Truly Perfect Samplers for Data Streams and Sliding Windows

Rajesh Jayaram
Carnegie Mellon University
Pittsburgh, USA
rkjayara@cs.cmu.edu

David P. Woodruff
Carnegie Mellon University
Pittsburgh, USA
dwoodruf@cs.cmu.edu

Samson Zhou
Carnegie Mellon University
Pittsburgh, USA
samsonzhou@gmail.com

ABSTRACT

In the G-sampling problem, the goal is to output an index i of a vector $f \in \mathbb{R}^n$, such that for all coordinates $j \in [n]$,

$$\mathbf{Pr}[i = j] = (1 \pm \epsilon)\frac{G(f_j)}{\sum_{k\in[n]} G(f_k)} + \gamma,$$

where $G : \mathbb{R} \to \mathbb{R}_{\geq 0}$ is some non-negative function. If $\epsilon = 0$ and $\gamma = 1/\text{poly}(n)$, the sampler is called *perfect*. In the data stream model, f is defined implicitly by a sequence of updates to its coordinates, and the goal is to design such a sampler in small space. Jayaram and Woodruff (FOCS 2018) gave the first perfect L_p samplers in turnstile streams, where $G(x) = |x|^p$, using polylog(n) space for $p \in (0, 2]$. However, to date all known sampling algorithms are not *truly perfect*, since their output distribution is only point-wise $\gamma = 1/\text{poly}(n)$ close to the true distribution. This small error can be significant when samplers are run many times on successive portions of a stream, and leak potentially sensitive information about the data stream.

In this work, we initiate the study of *truly perfect* samplers, with $\epsilon = \gamma = 0$, and comprehensively investigate their complexity in the data stream and sliding window models. We begin by showing that sublinear space truly perfect sampling is impossible in the turnstile model, by proving a lower bound of $\Omega\left(\min\left\{n, \log\frac{1}{\gamma}\right\}\right)$ for any G-sampler with point-wise error γ from the true distribution. We then give a general time-efficient sublinear-space framework for developing truly perfect samplers in the insertion-only streaming and sliding window models. As specific applications, our framework addresses L_p sampling for all $p > 0$, e.g., $\tilde{O}\left(n^{1-1/p}\right)$ space for $p \geq 1$, concave functions, and a large number of measure functions, including the $L_1 - L_2$, Fair, Huber, and Tukey estimators. The update time of our truly perfect L_p-samplers is $O(1)$, which is an exponential improvement over the running time of previous perfect L_p-samplers.

CCS CONCEPTS

• **Theory of computation → Streaming, sublinear and near linear time algorithms**.

KEYWORDS

streaming algorithms, sampling, sliding window model

PODS '22, June 12–17, 2022, Philadelphia, PA, USA

ACM ISBN 978-1-4503-9260-0/22/06.
https://doi.org/10.1145/3517804.3524139

ACM Reference Format:
Rajesh Jayaram, David P. Woodruff, and Samson Zhou. 2022. Truly Perfect Samplers for Data Streams and Sliding Windows. In *Proceedings of the 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS '22), June 12–17, 2022, Philadelphia, PA, USA.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3517804.3524139

1 INTRODUCTION

The streaming model of computation has emerged as an increasingly popular paradigm for analyzing massive data sets, such as network traffic monitoring logs, IoT sensor logs, financial market updates, e-commerce transaction logs, and scientific observations, as in computational biology, astronomy, or particle physics. In the (one-pass) streaming model, an underlying data set is implicitly defined through sequential updates that arrive one-by-one and can only be observed once, and the proposed algorithms are required to use space that is sublinear in the size of the input.

Sampling has proven to be a fundamental and flexible technique for the analysis of massive data. A significant line of active work has studied sampling techniques [19–23, 41, 44, 72] in big data applications such as network traffic analysis [39, 43, 49, 62, 70], database analysis [24, 42, 44–46, 59, 60], distributed computing [31, 32, 51, 71, 74], and data summarization [1, 35–37, 40, 61]. Given a non-negative weight function $G : \mathbb{R} \to \mathbb{R}_{\geq 0}$ and a vector $f \in \mathbb{R}^n$, the goal of a G-sampler is to return an index $i \in \{1, 2, \ldots, n\}$ with probability proportional to $G(f_i)$. In the data stream setting, the vector f, called the *frequency vector*, is given by a sequence of m updates (referred to as insertions or deletions) to its coordinates. More formally, in a data stream the vector f is initialized to 0^n, and then receives a stream of updates of the form $(i, \Delta) \in [n] \times \{-M, \ldots, M\}$ for some integer $M > 0$. The update (i, Δ) causes the change $f_i \leftarrow f_i + \Delta$. This is known as the *turnstile* model of streaming, as opposed to the *insertion-only* model where $\Delta \geq 0$ for all updates. A 1-pass G-sampler must return an index given only one pass through the updates of the stream.

The most well-studied weight functions are when $G(x) = |x|^p$ for some $p > 0$. Such samplers in their generality, known as L_p *samplers*, were introduced by [66]. L_p samplers have been used to develop efficient streaming algorithms for heavy hitters, L_p estimation, cascaded norm approximation, and finding duplicates [3, 15, 25, 53, 56, 66]. Formally, a G-sampler is defined as follows.

Definition 1.1 (G sampler). *Let $f \in \mathbb{R}^n$, $0 \leq \epsilon, \gamma < 1$, $0 < \delta < 1$ and $G : \mathbb{R} \to \mathbb{R}_{\geq 0}$ be a non-negative function satisfying $G(0) = 0$. An $(\epsilon, \gamma, \delta)$-approximate G-sampler is an algorithm that outputs an index $i \in [n]$ such that if $f \neq \vec{0}$, for each $j \in [n]$,*

$$\mathbf{Pr}\left[i = j\right] = (1 \pm \epsilon)\frac{G(f_j)}{\sum_{k\in[n]} G(f_k)} \pm \gamma \tag{1}$$

*and if $f = \vec{0}$, then the algorithm outputs a symbol $\perp$ with probability at least $1 - \delta$. The sampler is allowed to output **FAIL** with some probability δ, in which case it returns nothing.*

If $\epsilon = 0$ and $\gamma = n^{-c}$, where $c > 1$ is an arbitrarily large constant, then the sampler is called perfect. *If $\epsilon = 0$ and $\gamma = 0$, then the sampler is called truly perfect.*

In general, if $\epsilon > 0$ and $\gamma = n^{-c}$, where $c > 1$ is an arbitrarily large constant, an $(\epsilon, \gamma, \delta)$-sampler is commonly referred to as an ϵ-relative error *approximate sampler*. Notice that the guarantees on the distribution of a sampler are all conditioned on the sampler not outputting **FAIL**. In other words, conditioned on not outputting **FAIL**, the sampler must output a value in $[n] \cup \{\perp\}$ from a distribution which satisfies the stated requirements. When $f = \vec{0}$, the distribution in equation 1 is undefined; therefore the special symbol $\perp$ is needed to indicate this possibility.

In the case of L_p samplers with $p > 0$, the underlying distribution is given by $|f_j|^p / \|f\|_p^p$. Such samplers are particularly useful as subroutines for other streaming and data-analytic tasks. In the insertion-only model, the classical reservoir sampling technique of [72] gives an $O(\log n)$ space truly perfect L_1 sampling algorithm. However, when $p \neq 1$, or when negative updates are also allowed (i.e., the turnstile model), the problem becomes substantially more challenging. In fact, the question of whether such L_p samplers even exist using sublinear space was posed by Cormode, Murthukrishnan, and Rozenbaum [30].

Monemizadeh and Woodruff partially answered this question by showing the existence of an $(\epsilon, n^{-c}, 1/2)$-approximate L_p sampler for $p \in [1, 2]$ using poly $\left(\frac{c}{\epsilon}, \log n\right)$ bits of space in the turnstile model [66]. The space bounds were improved by Andoni, Krauthgamer, and Onak [3], and then later by Jowhari, Saglam, and Tardos [56], for $p \in (0, 2)$ to roughly $O\left(c\epsilon^{-\max(1,p)} \log^2 n\right)$ and $O\left(c\epsilon^{-2} \log^3 n\right)$ for $p = 2$. This matched the lower bound of $\Omega(\log^2 n)$ in terms of $\log n$ factors for $p < 2$, as shown in the same paper [56], but was loose in terms of ϵ. This gap was explained by Jayaram and Woodruff [53], who gave the first *perfect* $(0, n^{-c}, 1/2)$-F_p samplers in the streaming model, using $O\left(c \log^2 n\right)$ bits of space for $p \in (0, 2)$ and $O\left(c \log^3 n\right)$ bits of space for $p = 2$. For a further discussion on the development of L_p samplers in the streaming model, we direct the reader to the survey [28]. In addition to L_p sampling, [24] also considered samplers for certain classes of concave functions G in the insertion-only model of streaming.

Truly Perfect Sampling. Unfortunately, none of the aforementioned perfect samplers are *truly perfect*. Specifically, they have an additive error of $\gamma = n^{-c}$, and space depending linearly on c. While this may be acceptable for some purposes where only a small number of samples are required, this error can have significant downstream consequences when many samplers are run independently. For instance, a common usage of sampling for network monitoring and event detection is to generate samples on successive portions of the stream, which are reset periodically (e.g., minute by minute). Additionally, in a large, distributed database, many independent samplers can be run locally on disjoint portions of the dataset. These samples can be used as compact summaries of the database, providing informative statistics on the distribution of data across machines. While the samples generated on a single portion of the data may be accurate enough for that portion, the $1/\text{poly}(n)$ variation distance between the samples and true distribution accumulates over many portions. For large databases containing s distributed machines with $s \gg \text{poly}(n)$, or for samplers run on short portions of high throughput streams, the resulting gap in variation distance between the joint distributions of the samples and the true distribution can blow up to a *constant*. This results in large accuracy issues for sensitive tests run on the data.

