldsymbol{x}, h_0, \ldots, h_{k-1})$ when $\boldsymbol{x}$ is drawn uniformly at random from $\{0, 1\}^n$, denoted $\boldsymbol{x} \sim U_n$. Let $(U_n)^{b^k}$ denote the uniform distribution over length $n \cdot b^k$ bitstrings. The aim is to show, akin to Nisan's generator, for any "small-space computation", with high probability over $\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}$, the distribution $G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})$ is indistinguishable from the uniform distribution $(U_n)^{b^k}$.

Lemma 3.1 will be derived from this result: For any $\varepsilon > 0$, if all $\boldsymbol{h}_i^{(j)}$ are drawn independently from a 2-wise independent hash family,

$$
\mathbf{Pr}\left[\|Q(G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})) - Q((U_n)^{b^k})\| > \varepsilon\right] \leq \frac{(b^k - 1)^2 k}{(b-1)^2 \varepsilon^2} \cdot 2^{3w-n} \ . \tag{1}
$$

Note that each $\boldsymbol{h}_i$ can also be seen as function from $\{0, 1\}^n$ to $\{0, 1\}^{bn}$ defined as $\boldsymbol{h}_i(x) = \boldsymbol{h}_i^{(0)}(x) \circ \cdots \circ \boldsymbol{h}_i^{(b-1)}(x)$. We say $\boldsymbol{h}_i : \{0, 1\}^n \to \{0, 1\}^{bn}$ is drawn from the hash family $\mathcal{H}^b$. Since each $\boldsymbol{h}_i^{(j)}$ is sampled independently from a 2-wise independent hash family, $\mathcal{H}^b$ is a also 2-wise independent.

Definition 3.4. Let $A \subseteq \{0, 1\}^{bn}$, $h : \{0, 1\}^n \to \{0, 1\}^{bn}$, and $\varepsilon > 0$. We say that h is (ε, A)-*independent if* $|\mathbf{Pr}_{\boldsymbol{x}\sim U_n}[h(\boldsymbol{x}) \in A] - \varrho(A)| \leq \varepsilon$, where $\varrho(A) := |A|/2^{bn}$ is the density of set A.

The above definition corresponds to Definition 3 in [Nis92] but with some important differences. The differences leads to subtle changes in the whole analysis, but the overall structure of the analysis remain the same.

We now have the following lemma that corresponds to Lemma 1 of [Nis92].

Lemma 3.5. *Let $A \subseteq \{0, 1\}^{bn}$ and $\varepsilon > 0$. Then $\mathbf{Pr}_{h\sim\mathcal{H}^b}[h$ is not (ε, A)-independent$] < \frac{\varrho(A)}{\varepsilon^2 2^n}$.*

Proof. Consider a matrix M that has a row for each $x \in \{0,1\}^n$ and a column for each $h \in \mathcal{H}^b$. Let $M(x,h) = 1$ if $h(x) \in A$ and 0 otherwise. We define the function f that expresses the probability that a function h maps $x \sim U_n$ to a value in A:

$$f(h) = \mathop{\mathbf{E}}_{x \sim U_n}[M(x,h)] = \mathop{\mathbf{Pr}}_{x \sim U_n}[h(x) \in A].$$

For $h \sim \mathcal{H}^b$ we have

$$\mathop{\mathbf{E}}_{h \sim \mathcal{H}^b} f(h) = \mathop{\mathbf{E}}_{h \sim \mathcal{H}^b, x \sim U_n}[M(x,h)] = \mathop{\mathbf{E}}_{x \sim U_n} \mathop{\mathbf{Pr}}_{h \sim \mathcal{H}^b}[h(x) \in A] = \varrho(A),$$

where the last equality follows from $\mathbf{Pr}_{h \sim \mathcal{H}^b}[h(x) \in A] = \varrho(A)$, since $\mathcal{H}^b$ is 2-wise independent. We next bound the variance of $f(h)$ to show that $f(h)$ is close to $\varrho(A)$ with high probability:

$$\mathop{\mathbf{Var}}_{h \sim \mathcal{H}^b} f(h) = \mathop{\mathbf{E}}_{h \sim \mathcal{H}^b} (f(h) - \varrho(A))^2$$
$$= \mathop{\mathbf{E}}_{h \sim \mathcal{H}^b} (\mathop{\mathbf{E}}_{x \sim U_n}[M(x,h) - \varrho(A)])^2.$$

If $x_1 \sim U_n$ and $x_2 \sim U_n$ are independently drawn we can expand

$$(\mathop{\mathbf{E}}_{x \sim U_n}[M(x,h) - \varrho(A)])^2$$
$$= \mathop{\mathbf{E}}_{x_1 \sim U_n, x_2 \sim U_n}(M(x_1,h) - \varrho(A))(M(x_2,h) - \varrho(A)).$$

Using the fact that for every x, $\mathbf{E}_{h \sim \mathcal{H}^b}[M(x,h)] = \varrho(A)$ we obtain

$$\mathop{\mathbf{Var}}_{h \sim \mathcal{H}^b} f(h) = \mathop{\mathbf{E}}_{x_1 \sim U_n, x_2 \sim U_n} \mathop{\mathbf{E}}_{h \sim \mathcal{H}^b}[M(x_1,h)M(x_2,h) - \varrho(A)^2].$$

With probability $1 - 2^{-n}$ we have $x_1 \neq x_2$. Since h is drawn from a 2-wise independent hash family, $h(x_1)$ and $h(x_2)$ are independent in this case, so $\mathbf{E}_{h \sim \mathcal{H}^b}[M(x_1,h)M(x_2,h) \mid x_1 \neq x_2] = \varrho(A)^2$. With probability $1/2^n$, we have $x_1 = x_2$ and $\mathbf{E}_{h \sim \mathcal{H}^b}[M(x_1,h)M(x_2,h) \mid x_1 = x_2] = \mathbf{E}_{h \sim \mathcal{H}^b}[M(x_1,h)] = \varrho(A)$. Hence,

$$\mathop{\mathbf{Var}}_{h \sim \mathcal{H}^b} f(h) = (1 - 2^{-n})\varrho(A)^2 + 2^{-n}\varrho(A) - \varrho(A)^2 < \frac{\varrho(A)}{2^n}.$$

By Chebyshev's inequality,

$$\mathop{\mathbf{Pr}}_{h \sim \mathcal{H}^b}[|f(h) - \varrho(A)| \geq \varepsilon] \leq \frac{\varrho(A)}{2^n \varepsilon^2}. \qquad \square$$

We now have the following definition that corresponds to Definition 4 of [Nis92].

Definition 3.6. Let Q be a FSM on alphabet $\{0,1\}^n$ and pick $\varepsilon > 0$. For hash functions $h_0, \ldots, h_{k-1} : \{0,1\}^n \to \{0,1\}^{bn}$, we say $(h_0, \ldots, h_{k-1})$ is (ε, Q)-*good* if $\|Q(G_k(*, h_0, \ldots, h_{k-1})) - Q((U_n)^{b^k})\| \leq \varepsilon$.

For small ε the definition essentially says that the FSM Q cannot distinguish between the distributions $G_k(*, h_0, \ldots, h_{k-1})$ and $(U_n)^{b^k}$. The following lemma corresponds to Lemma 2 of [Nis92].

Lemma 3.7. *Let Q be a FSM of size 2^w over alphabet $\{0,1\}^n$ and k a nonnegative integer. Then for every $\varepsilon > 0$,*

$$\Pr_{\boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-1}\sim\mathcal{H}^b}[(\boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-1}) \text{ is not } ((b^k-1)\varepsilon, Q)\text{-good}]$$

$$\leq \frac{2^{3w}k}{(b-1)^2\varepsilon^2 2^n}.$$

Proof. The proof is a careful translation of the proof of Lemma 2 of [Nis92]. The proof is by induction on k. For $k = 0$, the statement is immediate since $G_0(x) \sim U_n$ which means that there is zero difference between the distributions. For the induction step assume that the statement holds for $k-1$. For every choice of $h_0,\ldots,h_{k-2} : \{0,1\}^n \to \{0,1\}^{bn}$ and every two states i,j of Q we define the set of seeds $y^{(0)},\ldots,y^{(b-1)}$ that when used one by one with G_{k-1} produces a string that takes Q from state i to state j:

$$B_{i,j}^{h_0,\ldots,h_{k-2}} = \{(y^{(0)},\ldots,y^{(b-1)}) \in \{0,1\}^{bn} \mid G_{k-1}(y^{(0)}, h_0,\ldots,h_{k-2}) \circ \cdots \circ G_{k-1}(y^{(b-1)}, h_0,\ldots,h_{k-2})$$
$$\text{takes } Q \text{ from } i \text{ to } j\}.$$

Now sample $\boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-1} \sim \mathcal{H}^b$ independently and consider the following events:

1. $(\boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-2})$ is $((b^{k-1}-1)\varepsilon, Q)$-good (see Definition 3.6), and

2. $\boldsymbol{h}_{k-1}$ is $(\varepsilon(b-1)/2^w, B_{i,j}^{h_0,\ldots,h_{k-2}})$-independent for all states i,j (see Definition 3.4).

We will see that these events happen simultaneously with probability at least $1 - \frac{2^{3w}k}{(b-1)^2\varepsilon^2 2^n}$. Furthermore, we will show that when this happens, $(\boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-1})$ is $((b^k-1)\varepsilon, Q)$-good, completing the induction step.

By the induction hypothesis, the probability of event 1 *not* happening is at most $\frac{2^{3w}(k-1)}{(b-1)^2\varepsilon^2 2^n}$. Lemma 3.5 tells us that for every choice of $h_0,\ldots,h_{k-2}$ and states i,j of Q, the probability that $\boldsymbol{h}_{k-1}$ is *not* $(\varepsilon(b-1)/2^w, B_{i,j}^{h_0,\ldots,h_{k-2}})$-independent is at most

$$\frac{\varrho(B_{i,j}^{h_0,\ldots,h_{k-2}})2^{2w}}{\varepsilon^2(b-1)^2 2^n} .$$

Using a union bound, the probability there exists a pair of states i and j such that $\boldsymbol{h}_{k-1}$ is not $(\varepsilon(b-1)/2^w, B_{i,j}^{h_0,\ldots,h_{k-2}})$-independent is at most

$$\sum_{i,j} \frac{\varrho(B_{i,j}^{h_0,\ldots,h_{k-2}})2^{2w}}{\varepsilon^2(b-1)^2 2^n} = \frac{2^{2w}}{\varepsilon^2(b-1)^2 2^n}\sum_{i,j}\varrho(B_{i,j}^{h_0,\ldots,h_{k-2}}) \leq \frac{2^{3w}}{\varepsilon^2(b-1)^2 2^n} .$$

The last inequality follows from the fact that for every fixed i, the sets $B_{i,j}^{h_0,\ldots,h_{k-2}}$, where j ranges over the states of Q, partition $\{0,1\}^{bn}$ and thus their ϱ-values sum up to 1. A union bound shows us that the probability of either event 1 or 2 not holding is at most $2^{3w}k/(b-1)^2\varepsilon^2 2^n$, as claimed.

Let $(G_{k-1}(*, \boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-2}))^b$ denote the distribution over bitstrings of length $n \cdot b^k$ obtained by concatenating b *independent* samples drawn from $G_{k-1}(*, \boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-2})$. Conditioning on events 1 and 2 we now bound $\|Q(G_k(*, \boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-1})) - Q((U_n)^{b^k})\|$. Using the triangle inequality,

$$\|Q(G_k(*, \boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-1})) - Q((U_n)^{b^k})\| \leq \|Q(G_{k-1}(*, \boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-1})) - Q((G_{k-1}(*, \boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-2}))^b)\|$$
$$+ \|Q((G_{k-1}(*, \boldsymbol{h}_0,\ldots,\boldsymbol{h}_{k-2}))^b) - Q((U_n)^{b^k})\|.$$

Below we will bound the first term by $(b-1)\varepsilon$ and second term by $(b^k - b)\varepsilon$, proving that the hash functions are $((b^k - 1)\varepsilon, Q)$-good as claimed.

Consider the matrix $Q(G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}))$. By definition, entry (i,j) of the matrix is equal to $\mathbf{Pr}_{\boldsymbol{x} \sim U_n}[\boldsymbol{h}_{k-1}(\boldsymbol{x}) \in B_{i,j}^{\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}}]$. Now consider $Q((G_{k-1}(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}))^b)$ where entry (i,j) is

$$\Pr_{\boldsymbol{y}^{(0)}, \ldots, \boldsymbol{y}^{(b-1)} \sim U_n}[G_{k-1}(\boldsymbol{y}^{(0)}, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}) \circ \cdots \circ G_{k-1}(\boldsymbol{y}^{(b-1)}, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}) \text{ takes } Q \text{ from } i \text{ to } j] .$$

By definition, the above quantity is exactly $\varrho(B_{i,j}^{\boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}})$. Since event 2 holds, we have that the absolute value of every entry of the matrix $Q(G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})) - Q((G_{k-1}(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}))^b)$ is at most $\varepsilon(b-1)/2^w$. Since each row has at most 2^w entries we obtain

$$\|Q(G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})) - Q((G_{k-1}(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}))^b)\| \leq (b-1)\varepsilon .$$

To bound the second term we define the transition matrices $M = Q(G_{k-1}(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}))$ and $N = Q((U_n)^{b^{k-1}})$ that describe transition probabilities using G_{k-1} and uniform inputs, respectively. Since the b parts of the input are independent we can express the matrices in the second term as powers of M and N:

$$Q((G_{k-1}(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1}))^b) = M^b \quad \text{and} \quad Q((U_n)^{b^k}) = N^b .$$

We can now invoke Lemma 3.3. Since event 1 holds we have $\|M - N\| \leq (b^{k-1} - 1)\varepsilon$, and thus

$$\|Q((G_{k-1}(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-2}))^b) - Q((U_n)^{b^k})\| \leq b(b^{k-1} - 1)\varepsilon = (b^k - b)\varepsilon.$$

Conditioning on events 1 and 2 we thus have $\|Q(G_k(*, \boldsymbol{h}_0, \ldots, \boldsymbol{h}_{k-1})) - Q((U_n)^{b^k})\| \leq (b^k - 1)\varepsilon$. $\square$

The proof of Lemma 3.1 now follows by choosing a small constant c and setting $w = cn$ and $\varepsilon = 2^{-cn}/(b^k - 1)$ in the above lemma. From Proposition 1 of [Nis92], it follows that any space(cn) algorithm that uses nb^k uniform random bits, for $b^k \leq 2^{cn}$, can use the bits from the pseudorandom generator and with probability at least $1 - 2^{-cn}$, the total variation distance of the final state of the algorithm using pseudorandom bits to the final state of the algorithm using uniform random bits is at most 2^{-cn}.

4 Moment Estimation for $p > 2$

Using min-stability of exponential random variables, Andoni [And17] gave an algorithm (Algorithm 1) to estimate F_p moments in the stream. We state their result below and describe their algorithm.

Consider the vector x obtained at the end of the stream applying all the updates sequentially to the starting vector 0. As Algorithm 1 is linear, we have that the final state of the algorithm depends only on the final vector and not on the order of the updates. To estimate the ℓ_p norm of a d-dimensional vector x, the algorithm first samples d independent standard exponential random variables $\boldsymbol{E}_1, \ldots, \boldsymbol{E}_d$ and creates a random vector $\boldsymbol{z} \in \mathbb{R}^d$ such that $\boldsymbol{z}_i = \boldsymbol{E}_i^{-1/p} x_i$. Andoni shows that $\|\boldsymbol{z}\|_\infty = \Theta(\|x\|_p)$ and therefore estimating the value of the coordinate in $\boldsymbol{z}$ with maximum absolute value gives a constant factor approximation for $\|x\|_p$. Andoni further shows that with high probability there are only at most $O(\log d)^p$ coordinates in $\boldsymbol{z}$ with absolute value $\geq \|x\|_p/(c \log d)$ for a constant c hence showing that the vector $\boldsymbol{z}$ has only a few coordinates with large values. The final property of $\boldsymbol{z}$ that Andoni uses is that with high probability $\|\boldsymbol{z}\|_2^2 = O(d^{1-2/p}\|x\|_p^2)$.

Conditioned on the above properties of z, Andoni argues that if S is a CountSketch matrix with $O(d^{1-2/p} \log d)$ rows formed using $O(\log d)$-wise independent hash functions, then with high probability $\|Sz\|_\infty = \Theta(\|z\|_\infty) = \Theta(\|x\|_p)$. To implement the streaming algorithm in sub-linear space we cannot store the exponential random variables and the vector z in the stream. So Andoni derandomizes his algorithm by using a pseudorandom generator of Nisan and Zuckerman [NZ96] and shows that exponential random variables E_i generated using the pseudorandom bits are sufficient to ensure that the algorithm produces a constant factor approximation to $\|x\|_p$. Also note that the algorithm does not need to explicitly store the d dimensional vector z; it is enough just to update the $O(d^{1-2/p} \log d)$ dimensional vector f in the stream.

Although Andoni's algorithm is space optimal for linear sketches up to constant factors, the update time using the Nisan-Zuckerman pseudorandom generator is $\text{poly}(d)$. We now show that Andoni's algorithm can be derandomized using our HashPRG to obtain an algorithm that is both space optimal for linear sketches and has an update time of $O(1)$ in the Word RAM model with a word size $\Omega(\log d)$.

Algorithm 1: Andoni's Algorithm with Independent Exponentials [And17]

Input: $p > 2$, $d \in \mathbb{N}$, a stream of updates $(i_1, v_1), \ldots, (i_m, v_m) \in [d] \times \{-M, \ldots, M\}$ for $m, M = \text{poly}(d)$

Output: An approximation to $\|x\|_p$ where $x \in \mathbb{R}^d$ is defined by the stream of updates

1 $E_1, \ldots, E_d \leftarrow$ Independent standard exponential random variables;
2 $d' \leftarrow O(d^{1-2/p} \log d)$;
3 $z \leftarrow 0_d$, $f \leftarrow 0_s$;
4 $h \leftarrow O(\log d)$-wise independent hash function from $[d]$ to $[d']$;
5 $\sigma \leftarrow O(\log d)$-wise independent hash function from $[d]$ to $\{-1, +1\}$;
6 **for** $j = 1, \ldots, m$ **do**
7 $\quad z_{i_j} \leftarrow z_{i_j} + E_i^{-1/p} v_j$;
8 $\quad f_{h(i_j)} \leftarrow f_{h(i_j)} + \sigma(i) E_i^{-1/p} v_j$;
9 **end**
10 **return** $\|f\|_\infty$;

4.1   Discretizing the Exponentials

We first show that we can replace exponential random variables in Andoni's algorithm with a simple discrete random variable and obtain all the guarantees we stated above that z satisfies even with this discrete random variable. For now, assume that $p = 1$. We will later generalize the guarantees for all p.

Suppose $x \in \mathbb{R}^d$ with all the coordinates of x being integers with absolute values $\leq \text{poly}_1(d)$. We can assume x has no nonzero coordinates. Consider the discrete random variable E that takes values in the set $\{1, 2, \ldots, 2^M\}$ for some $M = O(\log d)$ satisfying $2^M \geq d \, \text{poly}_1(d)$. Let

$$\mathbf{Pr}[E = 2^j] = \begin{cases} 1/2^{j+1} & 0 \leq j < M \\ 1/2^M & j = M. \end{cases}$$

We call the random variable E a *discrete exponential*[5]. Note that we can sample the random variable E quite easily from a uniform random bitstring of length M just based on the position of the first

---

[5]While E models the inverse of a continuous exponential random variable, we use the term "discrete exponential" for brevity.

appearance of 1 in the random bitstring. We mainly use the following properties of $\boldsymbol{E}$: for any $t > 0$, $\mathbf{Pr}[\boldsymbol{E} \geq t] \leq 1/t$. The statement clearly holds for $t \leq 1$. For $t \geq 1$, let 2^j be such that $t \leq 2^j < 2t$. Now,

$$\mathbf{Pr}[\boldsymbol{E} \geq t] = \mathbf{Pr}[\boldsymbol{E} = 2^j] + \mathbf{Pr}[\boldsymbol{E} = 2^{j+1}] + \cdots$$
$$= \frac{1}{2^{j+1}} + \frac{1}{2^{j+2}} + \cdots = \frac{1}{2^j} \leq \frac{1}{t}.$$

Similarly for any $t \leq 2^M$, we have $\mathbf{Pr}[\boldsymbol{E} \geq t] \geq \min(1, 1/(2t))$. We now prove the following lemma.