Moreover, this creates large issues for *privacy*, even when the identities of samples are anonymized. For instance, a non-truly perfect sampler may positively bias a certain subset $S \subset [n]$ of coordinates when a given entry is in the dataset (i.e, $x_i \neq 0$), and may negatively bias S if that entry is not present (i.e., $x_i = 0$). Given sufficiently many samples, an onlooker would be able to easily distinguish between these cases.

Another important application of truly perfect sampling is in situations where samples from previous portions of the stream influence future portions of the stream, causing cascading blow-ups in error. For instance, samples and sketches can be used to approximate gradient updates for gradient descent [50, 55, 67, 76], where a large number of future gradients naturally depend on the samples generated from prior ones. Unbiasedness is also important for interior point methods, since bias in estimates of the gradients can result in large drift, and therefore error, in the algorithm (see, e.g., Theorem 2 of [48]). Beyond non-adversarial adaptivity, we may also have a malicious attacker who uses adaptivity in an uncontrolled manner. For example, a malicious adversary can adaptively query a database for samples, with future queries depending on past samples. Such streams with adaptive updates are the focus of the field of *adversarial robust streaming* [4, 5, 7, 8, 11, 47, 65, 75]. Due to this adaptivity, the variation distance between the joint distributions can increase *exponentially*, causing large accuracy issues after only a small number of adaptive portions of the stream. Thus, even a perfect sampler would not avoid significant information leakage in such settings, and instead only a truly perfect sampler would be robust to drifts in the output distribution. Finally, truly perfect samplers are of fundamental importance in information-theoretic security.

The Sliding Window Model. While applicable for many situations in data analysis, the standard streaming model does not capture situations in which the data is considered time-sensitive. In applications such as network monitoring [26, 27, 29], event detection in social media [68], and data summarization [17, 38], recent data is considered more accurate and important than data that arrived prior to a certain time window. To address such settings, [33] introduced the *sliding window model*, where only the W most recent updates to the stream induce the underlying input data, for some window size parameter $W > 0$. The most recent W updates form the *active data*, whereas updates previous to the W most recent updates are *expired*. The goal is to aggregate information about the active data using space sublinear in W. We remark that, generally speaking, the sliding window model is insertion-only by definition. Hence, the sliding window model is a strict generalization of the standard insertion-only streaming model.

The sliding window model is more appropriate than the unbounded streaming model in a number of applications [6, 15, 63, 69, 73] and has been subsequently studied in a number of additional settings [9, 10, 12–15, 34, 57, 58]. To date, however, no truly perfect, perfect, or even approximate L_p samplers for the sliding window model are known, leaving a substantive gap in our understanding of sampling for these models.

1.1 Our Contributions

In this work, we initiate the study of *truly perfect samplers*, for general weight functions G in the data stream and sliding window models. We begin by studying the problem of truly perfect sampling in the *turnstile* model of streaming, where both positive and negative updates can be made to the coordinates in f. We demonstrate that in the turnstile model, the additive $1/\text{poly}(n)$ error in previous approximate and perfect L_p samplers is inherent [3, 53, 56, 66].

Theorem 1.2. *Fix constant $\epsilon_0 < 1$, integer $r \geq 1$, and let $2^{-n/2} \leq \gamma < \frac{1}{4}$. Let $G : \mathbb{R} \to \mathbb{R}_{\geq 0}$ be any function satisfying $G(x) > 0$ for $x \neq 0$, and $G(0) = 0$. Then any $(\epsilon_0, \gamma, \frac{1}{2})$-approximate G-sampler $\mathcal{A}$ in the r-pass turnstile streaming model must use $\Omega\left(\min\left\{n, \log\frac{1}{\gamma}\right\}\right)$ bits of space.*

Proof. Given such a sampler $\mathcal{A}$, we give an algorithm for the two-party, r-round equality problem as follows. Alice is given $x \in \{0,1\}^n$, and creates a stream with frequency vector given by $f = x$. Bob then adds the vector $-y$ into the stream so that the final state of the frequency vector induced by the stream is $f = x - y$. Alice and Bob each run $\mathcal{A}$ on their stream and repeatedly pass the state of the algorithm between each other over the course of r rounds. Bob then finally obtains the output of the streaming algorithm $\mathcal{A}(f)$ after r rounds. If the output is **FAIL**, or anything except for the symbol $\perp$, then Bob declares $\text{EQ}_n(x, y) = 0$. If the output is $\perp$, Bob declares $\text{EQ}_n(x, y) = 1$. Notice by definition of a $(\epsilon_0, \gamma, \frac{1}{2})$-$G$ sampler (Definition 1.1), if we actually had $x = y$, then $f = \vec{0}$, so if $\mathcal{A}$ does not output **FAIL**, then it must declare $\perp$ with probability at least $1 - \gamma$. Moreover, if $x \neq y$, then since $G((x-y)_i) > 0$ for some i, a correct sampler can output $\perp$ with probability at most γ. Furthermore, it can output **FAIL** in both cases with probability at most $\frac{1}{2}$.

The above protocol therefore satisfies that if $\text{EQ}_n(x, y) = 0$, Bob outputs 1 with probability at most γ, thus the refutation error is at most $\epsilon < \gamma$. Moreover, if $\text{EQ}_n(x, y) = 1$, then $\mathcal{A}$ outputs fail with probability $\frac{1}{2}$, and conditioned on not outputting fail it must output $\perp$ with probability at least $1 - \gamma$. Thus, the verification error is at most $\delta < 1/2 + \gamma < 3/4$. Then we have $n - \log(1-\delta) > n/2$, and $\log(\frac{(1-\delta)^2}{\epsilon}) > \log(\frac{1}{16\gamma})$. Thus the effective input size is given by

$$\hat{n} > \min\left\{\frac{n}{2}, \log\frac{1}{16\gamma}\right\}$$

Thus, by Theorem 2.1, we have

$$\begin{aligned} R^{(r),\text{ref}}_{\epsilon,\delta}(\text{EQ}_n) &\geq \frac{1}{8}(1-\delta)^2(\hat{n} + \log(1-\delta) - 5) \\ &\geq \frac{1}{8 \cdot 16}(\hat{n} - 7) \\ &= \Omega(\hat{n}) \end{aligned} \tag{2}$$

which completes the proof of the lower bound. □

Theorem 1.2 explains why all prior approximate and perfect samplers developed for the turnstile sliding window model paid a $1/\text{poly}(n)$ additive error in their variation distance. In particular, when $\gamma = n^{-c}$, our lower bound of $\Omega(c \log n)$ for a $(\epsilon, \gamma, \frac{1}{2})$-$F_p$ sampler for $p \in (0, 2]$ is nearly tight, given the upper bound of $O\left(c \log^2 n\right)$ of [53] for $p \in (0, 2)$ and $O\left(c \log^3 n\right)$ for $p = 2$, which achieve perfect sampling ($\epsilon = 0$). This demonstrates that $\log\frac{1}{\gamma}\,\text{polylog}\, n$ is the correct complexity of $(0, \gamma, \frac{1}{2})$-L_p sampling. Our lower bound is based on the fine-grained hardness of the EQUALITY problem from two-party communication complexity, demonstrating that a $(\epsilon, \gamma, 1/2)$-G sampler yields a communication protocol which solves EQUALITY with γ advantage.

Given the strong impossibility results for designing truly perfect samplers in the turnstile model, we shift our attention to the fundamental insertion-only model. Given a measure function $G : \mathbb{R} \to \mathbb{R}^{\geq 0}$ such that $G(x) = G(-x)$, $G(0) = 0$ and G is non-decreasing in $|x|$, we define $F_G = \sum_{i=1}^n G(f_i)$. We design a general framework for designing truly perfect G-samplers for a large number of useful functions G in insertion-only streams and sliding windows with insertion-only updates. The framework is developed in Section 3, wherein several instantiations of the framework are given for specific functions G. Our theorem in its most general form is as follows, although we remark that for several applications, such as for F_p estimation, significant additional work is needed to apply the theorem.

Framework 1.3. *Let G be a function such that $0 \leq G(x) - G(x-1) \leq \zeta$ for all $x \geq 1$. For insertion-only streams, there exists a perfect G sampler that succeeds with probability at least $1 - \delta$ and uses $O\left(\frac{\zeta m}{F_G} \log n \log\frac{1}{\delta}\right)$ bits of space. For the sliding window model with insertion-only updates, there exists a truly perfect G sampler that succeeds with probability at least $1 - \delta$ and uses $O\left(\frac{\zeta W}{F_G} \log^2 n \log\frac{1}{\delta}\right)$ bits of space. Further, the time to process each update is $O(1)$ in expectation. (See Theorem 3.1.)*

The main barrier to applying Framework 1.3 to any arbitrary measure function G is obtaining a "good" lower bound $\widehat{F_G}$ to $F_G = \sum_{i \in [n]} G(f_i)$. Moreover, this lower bound must be obtained correctly with probability 1, as any possibility of failure of a randomized algorithm would necessarily contribute to additive error to the distribution of the samples, resulting in only a perfect, but not truly perfect, sampler.

Interestingly, our samplers utilize a timestamp-based reservoir sampling scheme, as opposed to the common *precision sampling* framework used for other L_p samplers [3, 18, 51–53, 56]. This property of our samplers makes our framework particularly versatile, and for instance allows it to be applied to the sliding window model of streaming.

As specific applications of our framework, we obtain the first truly perfect samplers for many fundamental sampling problems, including L_p sampling, concave functions, and a large number of measure functions, including the $L_1 - L_2$, Fair, Huber, and Tukey estimators. For $p \geq 1$, our results for L_p sampling are as follows; we defer $p \in (0, 1)$ to Theorem 3.3.

Theorem 1.4. *For the insertion-only streaming model and $p \geq 1$, there exists a truly perfect L_p sampler that uses $O(1)$ update time and $O\left(n^{1-1/p} \operatorname{polylog}(n)\right)$ bits of space.*

Together, Theorem 1.2 and Theorem 1.4 show a strong separation between turnstile and insertion-only truly perfect L_p samplers; surprisingly, for every $p > 1$, a truly perfect L_p sampler exists with $O\left(n^{1-1/p} \operatorname{polylog}(n)\right)$ space in the insertion-only model, while in the turnstile model this requires $\Omega(n)$ space.

Another interesting feature of Theorem 1.4 is that the time to process each update is $O(1)$! In contrast, the perfect samplers of [53], which each are not truly perfect, have update time $n^{O(c)}$ to achieve variation distance $\frac{1}{n^c}$. Thus, we obtain an exponential improvement, and optimal running time.

Yet another interesting feature of our algorithm in Theorem 1.4 is that it is sampling-based rather than sketching-based, and thus if the indices have metadata associated with them then we can additionally return that metadata, whereas the sketching-based algorithm of [53] cannot. For example, the indices may be keys we sample by and each key is part of some document; in this case we sample by the keys but additionally can return the document sampled.

Notice also that the complexity in Theorem 1.4 above degrades as $p \to 1$. For instance, for $p = 1$, our bound degrades to $n^{1-1/p} \log n = \log n$, which matches the complexity of reservoir sampling in the insertion only model, or the sliding window truly perfect L_1 sampler by [15].

M-Estimators. In addition to Theorem 1.4, Framework 1.3 also implies truly perfect samplers for a number of M-estimators:

- Truly perfect G samplers for insertion-only streams that use $O\left(\log n \log \frac{1}{\delta}\right)$ bits of space when G is the $L_1 - L_2$ estimator, the Fair estimator, the Huber estimator, or the Tukey estimator. (See Corollary 3.4 and Theorem 5.4.)
- Truly perfect G samplers for sliding windows with insertion-only updates that use $O\left(\log^2 n \log \frac{1}{\delta}\right)$ bits of space when G is the $L_1 - L_2$ estimator, the Fair estimator, or the Huber estimator, or the Tukey estimator. (See Corollary 4.2 and Theorem 5.5.)

Matrix Norms. Framework 1.3 can also be extended to truly perfect sampling for matrix norms. That is, given a matrix $\mathbf{M} \in \mathbb{R}^{n \times d}$, the goal is to sample a row $\mathbf{m}_i$ of $\mathbf{M}$ with probability proportional to $G(\mathbf{m}_i)$ for some given function G. For example, when $G(\mathbf{x}) = \sqrt{\sum_{i \in [d]} x_i^2}$ is the L_2 norm of each row, then such a row sampling primitive would be equivalent to $L_{1,2}$ sampling, which has recently been used in adaptive sampling techniques (see [61] and references therein). See Theorem 3.5 for more details.

Turnstile Streams. Next, we show that Framework 1.3 can also be extended to strict turnstile streams, which combined with Theorem 1.2 shows the separation of general and strict turnstile streams. We give a general reduction as follows:

Theorem 1.5. *Suppose there exists a truly perfect L_p sampler in the one-pass insertion-only streaming model that uses S bits of space. Then there exists a truly perfect L_p sampler that uses $\tilde{O}(Sn^\gamma)$ space and $O\left(\frac{1}{\gamma}\right)$ passes over a strict turnstile stream, which induces intermediate frequency vectors with nonnegative coordinates.*

Truly Perfect F_0 Sampling. We give a truly perfect sampler for the important F_0 problem for both the insertion-only streaming model and the sliding window model (see Theorem 5.2 and Corollary 5.3). Our algorithm works by tracking the first $\sqrt{n}$ unique items in the stream to decide whether $F_0 > \sqrt{n}$. If $F_0 \leq \sqrt{n}$, then it suffices to output a random item among those collected. Otherwise, we simultaneously generate a set S of $\sqrt{n}$ random items, so that with constant probability, some item of S appears in the stream. We can then output any item of S that has appeared in the stream uniformly at random. Surprisingly, our algorithms use $O\left(\sqrt{n} \log n\right)$ space whereas there exists a truly perfect F_0 sampler in the random oracle model with space $O(\log n)$. We believe the complexity of truly perfect F_0 sampling without the assumption of a random oracle to be a fascinating open question.

Truly Perfect Sampling in the Random Order Model. In the full version of the paper [54], we demonstrate that for *random order* streams, we can design a truly perfect L_2 sampling using only $O\left(\log^2 n\right)$ bits of space. Since, as in our framework for truly perfect sampling, our algorithms are timestamp based, they also apply to the more challenging sliding window model of streaming. The complexity of our sampler is a $\log n$ factor smaller than the complexity of the best known previous samplers [53] in the adversarial order model, which had $\gamma = 1/\operatorname{poly}(n)$ additive error in their distribution, and which did not apply to the sliding window model. Our theorem for L_2 sampling on random order streams is as follows.

Theorem 1.6. *There exists a one-pass sliding window algorithm for random order insertion-only streams that outputs index $i \in [n]$ with probability $\frac{f_i^2}{F_2}$ and outputs **FAIL** with probability at most $\frac{1}{3}$, i.e., the algorithm is a truly perfect L_2 sampler, using $O\left(\log^2 n\right)$ bits of space and $O(1)$ update time.*

We generalize this approach to perfect L_p samplers on random order streams for integers $p > 2$.

Theorem 1.7. *Let $p > 2$ be a fixed integer. There exists a one-pass algorithm that outputs an index $i \in [n]$ with probability $\frac{f_i^p}{\sum_{j=1}^n f_j^p}$, and outputs **FAIL** with probability at most $\frac{1}{3}$ on a random-order insertion-only stream of length m , i.e., the algorithm is a truly perfect L_p sampler, using $O\left(m^{1-\frac{1}{p-1}} \log n\right)$ bits of space and $O(1)$ update time.*

For $p = 2$, intuitively our algorithm follows by tracking collisions between adjacent elements in the stream. Here, a collision occurs when to subsequent updates are made to the same coordinate $i \in [n]$. The probability that this occurs at a given timestep is $\frac{f_i(f_i-1)}{m(m-1)}$. Since this is not quite the right probability, we "correct" this distribution by a two part rejection sampling step to obtain a truly perfect sampler. For truly perfect L_p sampling on random order streams for integers $p > 2$, we store consecutive blocks of $m^{1-\frac{1}{p-1}}$ elements in the stream, along with their corresponding timestamps, and instead look for p-wise collisions within the block.

Fast Perfect L_p Sampling for $p < 1$. Lastly, we demonstrate that for insertion only streams, the runtime of the perfect L_p samplers of [53] can be significantly improved. Specifically, these prior samplers had update time which was a large polynomial in n: roughly, to obtain a $(0, \gamma, 1/2)$ L_p-sampler, these algorithms required $O\left(\frac{1}{\gamma}\right)$ runtime. This poses a serious barrier for usage of these perfect samplers in most applications. We demonstrate, however, that the tradeoff between runtime and distributional error in a perfect L_p sampler is not required for $p < 1$, by giving an algorithm with nearly tight space complexity, which achieves $\text{poly}(\log n)$ update time. See the full version of the paper [54] for more details.

1.2 Our Techniques

Truly Perfect L_p Samplers. We begin by describing our sampling framework for the case of L_p samplers. Specifically, to obtain a truly perfect L_p sampler for insertion-only streams of length m and for $p \geq 1$, we run $O\left(n^{1-1/p}\right)$ parallel instances of a single sampler, which uses only $\log n$ bits of space, but succeeds with probability $\Omega\left(\frac{1}{n^{1-1/p}}\right)$. Each instance of the single sampler first applies reservoir sampling to the updates in the stream to sample an item $s \in [n]$, along with a specific timestamp t_s when s was added to the reservoir sample, and keeps a counter c of how many times s appears in the stream after it is first sampled.

Inspired by a technique of Alon, Matias, and Szegedy for F_p-estimation [2], we then prove that if c occurrences of s appear afterwards and we output s with probability proportional to $c^p - (c-1)^p$, then by a telescoping argument, the probability of outputting each $i \in [n]$ is proportional to $|f_i|^p$. Thus a sampler that successfully outputs a coordinate must do so from the desired distribution. To implement the rejection sampling step, we need to obtain a good normalizing factor so that the resulting step forms a valid distribution. We demonstrate that it suffices to estimate $\|f\|_\infty$ to obtain a good normalizing factor, which results in acceptance probability of at least $\Omega\left(\frac{1}{n^{1-1/p}}\right)$. We carry out this step deterministically with the Misra-Gries sketch, which is necessary since failure of any randomized estimation algorithm would introduce additive error into the sampler. By repeating $O\left(n^{1-1/p}\right)$ times, we ensure that at least one sampler succeeds with constant probability, showing that we output **FAIL** with at most constant probability. We use a similar approach for $p \in (0, 1)$.

General Framework for Truly Perfect Sampler. The aforementioned telescoping argument for our truly perfect sampler for insertion-only updates can be viewed as using the identity $\sum_{c=1}^{f_i}(G(c) - G(c-1)) = G(f_i) - G(0)$ for $G(x) = x^p$. By the same reasoning, we can generalize the paradigm of "correcting" the sampling probability for any monotonic function G with $G(0) = 0$ by a subsequent rejection sampling step though the size of the acceptance probability depends on obtaining a good normalizing factor. For $F_G = \sum_{i=1}^{n} G(f_i)$, we show that we can similarly bound the probability each single sampler succeeds with probability roughly $\frac{F_G}{m}$. Thus we require roughly $\frac{m}{F_G}$ parallel instances of the single sampler.

For the case of the sliding window model, we can automatically expire the sample (s, t_s) as soon as t_s leaves the active window, causing an additional complexity in the sliding window model. However, by maintaining a number of "checkpoints", we can ensure that (s, t_s) is an active element with constant probability.