Lemma 4.1. *Let $x \in \mathbb{R}^d$ be an arbitrary vector with integer entries of absolute value at most $\mathrm{poly}_1(d)$. Let $\boldsymbol{E}_1, \ldots, \boldsymbol{E}_d$ be independent discrete exponential random variables. Then,*

1. With probability $\geq 95/100$,

$$\frac{\|x\|_1}{16} \leq \max_i |x_i| \boldsymbol{E}_i \leq 50 \|x\|_1.$$

2. For any T,

$$\mathbf{E}[|\{ i \mid |x_i| \boldsymbol{E}_i \geq \|x\|_1 / T \}|] \leq T.$$

Proof. Without loss of generality, we can assume that all the coordinates of x are nonzero. Using the distribution of the random variable $\boldsymbol{E}_i$, we obtain that

$$\mathbf{Pr}\left[\boldsymbol{E}_i \geq 50 \frac{\|x\|_1}{|x_i|}\right] \leq \frac{|x_i|}{50 \|x\|_1}.$$

Thus, by a union bound, with probability $\geq 1 - 1/50$, for all i, $\boldsymbol{E}_i \leq 50 \|x\|_1 / |x_i|$. Hence with probability $\geq 49/50$, $\max_i |x_i| \boldsymbol{E}_i \leq 50 \|x\|_1$. Similarly for all i,

$$\mathbf{Pr}\left[\boldsymbol{E}_i \geq \frac{\|x\|_1}{16 |x_i|}\right] \geq \min(8 |x_i| / \|x\|_1, 1).$$

If there exists an index i such that $|x_i| \geq \|x\|_1 / 8$, then we already have that with probability 1, $\max_i |x_i| \boldsymbol{E}_i \geq \|x\|_1 / 8$. Now assume that for all i, $|x_i| \leq \|x\|_1 / 8$. Using the independence of $\boldsymbol{E}_i$'s, we obtain that with probability $\geq 99/100$, there exists an index i such that $\boldsymbol{E}_i \geq \|x\|_1 / (16 |x_i|)$ and therefore with probability $\geq 99/100$, $\max_i |x_i| \boldsymbol{E}_i \geq \|x\|_1 / 16$. Hence, overall with probability $\geq 95/100$, we have

$$\frac{\|x\|_1}{16} \leq \max_i |x_i| \boldsymbol{E}_i \leq 50 \|x\|_1.$$

Let $T > 0$ be arbitrary. For a fixed i, we have $\mathbf{Pr}[\boldsymbol{E}_i \geq \|x\|_1 / (|x_i| T)] \leq |x_i| T / \|x\|_1$. Thus,

$$\mathbf{E}[|\{ i \mid |x_i| \boldsymbol{E}_i \geq \|x\|_1 / T \}|]$$
$$= \sum_i \mathbf{Pr}[\boldsymbol{E}_i$$
$$\geq \|x\|_1 / (|x_i| T)] \leq T \sum_i |x_i| / \|x\|_1 = T. \qquad \square$$

The following lemma extends the above properties to all $p > 2$ along with an additional property that Andoni uses in his proof.

Lemma 4.2. *Let $p > 2$ be arbitrary and let $x \in \mathbb{R}^d$ be an arbitrary vector of integer entries with absolute value at most $(\mathrm{poly}_1(d))^{1/p}$. Let $\boldsymbol{E}_1, \ldots, \boldsymbol{E}_d$ be independent discrete exponential random variables. Then,*

1. With probability $\geq 95/100$,

$$\frac{\|x\|_p}{16^{1/p}} \leq \max_i |x_i| \boldsymbol{E}_i^{1/p} \leq 50^{1/p} \|x\|_p.$$

2. For any $T > 0$, with probability $\geq 95/100$,

$$|\{ i \mid |x_i| \boldsymbol{E}_i^{1/p} \geq \|x\|_p / T^{1/p} \}| \leq 20T.$$

3.

$$\mathbf{E}[\sum_{i=1}^d (\boldsymbol{E}_i^{1/p} x_i)^2] = O_p(d^{1-2/p} \|x\|_p^2).$$

Proof. The first two properties follow from applying the previous lemma to the vector $x^{(p)} \in \mathbb{R}^d$ defined as $x_i^{(p)} = |x_i|^p$. To prove the last property, we have

$$\mathbf{E}[\sum_{i=1}^d (\boldsymbol{E}_i^{1/p} x_i)^2] = \sum_{i=1}^d x_i^2 \, \mathbf{E}[\boldsymbol{E}_i^{2/p}].$$

Now $\mathbf{E}[\boldsymbol{E}_i^{2/p}] = \sum_{j=0}^{M-1} 2^{2j/p}/2^{j+1} + 2^{2M/p}/2^M \leq (1/2) \sum_{j=0}^{\infty} 2^{j(2/p-1)} + 2^{M(2/p-1)} \leq 1/(2-2^{2/p}) + 1$. Hence

$$\mathbf{E}[\sum_{i=1}^d (\boldsymbol{E}_i^{1/p} x_i)^2] \leq \|x\|_2^2 \left(\frac{1}{2 - 2^{2/p}} + 1 \right)$$
$$= O_p(d^{1-2/p} \|x\|_p^2)$$

where the constant we are hiding blows up as $p \to 2$. $\qquad\square$

4.2 F_p **Estimation With HashPRG**

In the previous section, we discussed some of the properties satisfied by the random vector $\boldsymbol{z}$ given by $\boldsymbol{z}_i := \boldsymbol{E}_i^{1/p} x_i$ for any arbitrary vector x with integer entries of absolute value at most $(\mathrm{poly}_1(d))^{1/p}$. Now we show that the properties are still satisfied even when the random variables $\boldsymbol{E}_i$ are generated using HashPRG thereby showing that Algorithm 2 outputs a constant factor approximation to $\|x\|_p$ with probability $\geq 7/10$. The success probability can be increased to $1 - \delta$ by running $O(\log 1/\delta)$ independent copies of the algorithm and reporting the median.

Theorem 4.3. *Let $p > 2$ be a parameter and for $m = O(\mathrm{poly}(d))$ let the vector $x := 0 \in \mathbb{R}^d$ receive a stream of m updates $(i_1, v_1), (i_2, v_2), \ldots, (i_m, v_m)$ with $|v_j| \leq \mathrm{poly}(d)$ for all j. On receiving an update (i_j, v_j), the vector x is modified as $x_{i_j} \leftarrow x_{i_j} + v_j$. The Algorithm 2 uses $O_p(d^{1-2/p} \log(d))$ words of space and at the end of the stream outputs a constant factor approximation to $\|x\|_p$ with probability $\geq 7/10$. Further each update to x is processed by the algorithm in $O(1)$ time in the Word RAM model on a machine with a word size $\Omega(\log d)$.*

23

Proof. First we condition on the event that the hash families from which $\boldsymbol{h}, \boldsymbol{\sigma}$ are drawn are $O(\log d)$-wise independent. From Theorem 2.2, the event holds with probability $\geq 99/100$. From Theorem 3.2, HashPRG with parameters $n = O(\log d)$, $b = d^\epsilon$ and $k = O(1/\epsilon)$ fools a Finite State Machine with $\mathrm{poly}(d)$ states. Further the seed for HashPRG can be stored using $O((1/\epsilon)d^\epsilon \log d) = o(d^{1-2/p})$ bits if $\epsilon < 1 - 2/p$. Theorem 2.2 shows that $\boldsymbol{h}$ and $\boldsymbol{\sigma}$ can be stored using $O(d^\epsilon) = o(d^{1-2/p})$ bits of space and for any $i \in [d]$, $\boldsymbol{h}(i)$ and $\boldsymbol{\sigma}(i)$ can be evaluated in $O_{1/\epsilon}(1)$ time. Therefore, each update in the stream is processed in $O_{1/\epsilon}(1)$ time.

Let x be the vector at the end of the stream and for $i \in [d]$, $\boldsymbol{E}_i$ be the discrete exponential random variable computed by the algorithm for coordinate i. Let $T := (O(\log d))^p$. Define a FSM Q_x with a start state and other states being defined by the tuple (i, j, t, v) with $1 \leq i \leq d$, $0 \leq j \leq d, 1 \leq v \leq \mathrm{poly}(d)$ and $t \in \{\, \mathsf{less}, \mathsf{correct}, \mathsf{more}\,\}$. The machine Q_x clearly has $\mathrm{poly}(d)$ states. The FSM being in a state (i, j, less, v) denotes that it has processed the coordinates $x_1, \ldots, x_i$ until now and found that

$$|\{\, i' \leq i \mid |x_{i'}|\boldsymbol{E}_{i'}^{1/p} \geq \|a\|_p f^{1/p} \,\}| = j,$$

$$\max_{i' \leq i} |x_{i'}|\boldsymbol{E}_{i'}^{1/p} < \frac{\|x\|_p}{16^{1/p}}, \text{ and}$$

$$\sum_{i' \leq i} \mathrm{round}(\boldsymbol{E}_{i'}^{1/p})^2 \cdot (x_{i'})^2 = v.$$

Here $\mathrm{round}(x)$ denotes x rounded to the nearest integer. As $\boldsymbol{E}_{i'}^{1/p} \geq 1$ if $x_{i'} \neq 0$, we have

$$(\boldsymbol{E}_{i'}^{1/p} x_{i'})^2 \leq \mathrm{round}(\boldsymbol{E}_{i'}^{1/p})^2 \cdot (x_{i'})^2 \leq 4(\boldsymbol{E}_{i'}^{1/p} x_{i'})^2.$$

Note that using the bound on the absolute values of entries in x, the value v can be at most $\mathrm{poly}(d)$. Similarly, the state is in $(i, j, \mathsf{correct}, v)$ if a condition similar to above holds but instead we have

$$\frac{\|x\|_p}{16^{1/p}} \leq \max_{i' \leq i} |x_{i'}|\boldsymbol{E}_{i'}^{1/p} \leq 50^{1/p}\|x\|_p.$$

and in (i, j, more, v) if $\max_{i' \leq i} |x_{i'}|\boldsymbol{E}_{i'}^{1/p} > 50^{1/p}\|x\|_p$. Here the value of $\boldsymbol{E}_i$ is assigned based on the bitstring corresponding to an edge in the FSM. It is clear that we can construct such a Finite State Machine. Let $\mathcal{E}$ denote the event that $\max_i |x_i|\boldsymbol{E}_i^{1/p} \in [\|x\|_p/16^{1/p}, 50^{1/p}\|x\|_p]$, $|\{\, i \mid |x_i|\boldsymbol{E}_i^{1/p} \geq \|x\|_p/T^{1/p} \,\}| \leq 20T$ and $\sum_{i \leq i} \mathrm{round}((\boldsymbol{E}_i^{1/p})x_i)^2 = O_p(d^{1-2/p}\|x\|_p^2)$. By Lemma 4.2, the final state distribution of the FSM using uniform random edge at every state satisfies the event $\mathcal{E}$ with probability $\geq 85/100$. By Theorem 3.2, if the random variables $\boldsymbol{E}_1, \ldots, \boldsymbol{E}_d$ are generated using the random string sampled from HashPRG, then with probability $\geq 1 - 1/\mathrm{poly}(d)$, the final state of FSM Q_x satisfies the event $\mathcal{E}$ with probability $\geq 8/10$.

Thus, with probability $\geq 8/10$, the implicit vector $\boldsymbol{z} \in \mathbb{R}^d$ in the algorithm defined as $\boldsymbol{z}_i := \boldsymbol{E}_i^{1/p} x_i$ satisfies all the properties Andoni requires of the vector obtained by multiplying coordinates of x with scaled exponential random variables. Hence, with probability $\geq 75/100$, the maximum absolute value of the coordinate in $\boldsymbol{f}$ obtained by sketching $\boldsymbol{z}$ with a CountSketch matrix is a constant factor approximation to $\|x\|_p$.

Thus overall the algorithm outputs a constant factor approximation to $\|x\|_p$ with probability $\geq 7/10$. $\qquad\square$

Algorithm 2: F_p moment estimation using HashPRG

Input: $p > 2$, $d \in \mathbb{N}$, a stream of updates $(i_1, v_1), \ldots, (i_m, v_m) \in [d] \times \{-M, \ldots, M\}$ for $m, M = \text{poly}(d)$

Output: An approximation to $\|x\|_p$ where $x \in \mathbb{R}^d$ is defined by the stream of updates

1   $\epsilon \leftarrow$ A constant smaller than $1 - 2/p$;

2   $s \leftarrow$ Pseudorandom string constructed using HashPRG with parameters
    $n = O(\log d), b = d^\epsilon, k = O(1/\epsilon)$;
    `// The string s is only implicitly stored using the corresponding hash functions`
        `and the random seed to the generator`

3   $d' \leftarrow O(d^{1-2/p} \log d)$;

4   $h \leftarrow O(\log d)$-wise independent hash function from $[d] \to [d']$;

5   $\sigma \leftarrow O(\log d)$-wise independent hash function from $[d] \to \{+1, -1\}$;
    `// Both h and σ are drawn from hash family in Theorem 2.2 so that they can be`
        `stored using` $O_{1/\epsilon}(d^\epsilon)$ `bits and evaluated in` $O_{1/\epsilon}(1)$ `time on any input`

6   $f \leftarrow 0_m$;
    `// Stream Processing Begins`

7   **for** $j = 1, \ldots, m$ **do**

8      $E_{i_j} \leftarrow \text{DiscreteExponential}(i_j\text{-th chunk of } s)$;

9      $f_{h(i_j)} \leftarrow f_{h(i_j)} + \sigma(i_j) \cdot \text{round}(E_{i_j}^{1/p}) \cdot v_j$;

10 **end**

11 **return** $\|f\|_\infty$;

4.3   Comparison with Andoni's use of Nisan-Zuckerman PRG

Andoni argues that his algorithm can be run in $O(d^{1-2/p} \log d)$ words of space using the Nisan-Zuckerman pseudorandom generator, which shows that an S space algorithm using $\text{poly}(S)$ random bits can be run with just $O(S)$ random bits. Nisan-Zuckerman's algorithm takes an $O(S)$ length uniformly random string and stretches it by a factor of $O(S^\gamma)$ $(0 < \gamma < 1)$ by computing by computing $O(S^\gamma)$ blocks of $O(S)$ bits each. Each of the $O(S^\gamma)$ blocks of pseudorandom bits takes time $\text{poly}(S)$ time to compute which in our case is $\text{poly}(d^{1-2/p} \log d)$ and hence prohibitive.

4.4 ℓ_p sampling

As another application of HashPRG, we give a simple ℓ_p sampling algorithm for $p > 2$. Assume the same turnstile stream setting. At the end of the stream, ℓ_p sampling asks to output a coordinate i of the underlying vector x such that probability of sampling i is proportional to $|x_i|^p / \|x\|_p^p$. The problem has been widely studied (see [CJ19, JW18] and references therein). The perfect ℓ_p sampling algorithm of [JW18] for $p \in (0, 2)$ uses the following property of exponential random variables: if $E_1, \ldots, E_d$ are independent standard exponential random variables and $i^* = \arg\max_i x_i / E_i^{1/p}$, then $\mathbf{Pr}[i^* = i] = |x_i|^p / \|x\|_p^p$. This distribution exactly corresponds to what ℓ_p sampling asks. To implement the algorithm in a turnstile stream using a small amount of space, they first scale the coordinates with exponentials and then sketch the scaled vector using a data structure called *count-max* and show that the count-max data structure allows to recover the max coordinate in the vector obtained by scaling x with exponential random variables. Finally, they derandomize their construction using a half-space fooling pseudorandom generator.

    We show that using HashPRG, for $p > 2$, we obtain ℓ_p samplers that have a very fast update

time. For simplicity, we discuss an algorithm that samples from the following distribution:

$$\mathbf{Pr}[i \text{ is sampled}] \geq \frac{1}{1+\varepsilon}\frac{|x_i|^p}{\|x\|_p^p} \pm 1/\operatorname{poly}(d).$$

In the definition approximate perfect samplers, it is required that the $\mathbf{Pr}[i \text{ is sampled}]$ is $(1 \pm \varepsilon)|x_i|^p/\|x\|_p^p$ up to an additive $1/\operatorname{poly}(d)$ error. We discuss the simpler version as ℓ_p samplers are not our main focus.

For this, we work with a finer approximation of exponential random variables than what we used for approximating F_p moments. Assume that given a block of $O(\log d)$ uniform random bits, there is a fast way to convert the random bits into a fine enough discretization of the exponential random variables such that all the following property of the exponential random variables hold for this discretization:

1. The probability that $i^* = i$ and $\boldsymbol{E}_{i^*}^{-1/p}|x_{i^*}| \geq (1+\varepsilon)^{-1/p}\max_{i' \neq i}\boldsymbol{E}_i^{-1/p}|x_i|]$ is at least $|x_i|^p/(1+\varepsilon)\|x\|_p^p$, and

2. with probability at least $1 - 1/\operatorname{poly}(d)$,

$$\sum_{i=1}^{d}(\boldsymbol{E}_i^{-1/p}|x_i|)^2 \leq d^{1-2/p}(\boldsymbol{E}_{i^*}^{-1/p}|x_{i^*}|)^2 \operatorname{polylog}(d).$$

It can be easily shown that both of these properties hold for *continuous* exponential random variables and hence they hold for a suitable discretization of the continuous exponential random variables. The above properties crucially depend on $\boldsymbol{E}_1, \ldots, \boldsymbol{E}_d$ being independent but we can derandomize them by using HashPRG. Fix a vector x and design the following FSM for x: the FSM goes through each coordinate of x sequentially. The FSM tracks $\max_i \boldsymbol{E}_i^{-1/p}|x_i|$, the coordinate attaining the max, $\sum_i \boldsymbol{E}_i^{-2/p}x_i^2$. Clearly, all these statistics can be tracked using an FSM with $\operatorname{poly}(d)$ states similar to how we derandomized the properties of exponential random variables for approximating F_p moments. Using $n = O(\log d)$ and $b = d^c$ ($c < 1 - 2/p$) for HashPRG, thus we obtain that if the random variables $\boldsymbol{E}_i$ are generated using the pseudorandom string sampled from HashPRG, then for all i,

$$\Pr_{(\boldsymbol{E}_1, \ldots, \boldsymbol{E}_d) \sim \text{HashPRG}}[i^* = i \text{ and } \boldsymbol{E}_{i^*}^{-1/p}|x_{i^*}| \geq (1+\varepsilon)^{-1/p}\max_{i' \neq i}\boldsymbol{E}_i^{-1/p}|x_i|]$$
$$\geq \frac{|x_i|^p}{(1+\varepsilon)\|x\|_p^p} - 1/\operatorname{poly}(d).$$

and with probability $\geq 1 - 1/\operatorname{poly}(d)$, $\sum_{i=1}^{d}(\boldsymbol{E}_i^{-1/p}|x_i|)^2 \leq d^{1-2/p}(\boldsymbol{E}_{i^*}^{-1/p}|x_{i^*}|)^2 \operatorname{polylog}(d)$. Let $\boldsymbol{f}$ be a vector such that $\boldsymbol{f}_i = \boldsymbol{E}_i^{-1/p}x_i$. Condition on both the events. Then we have that the largest coordinate in $\boldsymbol{f}$ is at least a $1/(1+\varepsilon)$ factor larger than the second largest coordinate and that $\|\boldsymbol{f}\|_2^2 = O(d^{1-2/p}\|\boldsymbol{f}\|_\infty^2)$. Now hashing the coordinates of $\boldsymbol{f}$ into a CountSketch data structure with $O(d^{1-2/p}\operatorname{polylog}(d)/\varepsilon^2)$ rows using $O(\log d)$ wise preserves the large coordinate of $\boldsymbol{f}$ using the analysis in [And17]. Using $O(1)$ independent CountSketch data structure and finding the coordinate $i \in [d]$ which hashes to the max bucket of the CountSketch data structure in each of the $O(1)$ repetitions, we can extract the coordinate i and output it as the ℓ_p sample.

Note that the update time is $O(1)$ as a block of random bits from HashPRG with parameter $k = d^c$ can be obtained in $O(1)$ time and then the time to evaluate the hash functions of the CountSketch data structure is $O(1)$ when using the constructions from Theorem 2.2. The overall space complexity of the data structure is $O(d^{1-2/p}\operatorname{polylog}(d)/\varepsilon^2)$ bits. However, we note that the final step of computing which $i \in [d]$ hashes to the max coordinate in all $O(1)$ copies of the CountSketch data structure takes $O(d)$ time.

26

5 Moment Estimation for $0 < p < 2$

We assume as usual that a vector $x \in \mathbb{R}^d$ is being maintained in the stream. The vector x is initialized to 0 and then receives m updates of the form $(i, v) \in [d] \times \{-M, \ldots, M\}$ where upon receiving (i, v), we update $x_i \leftarrow x_i + v$. We assume that both $m, M \leq \text{poly}(d)$. At the end of the stream, we want to approximate $\|x\|_p^p$ up to a $1 \pm \varepsilon$ multiplicative factor. For ε such that $1/\sqrt{d} \leq \varepsilon \leq 1/d^c$ for a small enough constant, we show that the algorithm of [KNPW11] can be implemented in space of $O(\varepsilon^{-2} \log(d))$ *bits* of space and $O(\log d)$ update time per stream element. We measure the time complexity of the update algorithm in the Word RAM model with a word size of at least $\Omega(\log d)$. We first give a high level overview of the moment estimation algorithm of [KNPW11]. Throughout the section, we assume that $1/\sqrt{d} \leq \varepsilon \leq 1/d^c$.