Random-Order Sampling. Finally, we improve our perfect L_p samplers for random-order insertion-only streams by using distributional properties of each stream update. Namely, we modify our timestamp based sampling scheme to randomly sample a number of p-tuples and search for collisions among the p-tuples. For $p = 2$, the idea is to consider two adjacent elements and see if they collide. Intuitively, an arbitrary position in the window is item i with probability $\frac{f_i}{m}$ due to the random order of the stream, where m is the length of the stream. The next update in the stream is also item i with probability $\frac{f_i - 1}{m-1}$. Thus the probability that both the two positions are updates to coordinate i is $\frac{f_i(f_i-1)}{m(m-1)}$, which is not quite the right probability. Instead of using a telescoping argument as in the general framework, we instead "correct" this probability by sampling item i in a position with probability $\frac{1}{m}$. Otherwise, with probability $1 - \frac{1}{m}$, we sample item i if the item is in the next position as well. Now the probability of sampling i on the two adjacent elements is $\frac{1}{m}\frac{f_i}{m} + \frac{m-1}{m}\frac{f_i}{m}\frac{f_i-1}{m-1} = \frac{f_i^2}{m^2}$. We similarly generalize this argument to integer $p > 2$.

1.3 Preliminaries

We use $\mathbb{R}^{\geq 0}$ to denote a non-negative number. We use the notation $[n]$ to represent the set $\{1, \ldots, n\}$ for any positive integer n. We use $\text{poly}(n)$ to denote a fixed constant degree polynomial in n but we write $\frac{1}{\text{poly}(n)}$ to denote an arbitrary degree polynomial in n that can be determined from setting constants appropriately. When an event has probability $1 - \frac{1}{\text{poly}(n)}$ of occurring, we say the event occurs with high probability. We similarly use $\text{polylog}(n)$ to omit terms that are polynomial in $\log n$. We note that all space bounds in this paper are given in *bits*. In the insertion-only model of streaming, there is a vector $f \in \mathbb{R}^n$ which is initialized to the 0 vector. The stream then proceeds with a sequence of $m = \text{poly}(n)$ updates $i_1, i_2, \ldots, i_m$ (the exact value of m is unknown to the algorithm). After the first t time steps, the state of the vector, denoted $f^{(t)}$, is given by $f_j^{(t)} = \sum_{t' \leq t} \mathbf{1}(i_{t'} = j)$. In the sliding window model, at time step t only the most recent W updates form the active portion of the stream, so that in the sliding window model: $f_j^{(t)} \sum_{t' \in (t-W, t]} \mathbf{1}(i_{t'} = j)$.

2 LOWER BOUND FOR TRULY PERFECT SAMPLING IN TURNSTILE STREAMS

In this section, we demonstrate that truly perfect G samplers cannot exist in sublinear space in the turnstile model. Specifically, we show that any perfect sampler with additive error $\gamma = n^{-c}$ requires space at least $\Omega(c \log n)$. This demonstrates that no sublinear space truly perfect sampler can exist in the turnstile model, and demonstrates the tightness (up to $\log n$ factors), of the previously known perfect and approximate L_p samplers [3, 53, 56, 66].

Our lower bound is based on the fine-grained hardness of the EQUALITY problem from two-party communication complexity [16]. Specifically, consider the boolean function $\text{EQ}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ given by $\text{EQ}_n(x, y) = 1 \iff x = y$. In the two party, one way communication model, there are two parties: Alice and

Bob. Alice is given a input string $x \in \{0,1\}^n$ and Bob is given $y \in \{0,1\}^n$. Then Alice must send a single message M to Bob, who must then output whether $\text{EQ}_n(x,y)$ correctly with some probability. A communication protocol $\mathcal{P}$ is a randomized two-party algorithm which takes the input (x,y) and decides on a message M and output procedure out for Bob given (M,y). The communication cost of $\mathcal{P}$ is denoted $\text{cost}(\mathcal{P},x,y)$, and defined as the maximum number of bits sent in the message M over all coin flips of the algorithm, on inputs (x,y). We now define the randomized *refutation complexity* of a communication protocol for computing a Boolean function f. We define the *refutation cost*, *refutation error*, and *verification error* as:

$$\text{rcost}(\mathcal{P}) = \max_{(x,y)\in f^{-1}(0)} \text{cost}(\mathcal{P},x,y)$$
$$\text{rerr}(\mathcal{P}) = \max_{(x,y)\in f^{-1}(0)} \mathbf{Pr}\left[\text{out}(\mathcal{P}(x,y)) = 1\right]$$
$$\text{verr}(\mathcal{P}) = \max_{(x,y)\in f^{-1}(1)} \mathbf{Pr}\left[\text{out}(\mathcal{P}(x,y)) = 0\right].$$

We define the randomized refutation complexity of a function f for an integer $r \geq 1$ as

$$R^{(r),\text{ref}}_{\epsilon,\delta}(f) = \min_{\mathcal{P}} \{\text{rcost}(\mathcal{P}) : \text{rerr}(\mathcal{P}) \leq \epsilon, \text{verr}(\mathcal{P}) \leq \delta\}$$

where the minimum is restricted to r-round communication protocols $\mathcal{P}$. Finally, we define the effective instance size as

$$\hat{n} = \min\left\{n + \log(1-\delta),\ \log\left(\frac{(1-\delta)^2}{\epsilon}\right)\right\}.$$

Theorem 2.1 (Theorem 44 [16]). *We have* $R^{(r),\text{ref}}_{\epsilon,\delta}(\text{EQ}_n) \geq \frac{1}{8}(1-\delta)^2(\hat{n} + \log(1-\delta) - 5)$.

We show Theorem 1.2 by giving a reduction from EQUALITY and applying Theorem 2.1.

3 FRAMEWORK FOR TRULY PERFECT SAMPLING

In this section, we first give a framework for truly perfect sampling for some measure function $G : \mathbb{R} \to \mathbb{R}^{\geq 0}$ such that $G(x) = G(-x)$, $G(0) = 0$ and G is non-decreasing in $|x|$. If we define $F_G = \sum_{i=1}^n G(f_i)$, then we say that a truly perfect G sampler outputs index $i \in [n]$ with probability $\frac{G(f_i)}{F_G}$. We then show how to apply the framework to L_p sampling where $G(x) = |x|^p$ and to various M-estimators, such as the $L_1 - L_2$, Fair, Huber, and Tukey estimators.

3.1 Algorithmic Framework

Our algorithm is based on running parallel instances of a single sampler. Each instance uses $\log n$ bits of space, but only succeeds with small probability and thus we need to run multiple instances to ensure that with sufficiently high probability, some instance succeeds. Each instance uses reservoir sampling to sample an item s and keeps a counter c of how many times s appears in the stream after it is sampled.

We first describe the Sampler algorithm. Given a stream of elements $u_1, \ldots, u_m$, where each $u_i \in [n]$, Sampler selects an index $j \in [m]$ uniformly at random and outputs u_j as well as the number of instances of u_j that appear after time j. The algorithm uses reservoir sampling to ensure that each item is selected with probability $\frac{1}{m}$. A counter is also maintained to track the number of instances of the sample. Each time a new sample replaces the existing sample in the reservoir sampling procedure, the counter is reset to zero.

Algorithm 1 Sampler: Reservoir sampling, counting number of times item has appeared afterwards.

Input: A stream of updates $u_1, u_2, \ldots, u_m$, where each $u_i \in [n]$ represents a single update to a coordinate of the underlying vector f.
Output: Sample each coordinate u_i with probability $\frac{1}{m}$ and output the number of occurrences that appears afterwards.

$s \leftarrow \emptyset, c \leftarrow 0$
for each update u_r **do**
  $s \leftarrow u_r$ with probability $\frac{1}{r}$
  if s is updated to u_r **then**
    $c \leftarrow 0$ ▷Reset counter.
  if $u_r = s$ **then**
    $c \leftarrow c + 1$ ▷Increment counter.
return s and c.

By outputting s with probability $\frac{G(c)-G(c-1)}{\zeta}$, where ζ is a parameter such that $G(x) - G(x-1) \leq \zeta$ for all possible coordinates x in the frequency vector, i.e., $x \in \{f_1, \ldots, f_n\}$, then it can be shown by a telescoping argument that the probability of outputting each $i \in [n]$ is "corrected" to roughly $\frac{G(f_i)}{\zeta \cdot m}$, where m is the length of the stream. Hence if the sampler successfully outputs a coordinate, it follows the desired distribution.

Algorithm 2 Truly perfect G-sampler algorithm for insertion only streams.

Input: A stream of updates $u_1, u_2, \ldots, u_m$, where each $u_i \in [n]$ represents a single update to a coordinate of the underlying vector f, a measure function G.

Initialize an instance of Sampler. ▷Algorithm 1
for each update $u_t \in [n]$ **do**
  Update Sampler with u_t.
Let s be the sampled output of Sampler and let c be the number of times s has appeared afterwards.
Let ζ be a parameter such that $G(x) - G(x-1) \leq \zeta$ for all $x \geq 1$.
return s with probability $\frac{G(c+1)-G(c)}{\zeta}$.

Theorem 3.1. *Let G be a function such that $0 \leq G(x) - G(x-1) \leq \zeta$ for all $x \geq 1$. Given a lower bound $\widehat{F_G}$ on $F_G = \sum_{i\in[n]} G(f_i)$, then there exists a truly perfect G sampler for an insertion-only stream that outputs **FAIL** with probability at most δ and uses $O\left(\frac{\zeta m}{\widehat{F_G}} \log n \log \frac{1}{\delta}\right)$ bits of space. Further, the time to process each update is $O(1)$ in expectation.*

3.2 Applications in Data Streams

In this section, we show various applications of Algorithm 2 in the streaming model. The main barrier to applying Theorem 3.1 to any arbitrary measure function G is obtaining a "good" lower bound $\widehat{F_G}$ to $F_G = \sum_{i\in[n]} G(f_i)$.

3.2.1 Truly Perfect L_p Sampling on Insertion-Only Streams. We first consider truly perfect L_p sampling, where $G(x) = |x|^p$, for $p \geq 1$. Note that reservoir sampling is already a perfect L_1 sampler for $p = 1$ and it uses $O(\log n)$ bits of space on a stream of length $m = \text{poly}(n)$. For $p \in (1, 2]$, we first require the following norm estimation algorithm on insertion-only streams.

We now introduce an algorithm that for *truly perfect* L_p sampling using the above framework. We first describe the case of $p = 2$. Before describing our perfect L_2 sampler, we first recall the MisraGries data structure for finding heavy-hitters.

Theorem 3.2. *[64] There exists a deterministic one-pass streaming algorithm* MisraGries *that uses $O\left(\frac{1}{\epsilon}\log m\right)$ space on a stream of length m and outputs a list L of size $\frac{1}{2\epsilon}$ that includes all items i such that $f_i > 2\epsilon m$. Moreover, the algorithm returns an estimate $\widehat{f_i}$ for each $i \in L$ such that $f_i - \epsilon m \leq \widehat{f_i} \leq f_i$.*

Although we could obtain a perfect L_p sampler using any L_p estimation algorithm that succeeds with high probability, we can further remove the additive $\frac{1}{\text{poly}(n)}$ error of returning each coordinate $i \in [n]$ using the MisraGries data structure to obtain a *truly* perfect L_p sampler.

Theorem 3.3. *For a frequency vector f implicitly defined by an insertion-only stream, there exists an algorithm that returns each $i \in [n]$ with probability $\frac{f_i^p}{\sum_{j\in[n]} f_j^p}$, using space $O\left(m^{1-p}\log n\right)$ for $p \in (0, 1]$ and $O\left(n^{1-1/p}\log n\right)$ bits of space for $p \in [1, 2]$.*

3.2.2 M-estimators on Insertion-Only Streams. We generalize the paradigm of Algorithm 2 to sampling from general statistical M-estimator distributions. Recall that for a given a measure function $G : \mathbb{R} \to \mathbb{R}^{\geq 0}$ such that $G(x) = G(-x)$, $G(0) = 0$ and G is non-decreasing in $|x|$, we define $F_G = \sum_{i=1}^n G(f_i)$. Then a truly perfect M-estimator sampler outputs index $i \in [n]$ with probability $\frac{G(f_i)}{F_G}$.

Corollary 3.4. *There exist truly perfect G samplers for the insertion-only streaming model that succeed with probability at least $1 - \delta$ and use $O\left(\log n \log\frac{1}{\delta}\right)$ bits of space when G is the $L_1 - L_2$ estimator, the Fair estimator, or the Huber estimator.*

3.2.3 Matrix Norms on Insertion-Only Streams. We now consider the case of sampling row $\mathbf{m}_i$ from a matrix $\mathbf{M} \in \mathbb{R}^{n\times d}$ with rows $\mathbf{m}_1, \ldots, \mathbf{m}_n \in \mathbb{R}^d$ with probability $\frac{G(\mathbf{m}_i)}{F_G}$ for some function G, where we define $F_G = \sum_{j\in[n]} G(\mathbf{m}_j)$. We consider a insertion-only stream in the sense that each update in the stream is a non-negative update to some coordinate of $\mathbf{M}$.

Theorem 3.5. *Fix any non-negative function $G : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ satisfying $G(\vec{0}) = 0$. Let ζ be a parameter such that $G(\mathbf{x}) - G(\mathbf{x} - e_i) \leq \zeta$ for all $\mathbf{x} \in \left(\mathbb{R}^{\geq 0}\right)^d$, $i \in [d]$. Given a lower bound $\widehat{F_G}$ on F_G, then there exists a truly perfect G sampler for an insertion-only stream that succeeds with probability at least $1 - \delta$ and uses $O\left(\frac{\zeta dm}{\widehat{F_G}}\log n \log\frac{1}{\delta}\right)$ bits of space.*

For example, when $G(\mathbf{x}) = \sum_{i\in[d]} |x_i|$, then F_G is the $L_{1,1}$ norm. Then $F_G = m$, so that by Theorem 3.5, so we can sample a row $\mathbf{m}_i$ with probability proportional to its L_1 norm, using $O\left(d \log n \log\frac{1}{\delta}\right)$ bits of space. We can also apply Theorem 3.5 when $G(\mathbf{x}) = \sqrt{\sum_{i\in[d]} x_i^2}$ is the L_2 norm of each row, so that F_G is the $L_{1,2}$ norm crucially used in many adaptive sampling techniques (see [61] and references therein). Finally, we show in the full version of our paper [54] that our framework can be extended to strict turnstile streams, i.e., Theorem 1.5.

4 APPLICATIONS IN SLIDING WINDOWS

In this section, we give additional applications of our framework to truly perfect samplers on sliding windows. Recall that in the sliding window model, updates $u_1, \ldots, u_m$ to an underlying vector $f \in \mathbb{R}^n$ arrive sequentially as a data stream and the underlying vector f is determined by the most recent W updates $u_{m-W+1}, \ldots, u_m$, where $W > 0$ is the predetermined window size parameter. We assume that m and W are polynomially bounded in n, i.e., $O(\log m) = O(\log W) = O(\log n)$. For each update u_k, if $k < m - W + 1$, we say u_k is *expired*. Otherwise if $k \geq m - W + 1$, we say u_k is *active*.

4.1 M-estimators on Sliding Windows

In this section, we consider a general paradigm for sampling from general statistical M-estimator distributions. Recall that for a given a measure function $G : \mathbb{R} \to \mathbb{R}^{\geq 0}$ such that $G(x) = G(-x)$, $G(0) = 0$ and G is non-decreasing in $|x|$, we define $F_G = \sum_{i=1}^n G(f_i)$ so that F_G is also implicitly defined by only the most recent W updates. Then a truly perfect M-estimator sampler outputs index $i \in [n]$ with probability exactly $\frac{G(f_i)}{F_G}$.

The key argument in Theorem 1.4 was that if c instances of the sample s appeared after the initial sample, then s is output with probability proportional to $c^p - (c-1)^p$. By a telescoping argument, each index i is sampled with probability proportional to $\sum_{c=1}^{f_i} c^p - (c-1)^p = f_i^p$. Since G is non-decreasing in $|x|$, we can generalize to sampling each item with probability proportional to $G(c) - G(c-1)$, rather than $c^p - (c-1)^p$. This approach can be simulated in the sliding window model by checking whether s is an active element.

Theorem 4.1. *Let G be a function such that $G(x) - G(x-1) \leq \zeta$ for all $x \geq 1$. Then there exists a truly perfect G sampler for the sliding window model that succeeds with probability at least $1 - \delta$ and uses $O\left(\frac{\zeta W}{F_G}\log n \log\frac{1}{\delta}\right)$ bits of space.*

Using Theorem 4.1 and the properties of the M-estimators defined in Corollary 3.4, we have:

Corollary 4.2. *There exist truly perfect G samplers for the insertion-only sliding window model that succeed with probability at least $1 - \delta$ and use $O\left(\log n \log\frac{1}{\delta}\right)$ bits of space when G is the $L_1 - L_2$ estimator, the Fair estimator, or the Huber estimator.*

Algorithm 3 Truly perfect M-estimator sampler algorithm for the sliding window model on insertion only streams.

Input: A stream of updates $u_1, u_2, \ldots, u_m$, where each $u_i \in [n]$ represents a single update to a coordinate of the underlying vector f, a measure function G, and a size W for the sliding window.

Initialize instances of Sampler every W updates and keep the two most recent instances. ▷Algorithm 1
for each update $u_t \in [n]$ with $t \in [m]$ **do**
    Update each Sampler.
Let s be the sampled output of a Sampler and let c be the number of times s has appeared afterwards.
Let ζ be a parameter such that $G(x) - G(x-1) \leq \zeta$ for all $x \geq 1$.
if s was sampled within the last W updates **then**
    return s with probability $\frac{G(c+1)-G(c)}{\zeta}$.

We also give truly perfect L_p samplers on sliding windows. Specifically we prove Theorem 1.4, the details of which can be found in the full version of the paper [54].

5 F_0 SAMPLING

In this section, we give truly perfect F_0 samplers in the insertion-only streaming model. In the F_0 sampling problem, the goal is to sample a coordinate $i \in [n]$ from a frequency vector of length n such that $\mathbf{Pr}[i = j] = 0$ if $f_j = 0$ and $\mathbf{Pr}[i = j] = \frac{1}{F_0}$ otherwise, where $F_0 = |\{i \in [n],\ f_i \neq 0\}|$. We cannot immediately apply the framework of Algorithm 2 for F_0 sampling without trivializing the space complexity, due to the fact that F_0 can be substantially smaller than m.

We first remark that in the random oracle model, where an algorithm is given oracle access to a random hash function $h : [n] \to [0, 1]$, the well-known algorithm that outputs the nonzero coordinate $i \in [n]$ of f_i that minimizes $h(i)$ is truly perfect F_0 sampler, since each of the $|F_0|$ has probability $\frac{1}{|F_0|}$ of obtaining the minimal hash value for a random hash function.

Remark 5.1. *In the random oracle model, there exists a truly perfect F_0 sampler on insertion-only streams that uses $O(\log n)$ bits of space and constant update time.*

We now give a truly perfect F_0 sampler on insertion-only streams without the assumption of the random oracle model. We store the first $\sqrt{n}$ distinct items and in parallel sample a set S of $O(\sqrt{n})$ random items from the universe. We maintain the subset U of S that arrive in the stream. Now since we store the first $\sqrt{n}$ distinct items, we know whether $F_0 \leq \sqrt{n}$ or $F_0 > \sqrt{n}$. If $F_0 \leq \sqrt{n}$ then our algorithm has collected all items in the stream and can simply output an item uniformly at random. Otherwise if F_0 is larger than $\sqrt{n}$ at the end of the stream, then since S has size $O(\sqrt{n})$ and was generated at random, we have constant probability that some element of S has arrived in the stream. Our algorithm can then output a random element of U that has appeared in the stream. We give our algorithm in full in Algorithm 4.