5.1 Overview of Moment Estimation Algorithm of [KNPW11]

Their 1-pass algorithm is based on the Geometric Mean estimator of Li [Li08]. Li gives an estimator to compute F_p moment of a vector using p-stable random variables. Let x be a fixed d dimensional vector and $\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3 \in \mathbb{R}^d$ be random vectors with independent p-stable random variables. Then, Li showed that the estimator given by

$$\textbf{Est} = \frac{(|\langle \boldsymbol{v}_1, x \rangle||\langle \boldsymbol{v}_2, x \rangle||\langle \boldsymbol{v}_3, x \rangle|)^{p/3}}{\left(\frac{2}{\pi}\Gamma(p/3)\,\Gamma(2/3)\sin(\pi p/6)\right)^3} \tag{2}$$

satisfies $\mathbf{E}_{\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3}[\textbf{Est}] = \|x\|_p^p$ and $\mathbf{Var}_{\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3}[\textbf{Est}] = O(\|x\|_p^{2p})$. A $1 \pm \varepsilon$ approximate estimator can be obtained by averaging $O(1/\varepsilon^2)$ independent copies of the estimator but it makes the update time $\Omega(\varepsilon^{-2})$ which is prohibitive. To decrease the variance of the estimator, [KNPW11] hash the coordinates of x into buckets and estimate the contribution of each bucket to the F_p moment of x separately. When there are heavy coordinates in the vector x, variance of the estimator may still be too large. Therefore, they estimate the contribution of the *heavy* elements using a different data structure they call HighEnd and use the Li's estimator to only estimate the contribution of the light elements using a data structure they call LightEstimator. We show that when $(1/\sqrt{d}) \leq \varepsilon \leq 1/d^c$ for a constant $c < 1/2$, we can implement their algorithm using $O(\varepsilon^{-2} \log d)$ bits of space and an update time of $O(\log d)$ per each element in the stream in the Word RAM model.

5.2 The HighEnd Data Structure

As described above, the algorithm of [KNPW11] estimates the F_p moment of heavy entries and light entries separately. Their heavy entry moment estimation method, at the end of the stream, takes in $L \subseteq [d]$ satisfying the following conditions:

1. $L \supseteq \{i \in [d] \mid |x_i|^p \geq \alpha\|x\|_p^p\}$,

2. if $i \in L$, then $|x_i|^p \geq (\alpha/C)\|x\|_p^p$ for a constant $C \geq 1$, and

3. we know the sign of x_i for all $i \in L$.

We show in Appendix B how a set L satisfying the above properties can be computed in a turnstile stream using $O(\alpha^{-1} \log^2 d)$ bits of space and $O(\log d)$ update time per stream element. We use the CountSketch data structure [CCF04] along with the ExpanderSketch data structure [JST11] to obtain L. Note that the set L has size at most $O(1/\alpha)$.

Now we state the guarantees of the HighEnd data structure. We first define BasicHighEnd data structure and then define HighEnd by taking independent copies of the BasicHighEnd data structure. Let α be such that $1/\alpha = O(1/\varepsilon^2)$ and let $s = \Theta(1/\alpha)$ be a large enough power of 2. Let $\boldsymbol{h} : [d] \to [s]$ is picked at random from an r_h-wise independent hash family for $r_h = \Theta(\log 1/\alpha)$. Let $r = \Theta(\log 1/\varepsilon)$ be a sufficiently large power of 2. Let $\boldsymbol{g} : [d] \to [r]$ be drawn at random from an r_g-wise independent hash family for $r_g = r$. For each $i \in [d]$, we associate a random complex root of unity given by $\exp(\mathrm{i}2\pi\boldsymbol{g}(i)/r)$ where i denotes $\sqrt{-1}$. We initialize s counters $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_s$ to 0. Given an update of the form (i, v), we set $\boldsymbol{b}_{\boldsymbol{h}(i)} \leftarrow \boldsymbol{b}_{\boldsymbol{h}(i)} + \exp(\mathrm{i}2\pi\boldsymbol{g}(i)/r)v$.

The HighEnd data structure is defined by taking T independent copies of BasicHighEnd data structure with $T = O(\max(\log(1/\varepsilon), \log(1/\alpha))) = O(\log d)$. Each of the copies of the BasicHighEnd data structure is updated upon receiving an update (i, v) in the stream. Let $(\boldsymbol{h}^1, \boldsymbol{g}^1), \ldots, (\boldsymbol{h}^T, \boldsymbol{g}^T)$ be the hash functions corresponding to each of the BasicHighEnd data structures.

It is argued in [KNPW11] that storing the coefficients of complex numbers up to a precision of $\Theta(\log(d))$ bits suffices if the number of updates is poly(d), the magnitude of each update is bounded by poly(d) and $1/\sqrt{d} \leq \varepsilon \leq 1/d^c$. Thus the space complexity of HighEnd data structure (excluding the space required for storing the hash functions) is

$$O\left(\alpha^{-1}(\log d)^2\right) \text{ bits.}$$

By Theorem 2.2, for t a large enough constant we can construct random hash families $\mathcal{H} = \{h : [d] \to [s]\}$ and $\mathcal{G} = \{g : [d] \to [r]\}$ such that with probability $\geq 1 - 1/\text{poly}(d)$, the hash families $\mathcal{H}$ and $\mathcal{G}$ are r_h and r_g wise independent respectively. Now, if $\boldsymbol{h}^1, \ldots, \boldsymbol{h}^T$ are sampled independently from $\mathcal{H}$ and $\boldsymbol{g}^1, \ldots, \boldsymbol{g}^T$ are sampled independently from $\mathcal{G}$, they can be evaluated on any input in $O(1)$ time and therefore the update time per each stream element is $O(T) = O(\log d)$ in the Word RAM model. Each hash function $\boldsymbol{h}^i$ and $\boldsymbol{g}^i$ can be stored using $O(d^{c/2})$ bits and therefore the space required to store the hash functions necessary for the HighEnd data structure is $o(d^c) = o(\varepsilon^{-2})$ bits.

At the end of stream, by Theorem 11 of [KNPW11] we can use the HighEnd data structure to compute a value Ψ such that with probability $\geq 7/8$,

$$|\Psi - \|x_L\|_p^p| \leq O(\varepsilon)\|x\|_p^p.$$

We then have the following lemma:

Lemma 5.1. *Given $1/\sqrt{d} \leq \varepsilon \leq 1/d^c$ for a small enough constant, and α such that $1/\alpha = O(1/\varepsilon^2)$, there is a streaming algorithm that takes $O(\alpha^{-1}\log^2(d))$ bits of space and has an update time of $O(\log d)$ per stream element in the Word RAM model satisfying:*

1. *The algorithm outputs a set $L \subseteq [d]$ satisfying all the three properties stated above with probability $\geq 9/10$.*

2. *Conditioned on the list L satisfying those properties, the algorithm outputs a value Ψ such that with probability $\geq 65/100$,*

$$|\Psi - \|x_L\|_p^p| \leq O(\varepsilon)\|x\|_p^p.$$

By taking median of Ψ output by $O(1)$ independent instances of the HighEnd data structure, we have that conditioned on L satisfying all the properties, we have an estimate of $\|x_L\|_p^p$ with an additive error of $O(\varepsilon)\|x\|_p^p$ with probability $\geq 99/100$. Hence, by a union bound with probability $\geq 8/10$, the algorithm outputs both a good list L and a value Ψ satisfying $|\Psi - \|x_L\|_p^p| \leq O(\varepsilon)\|x\|_p^p$.

Note that for $\alpha = \varepsilon^2 \log(d)$, the algorithm uses $O(\varepsilon^{-2} \log d)$ *bits* of space. For this setting of α, we will now use the LightEstimator data structure of [KNPW11] to estimate $\|x_{[d]\setminus L}\|_p^p$. We will now describe the LightEstimator data structure and how it can be derandomized using HashPRG.

5.3 The LightEstimator Data Structure

Algorithm 3: LightEstimator using Independent p-stable random variables

---

Input: A parameter $p \in (0, 2)$, accuracy parameter ε such that $1/\sqrt{d} \leq \varepsilon \leq 1/d^c$ for a small enough constant c, a parameter α such that $1/\alpha = O(1/\varepsilon^2)$, a stream of updates $(i_1, v_1), \ldots, (i_m, v_m) \in [d] \times \{-M, \ldots, M\}$, a list $L \subseteq [d]$ of heavy coordinates revealed at the end of the stream

Output: An estimate of $\|x_{[d]\setminus L}\|_p^p$

1 $s \leftarrow O(1/\alpha)$;
2 $\boldsymbol{h} \leftarrow$ A hash function sampled from the construction in Theorem 2.3;
3 For $b \in [s]$ and $[j] \in [3]$, initialize $\boldsymbol{B}_{b,j} \leftarrow 0$;
4 For $i \in [d]$ and $[j] \in [3]$, let $\boldsymbol{A}_{i,j}$ be an independent p-stable random variable;
5 **for** $j = 1, \ldots, m$ **do**
6 $\quad$ $\boldsymbol{B}_{\boldsymbol{h}(i_j),1} \leftarrow \boldsymbol{B}_{\boldsymbol{h}(i_j),1} + \boldsymbol{A}_{i_j,1} v_j$;
7 $\quad$ $\boldsymbol{B}_{\boldsymbol{h}(i_j),2} \leftarrow \boldsymbol{B}_{\boldsymbol{h}(i_j),2} + \boldsymbol{A}_{i_j,2} v_j$;
8 $\quad$ $\boldsymbol{B}_{\boldsymbol{h}(i_j),3} \leftarrow \boldsymbol{B}_{\boldsymbol{h}(i_j),3} + \boldsymbol{A}_{i_j,3} v_j$;
9 **end**
$\quad$ // The set L is revealed to the algorithm
10 **for** $b = 1, \ldots, s$ **do**
11 $\quad$ **if** $b \notin \boldsymbol{h}(L)$ **then**
12 $\quad\quad$ $\mathbf{Est}(b) \leftarrow \frac{(|\boldsymbol{B}_{b,1}| \cdot |\boldsymbol{B}_{b,2}| \cdot |\boldsymbol{B}_{b,3}|)^{p/3}}{((2/\pi)\Gamma(p/3)\Gamma(2/3)\sin(\pi p/6))^3}$;
13 $\quad$ **else**
14 $\quad\quad$ $\mathbf{Est}(j) \leftarrow 0$;
15 $\quad$ **end**
16 **end**
17 $\Phi \leftarrow \frac{s}{s-|\boldsymbol{h}(L)|} \sum_{b=1}^{s} \mathbf{Est}(b)$;
18 **return** Φ;

As seen previously, the HighEnd data structure lets us compute the F_p moment of all the elements from L. We use the LightEstimator data structure to approximate the F_p moment of all the *light* elements i.e., the coordinates of x not in L.

Assume that at the end of processing the stream we are given a set $L \subseteq [d]$, $|L| \leq 2/\alpha$ and for all $i \notin L$, $|x_i|^p < \alpha \|x\|_p^p$.

Let $s = \Theta(1/\alpha) \geq 10|L|$ be a large enough power of 2. For $i \in [d]$ and $j \in [3]$, let $\boldsymbol{A}_{i,j}$ denote an independent p-stable random variable. Let $\boldsymbol{h} : [d] \rightarrow [s]$ be a hash function drawn using the construction of Theorem 2.3 with parameters $z = \lceil 2/\alpha + 2 \rceil$ and a large enough constant C. For $i \in [s]$ and $j \in [3]$, initialize the counters $\boldsymbol{B}_{i,j} = 0$. On receiving an update (i, v) in the stream, for each $j \in [3]$, update $\boldsymbol{B}_{\boldsymbol{h}(i),j} \leftarrow \boldsymbol{B}_{\boldsymbol{h}(i),j} + \boldsymbol{A}_{i,j} \cdot v$. As argued in [KNPW11], we need to store the values $\boldsymbol{B}_{i,j}$ only up to a precision of $\Theta(\log d)$ bits. Hence the space complexity of the LightEstimator data structure (excluding the space required to store/generate $\boldsymbol{A}_{i,j}$) is $O(\alpha^{-1} \log d)$ bits.

At the end of the stream, we receive the set $L \subseteq [d]$ of heavy hitters as specified in the previous

section and we will then compute an estimator for $\|x_{[d]\setminus L}\|_p^p$ as follows: for each $b \in [s] \setminus \boldsymbol{h}(L)$, define

$$\mathbf{Est}(b) := \frac{(|\boldsymbol{B}_{b,1}| \cdot |\boldsymbol{B}_{b,2}| \cdot |\boldsymbol{B}_{b,3}|)^{p/3}}{\left(\frac{2}{\pi}\Gamma(p/3)\,\Gamma(2/3)\sin(\pi p/6)\right)^3}.$$

Let $\boldsymbol{h}(L) = \{\,\boldsymbol{h}(i) \mid i \in L\,\}$ denote the buckets to which the elements of L are hashed into. Define the following estimator

$$\Phi = \frac{s}{s - |\boldsymbol{h}(L)|} \sum_{b \in [s]\setminus \boldsymbol{h}(L)} \mathbf{Est}(b).$$

It is shown in [NW10] for $p = 1$ and extended to all $0 < p < 2$ in [KNPW11] that

$$\underset{h,A}{\mathbf{E}}[\Phi] = (1 \pm \alpha^{10})\|x_{[d]\setminus L}\|_p^p. \tag{3}$$

We extend their analysis and show an upper bound on $\mathbf{Var}_{h,A}[\Phi]$.

Remark 5.2. Lemma 7 in [NW10] and Theorem 15 in [KNPW11] state that $\mathbf{E}_{h,A}[\Phi] = (1 \pm O(\varepsilon))\|x_{[d]\setminus L}\|_p^p$ for $p = 1$ and $0 < p < 2$ respectively. It can be seen from the proof of Lemma 7 in [NW10], we can obtain the above stronger result by just picking C large enough when constructing the hash family $\mathcal{H}$ using the construction in Theorem 2.3. A similar argument works to extend for $0 < p < 2$.

5.3.1 Analysis of the Estimator

Claim 5.3.

$$\underset{h,A}{\mathbf{Var}}[\Phi] \le O(\alpha)\|x\|_p^{2p}.$$

Proof. For simplicity, we define $\mathbf{Est}(b) = 0$ for all $b \in \boldsymbol{h}(L)$. We have

$$\underset{h,A}{\mathbf{E}}[\Phi^2] = \underset{h,A}{\mathbf{E}}\left[\left(\frac{s}{s - |\boldsymbol{h}(L)|}\sum_{b=1}^s \mathbf{Est}(b)\right)^2\right]$$

$$= \underset{h}{\mathbf{E}}\left[\left(\frac{s}{s - |\boldsymbol{h}(L)|}\right)^2\left(\sum_{b=1}^s \underset{A|h}{\mathbf{E}}(\mathbf{Est}(b))^2 + \sum_{b \ne b'} \underset{A|h}{\mathbf{E}}(\mathbf{Est}(b))(\mathbf{Est}(b'))\right)\right]$$

Now, for $b \notin \boldsymbol{h}(L)$, using the variance and mean of Li's estimator (2),

$$\underset{A|h}{\mathbf{E}}(\mathbf{Est}(b))^2 = \underset{A|h}{\mathbf{Var}}(\mathbf{Est}(b)) + \left(\underset{A|h}{\mathbf{E}}\mathbf{Est}(b)\right)^2 = O(1)\left(\sum_{j:\boldsymbol{h}(j)=b} |x_j|^p\right)^2 \tag{4}$$

and as p-stable random variables denoted by $\boldsymbol{A}$ are independent, for $b \ne b'$ with $b, b' \notin \boldsymbol{h}(L)$,

$$\underset{A|h}{\mathbf{E}}(\mathbf{Est}(b))(\mathbf{Est}(b')) = \underset{A|h}{\mathbf{E}}(\mathbf{Est}(b))\underset{A|h}{\mathbf{E}}(\mathbf{Est}(b'))$$

$$= \sum_{j:\boldsymbol{h}(j)=b} |x_j|^p \sum_{j:\boldsymbol{h}(j)=b'} |x_j|^p. \tag{5}$$

Hence,

$$
\mathop{\mathbf{E}}_{h,A}[\Phi^2] = \mathop{\mathbf{E}}_{h}\left[\left(\frac{s}{s-|\boldsymbol{h}(L)|}\right)^2 \left(O(1)\sum_{b\notin\boldsymbol{h}(L)}(\sum_{i:\boldsymbol{h}(i)=b}|x_i|^p)^2 + \sum_{b\neq b':b,b'\notin\boldsymbol{h}(L)}(\sum_{i:\boldsymbol{h}(i)=b}|x_i|^p)(\sum_{i:\boldsymbol{h}(i)=b'}|x_i|^p)\right)\right]
$$

$$
\leq \mathop{\mathbf{E}}_{h}\left[\left(\frac{s}{s-|\boldsymbol{h}(L)|}\right)^2 \left(O(1)\sum_{b\notin\boldsymbol{h}(L)}(\sum_{i:\boldsymbol{h}(i)=b}|x_i|^p)^2 + \left(\sum_{i:\boldsymbol{h}(i)\notin\boldsymbol{h}(L)}|x_i|^p\right)^2\right)\right]
$$

$$
\leq O(1)\mathop{\mathbf{E}}_{h}\left[\sum_{b\notin\boldsymbol{h}(L)}(\sum_{i:\boldsymbol{h}(i)=b}|x_i|^p)^2\right] + \mathop{\mathbf{E}}_{h}\left[\sum_{i\neq i'}\left(\frac{s}{s-|\boldsymbol{h}(L)|}\right)^2 \mathbf{1}[\boldsymbol{h}(i),\boldsymbol{h}(i')\notin\boldsymbol{h}(L)]|x_i|^p|x_{i'}|^p\right]
$$

where the last inequality follows from the fact that $|\boldsymbol{h}(L)| \leq |L| \leq s/10$. We first bound the second term. For any $i, i' \notin L$, with probability $1-\rho$, the hash function $\boldsymbol{h}$ is drawn from a hash family that is $|L|+2$-wise independent when restricted to the set $L\cup\{i,i'\}$ when restricted to $L\cup\{i,i'\}$. We can make $\rho \leq \varepsilon^C$ for any constant C by setting C large enough while sampling $\boldsymbol{h}$ from the hash family in Theorem 2.3. Let the event that the hash family is $|L|+2$-wise independent with respect to $L\cup\{i,i'\}$ be called Good. Conditioned on this event and the size $|\boldsymbol{h}(L)|$,

$$
\mathop{\mathbf{E}}_{h|\mathsf{Good},|\boldsymbol{h}(L)|}[\mathbf{1}[\boldsymbol{h}(i),\boldsymbol{h}(i')\notin\boldsymbol{h}(L)]] = \left(\frac{s-|\boldsymbol{h}(L)|}{s}\right)^2 \tag{6}
$$

which gives

$$
\mathop{\mathbf{E}}_{h}\left[\sum_{i\neq i'}\left(\frac{s}{s-|\boldsymbol{h}(L)|}\right)^2 \mathbf{1}[\boldsymbol{h}(i),\boldsymbol{h}(i')\notin\boldsymbol{h}(L)]|x_i|^p|x_{i'}|^p\right]
$$

$$
\leq \mathop{\mathbf{E}}_{h|\mathsf{Good}}\left[\sum_{i\neq i'}\left(\frac{s}{s-|\boldsymbol{h}(L)|}\right)^2 \mathbf{1}[\boldsymbol{h}(i),\boldsymbol{h}(i')\notin\boldsymbol{h}(L)]|x_i|^p|x_{i'}|^p\right]
$$

$$
+ \mathop{\mathbf{E}}_{h|\neg\mathsf{Good}}\left[\sum_{i\neq i'}\left(\frac{s}{s-|\boldsymbol{h}(L)|}\right)^2 \mathbf{1}[\boldsymbol{h}(i),\boldsymbol{h}(i')\notin\boldsymbol{h}(L)]|x_i|^p|x_{i'}|^p\right] \times \mathop{\mathbf{Pr}}_{h}[\neg\mathsf{Good}]
$$

$$
\leq \|x_{[d]\setminus L}\|_p^{2p} + 2\|x_{[d]\setminus L}\|_p^{2p}\rho = (1+2\rho)\|x_{[d]\setminus L}\|_p^{2p}. \tag{7}
$$

Here we used (6) to cancel out the $s^2/(s-|\boldsymbol{h}(L)|)^2$ factor in the expectation and we used that $s/(s-|\boldsymbol{h}(L)|) \leq 10/9$ with probability 1 as $|\boldsymbol{h}(L)| \leq |L| \leq s/10$. Thus

$$
\mathop{\mathbf{E}}_{h,A}[\Phi^2] \leq O(1)\mathop{\mathbf{E}}_{h}\left[\sum_{b\notin\boldsymbol{h}(L)}(\sum_{i:\boldsymbol{h}(i)=b}|x_j|^p)^2\right]
$$

$$
+ (1+2\rho)\|x_{[d]\setminus L}\|_p^{2p}. \tag{8}
$$

Now, we bound the first term.