Algorithm 4 Truly perfect F_0 sampler for insertion only streams.

Input: A stream of updates $u_1, u_2, \ldots, u_m$, where each $u_i \in [n]$ represents a single update to a coordinate of a underlying vector f.

Let S be a random subset of $[n]$ of size $2\sqrt{n}$.
Let T be the first unique $\sqrt{n}$ coordinate updates to f in the stream.
Let U be the subset of S that appears in the stream.
if $|T| < \sqrt{n}$ **then**
    return a random element of T
else if $|U| > 0$ **then**
    return a random element of U
else
    return FAIL

Theorem 5.2. *Given $\delta \in (0, 1)$, there exists a truly perfect F_0 sampler on insertion-only streams that uses space $O\left(\sqrt{n} \log n \log \frac{1}{\delta}\right)$ and constant update time and succeeds with probability at least $1 - \delta$. Moreover, the algorithm reports f_i along with the sampled index $i \in [n]$.*

Proof. Consider Algorithm 4. If $F_0 < \sqrt{n}$, then all nonzero coordinates of f will be stored in T in the insertion-only streaming model and so a random element of T is a truly perfect F_0 sample. Otherwise if $F_0 \geq \sqrt{n}$ and the algorithm does not return **FAIL**, then it must have output a random item in U, which is a subset of S. Since S is a subset of n chosen uniformly at random, then a random element of S is a truly perfect F_0 sample.

It remains to analyze the probability that Algorithm 4 returns **FAIL** when $F_0 \geq \sqrt{n}$, since it will never fail when $F_0 < \sqrt{n}$. Our algorithm will only fail if $|U| = 0$, so that none of the $2\sqrt{n}$ random items in S appeared in the stream. This can only occur with probability at most $\left(1 - \frac{2\sqrt{n}}{n}\right)^{\sqrt{n}}$, which is certainly at most $\frac{1}{e}$ for sufficiently large n. Hence by repeating $O\left(\log \frac{1}{\delta}\right)$ times, the algorithm has probability at least $1 - \delta$ of success. Then the space required is $O\left(\sqrt{n} \log n \log \frac{1}{\delta}\right)$.

Finally, note that the algorithm can track the frequency of each element in U and T, so the algorithm can also report the frequency f_i corresponding to the sampled index i. □

By modifying T to be the last unique $\sqrt{n}$ coordinate updates to f in the stream and storing timestamps for each element in U, Algorithm 4 extends naturally to the sliding window model.

Corollary 5.3. *Given $\delta \in (0, 1)$, there exists a truly perfect F_0 sampler in the sliding window model that uses space $O\left(\sqrt{n} \log n \log \frac{1}{\delta}\right)$ and constant update time and succeeds with probability at least $1 - \delta$.*

Recall that for the Tukey measure function, we have $G(x) = \frac{\tau^2}{6}\left(1 - \left(1 - \frac{x^2}{\tau^2}\right)^3\right)$ for $|x| \leq \tau$ and $G(x) = \frac{\tau^2}{6}$ otherwise, where $\tau > 0$ is some constant. We can now use our F_0 sampler of choice, say L0Sampler, to obtain a truly perfect sampler for the Tukey measure function by a similar approach to Algorithm 2. Each instance will use a subroutine of L0Sampler rather than Sampler. Now if c is the

number of instances of the index output by L0Sampler within the window, then we accept c with probability $\frac{G(c)}{G(\tau)}$.

Theorem 5.4. *Given $\delta \in (0,1)$, there exists a truly perfect G sampler for the insertion-only streaming model that succeeds with probability at least $1-\delta$ when G is the Tukey estimator. The algorithm uses $O\left(\log n \log \frac{1}{\delta}\right)$ bits of space in the random oracle model and $O\left(\sqrt{n}\log n \log \frac{1}{\delta}\right)$ otherwise.*

Proof. Note that only the indices that appear in the stream can be output. The probability that any index $i \in [n]$ that ap pears in the stream is output is $\frac{1}{F_0}\cdot\frac{G(f_i)}{G(\tau)}$. Then a single instance of the algorithm returns an output with probability $\frac{F_G}{F_0\cdot G(\tau)} \geq \frac{G(1)}{G(\tau)}$. Hence repeating $O\left(\log \frac{1}{\delta}\right)$ times outputs an instance with probability at least $1-\delta$. Since L0Sampler requires $O(\log n)$ space in the random oracle model, then the total space used is $O\left(\log n \log \frac{1}{\delta}\right)$ bits of space. Otherwise by Theorem 5.2, L0Sampler requires $O\left(\sqrt{n}\log n\right)$ space so that the total space used is $O\left(\sqrt{n}\log n \log \frac{1}{\delta}\right)$. □

We also obtain a truly perfect G sampler for the insertion-only sliding window model when G is the Tukey estimator by a similar approach to Algorithm 3. We again use a truly perfect F_0 sampling algorithm L0Sampler of choice and, if c is the number of instances of the index s output by L0Sampler within the window, then we accept s with probability $\frac{G(c)}{G(\tau)}$.

Theorem 5.5. *Given $\delta \in (0,1)$, there exists a truly perfect G sampler for the sliding window model that succeeds with probability at least $1-\delta$ when G is the Tukey estimator. The algorithm uses $O\left(\log n \log \frac{1}{\delta}\right)$ bits of space in the random oracle model and $O\left(\sqrt{n}\log n \log \frac{1}{\delta}\right)$ otherwise.*

Finally, we show in the full version of the paper [54] that the same guarantees of Theorem 5.2 can be extended to strict turnstile streams.

6 CONCLUSION

Our work shows that surprisingly, truly perfect samplers exist in the insertion-only model with a sublinear amount of memory for a large class of loss functions, a large class of objects (vectors, matrices), and several different models (data stream and sliding window). Establishing better upper bounds, or proving lower bounds is a very intriguing open question. In particular, the usual communication complexity lower bound arguments do not seem to apply, and already for our turnstile lower bound we needed to use fine-grained properties of the equality function.

ACKNOWLEDGEMENTS

This work was supported in part by a Simons Investigator Award and by the National Science Foundation under Grant No. CCF-1815840.

REFERENCES

[1] Ankit Aggarwal, Amit Deshpande, and Ravi Kannan. 2009. Adaptive Sampling for k-Means Clustering. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, APPROX and 13th International Workshop, RANDOM. Proceedings*. 15–28.

[2] Noga Alon, Yossi Matias, and Mario Szegedy. 1999. The Space Complexity of Approximating the Frequency Moments. *J. Comput. Syst. Sci.* 58, 1 (1999), 137–147.

[3] Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. 2011. Streaming Algorithms via Precision Sampling. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS*. 363–372.

[4] Idan Attias, Edith Cohen, Moshe Shechner, and Uri Stemmer. 2021. A Framework for Adversarial Streaming via Differential Privacy and Difference Estimators. *CoRR* abs/2107.14527 (2021).

[5] Dmitrii Avdiukhin, Slobodan Mitrovic, Grigory Yaroslavtsev, and Samson Zhou. 2019. Adversarially Robust Submodular Maximization under Knapsack Constraints. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD*. 148–156.

[6] Brian Babcock, Shivnath Babu, Mayur Datar, Rajeev Motwani, and Jennifer Widom. 2002. Models and Issues in Data Stream Systems. In *Proceedings of the Twenty-first ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*. 1–16.

[7] Omri Ben-Eliezer, Rajesh Jayaram, David P Woodruff, and Eylon Yogev. 2020. A Framework for Adversarially Robust Streaming Algorithms. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 63–80.

[8] Omri Ben-Eliezer and Eylon Yogev. 2020. The adversarial robustness of sampling. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 49–62.

[9] Vladimir Braverman, Petros Drineas, Cameron Musco, Christopher Musco, Jalaj Upadhyay, David P. Woodruff, and Samson Zhou. 2020. Near Optimal Linear Algebra in the Online and Sliding Window Models. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS*. 517–528.

[10] Vladimir Braverman, Elena Grigorescu, Harry Lang, David P. Woodruff, and Samson Zhou. 2018. Nearly Optimal Distinct Elements and Heavy Hitters on Sliding Windows. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*. 7:1–7:22.

[11] Vladimir Braverman, Avinatan Hassidim, Yossi Matias, Mariano Schain, Sandeep Silwal, and Samson Zhou. 2021. Adversarial Robustness of Streaming Algorithms through Importance Sampling. *CoRR* abs/2106.14952 (2021).

[12] Vladimir Braverman, Harry Lang, Keith Levin, and Morteza Monemizadeh. 2015. Clustering on Sliding Windows in Polylogarithmic Space. In *35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS*. 350–364.

[13] Vladimir Braverman, Harry Lang, Keith Levin, and Morteza Monemizadeh. 2016. Clustering Problems on Sliding Windows. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. 1374–1390.

[14] Vladimir Braverman and Rafail Ostrovsky. 2007. Smooth Histograms for Sliding Windows. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS), Proceedings*. 283–293.

[15] Vladimir Braverman, Rafail Ostrovsky, and Carlo Zaniolo. 2012. Optimal sampling from sliding windows. *J. Comput. Syst. Sci.* 78, 1 (2012), 260–272.

[16] Joshua Brody, Amit Chakrabarti, Ranganath Kondapally, David P Woodruff, and Grigory Yaroslavtsev. 2014. Certifying equality with limited interaction. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014)*.

[17] Jiecao Chen, Huy L. Nguyen, and Qin Zhang. 2016. Submodular Maximization over Sliding Windows. *CoRR* abs/1611.00129 (2016).

[18] Xi Chen, Rajesh Jayaram, Amit Levi, and Erik Waingarten. 2020. An Improved Analysis of the Quadtree for High Dimensional EMD.

[19] Edith Cohen. 2018. Stream Sampling Framework and Application for Frequency Cap Statistics. *ACM Trans. Algorithms* 14, 4 (2018), 52:1–52:40.

[20] Edith Cohen, Graham Cormode, and Nick G. Duffield. 2011. Structure-Aware Sampling: Flexible and Accurate Summarization. *Proc. VLDB Endow.* 4, 11 (2011), 819–830.

[21] Edith Cohen, Graham Cormode, and Nick G. Duffield. 2012. Don't let the negatives bring you down: sampling from streams of signed updates. In *ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS*. 343–354.

[22] Edith Cohen, Nick G. Duffield, Haim Kaplan, Carsten Lund, and Mikkel Thorup. 2011. Efficient Stream Sampling for Variance-Optimal Estimation of Subset Sums. *SIAM J. Comput.* 40, 5 (2011), 1402–1431.

[23] Edith Cohen, Nick G. Duffield, Haim Kaplan, Carsten Lund, and Mikkel Thorup. 2014. Algorithms and estimators for summarization of unaggregated data streams. *J. Comput. Syst. Sci.* 80, 7 (2014), 1214–1244.

[24] Edith Cohen and Ofir Geri. 2019. Sampling Sketches for Concave Sublinear Functions of Frequencies. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems, NeurIPS*. 1361–1371.

[25] Edith Cohen, Rasmus Pagh, and David P. Woodruff. 2020. WOR and p's: Sketches for ℓ_p-Sampling Without Replacement. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems, NeurIPS*.

[26] Graham Cormode. 2013. The continuous distributed monitoring model. *SIGMOD Record* 42, 1 (2013), 5–14.

[27] Graham Cormode and Minos N. Garofalakis. 2008. Streaming in a connected world: querying and tracking distributed data streams. In *EDBT 2008, 11th International Conference on Extending Database Technology, Proceedings*. 745.

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

[29] Graham Cormode and S. Muthukrishnan. 2005. What's new: finding significant differences in network data streams. *IEEE/ACM Transactions on Networking* 13, 6 (2005), 1219–1232.

[30] Graham Cormode, S. Muthukrishnan, and Irina Rozenbaum. 2005. Summarizing and Mining Inverse Distributions on Data Streams via Dynamic Inverse Sampling. In *Proceedings of the 31st International Conference on Very Large Data Bases, Trondheim, Norway, August 30 - September 2, 2005*. ACM, 25–36.

[31] Graham Cormode, S. Muthukrishnan, Ke Yi, and Qin Zhang. 2010. Optimal sampling from distributed streams. In *Proceedings of the Twenty-Ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS*. 77–86.

[32] Graham Cormode, S. Muthukrishnan, Ke Yi, and Qin Zhang. 2012. Continuous sampling from distributed streams. *J. ACM* 59, 2 (2012), 10:1–10:25.

[33] Mayur Datar, Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 2002. Maintaining Stream Statistics over Sliding Windows. *SIAM J. Comput.* 31, 6 (2002), 1794–1813.

[34] Mayur Datar and Rajeev Motwani. 2007. The Sliding-Window Computation Model and Results. In *Data Streams - Models and Algorithms*. Springer, 149–167.

[35] Amit Deshpande, Luis Rademacher, Santosh S. Vempala, and Grant Wang. 2006. Matrix Approximation and Projective Clustering via Volume Sampling. *Theory of Computing* 2, 12 (2006), 225–247.

[36] Amit Deshpande and Kasturi R. Varadarajan. 2007. Sampling-based dimension reduction for subspace approximation. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*. 641–650.

[37] Amit Deshpande and Santosh S. Vempala. 2006. Adaptive Sampling and Fast Low-Rank Matrix Approximation. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 9th International Workshop, APPROX and 10th International Workshop, RANDOM. Proceedings*. 292–303.

[38] Alessandro Epasto, Silvio Lattanzi, Sergei Vassilvitskii, and Morteza Zadimoghaddam. 2017. Submodular Optimization Over Sliding Windows. In *Proceedings of the 26th International Conference on World Wide Web, WWW*. 421–430.

[39] Cristian Estan and George Varghese. 2003. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. *ACM Trans. Comput. Syst.* 21, 3 (2003), 270–313.

[40] Alan M. Frieze, Ravi Kannan, and Santosh S. Vempala. 2004. Fast monte-carlo algorithms for finding low-rank approximations. *J. ACM* 51, 6 (2004), 1025–1041.

[41] Rainer Gemulla, Wolfgang Lehner, and Peter J. Haas. 2008. Maintaining bounded-size sample synopses of evolving datasets. *VLDB J.* 17, 2 (2008), 173–202.

[42] Phillip B. Gibbons and Yossi Matias. 1998. New Sampling-Based Summary Statistics for Improving Approximate Query Answers. In *SIGMOD, Proceedings ACM SIGMOD International Conference on Management of Data*. 331–342.

[43] Anna C Gilbert, Yannis Kotidis, S Muthukrishnan, and Martin Strauss. 2001. Quicksand: Quick summary and analysis of network data. Technical report.

[44] Peter J. Haas. 2016. Data-Stream Sampling: Basic Techniques and Results. In *Data Stream Management - Processing High-Speed Data Streams*. Springer, 13–44.

[45] Peter J. Haas, Jeffrey F. Naughton, S. Seshadri, and Arun N. Swami. 1996. Selectivity and Cost Estimation for Joins Based on Random Sampling. *J. Comput. Syst. Sci.* 52, 3 (1996), 550–569.

[46] Peter J. Haas and Arun N. Swami. 1992. Sequential Sampling Procedures for Query Size Estimation. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. 341–350.

[47] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, Yossi Matias, and Uri Stemmer. 2020. Adversarially Robust Streaming Algorithms via Differential Privacy. In *Advances in Neural Information Processing Systems*.

[48] Xiaowei Hu, LA Prashanth, András György, and Csaba Szepesvari. 2016. (Bandit) convex optimization with biased noisy gradient oracles. In *Artificial Intelligence and Statistics*. PMLR, 819–828.

[49] Ling Huang, XuanLong Nguyen, Minos N. Garofalakis, Joseph M. Hellerstein, Michael I. Jordan, Anthony D. Joseph, and Nina Taft. 2007. Communication-Efficient Online Detection of Network-Wide Anomalies. In *INFOCOM. 26th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies*. 134–142.

[50] Nikita Ivkin, Daniel Rothchild, Enayat Ullah, Ion Stoica, Raman Arora, et al. 2019. Communication-efficient distributed sgd with sketching. In *Advances in Neural Information Processing Systems*. 13144–13154.

[51] Rajesh Jayaram, Gokarna Sharma, Srikanta Tirthapura, and David P. Woodruff. 2019. Weighted Reservoir Sampling from Distributed Streams. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS*. 218–235.

[52] Rajesh Jayaram and David P. Woodruff. 2018. Data Streams with Bounded Deletions. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database System*. 341–354.

[53] Rajesh Jayaram and David P. Woodruff. 2018. Perfect L_p Sampling in a Data Stream. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS*. 544–555.

[54] Rajesh Jayaram, David P. Woodruff, and Samson Zhou. 2021. Truly Perfect Samplers for Data Streams and Sliding Windows. *CoRR* abs/2108.12017 (2021).

[55] Rie Johnson and Tong Zhang. 2013. Accelerating Stochastic Gradient Descent using Predictive Variance Reduction. In *27th Annual Conference on Neural Information Processing Systems*. 315–323.

[56] Hossein Jowhari, Mert Saglam, and Gábor Tardos. 2011. Tight bounds for L_p samplers, finding duplicates in streams, and related problems. In *Proceedings of the 30th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS*. 49–58.

[57] Lap-Kei Lee and H. F. Ting. 2006. Maintaining significant stream statistics over sliding windows. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. 724–732.

[58] Lap-Kei Lee and H. F. Ting. 2006. A simpler and more efficient deterministic scheme for finding frequent items over sliding windows. In *Proceedings of the Twenty-Fifth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*. 290–297.

[59] Richard J. Lipton and Jeffrey F. Naughton. 1995. Query Size Estimation by Adaptive Sampling. *J. Comput. Syst. Sci.* 51, 1 (1995), 18–25.

[60] Richard J. Lipton, Jeffrey F. Naughton, and Donovan A. Schneider. 1990. Practical Selectivity Estimation through Adaptive Sampling. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. 1–11.

[61] Sepideh Mahabadi, Ilya P. Razenshteyn, David P. Woodruff, and Samson Zhou. 2020. Non-adaptive adaptive sampling on turnstile streams. In *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC*. 1251–1264.

[62] Jianning Mai, Chen-Nee Chuah, Ashwin Sridharan, Tao Ye, and Hui Zang. 2006. Is sampled data sufficient for anomaly detection?. In *Proceedings of the 6th ACM SIGCOMM Internet Measurement Conference, IMC*. 165–176.

[63] Gurmeet Singh Manku and Rajeev Motwani. 2012. Approximate Frequency Counts over Data Streams. *PVLDB* 5, 12 (2012), 1699.

[64] Jayadev Misra and David Gries. 1982. Finding Repeated Elements. *Sci. Comput. Program.* 2, 2 (1982), 143–152.

[65] Slobodan Mitrovic, Ilija Bogunovic, Ashkan Norouzi-Fard, Jakub Tarnawski, and Volkan Cevher. 2017. Streaming Robust Submodular Maximization: A Partitioned Thresholding Approach. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*. 4557–4566.

[66] Morteza Monemizadeh and David P. Woodruff. 2010. 1-Pass Relative-Error L_p-Sampling with Applications. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. 1143–1160.

[67] Deanna Needell, Nathan Srebro, and Rachel Ward. 2016. Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm. *Math. Program.* 155, 1-2 (2016), 549–573.