$$
\mathop{\mathbf{E}}_{h}\left[\sum_{b\notin\boldsymbol{h}(L)}\left(\sum_{i:\boldsymbol{h}(i)=b}|x_j|^p\right)^2\right] = \mathop{\mathbf{E}}_{h}\left[\sum_{b\notin\boldsymbol{h}(L)}\sum_{i:\boldsymbol{h}(i)=b}|x_i|^{2p} + \sum_{b\notin\boldsymbol{h}(L)}\sum_{i\neq i':\boldsymbol{h}(i)=\boldsymbol{h}(i')=b}|x_i|^p|x_{i'}|^p\right]
$$

$$
\leq \|x_{[d]\setminus L}\|_{2p}^{2p} + \mathop{\mathbf{E}}_{h}\left[\sum_{b}\sum_{i\neq i'}\mathbf{1}[b\notin\boldsymbol{h}(L),\boldsymbol{h}(i)=\boldsymbol{h}(i')=b]|x_i|^p|x_{i'}|^p\right].
$$

31

Again, for $i, i' \notin L$, with probability $1 - \rho$, the hash function $\boldsymbol{h}$ is drawn from a hash family that is $|L| + 2$-wise independent on the set $L \cup \{i, i'\}$. Conditioning on that event, for any $b \in [s]$,

$$\mathbf{Pr}[\mathbf{1}[b \notin \boldsymbol{h}(L), \boldsymbol{h}(j) = b, \boldsymbol{h}(j') = b]] \leq O(1/s^2).$$

Hence using an argument similar to the one in proving (7), we have

$$\mathbf{E}_{\boldsymbol{h}} \left[\sum_{b \notin \boldsymbol{h}(L)} \left(\sum_{j : \boldsymbol{h}(j) = b} |x_j|^p \right)^2 \right]$$
$$= \|x_{[d] \setminus L}\|_{2p}^{2p} + O(1/s + \rho)\|x_{[d] \setminus L}\|_p^{2p}.$$

Therefore

$$\mathbf{Var}_{\boldsymbol{h}, \boldsymbol{A}}[\Phi] \leq O(1)\|x_{[d] \setminus L}\|_{2p}^{2p} + O(1/s + \rho)\|x_{[d] \setminus L}\|_p^{2p} + (1 + O(\rho))\|x_{[d] \setminus L}\|_p^{2p} - (1 - \alpha^{10})^2\|x_{[d] \setminus L}\|_p^{2p}$$

$$\leq O(1)\|x_{[d] \setminus L}\|_{2p}^{2p} + O(1/s + \rho)\|x_{[d] \setminus L}\|_p^{2p} + O(\rho + \alpha^{10})\|x_{[d] \setminus L}\|_p^{2p}.$$

Now, $\|x_{[d] \setminus L}\|_{2p}^{2p} \leq \alpha\|x\|_p^p\|x_{[d] \setminus L}\|_p^p$ using the fact that all coordinates in $x_{[d] \setminus L}$ have pth power at most $\alpha\|x\|_p^p$. As $\rho \leq \varepsilon^C$ and $s = \Theta(1/\alpha)$,

$$\mathbf{Var}_{\boldsymbol{h}, \boldsymbol{A}}[\Phi] \leq O(\alpha)\|x\|_p^{2p}. \qquad \square$$

Let $\Phi_1, \ldots, \Phi_T$ be T independent copies of the estimator Φ for $T = \Theta(\log(1/\varepsilon)) = \Theta(\log d)$. As $\alpha = \varepsilon^2 \log(d)$, $\mathbf{Var}[(\Phi_1 + \cdots + \Phi_T)/T] \leq O(\varepsilon^2 \log(d)/T)\|x\|_p^{2p} \leq (\varepsilon^2/100)\|x\|_p^{2p}$. By Chebyshev inequality, with probability $\geq 99/100$, $\bar{\Phi} := \frac{\Phi_1 + \cdots + \Phi_T}{T}$ lies in the interval

$$[(1 - \alpha^{10})\|x_{[d] \setminus L}\|_p^p - (\varepsilon/10)\|x\|_p^p, (1 + \alpha^{10})\|x_{[d] \setminus L}\|_p^p + (\varepsilon/10)\|x\|_p^p].$$

Hence, using $O(\log d)$ independent instantiations of the LightEstimator data structure, we can obtain an estimate $\bar{\Phi}$ for $\|x_{[d] \setminus L}\|_p$ so that with probability $\geq 7/10$, $\Psi + \bar{\Phi} \in [(1 - O(\varepsilon))\|x\|_p^p, (1 + O(\varepsilon))\|x\|_p^p]$. Throughout the analysis we assumed that the p-stable random variables $\boldsymbol{A}_{i,j}$ are independent. To implement the algorithm in sub-linear space, we derandomize p-stable random variables using HashPRG.

5.3.2 Derandomizing p-stable random variables using HashPRG

Lemma 5.4. *Given $p \in (0, 2)$, an accuracy parameter ε such that $1/\sqrt{d} \leq \varepsilon \leq 1/d^c$ for a constant c, a parameter α such that $1/\alpha \leq O(1/\varepsilon^2)$ and a stream of updates $(i_1, v_1), \ldots (i_m, v_m) \in [d] \times \{-M, \ldots, M\}$ for $m, M \leq \mathrm{poly}(d)$ to the vector x, there is a streaming algorithm that uses $O(\alpha^{-1} \log d)$ bits of space and an update time of $O(\log d)$ per stream element in the Word RAM model. At the end of processing the stream, the algorithm takes in a set $L \subseteq [d]$ of heavy coordinates satisfying for all $i \notin L$, $|x_i|^p \leq \alpha\|x\|_p^p$ and outputs a value Φ satisfying*

$$\mathbf{E}[\Phi] = (1 \pm \varepsilon^C)\|x_{[d] \setminus L}\|_p^p \pm \frac{1}{\mathrm{poly}(d)}.$$

and

$$\mathbf{Var}[\Phi] \leq O(\alpha)\|x\|_p^p.$$

The algorithm in above lemma is given by a modified version of Algorithm 3. Instead of using independent p-stable random variables $\boldsymbol{A}_{i,j}$, the algorithm uses HashPRG to obtain a pseudorandom string and uses the pseudorandom bits to compute the p-stable random variables.

Proof. Consider one instance of LightEstimator data structure as in Algorithm 3. It consists of the following objects:

1. A hash function $\boldsymbol{h} : [d] \to [s]$ for $s = O(1/\alpha)$ drawn from a hash family as stated in Theorem 2.3 with parameter $z = O(1/\alpha)$ and C being a large enough constant.

2. There is a table of $s = O(1/\alpha)$ counters each maintained with a precision of $O(\log d)$ bits.

3. The p-stable random variables $\boldsymbol{A}_{i,j}$ for $i \in [d]$ and $j = 1, 2, 3$.

The hash function $\boldsymbol{h}$ can be stored using $O(\alpha^{-1} \log d)$ bits by Theorem 2.3. The counters can be maintained using $O(\alpha^{-1} \log d)$ bits as well. So, we are left with derandomizing the p-stable random variables.

The algorithm overall needs $O(d \log d)$ uniform random bits to generate three p-stable random variables for each of the coordinates. We use HashPRG to obtain the pseudorandom bits and use them to generate p-stable random variables. We critically use the fact that our analysis of the estimator Φ as described in the previous section needs to use only the mean and variance of Φ to show that we only have to "fool" multiple $O(\log d)$ space algorithms and hence using HashPRG as described in Theorem 3.2 is enough to generate the pseudorandom bits to compute the p-stable random variables. Further, for each update in the stream, the necessary block of pseudorandom bits can be generated in $O(1)$ time for each update. The space required to store the randomness necessary for the pseudo random generator is $O(d^\epsilon) = O(1/\varepsilon^2)$ bits when $\epsilon \leq 2c$ where $\varepsilon \leq 1/d^c$.

So, the overall algorithm on each update (i, v) is as follows: We use $\boldsymbol{h}$ to hash i into one of the $O(1/\alpha)$ buckets. Note that $\boldsymbol{h}(i)$ can be computed in $O(1)$ time. Using the value of i, we generate a block of $O(\log d)$ pseudorandom bits from the pseudorandom generator and then use the bits to compute 3 samples from a p-stable distribution. Let the samples be $\boldsymbol{A}_{i,1}, \boldsymbol{A}_{i,2}, \boldsymbol{A}_{i,3}$. We update the counters $\boldsymbol{B}_{h(i),j} \leftarrow \boldsymbol{B}_{h(i),j} + \boldsymbol{A}_{i,j} v$ for $j \in [3]$. As discussed, generating the pseudorandom bits and updating the counters can be performed in $O(1)$ time (assuming the pseudorandom bits can be converted to samples from a p-stable distribution in $O(1)$ time).

Let $\boldsymbol{\gamma} \sim$ HashPRG be a string sampled from the pseudorandom generator. Let $\boldsymbol{A}_{\boldsymbol{\gamma}}$ denote the p-stable random variables generated using $\boldsymbol{\gamma}$. Let $\boldsymbol{A}$ denote p-stable random variables generated using a uniform random string of bits. Hence the random variables $\boldsymbol{A}_{i,j}$ are independent. Fix a hash function $\boldsymbol{h}$, a vector x and a set of heavy elements L. Now consider the estimator we use to estimate the F_p moments of the light elements $[d] \setminus L$:

$$\Phi_{\boldsymbol{h},\boldsymbol{A}} = \frac{s}{(s - |\boldsymbol{h}(L)|)\theta} \times \sum_{b \in [s] \setminus \boldsymbol{h}(L)} (|\sum_{i:\boldsymbol{h}(i)=b} \boldsymbol{A}_{i,1} x_i| \cdot |\sum_{i:\boldsymbol{h}(i)=b} \boldsymbol{A}_{i,2} x_i| \cdot |\sum_{i:\boldsymbol{h}(i)=b} \boldsymbol{A}_{i,3} x_i|)^{p/3}$$

where we use θ to denote the denominator in (2). Fix some bucket $b \in [s] \setminus \boldsymbol{h}(L)$. Define

$$\Phi_{\boldsymbol{h},\boldsymbol{A}}^{(b)} := \frac{s}{(s - |\boldsymbol{h}(L)|)\theta} (|\sum_{i:\boldsymbol{h}(i)=b} \boldsymbol{A}_{i,1} x_i| \cdot |\sum_{i:\boldsymbol{h}(i)=b} \boldsymbol{A}_{i,2} x_i| \cdot |\sum_{i:\boldsymbol{h}(i)=b} \boldsymbol{A}_{i,3} x_i|)^{p/3}.$$

For $b \in \boldsymbol{h}(L)$, we define $\Phi_{\boldsymbol{h},\boldsymbol{A}}^{(b)} = 0$. The quantity $\Phi_{\boldsymbol{h},\boldsymbol{A}}^{(b)}$ can be computed by an $O(\log d)$ space algorithm in a single pass over the uniform random string of bits used to generate the p-stable random variables $\boldsymbol{A}$ by going over the values $i = 1, \ldots, d$ and ignoring the random bits that correspond to all i such that $\boldsymbol{h}(i) \neq b$. We formalize the algorithm by constructing an FSM $Q_{x,L,\boldsymbol{h},b}$

with poly(d) states. Note that we fixed $x, L, \boldsymbol{h}$ and b. Let the state of automaton be of the form (i, c_1, c_2, c_3) where $i \in [d]$ and c_1, c_2, c_3 denote the counters. The FSM being in state (i, c_1, c_2, c_3) denotes that it has processed $x_1, \ldots, x_i$ and found that for $j \in [3]$

$$c_j = \sum_{i' \leq i : h(i') = b} \boldsymbol{A}_{i',j} x_{i'}.$$

When in state (i, c_1, c_2, c_3), if $\boldsymbol{h}(i+1) \neq b$, the FSM directly transitions to the state $(i+1, c_1, c_2, c_3)$ ignoring the alphabet in the input to the FSM. If $\boldsymbol{h}(i) = b$, then the FSM reads the alphabet in the input string. Uses the $\{0, 1\}^{O(\log d)}$ size bit string that it reads to construct three p-stable random variables $\boldsymbol{A}_{i+1,1}, \boldsymbol{A}_{i+1,2}, \boldsymbol{A}_{i+1,3}$ and then transitions to the state $(i + 1, c'_1, c'_2, c'_3)$ where for $j \in [3]$

$$c'_j = c_j + \boldsymbol{A}_{i+1,j} x_{i+1}.$$

Note that all the above operations are performed only with a precision of $O(\log d)$ bits. Hence the Finite State Machine has only poly(d) states. From the state (d, c_1, c_2, c_3), the algorithm transitions to $(\text{final}, (s/(s - |\boldsymbol{h}(L)|)\theta)(|c_1||c_2||c_3|)^{p/3})$.

Given a uniform random string as input, the final state of FSM $Q_{c,L,\boldsymbol{h},b}$ encodes the value $\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}}$ and given $\boldsymbol{\gamma} \sim \text{HashPRG}$ as input, the final state of FSM $Q_{c,L,\boldsymbol{h},b}$ encodes the value of $\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}_{\boldsymbol{\gamma}}}$.

Let $d^b_{\boldsymbol{h}}$ be the distribution of the value of $\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}}$ conditioned on $\boldsymbol{h}$. Now define $(d^b_{\boldsymbol{h}})'$ to be the distribution of $\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}_{\boldsymbol{\gamma}}}$ i.e., the value of the estimator for bth bucket computed using p-stable random variables generated from a random $\boldsymbol{\gamma}$ HashPRG. As FSM $Q_{x,L\boldsymbol{h},b}$ has only poly(d) states, we obtain using Theorem 3.2 that

$$d_{\text{TV}}(d^b_{\boldsymbol{h}}, ((d_{\boldsymbol{h}})^b)') \leq 1/\text{poly}(d).$$

The above is true for any fixing of $\boldsymbol{h}$, x, L and b. As for any values of $\boldsymbol{h}$ and $\boldsymbol{A}$, we have $|\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}}| \leq$ poly(d), we obtain that for any fixed $\boldsymbol{h}$, x, L and b,

$$\left| \mathbf{E}_{\boldsymbol{A}}[\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}}] - \mathbf{E}_{\boldsymbol{A}_{\boldsymbol{\gamma}}}[\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}_N}] \right| \leq \frac{1}{\text{poly}(d)}$$

which implies that

$$\left| \mathbf{E}_{\boldsymbol{A}}[\Phi_{\boldsymbol{h},\boldsymbol{A}}] - \mathbf{E}_{\boldsymbol{A}_{\boldsymbol{\gamma}}}[\Phi_{\boldsymbol{h},\boldsymbol{A}_{\boldsymbol{\gamma}}}] \right| \leq s \cdot \frac{1}{\text{poly}(d)} \leq \frac{1}{\text{poly}(d)}.$$

Therefore, using (3) we have

$$\mathbf{E}_{\boldsymbol{h},\boldsymbol{A}_{\boldsymbol{\gamma}}}[\Phi_{\boldsymbol{h},\boldsymbol{A}_{\boldsymbol{\gamma}}}] = \mathbf{E}_{\boldsymbol{h},\boldsymbol{A}}[\Phi_{\boldsymbol{h},\boldsymbol{A}}] \pm 1/\text{poly}(d) = (1 \pm \varepsilon^C)\|x_{[d] \setminus L}\|^p_p \pm 1/\text{poly}(d). \tag{9}$$

Similarly, we have

$$(\Phi_{\boldsymbol{h},\boldsymbol{A}})^2 = \sum_{b,b' \in [s]} \Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}} \cdot \Phi^{(b')}_{\boldsymbol{h},\boldsymbol{A}}.$$

Again, for any fixed pair b, b', we can compute the product $\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}} \cdot \Phi^{(b')}_{\boldsymbol{h},\boldsymbol{A}}$ in $O(\log d)$ space using one pass over the uniform random string used to generate the p-stable random variables. We can construct a Finite State Machine very similar to the one above to show that the value $\Phi^{(b)}_{\boldsymbol{h},\boldsymbol{A}} \cdot \Phi^{(b')}_{\boldsymbol{h},\boldsymbol{A}}$ can be computed by a machine with poly(d) states. Now, we have that the total variation distance

between the distributions of the product when using a uniform random string to generate p-stable random variables and HashPRG to generate the p-stable random variables is at most $1/\operatorname{poly}(d)$ and hence we obtain that

$$|\mathop{\mathbf{E}}_{A}(\Phi^{(b)}_{h,A} \cdot \Phi^{(b')}_{h,A}) - \mathop{\mathbf{E}}_{A_\gamma}(\Phi^{(b)}_{h,A_\gamma} \cdot \Phi^{(b')}_{h,A_\gamma})| \leq 1/\operatorname{poly}(d).$$

Summing over all the pairs (b, b'), we obtain for any $x, L, \boldsymbol{h}$ that,

$$|\mathop{\mathbf{E}}_{A}[(\Phi_{h,A})^2] - \mathop{\mathbf{E}}_{A_\gamma}[(\Phi_{h,A_\gamma})^2]| \leq \frac{1}{\operatorname{poly}(d)}$$

which implies

$$\mathop{\mathbf{E}}_{h,A_\gamma}[(\Phi_{h,A_\gamma})^2] \leq \mathop{\mathbf{E}}_{h,A}[(\Phi_{h,A})^2] + \frac{1}{\operatorname{poly}(d)}.$$

We then have

$$\mathop{\mathbf{Var}}_{h,A_\gamma}[\Phi_{h,A_\gamma}] \leq \mathop{\mathbf{E}}_{h,A}[(\Phi_{h,A})^2] + \frac{1}{\operatorname{poly}(d)} - (\mathop{\mathbf{E}}_{h,A}[\Phi_{h,A}])^2 + \frac{\mathbf{E}_{h,A}[\Phi_{h,A}]}{\operatorname{poly}(d)}.$$

As the $\operatorname{poly}(d)$ term can be made $\geq d^C$ for a large enough constant C, we obtain that

$$\mathop{\mathbf{Var}}_{h,A_\gamma}[\Phi_{h,A_N}] \leq \mathop{\mathbf{Var}}_{h,A}[\Phi_{h,A}] + \frac{1}{\operatorname{poly}(d)} \leq O(\alpha)\|x\|_p^{2p}.$$

by Claim 5.3 and using the fact that x is a nonzero vector with integer coordinates and $\alpha \geq 1/\operatorname{poly}(d)$. $\qquad\square$

5.4 Wrap-up

Theorem 5.5. *Given $p \in (0, 2)$, an accuracy parameter ε such that $1/\sqrt{d} \leq \varepsilon \leq 1/d^c$ for a constant $0 < c < 1/2$ and a stream of updates $(i_1, v_1), \dots (i_m, v_m) \in [d] \times \{-M, \dots, M\}$ for $m, M \leq \operatorname{poly}(d)$ to the vector x, there is a streaming algorithm that uses $O(\varepsilon^{-2} \log d)$ bits of space and has an update time of $O(\log d)$ per stream element that outputs with probability $\geq 7/10$, a value v such that*

$$v = (1 \pm \varepsilon)\|x\|_p^p.$$

Proof. Setting $\alpha = \varepsilon^2 \log(d)$, we have that the set of α heavy hitters L can be computed in $O(\alpha^{-1} \log^2(d)) = O(\varepsilon^{-2} \log d)$ bits of space using Lemma B.1 and has an update time of $O(\log d)$ per stream element. The set L satisfies all the properties in Lemma B.1 with probability $\geq 9/10$. By Lemma 5.1, the HighEnd data structure can be maintained in $O(\alpha^{-1} \log^2(d)) = O(\varepsilon^{-2} \log d)$ bits of space and has an update time of $O(\log d)$ per stream element. Conditioned on L satisfying all the properties, we have that the value Ψ output by HighEnd data structure satisfies with probability $\geq 9/10$,

$$\Psi = (1 \pm \varepsilon)\|x_L\|_p^p.$$

By Lemma 5.4, the LightEstimator data structure can be maintained in $O(\alpha^{-1} \log d) = O(\varepsilon^{-2})$ bits of space. We also have that the data structure can be update in $O(1)$ time per stream element and conditioned on the set L having all the properties, the value Φ output by the algorithm satisfies

$$\mathbf{E}[\Phi] = (1 \pm \varepsilon^C)\|x_{[d] \setminus L}\|_p^p + 1/\operatorname{poly}(d)$$

and

$$\mathbf{Var}[\Phi] = O(\varepsilon^2 \log d)\|x\|_p^{2p}.$$

Maintaining $r = O(\log d)$ independent copies of LightEstimator in the stream and considering their outputs $\Phi_1, \ldots, \Phi_r$, conditioned on L having all the properties, we obtain using Chebyshev's inequality that with probability $\geq 99/100$,

$$\bar{\Phi} = \frac{\Phi_1 + \cdots + \Phi_r}{r} = (1 \pm \varepsilon^C)\|x_{[d]\setminus L}\|_p^p + \frac{1}{\text{poly}(d)} + \varepsilon\|x\|_p^p.$$

Thus, by a union bound, with probability $\geq 7/10$,

$$\Psi + \bar{\Phi} = (1 \pm O(\varepsilon))\|x\|_p^p + \frac{1}{\text{poly}(d)} = (1 \pm O(\varepsilon))\|x\|_p^p$$

using the fact that $\varepsilon > (1/\sqrt{d})$ and there is at least one nonzero integer coordinate in x. $\qquad\square$

6 Derandomizing CountSketch with HashPRG

CountSketch [CCF04] is a random linear map of a vector $x \in \mathbb{R}^d$ to $Ax \in \mathbb{R}^D$. For parameters r, t such that $D = rt$, the CountSketch $\text{CS}(x)$ is defined by two sequences of random independent hash functions: $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_r : \{0, \ldots, d-1\} \to [t]$ and $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_r : \{0, \ldots, d-1\} \to \{-1, +1\}$. To simplify our exposition we will assume that t is a power of two and that r is odd. We additionally assume that the coordinates of x are 0-indexed so that $x = (x_0, \ldots, x_{d-1})$. Indexing $\text{CS}(x) \in \mathbb{R}^D$ by $(i, j) \in [r] \times [t]$ the entries are defined as:

$$\text{CS}(x)_{i,j} := \sum_{\ell=0}^{d-1} \boldsymbol{s}_i(\ell)\, x_\ell\, [\boldsymbol{g}_i(\ell) = j].$$

For $x \in \mathbb{R}^d$ and $\ell \in \{0, \ldots, d-1\}$ we use $\text{CS}(x)$ to approximate x_ℓ with the following estimator:

$$\hat{x}_\ell = \text{median}(\{\boldsymbol{s}_i(\ell) \cdot \text{CS}(x)_{i,\boldsymbol{g}_i(\ell)} \mid i \in [r]\})\ . \tag{10}$$

Charikar, Chen, and Farach-Colton [CCF04] upper bounded the estimation error $|\hat{x}_\ell - x_\ell|$ in terms of the norm of the vector x and the parameters r, t. Their analysis only relies on using *pairwise independent* hash functions (independently for each repetition), which require $O(r \log d)$ bits of storage and allow the estimator to be computed in $O(r)$ time assuming constant time arithmetic operations.

Minton and Price [MP14] presented a tighter analysis of the distribution of the estimation error, focusing on $r = \Theta(\log d)$ repetitions, under the assumption that the hash function values of $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_r$ and $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_r$ are *fully independent*. This assumption is used in order to argue about the Fourier transform of the error distribution. In our notation they show the following lemma:

Lemma 6.1. *For every* $\alpha \in [0, 1]$, $\ell \in [d]$, $\mathbf{Pr}\left[|\hat{x}_\ell - x_\ell| > \alpha\,\Delta\right] < 2\exp\left(-\Omega\left(\alpha^2 r\right)\right)$, *where* $\Delta = \|tail_t(x)\|_2/\sqrt{t}$.

Literally storing fully random hash functions would require $O(rd \log t)$ bits, so this is not attractive when $d \gg t$, which is the setting where using CountSketch is of interest. Minton and Price note that for integer vectors $x \in \{-M, \ldots, +M\}^d$ where M is polynomial in d, it is possible to use the pseudorandom generator of Nisan [Nis92] to replace the fully independent hash

functions, keeping the tail bound of Lemma 6.1 up to a $1/\operatorname{poly}(d)$ additive term. However, this comes with considerable overhead: The space complexity increases by an $\Omega(\log(dt))$ factor, and the time per update/query increases by a factor $\Omega(rt)$. Jayaram and Woodruff [JW18] later considered a modification of CountSketch (with the same space and error guarantees) and showed that the multiplicative space overhead can be reduced to $O((\log \log d)^2)$ using a pseudo-random generator for fooling halfspaces. The time complexity increases by an unspecified polylogarithmic factor compared to the fully random setting. Though it is technically not accurate we will still refer to their sketch as CountSketch.

In this section we present an alternative derandomization of [MP14] using hash functions computed using HashPRG. Specifically, for $i \in [r]$ and $\rho \in \{0, \ldots, d-1\}$ we use block number $(i-1)d + \rho$ from the output of HashPRG to get the random bits for $\boldsymbol{s}_i(\rho)$ and $\boldsymbol{g}_i(\rho)$. Since a given output block of HashPRG can be computed efficiently, these hash functions can be efficiently evaluated.

Theorem 6.2. *Let d be the dimension of the vectors and M be the maximum absolute value of coordinates in the vector. Let t and r be the parameters of the CountSketch map as defined above. Let $b \geq 2$ be an integer denoting the branching factor of HashPRG. Let $w = \Omega(\log d + \log M)$. There exists a randomized linear sketch $\mathrm{CS}_{HashPRG} : \{-M, \ldots, M\}^d \to \{-2^w, \ldots, 2^w\}^{tr}$ that can be implemented on a word RAM with word size w with the following properties:*

- *The parameters required to defined the map $\mathrm{CS}_{HashPRG}$ can be stored in $O(b \log_b d)$ words of space and given a vector $x \in \{-M, \ldots, M\}^d$, the resulting vector $\mathrm{CS}_{HashPRG}(x)$ can be stored in $O(rt)$ words of space.*

- *Given $\mathrm{CS}_{HashPRG}(x)$ and an update (ℓ, u_ℓ) corresponding to a vector u with a single nonzero entry u_ℓ, we can compute $\mathrm{CS}_{HashPRG}(x + u)$ in time $O(r \log_b d)$.*

- *For every $x \in \{-M, \ldots, M\}^d$, $\alpha \in [0, 1]$, and $\ell \in [d]$, we can compute an estimator $\hat{x}_\ell$ from $\mathrm{CS}_{HashPRG}$ in time $O(r \log_b d)$ such that $\mathbf{Pr}\left[|\hat{x}_\ell - x_\ell| > \alpha\,\Delta\right] < 2\exp\left(-\Omega\left(\alpha^2 r\right)\right) + 2^{-Cw}$, where $\Delta = \|tail_t(x)\|_2/\sqrt{t}$.*

Figure 1 compares Theorem 6.2 to previously known ways of choosing the hash functions for CountSketch. Our construction is the first one that is able to match the space usage of CountSketch with pairwise independent hash functions (for $r = O(\log d)$ repetitions) while showing the strong concentration known for fully random hash functions. With CountSketch table size $t = d^{\Omega(1)}$ and word length $w = O(\log d)$ we also match the update time of pairwise independence on the Word RAM.

6.1 PRGs for Space-bounded Computation and CountSketch

Like Minton and Price [MP14] we will consider vectors $x \in \{-M, \ldots, M\}^d$ for a positive integer M. For concreteness we consider CountSketch with entries that are w-bit machine words. We can relax the requirement from [MP14] that M is polynomial in d, and instead assume that $M < 2^{w-1}/d$, which is also necessary to ensure that there are no overflows when computing $\mathrm{CS}(x)$.

To derandomize CountSketch, we describe a small-space algorithm for any fixed input vector x, query ℓ and threshold $\alpha\Delta$. The algorithm makes a single pass over the output from HashPRG and determines whether the estimator $\hat{x}_\ell$ computed using HashPRG has error exceeding $\alpha\Delta$. This is done *without* computing $\mathrm{CS}(x)$, and in fact even without computing $\hat{x}_\ell$. The algorithm makes critical use of the *symmetry* of HashPRG, namely, that the distribution of hash values is unchanged by permuting the inputs using a mapping of the form $\rho \mapsto \rho \oplus \ell$. We stress that Nisan's generator

Hash function	Space in words	Bounds small error	Update time
Pairwise independent [CCF04]	D	No	$\log d$
Fully random [MP14]	$d \log d$	Yes	$\log d$
Nisan's generator [MP14, Nis92]	$D \log(d)$	Yes	$t \log^3(d)$
Halfspace Fooling PRGs [JW18]	$D(\log \log d)^2$	Yes	$(\log d)^{O(1)}$
HashPRG $(b = t)$	D	Yes	$\log^2(d)/\log t$
HashPRG $(b = d^{\Omega(1)})$	$D + b$	Yes	$\log d$

Figure 1: Overview of CountSketch guarantees with different kinds of random hash functions. For simplicity we focus on the case of $r = O(\log d)$ repetitions and d-dimensional input vectors that contain $O(\log d)$-bit integers such that the CountSketch itself (without hash functions) uses space $D = O(t \log d)$ words. With pairwise independence we can only tightly bound the probability of exceeding error $\Delta = \|\text{tail}_t(x)\|_2/\sqrt{t}$, while the other hash functions allow us to bound the probability of smaller errors. Time bounds are for implementation on a Word RAM with word size $w = O(\log d)$. Parameters with a particularly bad impact on space or time are highlighted in red color.

does not have this symmetry property, and that we are not aware of an equally space-efficient finite state machine for evaluating the error of CountSketch using Nisan's generator.

The finite state machine. Consider $x \in \{-M, \ldots, M\}^d$, $\ell \in \{0, \ldots, d-1\}$, and a given error threshold $\alpha\Delta \in \mathbb{R}$. A choice of hash functions $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_r : \{0, \ldots, d-1\} \to [t]$ and $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_r : \{0, \ldots, d-1\} \to \{-1, +1\}$ can be represented as a binary string $\boldsymbol{\gamma} \in \{0,1\}^{rdw}$ where a block of w consecutive bits encodes hash values $\boldsymbol{s}_i(\rho)$ and $\boldsymbol{g}_i(\rho)$ for $i \in [r]$ and $\rho \in \{0, \ldots, d-1\}$. We order the blocks such that hash values with $i = 1$ come first, then $i = 2$ and so on. Concretely we may take $\boldsymbol{s}_i(\rho) = 2\gamma_{(i-1)dw+\rho w} - 1$ and $\boldsymbol{g}_i(j) = 1 + \sum_{k=1}^{\log(t)} \gamma_{(i-1)dw+(\rho+1)w-k} 2^{k-1}$ such that the hash values can be extracted from a block in constant time. We consider two ways of sampling the string $\boldsymbol{\gamma}$:

- First, we may choose $\boldsymbol{\gamma} \sim (U_2)^{rdw}$ with independent, random bits. It is easy to see that this is the same as choosing the hash functions with full independence, so Lemma 6.1 holds for this choice of $\boldsymbol{\gamma}$.

- Second, for every b, k such that $dr \leq b^k < 2^{cw}$ we can use HashPRG with block size $n = w$ to generate $\boldsymbol{\gamma} \sim G_k(*, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_k)$. This corresponds to the hash functions we use to derandomize CountSketch.

As we noted in the introduction, the distribution of $\boldsymbol{\gamma}^{\oplus \ell}$ is the same as the distribution of $\boldsymbol{\gamma}$. Thus, we can assume that an algorithm that makes a single pass over the string $\boldsymbol{\gamma}$ reads the rd blocks of bits in the order $0 \oplus \ell, 1 \oplus \ell, \ldots, (rd-1) \oplus \ell$ so that for each repetition $i \in [t]$, the algorithm gets to know the value $\boldsymbol{g}_i(\ell)$ before reading the pseudorandom bits corresponding to other blocks. Thus, for each repetition $i \in [t]$, we can assume that an algorithm making a single pass over the string $\boldsymbol{\gamma}$ *sees* the values $\boldsymbol{g}_i(0 \oplus \ell), \boldsymbol{g}_i(1 \oplus \ell), \ldots, \boldsymbol{g}_i((d-1) \oplus \ell)$ in that order and similarly the values of $\boldsymbol{s}_i$.

To analyze the properties of the CountSketch data structure constructed using the string $\boldsymbol{\gamma}$, we consider an FSM $Q = Q_{x,\ell,\alpha\Delta}$ with states

$$\{0, \ldots, r+1\}^3 \times \{-1, +1\} \times [t] \times \{0, \ldots, d\} \times \{-2^w, \ldots, 2^w\},$$

plus a special start state $\perp$. We use $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_r$ and $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_r$ to refer to the hash functions encoded by $\boldsymbol{\gamma}$. The idea is that the FSM Q computes the r simple estimators in (10) one at a time, and keeps track of the number of these estimators that deviate from x_ℓ by more than $\alpha\Delta$ (with separate accounting for overestimates and underestimates). More precisely, when Q is in state $(\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7)$ it signifies that:

- It has fully processed the hash functions $\boldsymbol{s}_i$, $\boldsymbol{g}_i$ for $i < \beta_1$, that β_2 of these hash function pairs produced an estimate $\boldsymbol{s}_i(\ell) \cdot \mathrm{CS}(x)_{i,\boldsymbol{g}_i(\ell)} < x_\ell - \alpha\Delta$, and β_3 pairs produced an estimate $\boldsymbol{s}_i(\ell) \cdot \mathrm{CS}(x)_{i,\boldsymbol{g}_i(\ell)} > x_\ell + \alpha\Delta$,

- $\boldsymbol{s}_{\beta_1}(\ell) = \beta_4$ and $\boldsymbol{g}_{\beta_1}(\ell) = \beta_5$, and

- $\beta_4 \sum_{0 \le z < \beta_6} x_{z \oplus \ell} \boldsymbol{s}_{\beta_1}(z \oplus \ell)[\boldsymbol{g}_{\beta_1}(z \oplus \ell) = \boldsymbol{g}_{\beta_1}(\ell)] = \beta_7$.

From state $\perp$, the FSM Q transitions to $(1, 0, 0, \boldsymbol{s}_1(\ell), \boldsymbol{g}_1(\ell), 0, 0)$, where the values $\boldsymbol{s}_1(\ell), \boldsymbol{g}_1(\ell)$ are determined by the zeroth block since we assume that the FSM sees the blocks as they are ordered in the string $\boldsymbol{\gamma}^{\oplus \ell}$. From this point on, when in state $(\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7)$:

- If $\beta_6 = d$ we have $\beta_7 = \boldsymbol{s}_{\beta_1}(\ell)\mathrm{CS}(x)_{\beta_1, \boldsymbol{g}_{\beta_1}(\ell)}$, so we can decide whether simple estimator number β_1 has an error above $\alpha\Delta$ and update β_2 or β_3 accordingly. Finally we can increment β_1, set $\beta_6 = 0$, and update β_4, β_5 to reflect the values of the new hash values, available in the next block.

- Otherwise, when $\beta_1 \le r$, Q has access to $\boldsymbol{s}_{\beta_1}(\beta_6 \oplus \ell)$ and $\boldsymbol{g}_{\beta_1}(\beta_6 \oplus \ell)$ from the next block of bits it reads. This allows us to increment β_6 when simultaneously increasing β_7 by $\beta_4 \boldsymbol{s}_{\beta_1}(\beta_6 \oplus \ell)x_{\beta_6 \oplus \ell}$ if $\boldsymbol{g}_{\beta_1}(\beta_6 \oplus \ell) = \beta_5$.

- Finally, when $\beta_1 = r + 1$ we ignore the rest of the input, remaining in the same state.

After reaching the end, the values β_2, β_3 determine whether the estimator $\hat{x}_\ell$ in (10) has an error of more than $\alpha\Delta$: If $\beta_2 \ge \lceil r/2 \rceil$ then $\hat{x}_\ell < x_\ell - \alpha\Delta$, if $\beta_3 \ge \lceil r/2 \rceil$ then $\hat{x}_\ell > x_\ell + \alpha\Delta$, and otherwise $|\hat{x}_\ell - x_\ell| \le \alpha\Delta$.

The number of states in Q is $O(r^3 t d 2^w)$ and the number of blocks of bits read is rd, both of which are $2^{O(w)}$. Theorem 3.1 with $n = O(w)$ now implies that the total variation distance $\|Q(G_k(*, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_k)) - Q((U_n)^{2^k})\|$ is at most 2^{-cw} for *some* constant $c > 0$. The additive term 2^{-cw} can be made smaller than 2^{-Cw} for *any* constant $C > 1$ by adjusting the parameter n since a Word RAM with a constant factor larger word size can be simulated with a constant factor overhead in time. In particular, error probabilities grow by at most 2^{-Cw} when switching from fully random hash functions to hash functions defined by to HashPRG. Finally, note that the space usage of HashPRG is $O(w^2 b / \log b)$ bits, and that we can compute the hash function values $\boldsymbol{g}_1(\ell), \ldots, \boldsymbol{g}_r(\ell)$ and $\boldsymbol{s}_1(\ell), \ldots, \boldsymbol{s}_r(\ell)$ in time $O(rw / \log b)$.

6.2 Alternative Derandomizations of CountSketch with Nisan's PRG

We note that there are alternative derandomizations using Nisan's PRG that do not incur the naïve $O(\log d)$ factor overhead in terms of space in some regimes. We can use the fact that Nisan's PRG also "fools" small space algorithms that make multiple passes over the pseudorandom string. Specifically, Lemma 2.3 of [DPS11] shows that Nisan's PRG fools a small space algorithm that makes 2 passes over the pseudorandom string. We derandomize each repetition $i \in [r]$ separately. Fix an index $\ell \in [d]$ and repetition $i \in [r]$ and consider the string used to compute hash functions $\boldsymbol{g}_i$ and $\boldsymbol{s}_i$ for repetition i. An algorithm in the first pass over the string computes the bucket

into which the index ℓ gets hashed into and in the second pass over the string computes the value, denoted by $\hat{x}_{i,g_i(\ell)}$, of the bucket into which the ℓ-th coordinate gets hashed into in the i-th repetition. Finally the algorithm terminates with the value $s_i(\ell)\hat{x}_{i,g_i(\ell)}$. As this algorithm overall takes $O(\log d)$ space, it is "fooled" by Nisan's PRG with a seed length of $O(\log^2 d)$. Thus, we obtain $\mathbf{Pr}[|s_i(\ell)\hat{x}_{i,g_i(\ell)} - x_\ell| \leq \alpha\Delta] \gtrsim \alpha$ as in proof of Theorem 4.1 of [MP14] even when each repetition of CountSketch is independently derandomized using a string sampled from Nisan's PRG. Overall, if each repetition is derandomized using an independent pseudorandom string, we obtain that $\mathbf{Pr}[|\hat{x}_\ell - x_\ell| > \alpha\Delta] \leq 2\exp(-\Omega(\alpha^2 r))$. While this derandomization has a fast update time, a drawback is that the overall seed length is $O(r\log^2 d)$ (a string of $O(\log^2 d)$ bits for each $i \in [r]$) which can be larger than the space complexity of CountSketch ($O(rt\log d)$ bits) when $t = o(\log d)$.

There is also another derandomization using Nisan's PRG which avoids space blow-up in the case of $rt = \omega(\log d)$. Instead of using independent samples from Nisan's PRG for each repetition of CountSketch, we derandomize the construction all-at-once. Consider the following algorithm that uses a single sample from Nisan's PRG to construct all hash functions g_i and s_i. Fix a vector x and coordinate ℓ. The algorithm makes a first pass over the string to determine into which bucket the index ℓ gets hashed into in each of the repetitions. The information can be stored using $O(r\log t)$ bits. In the second pass over the random string, the algorithm can then determine if the estimate $\hat{x}_\ell$ satisfies $|\hat{x}_\ell - x_\ell| \leq \alpha\Delta$. Overall the algorithm uses a space of $O(\log d + r\log t)$ bits. Using this fact, one can obtain a derandomization of CountSketch with the above estimation error guarantee and the CountSketch derandomized using Nisan's generator can be stored in $O(r \cdot t + \log d)$ words of space, which is asymptotically the same as the space complexity of CountSketch derandomized using HashPRG. However, Nisan's generator needs to fool an $O(r \cdot \log t + \log d)$ space algorithm which makes the update time much slower, on a machine with $O(\log d)$ word size, compared to the derandomization using HashPRG which only has to fool an $O(\log d)$ space algorithm.

7   Private CountSketch

Pagh and Thorup [PT22] recently analyzed the estimation error of CountSketch data structure made private using the Gaussian Mechanism. Their analysis requires the hash functions $g_1, \ldots, g_r : \{0, \ldots, d-1\} \to [t]$ and $s_1, \ldots, s_r : \{0, \ldots, d-1\} \to \{+1, -1\}$ to be fully random. They define

$$\mathrm{PCS}(x) = \mathrm{CS}(x) + \nu$$

where ν is a $D = r \cdot t$ dimensional vector with independent Gaussian random variables of mean 0 and variance σ^2. By taking σ to be appropriately large, we obtain that $\mathrm{PCS}(x)$ is (ε, δ)-differentially private. For $\ell \in \{0, 1, \ldots, d-1\}$, we can define the estimator $\hat{x}_\ell$ as

$$\hat{x}_\ell = \mathrm{median}(\{\, s_i(\ell) \cdot \mathrm{PCS}(x)_{i,g_i(\ell)} \mid i \in [r] \,\}).$$

As we saw in Section 6, if σ were 0, then the hash functions g_i and s_i can be derandomized using HashPRG while obtaining tail bounds on the estimation error $|x_\ell - \hat{x}_\ell|$. A similar argument which crucially uses the symmetry property of HashPRG shows that PCS can also be derandomized using HashPRG. For the case of Private CountSketch with fully random hash functions the following theorem is shown in [PT22]:

Theorem 7.1 ([PT22]). *For every $\alpha \in [0,1]$ and every $\ell \in [d]$, the estimation error of private CountSketch with r repetitions, table size t and $\nu \sim N(0, \sigma^2)^D$, then*

$$\mathbf{Pr}[|\hat{x}_\ell - x_\ell| \geq \alpha\max(\Delta, \sigma)] \leq 2\exp(-\Omega(\alpha^2 r))$$

where $\Delta = \|tail_t(x)\|_2/\sqrt{t}$.