[68] Miles Osborne, Sean Moran, Richard McCreadie, Alexander Von Lunen, Martin Sykora, Elizabeth Cano, Neil Ireson, Craig MacDonald, Iadh Ounis, Yulan He, Tom Jackson, Fabio Ciravegna, and Ann O'Brien. 2014. Real-Time Detection, Tracking and Monitoring of Automatically Discovered Events in Social Media. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*.

[69] Odysseas Papapetrou, Minos N. Garofalakis, and Antonios Deligiannakis. 2015. Sketching distributed sliding-window data streams. *VLDB J.* 24, 3 (2015), 345–368.

[70] Marina Thottan, Guanglei Liu, and Chuanyi Ji. 2010. Anomaly Detection Approaches for Communication Networks. In *Algorithms for Next Generation Networks*. Springer, 239–261.

[71] Srikanta Tirthapura and David P. Woodruff. 2011. Optimal Random Sampling from Distributed Streams Revisited. In *Distributed Computing - 25th International Symposium, DISC. Proceedings*, Vol. 6950. 283–297.

[72] Jeffrey Scott Vitter. 1985. Random Sampling with a Reservoir. *ACM Trans. Math. Softw.* 11, 1 (1985), 37–57.

[73] Zhewei Wei, Xuancheng Liu, Feifei Li, Shuo Shang, Xiaoyong Du, and Ji-Rong Wen. 2016. Matrix Sketching Over Sliding Windows. In *Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference*. 1465–1480.

[74] David P. Woodruff and Peilin Zhong. 2016. Distributed low rank approximation of implicit functions of a matrix. In *32nd IEEE International Conference on Data Engineering, ICDE*. 847–858.

[75] David P. Woodruff and Samson Zhou. 2020. Tight Bounds for Adversarially Robust Streams and Sliding Windows via Difference Estimators. *CoRR* abs/2011.07471 (2020).

[76] Peilin Zhao and Tong Zhang. 2015. Stochastic Optimization with Importance Sampling for Regularized Loss Minimization. In *Proceedings of the 32nd International Conference on Machine Learning, ICML*, Vol. 37. 1–9.

A APPENDIX

Proof of Theorem 3.1 : The probability that s is the j^{th} particular instance of item i inside the stream is $\frac{1}{m}$. Since the number of instances of i appearing after j is $f_i - j$ then the probability that i is output is

$$\sum_{j=1}^{f_i} \frac{1}{m} \frac{G(f_i - j + 1) - G(f_i - j)}{\zeta} = \frac{G(f_i)}{\zeta m}.$$

We note that $G(f_i - j + 1) - G(f_i - j) \leq \zeta$ for all $j \in [f_i]$, so returning s with probability $\frac{G(c+1)-G(c)}{\zeta}$ is valid.

Hence the probability that some index is returned by Algorithm 2 is $\sum_{i\in[n]} \frac{G(f_i)}{\zeta m} = \frac{F_G}{\zeta m}$, where $F_G = \sum G(f_i)$. Thus by repeating the sampler $O\left(\frac{\zeta m}{F_G} \log \frac{1}{\delta}\right)$ times, the algorithm will output a sample s with probability at least $1 - \delta$. Although the algorithm does not actually have the value of F_G, given a lower bound $\widehat{F_G}$ on F_G, then it suffices to repeat the sampler $O\left(\frac{\zeta m}{\widehat{F_G}} \log \frac{1}{\delta}\right)$ times. Moreover, the sample s will output each index $i \in [n]$ with probability $\frac{G(f_i)}{F_G}$. Each instance only requires $O(\log n)$ bits to maintain the counter c, assuming $\log m = O(\log n)$. Thus the total space used is $O\left(\frac{\zeta m}{\widehat{F_G}} \log n \log \frac{1}{\delta}\right)$ bits of space.

Finally, we remark that the runtime of the algorithm can be optimized to constant time per update by storing a hash table containing a count and a list of offsets. Specifically, when item i is first sampled by some repetition of the algorithm, then we start counting the number of subsequent instances of i in the stream. If i is subsequently sampled by another independent instance of the reservoir sampling at some time t, then it suffices to store the value of the counter at the time t as an offset. This value does not change and can now be used to correctly recover the correct count of the number of instances of i after time t by subtracting this offset from the largest count. Finally, we can maintain a hash table with pointers to the head and tail of the list, so that when an item i is sampled, we can efficiently check whether that item is already being tracked by another repetition of the sampler. Hence the update time is $O(1)$ worst-case once the hash bucket for i is determined and $O(1)$ in expectation overall given the assignment of the bucket by the hash function. Finally, note that by design of the offsets, we can build the correct counters at the end of the stream to determine the corresponding sampling probabilities. □

Theorem A.1. *For $p \in [1, 2]$ and a frequency vector f implicitly defined by an insertion-only stream, there exists an algorithm that returns each $i \in [n]$ with probability $\frac{f_i^p}{\sum_{j\in[n]} f_j^p}$, using $O\left(n^{1-1/p} \log n\right)$ bits of space.*

Proof. By Theorem 3.2, using a single MisraGries data structure with $O\left(n^{1-1/p} \log n\right)$ bits of space allows us to obtain a number Z such that

$$\|f\|_\infty \leq Z \leq \|f\|_\infty + \frac{m}{n^{1-1/p}}.$$

Note that for $p \in [1, 2]$, the function $x^p - (x-1)^p$ is maximized at $p = 2$, equal to $2x^{p-1}$, by the generalized binomial theorem. Since $f_i \leq \|f\|_\infty$, then we have $(f_i)^p - (f_i - 1)^p \leq 2\|f\|_\infty^{p-1}$ for any $i \in [n]$, so that $\zeta = 2Z^{p-1}$ induces a valid sampling procedure. Hence each instance outputs some index $i \in [n]$ with probability at least $\frac{F_p}{2Z^{p-1}m}$. If $\|f\|_\infty \geq \frac{m}{n^{1-1/p}}$, then we have $2Z \leq 4\|f\|_\infty \leq 4\|f\|_p$, so that

$$\frac{F_p}{2Z^{p-1}m} \geq \frac{F_p}{4L_p^{p-1} \cdot m} = \frac{L_p}{4F_1} \geq \frac{1}{4n^{1-1/p}}.$$

On the other hand, if $\|f\|_\infty \leq \frac{m}{n^{1-1/p}}$, then we have $2Z \leq \frac{4m}{n^{1-1/p}}$, so that

$$\frac{F_p}{2Z^{p-1}m} \geq \frac{F_p \cdot n^{(p-1)^2/p}}{4m^p} = \frac{F_p \cdot n^{(p-1)^2/p}}{4F_1^p} \geq \frac{F_p \cdot n^{(p-1)^2/p}}{4F_p \cdot n^{(p-1)}}$$

$$= \frac{1}{4n^{1-1/p}}.$$

Therefore, the probability that an instance outputs some index $i \in [n]$ is at least $\frac{1}{4n^{1-1/p}}$, and it suffices to use $O\left(n^{1-1/p}\right)$ such instances, with total space $O\left(n^{1-1/p} \log n\right)$ bits of space. By Theorem 3.2, conditioned on an index being returned by the algorithm, the probability that each coordinate $i \in [n]$ is output is $\frac{f_i^p}{F_p}$. □

Theorem A.2. *For $p \in (0, 1]$ and a frequency vector f implicitly defined by an insertion-only stream, there exists an algorithm that returns each $i \in [n]$ with probability $\frac{f_i^p}{\sum_{j\in[n]} f_j^p}$, using $O\left(m^{1-p} \log n\right)$ bits of space.*

Proof. Note that for $p \in (0, 1]$, we have $(f_i)^p - (f_i - 1)^p \leq 1$ for any $i \in [n]$, so that $\zeta = 1$ induces a valid sampling procedure. Hence each instance outputs some index $i \in [n]$ with probability at least $\frac{F_p}{m} \geq \frac{1}{m^{1-p}}$. Therefore, it suffices to use $O\left(m^{1-p}\right)$ such instances, with total space $O\left(m^{1-p} \log n\right)$ bits of space and conditioned on an index being returned by the algorithm, the probability that each coordinate $i \in [n]$ is output is $\frac{f_i^p}{F_p}$. □

Together, Theorem A.1 and Theorem A.2 give Theorem 3.3. We now prove Corollary 3.4.

Corollary 3.4. *There exist truly perfect G samplers for the insertion-only streaming model that succeed with probability at least $1 - \delta$ and use $O\left(\log n \log \frac{1}{\delta}\right)$ bits of space when G is the $L_1 - L_2$ estimator, the Fair estimator, or the Huber estimator.*

Proof. For the $L_1 - L_2$ estimator, $G(x) = 2\left(\sqrt{1 + \frac{x^2}{2}} - 1\right)$ so that $G(x) - G(x-1) < 3$ for $x \geq 1$. Moreover, $G(x) > |x|$ so $F_G > m$. Hence by Theorem 3.1, there exists a perfect G sampler that uses $O(\log n)$ bits of space when G is the $L_1 - L_2$ estimator.

For the Fair estimator, we have $G(x) = \tau|x| - \tau^2 \log\left(1 + \frac{|x|}{\tau}\right)$ for some constant $\tau > 0$ so that $G(x) - G(x-1) < \tau$ for $x \geq 1$. Since $G(x) > \tau|x|$ and thus $F_G > \tau m$, then by Theorem 3.1, there exists a perfect G sampler that uses $O(\log n)$ bits of space when G is the Fair estimator.

For the Huber measure function, we have $G(x) = \frac{x^2}{2\tau}$ for $|x| \leq \tau$ and $G(x) = |x| - \frac{\tau}{2}$ otherwise, where $\tau > 0$ is some constant parameter. Hence, $G(x) - G(x-1) < 1$ and $G(x) > \frac{\tau}{2} \cdot m$, so there exists a perfect G sampler that uses $O(\log n)$ bits of space when G is the Huber estimator by Theorem 3.1. □