We now derandomize the requirement that the hash functions g_i and s_i be fully random. The proof is an extension of the proof in Section 6 and proceeds very similarly. We detail it below for completeness. Instead of constructing an FSM, we describe a small space algorithm which makes a single pass over the randomness while updating its state after reading a block of random bits. The algorithm can be easily converted to an FSM.

Fix a vector $x \in \mathbb{R}^d$ (corresponds to the final value of the vector in the stream), $\nu \in \mathbb{R}^D$ (corresponds to the Gaussian vector we add to CountSketch to make it private), a coordinate $\ell \in [d]$ and a parameter α. We initialize three variables $\mathsf{acc} = \mathsf{deficit} = \mathsf{excess} = 0$. The algorithm reads a block of w bits from the input string. Using the symmetry of HashPRG we again assume that the FSM sees the blocks of γ as they are in $\gamma^{\oplus \ell}$. Thus, for each repetition, the block that the FSM sees first can be used to determine $g_i(\ell)$ and $s_i(\ell)$. For the first repetition, using the zeroth block of bits, the FSM computes and stores the values $g_1(\ell)$ and $s_1(\ell)$ locally and update the accumulator acc to $s_1(\ell)x_{0\oplus\ell}$ (note that $x_{0\oplus\ell} = x_\ell$) and move to the next block of bits in the input string.

Suppose we are reading the j-th block of random bits. Again, we can determine $g_i(j \oplus \ell)$ and $s_i(j \oplus \ell)$. If $h_1(j \oplus \ell) \neq g_1(\ell)$, we move to reading the $(j+1)$-th block. If $g_1(j \oplus \ell) = g_1(\ell)$, we add $s_1(j \oplus \ell)x_{j\oplus\ell}$ to the accumulator and move to the $(j+1)$-th block. After reading d blocks, if $s_1(\ell)(\mathsf{acc} + \nu_{1,g_1(\ell)}) \geq x_\ell + \alpha \max(\Delta, \sigma)$ we increase the variable excess by 1. If $s_1(\ell)(\mathsf{acc} + \nu_{1,g_1(\ell)}) \leq x_\ell - \alpha \max(\Delta, \sigma)$, we increase $\mathsf{deficit}$ by 1. Then we zero out the variable acc, remove the stored values $g_1(\ell)$ and $s_1(\ell)$ and repeat the process by reading the next block of bits. We again determine $g_2(\ell)$ and $s_2(\ell)$ and set acc to $s_2(\ell)x_\ell$. We then move to the next block and so on. We do the process in total for r of times. Finally, if $\mathsf{excess} < r/2$ and $\mathsf{deficit} < r/2$, we set the variable Status to $\mathsf{Success}$ and otherwise set Status to $\mathsf{Failure}$.

The algorithm, at any point of time needs to store only $O(\log d)$ bits. As the algorithm needs only "read" access to x, the entire algorithm can be converted to an FSM, which we call $Q_{x,\nu,\ell,\alpha}$, that has poly(d) states over the alphabet $\{0,1\}^w$. The variable Status determines if the FSM ends in a state $\mathsf{Success}$ or $\mathsf{Failure}$.

If the input string to the FSM is uniformly random, then the hash functions g_i and s_i are fully random. Therefore by Theorem 7.1, we obtain that

$$\mathbf{E}_{\nu} \mathbf{Pr}_{\gamma \sim U}[\text{Final State} = \mathsf{Success}] \geq 1 - 2\exp(-\Omega(\alpha^2 r)).$$

Now consider HashPRG with a block size $w = \Omega(\log d)$. By Theorem 3.2, for any fixed x, ν, ℓ and α that

$$\mathbf{Pr}_{\gamma \sim \text{HashPRG}}[\text{Final State of FSM } Q_{x,\nu,\ell,\alpha} \text{ on input } \gamma]$$
$$\geq \mathbf{Pr}_{\gamma \sim U}[\text{Final State of FSM } Q_{x,\nu,\ell,\alpha} \text{ on input } \gamma] - O(2^{-cw}).$$

By taking an expectation over $\nu \sim N(0, \sigma^2)^D$, we have

$$\mathbf{E}_{\nu} \mathbf{Pr}_{\gamma \sim \text{HashPRG}}[\text{Final State of FSM } Q_{x,\nu,\ell,\alpha} \text{ on input } \gamma] \geq 1 - 2\exp(-\Omega(\alpha^2 r)) - O(2^{-cw}).$$

We therefore obtain that with probability $\geq 1 - 2\exp(-\Omega(\alpha^2 r)) - 2 \cdot 2^{-cw}$ over $\nu \sim N(0, \sigma^2)^D$ and $\gamma \sim \text{HashPRG}$, if g_i and s_i are hash functions constructed as a function of γ as described in the above algorithm, then with probability $\geq 1 - \exp(-\Omega(\alpha^2 r)) - O(2^{-cw})$ over ν and γ,

$$|\text{median}_{i \in [r]}(s_i(\ell)(\sum_{j=0}^{d-1}[g_i(j) = g_i(\ell)]s_j(i)x_j + \nu_{i,g_i(\ell)})) - x_\ell| \leq \alpha \max(\Delta, \sigma).$$

Thus we have the following theorem.

Theorem 7.2. *For every $\alpha \in [0,1]$ and every $\ell \in \{0, \ldots, d-1\}$, the estimation error of private CountSketch with r repetitions, table size t derandomized using HashPRG with a block size of w and $\nu \sim N(0, \sigma^2)^D$, then*

$$\mathbf{Pr}[|\hat{x}_\ell - x_\ell| \geq \alpha \max(\Delta, \sigma)] \leq 2 \exp(-\Omega(\alpha^2 r)) + O(2^{-cw})$$

where $\Delta = \|tail_t(x)\|_2 / \sqrt{t}$.

8 Estimating $\|x\|_\infty$

Let $x \in \mathbb{R}^d$ be the underlying vector we are maintaining in a turnstile stream. Assume that the coordinates of x are integers bounded in absolute value by $\text{poly}(d)$. We give a simple algorithm that uses only $O(\varepsilon^{-2} \log d \log 1/\varepsilon)$ bits of space and approximates $\|x\|_\infty$ up to an additive error of $\varepsilon \|x\|_2$. We give a matching lower bound and show that our algorithm is tight up to constant factors.

Let $t \leq d$ be a parameter that we set later. Let $\boldsymbol{L} : \mathbb{R}^d \to \mathbb{R}^t$ be a randomized linear map defined as

$$(\boldsymbol{L}x)_i = \sum_{j : \boldsymbol{h}(j) = i} \boldsymbol{s}(i) x_i.$$

Here we assume that the hash function $\boldsymbol{h}$ is drawn from a 2-wise independent hash family $\mathcal{H} = \{\, h : [d] \to [t] \,\}$ and the sign function $\boldsymbol{s}$ is drawn from a 4-wise independent hash family $\mathcal{S} = \{\, s : [d] \to \{+1, -1\} \,\}$. We note that there exist families $\mathcal{H}$ and $\mathcal{S}$ such that $\boldsymbol{h}$ and $\boldsymbol{s}$ can be sampled from their respective families and stored using $O(\log d)$ bits. We prove the following lemma:

Lemma 8.1. *Given a parameter α, if $t \geq 1/2\alpha^4\delta$, then with probability $\geq 1 - 3\delta$, the following simultaneously hold:*

1. *$\|\boldsymbol{L}\|_\infty = \|x\|_\infty \pm (2\sqrt{\alpha}/\delta^{1/4})\|x\|_2$ and*

2. *$\|\boldsymbol{L}x\|_2^2 \leq (1 + 2\alpha^2)\|x\|_2^2$.*

Proof. Define $\textsc{Large} := \{\, j \in [d] \mid |x_j| \geq \alpha\|x\|_2 \,\}$ and $\textsc{Small} := [d] \setminus \textsc{Large}$. Note that $|\textsc{Large}| \leq 1/\alpha^2$. Let $x_{\textsc{Large}} \in \mathbb{R}^d$ be the d-dimensional vector with only the $\textsc{Large}$ coordinates of vector x and define $x_{\textsc{Small}} = x - x_{\textsc{Large}}$. The following result is now a simple consequence of the 2-wise independence of $\boldsymbol{h}$.

Lemma 8.2. *If $t \geq |\textsc{Large}|^2/(2\delta)$, then with probability $\geq 1 - \delta$, for all $j, j' \in \textsc{Large}$ with $j \neq j'$, we have $\boldsymbol{h}(j) \neq \boldsymbol{h}(j')$.*

Proof. For $j, j' \in \textsc{Large}$ with $j \neq j'$, let $\boldsymbol{X}_{j,j'} = 1$ if $\boldsymbol{h}(j) = \boldsymbol{h}(j')$ and 0 otherwise. Using the 4-wise independence of $\mathcal{H}$, we have $\mathbf{E}[\boldsymbol{X}_{j,j'}] = 1/t$. Hence, $\mathbf{E}[\sum_{j<j'} \boldsymbol{X}_{j,j'}] \leq |\textsc{Large}|^2/(2t)$. By Markov's inequality, with probability $\geq 1 - \delta$, $\sum_{j<j'} \boldsymbol{X}_{j,j'} \leq |\textsc{Large}|^2/(2t\delta) \leq 1/2$ if $t \geq |\textsc{Large}|^2/\delta$. By definition, the random variable $\sum_{j<j'} \boldsymbol{X}_{j,j'}$ takes only non-negative integer values. Hence, we obtain that with probability $\geq 1 - \delta$, $\sum_{j<j'} \boldsymbol{X}_{j,j'} = 0$ which implies that the hash function $\boldsymbol{h}$ hashes each of the coordinates in the set $\textsc{Large}$ to distinct locations. $\square$

As $t \geq 1/2\alpha^4\delta \geq |\textsc{Large}|^2/2\delta$, we get that all the coordinates in the set $\textsc{Large}$ are hashed to distinct buckets by the hash function $\boldsymbol{h}$ with probability $1 - \delta$. We now bound $\|\boldsymbol{L}x_{\textsc{Small}}\|_\infty$. By

42

definition of the set SMALL, we have $\|x_{\text{SMALL}}\|_\infty \le \alpha\|x\|_2$ and $\|x_{\text{SMALL}}\|_2 \le \|x\|_2$. Let $i \in [t]$ be arbitrary. We have $(\boldsymbol{L}x_{\text{SMALL}})_i = \sum_{j\in\text{SMALL}}[\boldsymbol{h}(j) = i]\boldsymbol{s}(j)x_j$ and

$$\mathbf{E}[(\boldsymbol{L}x_{\text{SMALL}})_i^2] = \sum_{j,j'\in\text{SMALL}} \mathbf{Pr}[\boldsymbol{h}(j) = i, \boldsymbol{h}(j') = i]\,\mathbf{E}[\boldsymbol{s}(j)\boldsymbol{s}(j')]x_j x_{j'}$$

using the independence of $\boldsymbol{h}$ and $\boldsymbol{s}$. For $j \ne j' \in \text{SMALL}$, using the 4-wise independence of $\boldsymbol{s}$, we get $\mathbf{E}[\boldsymbol{s}(j)\boldsymbol{s}(j')] = 0$ and therefore, the above expression simplifies to

$$\mathbf{E}[(\boldsymbol{L}x_{\text{SMALL}})_i^2] = \sum_{j\in\text{SMALL}} \mathbf{Pr}[\boldsymbol{h}(j) = i]x_j^2 = \|x_{\text{SMALL}}\|_2^2/t.$$

Similarly, we have

$$\mathbf{E}[(\boldsymbol{L}x_{\text{SMALL}})_i^4] = \sum_{j_1,j_2,j_3,j_4\in\text{SMALL}} \mathbf{Pr}[\bigwedge_{k=1}^4 (\boldsymbol{h}(j_k) = i)]\,\mathbf{E}[\prod_{k=1}^4 \boldsymbol{s}(j_k)] \prod_{k=1}^4 x_{j_k}.$$

We now see using the 4-wise independence of $\boldsymbol{s}$ that $\mathbf{E}[\boldsymbol{s}(j_1)\boldsymbol{s}(j_2)\boldsymbol{s}(j_3)\boldsymbol{s}(j_4)]$ is nonzero only when all the indices j_1, j_2, j_3, j_4 are equal or we can pair the indices j_1, j_2, j_3, j_4 into two groups taking same values. Hence,

$$\mathbf{E}[(\boldsymbol{L}x_{\text{SMALL}})_i^4] = \frac{\|x_{\text{SMALL}}\|_4^4}{t} + \frac{6}{t^2} \sum_{j_1<j_2\in\text{SMALL}} x_{j_1}^2 x_{j_2}^2$$

$$= \frac{\|x_{\text{SMALL}}\|_4^4}{t} + \frac{3}{t^2}\left(\|x_{\text{SMALL}}\|_2^4 - \|x_{\text{SMALL}}\|_4^4\right)$$

which then implies

$$\mathbf{Var}[(\boldsymbol{L}x_{\text{SMALL}})_i^2] \le (1/t)\|x_{\text{SMALL}}\|_4^4 + (2/t^2)\|x_{\text{SMALL}}\|_2^4.$$

Since $\|x_{\text{SMALL}}\|_4^4 \le \|x_{\text{SMALL}}\|_\infty^2\|x_{\text{SMALL}}\|_2^2 \le \alpha^2\|x\|_2^4$, we get $\mathbf{Var}[(\boldsymbol{L}x_{\text{SMALL}})_i^2] \le (2/t^2 + \alpha^2/t)\|x\|_2^4$. By Chebyshev's inequality,

$$\mathbf{Pr}[(\boldsymbol{L}x_{\text{SMALL}})_i^2 \ge \|x_{\text{SMALL}}\|_2^2/t + \gamma] \le \frac{(2/t^2 + \alpha^2/t)\|x\|_2^4}{\gamma^2}.$$

By a union bound over all t values of i, we get that $\|\boldsymbol{L}x_{\text{SMALL}}\|_\infty^2 \le \|x_{\text{SMALL}}\|_2^2/t + \gamma$ with probability $\ge 1 - (2/t + \alpha^2)\|x\|_2^4/\gamma^2$. For $t \ge 1/\alpha^2$ and $\gamma = (2\alpha/\sqrt{\delta})\|x\|_2^2$, we have that with probability $\ge 1 - \delta$, $\|\boldsymbol{L}x_{\text{SMALL}}\|_\infty^2 \le (\alpha^2 + 2\alpha/\sqrt{\delta})\|x\|_2^2 \le (3\alpha/\sqrt{\delta})\|x\|_2^2$. We thus finally have that if $t \ge 1/\alpha^2$, then with probability $\ge 1 - \delta$,

$$\|\boldsymbol{L}x_{\text{SMALL}}\|_\infty \le \frac{2\sqrt{\alpha}}{\delta^{1/4}}\|x\|_2.$$

Hence, by a union bound, with probability $\ge 1 - 2\delta$, $\|\boldsymbol{L}x_{\text{LARGE}}\|_\infty = \|x_{\text{LARGE}}\|_\infty$ and $\|\boldsymbol{L}x_{\text{SMALL}}\|_\infty \le (2\sqrt{\alpha}/\delta^{1/4})\|x\|_2$. Condition on this event. If $\|x\|_\infty \ge \alpha\|x\|_2$, then $\|x_{\text{LARGE}}\|_\infty = \|x\|_\infty$ and by triangle inequality, we get

$$\|\boldsymbol{L}x\|_\infty = \|\boldsymbol{L}x_{\text{LARGE}} + \boldsymbol{L}x_{\text{SMALL}}\|_\infty$$

$$= \|\boldsymbol{L}x_{\text{LARGE}}\|_\infty \pm \|\boldsymbol{L}x_{\text{SMALL}}\|_\infty$$

$$= \|x\|_\infty \pm (2\sqrt{\alpha}/\delta^{1/4})\|x\|_2.$$

If $\|x\|_\infty < \alpha\|x\|_2$, then $\text{LARGE} = \emptyset$ and $\|\boldsymbol{L}x\|_\infty = \|\boldsymbol{L}x_{\text{SMALL}}\|_\infty \leq (2\sqrt{\alpha}/\delta^{1/4})\|x\|_2$ which clearly satisfies $\|\boldsymbol{L}x\|_\infty = \|x\|_\infty \pm (2\sqrt{\alpha}/\delta^{1/4})\|x\|_2$.

We will now bound $\|\boldsymbol{L}x\|_2^2$. First we have,

$$\|\boldsymbol{L}x\|_2^2 = \sum_{i \in [t]} \left(\sum_{j \in [d]: \boldsymbol{h}(j)=i} \boldsymbol{s}(j)x_j \right)^2$$

$$= \sum_{i \in [t]} \left(\sum_{\substack{j \in [d] \\ \boldsymbol{h}(j)=i}} x_j^2 + \sum_{\substack{j_1 \neq j_2 \in [d]: \\ \boldsymbol{h}(j_1)=\boldsymbol{h}(j_2)=i}} \boldsymbol{s}(j_1)\boldsymbol{s}(j_2)x_{j_1}x_{j_2} \right)$$

$$= \|x\|_2^2 + 2\sum_{i \in [t]} \sum_{j_1 < j_2 \in [d]} [\boldsymbol{h}(j_1) = \boldsymbol{h}(j_2) = i]\boldsymbol{s}(j_1)\boldsymbol{s}(j_2)x_{j_1}x_{j_2}.$$

By 4-wise independence of $\boldsymbol{s}$, we get $\mathbf{E}[\boldsymbol{s}(j_1)\boldsymbol{s}(j_2)] = 0$ for $j_1 \neq j_2$ and therefore get $\mathbf{E}[\|\boldsymbol{L}x\|_2^2] = \|x\|_2^2$. We now bound $\mathbf{Var}(\|\boldsymbol{L}x\|_2^2)$.

$$\mathbf{Var}(\|\boldsymbol{L}x\|_2^2) = \mathbf{E}[(\|\boldsymbol{L}x\|_2^2 - \|x\|_2^2)^2]$$

$$= 4\sum_{i_1,i_2} \sum_{\substack{j_1 < j_2 \\ j_3 < j_4}} \mathbf{Pr}[\boldsymbol{h}(j_1) = \boldsymbol{h}(j_2) = i_1, \boldsymbol{h}(j_3) = \boldsymbol{h}(j_4) = i_2] \times \mathbf{E}[\boldsymbol{s}(j_1)\boldsymbol{s}(j_2)\boldsymbol{s}(j_3)\boldsymbol{s}(j_4)]x_{j_1}x_{j_2}x_{j_3}x_{j_4}.$$

Note that by 4-wise independence of the hash family from which the function $\boldsymbol{s}$ is drawn, we get that $\mathbf{E}[\boldsymbol{s}(j_1)\boldsymbol{s}(j_2)\boldsymbol{s}(j_3)\boldsymbol{s}(j_4)]$ is 0 unless $j_1 = j_3$ and $j_2 = j_4$ (since $j_1 < j_2$ and $j_3 < j_4$). If $j_1 = j_3$ and $j_2 = j_4$, we additionally have that $\mathbf{Pr}[\boldsymbol{h}(j_1) = \boldsymbol{h}(j_2) = i_1, \boldsymbol{h}(j_3) = \boldsymbol{h}(j_4) = i_2] = 0$ unless $i_1 = i_2$. Thus, the above expression simplifies to

$$\mathbf{Var}(\|\boldsymbol{L}x\|_2^2) = 4\sum_{i} \sum_{j_1 < j_2} \mathbf{Pr}[\boldsymbol{h}(j_1) = \boldsymbol{h}(j_2) = i]x_{j_1}^2 x_{j_2}^2$$

$$= 4\sum_{i}(1/t^2) \sum_{j_1 < j_2} x_{j_1}^2 x_{j_2}^2$$

$$= (2/t)(\|x\|_2^4 - \|x\|_4^4).$$

For $t \geq 1/2\alpha^4\delta$, we have $\mathbf{Var}(\|\boldsymbol{L}x\|_2^2) \leq 4\alpha^4\delta\|x\|_2^4$ and by Chebyshev inequality, we get that $\mathbf{Pr}[\|\boldsymbol{L}x\|_2^2 \geq \|x\|_2^2 + 2\alpha^2\|x\|_2^2] \leq \delta$. By the union bound, we obtain the result. $\square$

We now prove the following theorem.

Theorem 8.3. *There is a turnstile stream algorithm using $O(\varepsilon^{-2}\log(1/\varepsilon)\log d)$ bits of space and outputs an estimate to $\|x\|_\infty$ up to an additive error of $\varepsilon\|x\|_2$ with probability $\geq 9/10$.*

Proof. In the above lemma, setting $\delta = 1/100$, $\alpha = \varepsilon^2/160$, we get that for $t = C/\varepsilon^8$ for a large enough constant C, with probability $\geq 97/100$, $\|\boldsymbol{L}x\|_\infty = \|x\|_\infty \pm (\varepsilon/2)\|x\|_2$ and $\|\boldsymbol{L}x\|_2^2 \leq (1 + 4\varepsilon^4)\|x\|_2^2$. Condition on this event. From [CCF04], if a vector $y \in \mathbb{R}^m$ is being updated in a turnstile stream, the CountSketch data structure with parameters $t = O(1/\varepsilon^2)$ and $r = O(\log m)$ can be used to recover a vector $\hat{y}$ such that with probability $\geq 99/100$,

$$\|y - \hat{y}\|_\infty \leq (\varepsilon/3)\|y\|_2$$

and therefore have that $\|\hat{y}\|_\infty = \|y\|_\infty \pm (\varepsilon/3)\|y\|_2$. Note that the hash functions for the CountSketch data structure can be stored using $O(r \cdot \log m) = O(\log^2 m)$ bits. If the vector y has entries bounded

by poly(d), then the CountSketch data structure can be stored in $O(t \cdot r \cdot \log d)$ bits. We now note that $\boldsymbol{L}x$ is a C/ε^8 dimensional vector with entries bounded by poly(d). Hence using a CountSketch data structure, we can obtain an estimate **Est** that satisfies

$$\textbf{Est} = \|\boldsymbol{L}x\|_\infty \pm (\varepsilon/3)\|\boldsymbol{L}x\|_2 = \|x\|_\infty \pm \varepsilon\|x\|_2.$$

The two stage sketching procedure can be implemented in a turnstile stream as follows: when we receive an update (i, Δ) to the vector x, we supply the update $(\boldsymbol{h}(i), \boldsymbol{s}(i)\Delta)$ to the CountSketch data structure. The overall space usage of the algorithm is $O(\log d + (\log 1/\varepsilon)^2 + \varepsilon^{-2} \log(1/\varepsilon) \log d)$ bits where we use $O(\log d)$ bits to store the hash functions corresponding to $\boldsymbol{L}$, $O(\log 1/\varepsilon)^2$ bits to store the hash functions corresponding to the CountSketch data structure and $O(\varepsilon^{-2} \log(1/\varepsilon) \log d)$ bits to store the CountSketch table itself.

To process each update in the stream, we require $O(\log 1/\varepsilon)$ time in the Word RAM model with a word size $O(\log d)$ as each update only involves evaluating $O(\log 1/\varepsilon)$ constant wise independent hash functions. $\qquad \square$

We will now show that the above is tight up to constant factors.

8.1 Space Lower Bound for estimating $\|x\|_\infty$ in a turnstile stream

To lower bound the space complexity of the turnstile streaming algorithm, we reduce from the Augmented Sparse Set-Disjointness problem. We define this communication problem as a combination of the so-called Augmented INDEX problem [CW09] and Sparse Set-Disjointness problem [DKS12]. In this problem, Alice is given sets $A_1, \ldots, A_t \subseteq [n]$ and Bob is given the sets $B_1, \ldots, B_t \subseteq [n]$. Assume that for all $j \in [t]$, $|A_j| = |B_j| = k$. Bob is given an index j and the sets $A_1, \ldots, A_{j-1}$ and has to output using a one-way message M from Alice if $A_j \cap B_j = \emptyset$ or not. Suppose that Alice and Bob have access to a shared random string. We say that a randomized one-way protocol has δ error if for any instance of the Augmented Sparse Set-Disjointness problem, when Alice and Bob run the protocol Π, Bob outputs the correct answer with probability $\geq 1 - \delta$. Note that in the one-way protocol, Alice can only send a single message M (possibly randomized using the shared random string) to Bob. The δ-error communication complexity of the Augmented Sparse Set-Disjointness problem is then defined as the minimum over all δ-error protocols of the maximum length, measured in terms of number of bits, of the message sent by Alice over all the inputs. By Yao's minimax principle, we can lower bound the communication complexity of the problem by exhibiting a distribution over the inputs such that any *deterministic* protocol must have a large communication complexity for Bob to output correct answer with probability $\geq 1 - \delta$ (over the distribution of inputs). We now show that there is a small enough constant δ for which the δ-error communication complexity of the Augmented Sparse Set-Disjointness problem is $\Omega(tk \log k)$.

Theorem 8.4. *Let t be arbitrary. If $n \geq k^2$, there exists a small enough universal constant δ (independent of t, k and n) such that the δ error randomized communication complexity of the Augmented Sparse Set-Disjointness problem is $\Omega(tk \log k)$.*

Proof. Given parameters k and $\alpha < 1/2$, [DKS12] show that there exists a family $\mathcal{X}$ of $2^{\alpha k \log k}$ subsets of $[k^2]$ such that (i) $|X| = k$ for all $X \in \mathcal{X}$ and (ii) for all $X \neq X' \in \mathcal{X}$, we have $|X \cap X'| \leq \alpha k$. They also show that corresponding to the family $\mathcal{X}$, there is a family $\mathcal{Y}$ of subsets of $[k^2]$ such that (i) $|Y| = k$ for all $Y \in \mathcal{Y}$, (ii) $|\mathcal{Y}| \leq ak \log k$ for an absolute constant $a = a(\alpha)$ (independent of k) and (iii) for $X \neq X' \in \mathcal{X}$, there is at least one set $Y \in \mathcal{Y}$ such that exactly one of the sets $X \cap Y$ and $X' \cap Y$ is non-empty. The last property implies that the sequence $(\text{DISJ}(X, Y))_{Y \in \mathcal{Y}}$ is distinct for each $X \in \mathcal{X}$. Here $\text{DISJ}(X, Y) = 1$ if $X \cap Y = \emptyset$ and 0 otherwise.

We will now define a distribution over the instances of the Augmented Set-Disjointness problem such that any deterministic protocol must have Alice sending a message with $\Omega(tk \log k)$ bits which will prove the theorem by Yao's minimax principle. Let $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_t \sim \mathcal{X}$ and $\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_t \sim \mathcal{Y}$ be drawn independently. We let $\boldsymbol{X} = (\boldsymbol{X}_1, \ldots, \boldsymbol{X}_t)$ denote the sets given to Alice and $\boldsymbol{Y} = (\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_t)$ denote the sets given to Bob. Let $\boldsymbol{i} \sim [t]$ drawn uniformly at random be the index given to Bob. Let $M(\boldsymbol{X})$ be the message sent by Alice when running an arbitrary deterministic protocol with an error δ over the distribution defined by the random variables $\boldsymbol{X}$, $\boldsymbol{Y}$, and $\boldsymbol{i}$. By the chain rule of entropy,

$$H(\boldsymbol{X} \mid M(\boldsymbol{X})) = \sum_{i \in [t]} H(\boldsymbol{X}_i \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}).$$

As $\boldsymbol{X}_i$ is uniquely identifiable by the sequence $(\text{DISJ}(\boldsymbol{X}_i, Y))_{Y \in \mathcal{Y}}$, we have

$$\begin{aligned}
H(\boldsymbol{X} \mid M(\boldsymbol{X})) &= \sum_{i \in [t]} H((\text{DISJ}(\boldsymbol{X}_i, Y))_{Y \in \mathcal{Y}} \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}) \\
&\leq \sum_{i \in [t]} \sum_{Y \in \mathcal{Y}} H(\text{DISJ}(\boldsymbol{X}_i, Y) \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}) \qquad \text{(sub-additivity)} \\
&= \sum_{i \in [t]} \sum_{y \in \mathcal{Y}} H(\text{DISJ}(\boldsymbol{X}_i, \boldsymbol{Y}_i) \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}, \boldsymbol{Y}_i = y) \\
&= \sum_{i \in [t]} |\mathcal{Y}| H(\text{DISJ}(\boldsymbol{X}_i, \boldsymbol{Y}_i) \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}, \boldsymbol{Y}_i) \qquad \text{(since } \boldsymbol{Y}_i \text{ is uniform over } \mathcal{Y}) \\
&= |\mathcal{Y}| \sum_{i \in [t]} H(\text{DISJ}(\boldsymbol{X}_i, \boldsymbol{Y}_i) \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}, \boldsymbol{Y}).
\end{aligned}$$

Here the last equality follows from the fact that $\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_t$ are mutually independent and are also independent of $\boldsymbol{X}_{<i}$ and $M(\boldsymbol{X})$. As the *deterministic* protocol has δ error over the distribution of the instances we defined above, we have that there is a deterministic function f (which Bob runs to output his answer) such that

$$\Pr_{X,Y,i}[f(M(\boldsymbol{X}), \boldsymbol{i}, \boldsymbol{Y}, \boldsymbol{X}_{<i}) = \text{DISJ}(\boldsymbol{X}_i, \boldsymbol{Y}_i)] \geq 1 - \delta.$$

We then have that with probability $1 - \sqrt{\delta}$ over $\boldsymbol{i}$ that

$$\Pr_{X,Y}[f(M(\boldsymbol{X}), \boldsymbol{i}, \boldsymbol{Y}, \boldsymbol{X}_{<i}) = \text{DISJ}(\boldsymbol{X}_i, \boldsymbol{Y}_i)] \geq 1 - \sqrt{\delta}.$$

Let $\text{GOOD} \subseteq [t]$ denote all the indices i for which the above holds. We have $|\text{GOOD}| \geq (1 - \sqrt{\delta})t$. We now note that $\text{DISJ}(\boldsymbol{X}_i, \boldsymbol{Y}_i)$ is a 0/1 random variable. By Fano's inequality, for all $i \in \text{GOOD}$, we obtain

$$H(\text{DISJ}(\boldsymbol{X}_i, \boldsymbol{Y}_i) \mid M(\boldsymbol{X}), \boldsymbol{Y}, \boldsymbol{X}_{<i}) \leq H(\sqrt{\delta}).$$

For $i \notin \textsc{Good}$, we simply have $H(\textsc{Disj}(\boldsymbol{X}_i, \boldsymbol{Y}_i) \mid M(\boldsymbol{X}), \boldsymbol{Y}, \boldsymbol{X}_{<i}) \leq H(\textsc{Disj}(\boldsymbol{X}_i, \boldsymbol{Y}_i)) \leq 1$. Thus,

$$
\begin{aligned}
H(\boldsymbol{X} \mid M(\boldsymbol{X})) &\leq |\mathcal{Y}| \sum_{i \in [t]} H(\textsc{Disj}(\boldsymbol{X}_i, \boldsymbol{Y}_i) \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}, \boldsymbol{Y}) \\
&= |\mathcal{Y}| \sum_{i \in \textsc{Good}} H(\textsc{Disj}(\boldsymbol{X}_i, \boldsymbol{Y}_i) \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}, \boldsymbol{Y}) \\
&\quad + |\mathcal{Y}| \sum_{i \notin \textsc{Good}} H(\textsc{Disj}(\boldsymbol{X}_i, \boldsymbol{Y}_i) \mid M(\boldsymbol{X}), \boldsymbol{X}_{<i}, \boldsymbol{Y}) \\
&\leq |\mathcal{Y}| t (1 - \sqrt{\delta}) H(\sqrt{\delta}) + |\mathcal{Y}| t \sqrt{\delta} \\
&= |\mathcal{Y}| t (\sqrt{\delta} + (1 - \sqrt{\delta}) H(\sqrt{\delta})).
\end{aligned}
$$

Now, $H(\boldsymbol{X} \mid M(\boldsymbol{X})) \geq H(\boldsymbol{X}) - H(M(\boldsymbol{X}))$ by the chain rule and sub-additivity which implies from the above inequality that

$$
H(M(\boldsymbol{X})) \geq H(\boldsymbol{X}) - |\mathcal{Y}| t (\sqrt{\delta} + (1 - \sqrt{\delta}) H(\sqrt{\delta})).
$$

As $H(\boldsymbol{X}) = H((\boldsymbol{X}_1, \ldots, \boldsymbol{X}_t)) = t H(\boldsymbol{X}_1) = t \alpha k \log k$, we have

$$
\begin{aligned}
H(M(\boldsymbol{X})) &\geq t \alpha k \log k - t a k \log k (\sqrt{\delta} + (1 - \sqrt{\delta}) H(\sqrt{\delta})) \\
&\geq t k \log k (\alpha - a(\sqrt{\delta} + (1 - \sqrt{\delta}) H(\sqrt{\delta}))).
\end{aligned}
$$

Now we note that $H(\sqrt{\delta}) \leq 2\delta^{1/4}$ and get $H(M(\boldsymbol{X})) \geq t k \log k (\alpha - a(\sqrt{\delta} + 2\delta^{1/4}))$. As $a = a(\alpha)$ is purely a function of α independent of k, by picking δ to be a small enough function of α, we get $H(M(\boldsymbol{X})) \geq (\alpha/2) t k \log k$. Finally, this implies that $\max_{\boldsymbol{X}} |M(\boldsymbol{X})| \geq H(M(\boldsymbol{X})) \geq (\alpha t k \log k)/2$. Picking $\alpha = 1/2$, we obtain that there is a small enough constant δ and a product distribution $\mathcal{D}$ over $\mathcal{X} \otimes \mathcal{Y} \otimes [t]$ such that any *deterministic* one-way protocol that solves the Augmented Sparse Set-Disjointness problem with probability $\geq 1 - \delta$ over the distribution $\mathcal{D}$ must have a communication complexity of $\Omega(t k \log k)$ bits. $\square$

Using the above lower bound, we can now show that any turnstile streaming algorithm that approximates the ℓ_∞ norm of a d dimensional vector x with integer coordinates bounded in absolute value by $\mathrm{poly}(d)$ up to an additive error of $\varepsilon \|x\|_2$ must use $O(\varepsilon^{-2} \log(1/\varepsilon) \log d)$ bits of space.

Theorem 8.5. *There exists a small enough constant δ such that if $\varepsilon \geq 6((\log d)/d)^{1/4}$, any turnstile streaming algorithm that estimates $\|x\|_\infty$ up to an additive error of $\varepsilon \|x\|_2$ with probability $\geq 1 - \delta$, of a d-dimensional vector x with integer entries bounded in absolute value by $\mathrm{poly}(d)$, must use $\Omega(\varepsilon^{-2} \log(1/\varepsilon) \log(d))$ bits of space.*

Proof. Let $t = \log d$, $k = 1/\varepsilon^2$ and $n = 1/\varepsilon^4$. From the above theorem, the Augmented Sparse Set-Disjointness problem with these parameters has a randomized communication complexity of $\Omega(\varepsilon^{-2} \log(1/\varepsilon) \log d)$ bits.

Suppose given an instance of the Augmented Sparse Set-Disjointness problem, Alice computes a vector $x \in \mathbb{R}^{(\log d) \cdot 1/\varepsilon^4}$ as follows: the vector x is divided into $\log d$ blocks—one for each of the sets $A_1, \ldots, A_{\log d}$. The i-th block of vector x, for $i = 1, \ldots, t$, is defined to be $10^{t-i} \cdot a_i$ where a_i is a binary-vector representation of the set A_i. As $t = \log d$, we obtain that $\|x\|_\infty \leq \mathrm{poly}(d)$. Let $\boldsymbol{M}$ be a randomized turnstile streaming algorithm for estimating $\|x\|_\infty$. Let $\boldsymbol{M}(x)$ be the state of the turnstile streaming algorithm after feeding the coordinates of the vector x to $\boldsymbol{M}$. Alice transmits the state $\boldsymbol{M}(x)$ to Bob. As Bob has an index i and the sets $A_1, \ldots, A_{i-1}$, Bob can construct a vector $y \in \mathbb{R}^{(\log d) \cdot 1/\varepsilon^4}$ such that the j-th block of vector y for $j = 1, \ldots, i - 1$ is the same as the

j-th block of the vector x. The rest of the blocks of y are set to be 0. We now note that the only non-zero blocks of the vector $x - y$ are $i, i + 1, \ldots, t$.

Bob feeds the updates corresponding to the vector $-y$ to the streaming algorithm $\boldsymbol{M}$ starting with the state $\boldsymbol{M}(x)$ to obtain $\boldsymbol{M}(x - y)$. Finally, Bob defines a vector z with the i-th block being the vector $10^{t-i} \cdot b_i$ where b_i is the binary vector corresponding to the set B_i. The rest of the blocks of z are set to be 0. Bob finally updates the state of the streaming algorithm to obtain $\boldsymbol{M}(x - y + z)$. If $A_i \cap B_i = \emptyset$, we have $\|x - y + z\|_\infty = 10^{t-i}$ and if $A_i \cap B_i \neq \emptyset$, then $\|x - y + z\|_\infty = 2 \cdot 10^{t-i}$. Additionally, we have $\|x - y + z\|_2^2 \leq 4 \cdot 10^{2(t-i)} \cdot k + \sum_{j=i+1}^{t} 10^{2(t-j)} \cdot k \leq 5 \cdot 10^{2(t-i)} \cdot k$. Thus, $\|x - y + z\|_2 \leq 3 \cdot 10^{t-i} \cdot (1/\varepsilon)$ since $k = 1/\varepsilon^2$. Thus an approximation of $\|x - y + z\|_\infty$ up to an additive error of $(\varepsilon/6)\|x - y + z\|_2$ lets Bob output the correct answer for the instance of Augmented Sparse Set-Disjointness problem.

By the lower bound on communication complexity of the Augmented Sparse Set-Disjointess problem, we obtain that any turnstile streaming algorithm that outputs, with probability $1 - \delta$ for a small enough constant δ, an approximation to $\|x\|_\infty$ up to an additive error of $(\varepsilon/6)\|x\|_\infty$, for a $(1/\varepsilon^4) \log d$ dimensional vector x with coordinates of absolute values bounded by $\mathrm{poly}(d)$, must use $\Omega(\varepsilon^{-2} \log(1/\varepsilon) \log d)$ bits. As $\varepsilon \geq ((\log d)/d)^{1/4}$ implies $d \geq \varepsilon^{-4} \log d$, we obtain the result. $\qquad\square$

Note that the above lower bound crucially uses that the algorithm is a turnstile streaming algorithm and does not hence lower bound the space complexity of the algorithms in the insertion-only streams where only nonnegative updates are allowed to the vector x being maintained in the stream.

8.2 Tighter bounds for vectors with large $\|x\|_\infty$

The lower bound in the previous section shows that $\Theta(\varepsilon^{-2} \log(d) \log(1/\varepsilon))$ bits of space is both necessary and sufficient to approximate $\|x\|_\infty$ up to an additive error $\varepsilon\|x\|_2$. The hard instance in the lower bound has the property that $\|x\|_\infty = O(\varepsilon\|x\|_2)$. We show that assuming the vector x satisfies, $\|x\|_\infty \geq c\|x\|_2$ for a constant c, we can beat the lower bound and obtain a better than $\varepsilon\|x\|_2$ additive error using $O(\varepsilon^{-2} \log d)$ bits of space. Note that the condition $\|x\|_\infty \geq c\|x\|_2$ is natural in certain settings where the coordinates of the vector x follow the Zipf's law. We prove the following theorem.

Theorem 8.6. *Suppose a d dimensional vector x is being maintained in a turnstile stream. Assume that the coordinates of x are integers and are bounded in absolute value by $\mathrm{poly}(d)$. If x is such that $\|x\|_\infty \geq c\|x\|_2$ for a universal constant c, then there is an algorithm that uses $O(\varepsilon^{-2} \log d)$ bits of space and outputs an estimate $\boldsymbol{Est}$ that satisfies*

$$\boldsymbol{Est} = \|x\|_\infty \pm \varepsilon\|x\|_2$$

with probability $\geq 9/10$. The update time of the algorithm is $O(\log 1/\varepsilon)$ in the Word RAM model with a word size $\Omega(\log d)$.

Proof. Without loss of generality, assume $\varepsilon \leq c/10$. Let $\textsc{Large} = \{ i \mid |x_i| \geq (c/2)\|x\|_2 \}$. We have $|\textsc{Large}| \leq 4/c^2$. If $\boldsymbol{L} : \mathbb{R}^d \to \mathbb{R}^{d'}$ is a randomized linear map to $d' = \mathrm{poly}(1/\varepsilon)$ dimension constructed using a 2-wise independent hash function $\boldsymbol{h}$ and a 4-wise independent sign function $\boldsymbol{s}$ as in Lemma 8.1, then we have that

1. all the coordinates in $\textsc{Large}$ are hashed to different coordinates,

2. $\|\boldsymbol{L}x\|_2^2 \leq (1 + \varepsilon^8)\|x\|_2^2$,

3. for all $i \notin \text{LARGE}$, $|(\boldsymbol{L}x)_{\boldsymbol{h}(i)}| \leq (c/2 + \varepsilon^2/8)\|x\|_2 \leq (3c/5)\|x\|_2$, and

4. for all $i \in \text{LARGE}$, $|(\boldsymbol{L}x)_{\boldsymbol{h}(i)}| = |x_i| \pm (\varepsilon^2/8)\|x\|_2$.

Now, let $r = O(\log 1/\varepsilon)$, $t = O(1/\varepsilon^2)$ and $b = 1/\varepsilon$. Instantiate a CountSketch data structure CS : $\mathbb{R}^{d'} \rightarrow \mathbb{R}^{rt}$ with these parameters and derandomized using HashPRG as in Theorem 6.2 with a word size $w = \Omega(\log d)$. From Theorem 6.2, the parameters (random seed for HashPRG) of the map CS and the value $\text{CS}(\boldsymbol{L}x)$ can be stored using $O(rt + b\log_b d') = O(\varepsilon^{-2}\log 1/\varepsilon + \varepsilon^{-1}\log_{\varepsilon^{-1}}\text{poly}(1/\varepsilon)) = O(\varepsilon^{-2}\log 1/\varepsilon)$ *words* of space. We also have that the update time of the CountSketch data structure instantiated with these parameters is $O(\log 1/\varepsilon)$ in the Word RAM model with a word size $\Omega(\log d)$. If x receives a turnstile update (i, Δ), then

$$\boldsymbol{L}(x + \Delta e_i) = \boldsymbol{L}x + \Delta \cdot (\boldsymbol{L}e_i).$$

By definition of the map $\boldsymbol{L}$, the vector $\boldsymbol{L}e_i$ is nonzero in exactly one coordinate $\boldsymbol{h}(i)$. Thus, we further obtain

$$\text{CS}(\boldsymbol{L}(x + \Delta e_i)) = \text{CS}(\boldsymbol{L}x + \boldsymbol{s}(i)\Delta e_{\boldsymbol{h}(i)}).$$

Now, by Theorem 6.2, the vector $\text{CS}(\boldsymbol{L}(x + \Delta e_i))$ can be computed using the value of $\text{CS}(\boldsymbol{L}x)$ in time $O(r\log_b d') = O(\log 1/\varepsilon)$ time in Word RAM model. Thus, the randomized two-level sketch $\text{CS} \circ \boldsymbol{L}$ can be applied to the underlying vector x in a turnstile stream using a total space of $O(\varepsilon^{-2}\log 1/\varepsilon)$ *words* of space and each turnstile update can be processed in $O(\log 1/\varepsilon)$ time on a Word RAM machine with a word size $\Omega(\log d)$.

Theorem 6.2 also gives the following recovery guarantees: for any $i \in [d']$ and $\alpha < 1$, we can recover a value $\widehat{(\boldsymbol{L}x)_i}$ such that

$$\Pr_{\text{CS}}[|\widehat{(\boldsymbol{L}x)_i} - (\boldsymbol{L}x)_i| \geq \alpha\varepsilon\|\boldsymbol{L}x\|_2] \leq \exp(-\alpha^2 r) + 1/\text{poly}(d).$$

Setting $\alpha = 1/\sqrt{\log 1/\varepsilon}$, a union bound over the indices in the set $h(\text{LARGE}) \subseteq [d']$ gives that with probability $\geq 99/100$ over CS, for all $i \in h(\text{LARGE})$, $|\widehat{(\boldsymbol{L}x)_i} - (\boldsymbol{L}x)_i| \leq (\varepsilon/\sqrt{\log 1/\varepsilon})\|\boldsymbol{L}x\|_2$ and setting $\alpha = 1$ and a union bound over all the coordinates $i \in [d']$ gives that with probability $\geq 99/100$, for all $i \in [d]$, $|\widehat{(\boldsymbol{L}x)_i} - (\boldsymbol{L}x)_i| \leq \varepsilon\|\boldsymbol{L}x\|_2$. Conditioned on the properties of the map $\boldsymbol{L}$ above, overall, we obtain that with a probability $\geq 9/10$,

$$\|\boldsymbol{L}x\|_\infty = \|x\|_\infty \pm \frac{2\varepsilon}{\sqrt{\log 1/\varepsilon}}\|x\|_2.$$

Hence, $\|x\|_\infty$ can be estimated to an additive error of $(\varepsilon/\sqrt{\log 1/\varepsilon})\|x\|_2$ using only $O(\varepsilon^{-2}\log 1/\varepsilon)$ words of space and the time to update the state in a turnstile stream is $O(\log 1/\varepsilon)$ in the Word RAM model with a word size $\Omega(\log d)$. Setting $\varepsilon' = 2\varepsilon/\sqrt{\log 1/\varepsilon}$, we obtain the result. $\qquad\square$

Acknowledgements

References

[AGM12]   Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Graph sketches: Sparsification, Spanners, and Subgraphs. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, pages 5–14, 2012.

[AKO11]   Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Streaming algorithms via precision sampling. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 363–372. IEEE, 2011.

[And17]   Alexandr Andoni. High frequency moments via max-stability. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6364–6368. IEEE, 2017.

[ANPW13]   Alexandr Andoni, Huy L Nguyen, Yury Polyanskiy, and Yihong Wu. Tight lower bound for linear sketches of moments. In *International Colloquium on Automata, Languages, and Programming*, pages 25–32. Springer, 2013.

[BCIW16]   Vladimir Braverman, Stephen R. Chestnut, Nikita Ivkin, and David P. Woodruff. Beating CountSketch for heavy hitters in insertion streams. In *STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 740–753. ACM, New York, 2016.

[BGKS06]   Lakshminath Bhuvanagiri, Sumit Ganguly, Deepanjan Kesh, and Chandan Saha. Simpler algorithm for estimating frequency moments of data streams. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 708–713. ACM Press, 2006.

[BGW20]   Mark Braverman, Sumegha Garg, and David P. Woodruff. The coin problem with applications to data streams. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science*, pages 318–329. IEEE Computer Soc., Los Alamitos, CA, 2020.

[BO10]   Vladimir Braverman and Rafail Ostrovsky. Zero-one frequency laws. In *STOC'10—Proceedings of the 2010 ACM International Symposium on Theory of Computing*, pages 281–290. ACM, New York, 2010.

[BYJKS04]   Ziv Bar-Yossef, Thathachar S Jayram, Ravi Kumar, and D Sivakumar. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences*, 68(4):702–732, 2004.

[CCF04]   Moses Charikar, Kevin C. Chen, and Martin Farach-Colton. Finding frequent items in data streams. *Theoretical Computer Science*, 312(1):3–15, 2004.

[Cha02]   Moses S Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 380–388, 2002.

[CIW23]   Justin Y. Chen, Piotr Indyk, and David P. Woodruff. Space-optimal profile estimation in data streams with applications to symmetric functions. Unpublished, 2023.

[CJ19]   Graham Cormode and Hossein Jowhari. L_p samplers and their applications: A survey. *ACM Computing Surveys (CSUR)*, 52(1):1–31, 2019.

[CKS03]   Amit Chakrabarti, Subhash Khot, and Xiaodong Sun. Near-optimal lower bounds on the multi-party communication complexity of set disjointness. In *18th IEEE Annual Conference on Computational Complexity, 2003. Proceedings.*, pages 107–117. IEEE, 2003.

[CPT15]   Tobias Christiani, Rasmus Pagh, and Mikkel Thorup. From independence to expansion and back again. In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pages 813–820. ACM, New York, 2015.

[CW09]   Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, page 205–214, New York, NY, USA, 2009. Association for Computing Machinery.

[Die96]   Martin Dietzfelbinger. Universal hashing and k-wise independent random variables via integer arithmetic without primes. In *STACS 96, 13th Annual Symposium on Theoretical Aspects of Computer Science, Grenoble, France, February 22-24, 1996, Proceedings*, volume 1046 of *Lecture Notes in Computer Science*, pages 569–580. Springer, 1996.

[DKS10]   Anirban Dasgupta, Ravi Kumar, and Tamás Sarlós. A sparse Johnson-Lindenstrauss transform. In *STOC'10—Proceedings of the 2010 ACM International Symposium on Theory of Computing*, pages 341–350. ACM, New York, 2010.

[DKS12]   Anirban Dasgupta, Ravi Kumar, and D. Sivakumar. Sparse and lopsided set disjointness via information theory. In *Approximation, randomization, and combinatorial optimization*, volume 7408 of *Lecture Notes in Comput. Sci.*, pages 517–528. Springer, Heidelberg, 2012.

[DPS11]   Matei David, Periklis A Papakonstantinou, and Anastasios Sidiropoulos. How strong is Nisan's pseudo-random generator? *Information processing letters*, 111(16):804–808, 2011.

[FIS05]   Gereon Frahling, Piotr Indyk, and Christian Sohler. Sampling in dynamic data streams and applications. In *Proceedings of the twenty-first annual symposium on Computational geometry*, pages 142–149, 2005.

[Gan15]   Sumit Ganguly. Taylor polynomial estimator for estimating frequency moments. In *Automata, languages, and programming. Part I*, volume 9134 of *Lecture Notes in Comput. Sci.*, pages 542–553. Springer, Heidelberg, 2015.

[GKM18]   Parikshit Gopalan, Daniel M. Kane, and Raghu Meka. Pseudorandomness via the discrete fourier transform. *SIAM Journal on Computing*, 47(6):2451–2487, 2018.

[GKMS03]   Anna C. Gilbert, Yannis Kotidis, S Muthukrishnan, and Martin J Strauss. One-pass wavelet decompositions of data streams. *IEEE Transactions on knowledge and data engineering*, 15(3):541–554, 2003.

[GW18]   Sumit Ganguly and David P. Woodruff. High probability frequency moment sketches. In *45th International Colloquium on Automata, Languages, and Programming*, volume 107 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 58, 15. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2018.

[HNO08]   Nicholas JA Harvey, Jelani Nelson, and Krzysztof Onak. Sketching and streaming entropy via approximation theory. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 489–498. IEEE, 2008.

[Ind06]   Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *Journal of the ACM (JACM)*, 53(3):307–323, 2006.

[IW05]   Piotr Indyk and David Woodruff. Optimal approximations of the frequency moments of data streams. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 202–208. ACM, New York, 2005.

[Jay09]   T. S. Jayram. Hellinger strikes back: a note on the multi-party information complexity of AND. In *Approximation, randomization, and combinatorial optimization*, volume 5687 of *Lecture Notes in Comput. Sci.*, pages 562–573. Springer, Berlin, 2009.

[JST11]   Hossein Jowhari, Mert Sağlam, and Gábor Tardos. Tight bounds for ℓ_p samplers, finding duplicates in streams, and related problems. In *Proceedings of the thirtieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 49–58, 2011.

[JW18]   Rajesh Jayaram and David P. Woodruff. Perfect L_p sampling in a data stream. In *Symposium on Foundations of Computer Science (FOCS)*, pages 544–555, 2018.

[KNPW11]   Daniel M. Kane, Jelani Nelson, Ely Porat, and David P. Woodruff. Fast moment estimation in data streams in optimal space. In *STOC'11—Proceedings of the 43rd ACM Symposium on Theory of Computing*, pages 745–754. ACM, New York, 2011.

[KNW10]   Daniel M. Kane, Jelani Nelson, and David P. Woodruff. On the exact space complexity of sketching and streaming small norms. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1161–1178. SIAM, Philadelphia, PA, 2010.

[Li08]   Ping Li. Estimators and tail bounds for dimension reduction in l_α ($0 < \alpha \leq 2$) using stable random projections. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 10–19. ACM, New York, 2008.

[LNNT16]   Kasper Green Larsen, Jelani Nelson, Huy L. Nguyen, and Mikkel Thorup. Heavy hitters via cluster-preserving clustering. In *57th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2016*, pages 61–70. IEEE Computer Soc., Los Alamitos, CA, 2016.

[MP14]   Gregory T. Minton and Eric Price. Improved concentration bounds for Count-Sketch. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 669–686. ACM, New York, 2014.

[MW10]   Morteza Monemizadeh and David P. Woodruff. 1-pass relative-error L_p-sampling with applications. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1143–1160. SIAM, Philadelphia, PA, 2010.

[Nis92]   Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.

[NW10]   Jelani Nelson and David P Woodruff. Fast Manhattan sketches in data streams. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 99–110, 2010.

[NZ96]   Noam Nisan and David Zuckerman. Randomness is linear in space. *J. Comput. System Sci.*, 52(1):43–52, 1996.

[PP08]   Anna Pagh and Rasmus Pagh. Uniform hashing in constant time and optimal space. *SIAM Journal on Computing*, 38(1):85–96, 2008.

[PT22]   Rasmus Pagh and Mikkel Thorup. Improved utility analysis of private CountSketch. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

A   Nisan's Pseudorandom Generator

We say that a *randomized* program uses space w with a block size n if it accepts its random bits as an n bit block at a time and uses at most space w between the different blocks of random bits. Such programs can be modeled as a finite state machine over at most 2^w states, taking an input string over the alphabet $\{0,1\}^n$. Nisan [Nis92] constructed a pseudorandom generator which requires only a small uniform random seed that "fools" a space w program with a block size n.

Let $\boldsymbol{h}_1,\ldots,\boldsymbol{h}_k$ be independent hash functions drawn from a 2-wise independent hash family $\mathcal{H} = \{\, h : \{0,1\}^n \to \{0,1\}^n \,\}$. These hash functions together with $\boldsymbol{x} \in \{0,1\}^n$, sampled uniformly at random, serve as the *seed* of the generator $G_k : \{0,1\}^n \to \{0,1\}^{2^k \cdot n}$, defined recursively as follows:

$$G_0(x) := x$$
$$G_k(x, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_k) := G_{k-1}(x, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_{k-1}) \circ G_{k-1}(\boldsymbol{h}_k(x), \boldsymbol{h}_1, \ldots, \boldsymbol{h}_{k-1}),$$

where $\circ$ denotes the string concatenation. For a given choice of $\boldsymbol{h}_1,\ldots,\boldsymbol{h}_k$, define the distribution $G_k(*, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_k)$ over bitstrings of length $2^k \cdot n$ to be the distribution of $G_k(\boldsymbol{x}, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_k)$ for random $\boldsymbol{x} \in \{0,1\}^n$. Nisan showed that for any fixed FSM with at most 2^w states over alphabet $\{0,1\}^n$, with high probability over the hash functions $\boldsymbol{h}_1,\ldots,\boldsymbol{h}_k$, the distribution $G_k(*, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_k)$ is indistinguishable from the uniform distribution over $\{0,1\}^{2^k \cdot n}$. The power of Nisan's generator is summarized by the following lemma (using notation from section 3):

Lemma A.1. *There exists a constant $c > 0$ such that given integers n and $w \le cn$ and parameter $k \le cn$, for any FSM Q with 2^w states, if $\boldsymbol{h}_1, \ldots, \boldsymbol{h}_k : \{0,1\}^n \to \{0,1\}^n$ are drawn independently from a 2-wise independent hash family, then with probability $\ge 1 - 2^{-cn}$,*

$$\|Q(G_k(*, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_k)) - Q((U_n)^{2^k})\| \le 2^{-cn}$$

where $\|M\| := \max_i \sum_j |M_{ij}|$.

Note that $\|M\| = \max_{x \ne 0} \|Mx\|_\infty / \|x\|_\infty$ where $\|x\|_\infty = \max_i |x_i|$. We therefore have that for any two matrices A and B, $\|A + B\| \le \|A\| + \|B\|$ and $\|AB\| \le \|A\|\|B\|$. Therefore, we obtain that with probability $\ge 1 - 2^{-cn}$ over the hash functions $\boldsymbol{h}_1, \ldots, \boldsymbol{h}_k$, we have that the total variation distance between the distribution of final state using a random string drawn from $(U_n)^{2^k}$ and a random string drawn from $G_k(*, \boldsymbol{h}_1, \ldots, \boldsymbol{h}_k)$ is at most 2^{-cn}.

Specifically, for a $w = O(\log d)$ space algorithms using $\mathrm{poly}(d)$ random bits, we have that we can use Nisan's Generator with $n, k = O(\log d)$. The time to evaluate a block of n random bits is then

$\Omega(k) = \Omega(\log d)$ in the Word RAM model as the k hash functions have to be applied sequentially to the random seed. Using our new pseudorandom generator, which we call HashPRG, we show that we can set $k = O(1)$ at the expense of using more space to store the hash functions.

B   Finding Heavy Entries

See Section 6 for the definition of CountSketch data structure. Note that $[t]$ denotes the range of locations the coordinate gets hashed into and r denotes the number of repetitions. Further for each $\ell \in [d]$, $\hat{x}_\ell$ defined in (10) denotes our estimate for the value of coordinate x_ℓ.

Jowhari, Sağlam and Tardos [JST11] show that if $r = O(\log d)$, then with probability $\geq 1 - 1/\operatorname{poly}(d)$, for all ℓ,

$$|x_\ell - \hat{x}_\ell| \leq \frac{\|x\|_p}{t^{1/p}}.$$

By picking $t = (\phi/10)^{-p}$ we obtain that with probability $1 - 1/\operatorname{poly}(d)$, for all $\ell \in [d]$, $|x_\ell - \hat{x}_\ell| \leq (\phi/10)\|x\|_p$. The algorithm uses $O((\phi/10)^{-p} \log^2 d)$ bits of space and has an update time of $O(\log d)$ per stream element. Condition on the event that for all $\ell \in [d]$, $|x_\ell - \hat{x}_\ell| \leq (\phi/10)\|x\|_p$ for all $\ell \in [d]$.

Concurrently, run the algorithm of [KNW10] with $\varepsilon = 1/4$ to obtain a value v such that with probability $\geq 99/100$

$$(9/10)\|x\|_p \leq v \leq (11/10)\|x\|_p.$$

Note that for constant ε, their algorithm uses $O(\log d)$ bits of space and has an update time of $O(1)$ per stream element in the Word RAM model. Condition on this event as well.

Now, let L' be the set returned by heavy-hitters algorithm of [LNNT16] with parameter ϕ. Their algorithm uses $O(\phi^{-p} \log^2(d))$ bits of space and has an update time of $O(\log d)$ per stream element. At the end of processing the stream, in time $O(\phi^{-p} \operatorname{poly}(\log d))$, they return a set L' satisfying $|L'| = O(\phi^{-p})$ and

$$L' \supseteq \{\, \ell \mid |x_\ell| \geq \phi\|x\|_p \,\}.$$

The set L' contains all the heavy-hitters and may contain additional coordinates as well. To filter the list L', we use the estimates $\hat{x}_\ell$ given by the CountSketch data structure. Define

$$L = \{\, i \in L' \mid |\hat{x}_i| \geq (8/10)\phi v \,\}.$$

Conditioned on the correctness of $L', \hat{x}_\ell$ and the estimate v, we prove properties about the set L. If $|x_\ell| \geq \phi\|x\|_p$, then $|\hat{x}_\ell| \geq (9\phi/10)\|x\|_p \geq (9\phi/11)v \geq (8\phi/10)v$. Therefore, $\ell \in L$. On the other hand, if $\ell \in L$ then

$$|x_\ell| \geq |\hat{x}_\ell| - (\phi/10)\|x\|_p \geq (8/10)\phi v - (\phi/10)\|x\|_p \geq (6\phi/10)\|x\|_p.$$

As $|x_\ell - \hat{x}_\ell| \leq (\phi/10)\|x\|_p$ for all ℓ and $|x_\ell| \geq (6\phi/10)\|x\|_p$ for all $\ell \in L$, we also obtain that for all $\ell \in L$, $\operatorname{sign}(x_\ell) = \operatorname{sign}(\hat{x}_\ell)$. As the list L' has size at most $O(\phi^{-p})$, the post processing can be performed in time $O(\phi^{-p} \operatorname{poly}(\log d))$. Thus, we over all have the following lemma.

Lemma B.1. *Given a stream of updates $(i_1, v_1), \ldots, (i_m, v_m) \in [d] \times \{-M, \ldots, M\}$ for $m, M \leq \operatorname{poly}(d)$, a parameter ϕ and $p \in (0, 2)$, there is a streaming algorithm that uses $O(\phi^{-p} \log^2(d))$ bits of space and has an update time of $O(\log d)$ per stream element and outputs a set $L \subseteq [d]$ at the end of the stream that with probability $\geq 9/10$ satisfies:*

1. $L \supseteq \{ \ell \in [d] \mid |x_\ell| \geq \phi \|x\|_p \}$.

2. For all $\ell \in L$, $|x_\ell| \geq (6\phi/10)\|x\|_p$.

3. For all $\ell \in L$, the algorithm also outputs $\text{sign}(x_\ell)$.

At the end of the stream, the algorithm takes only $O(\phi^{-p} \operatorname{poly}(\log d))$ time to compute the set L.