
Fast White-Box Adversarial Streaming Without a Random Oracle

Ying Feng¹ Aayush Jain¹ David P. Woodruff¹

Abstract

Recently, the question of adversarially robust streaming, where the stream is allowed to depend on the randomness of the streaming algorithm, has gained a lot of attention. In this work, we consider a strong white-box adversarial model (Ajtai et al. PODS 2022), in which the adversary has access to all past random coins and the parameters used by the streaming algorithm. We focus on the sparse recovery problem and extend our result to other tasks such as distinct element estimation and low-rank approximation of matrices and tensors. The main drawback of previous work is that it requires a *random oracle*, which is especially problematic in the streaming model since the amount of randomness is counted in the space complexity of a streaming algorithm. Also, the previous work suffers from large update time. We construct a near-optimal solution for the sparse recovery problem in white-box adversarial streams, based on the subexponentially secure Learning with Errors assumption. Importantly, our solution does not require a random oracle and has a polylogarithmic per item processing time. We also give results in a related white-box adversarially robust distributed model. Our constructions are based on homomorphic encryption schemes satisfying very mild structural properties that are currently satisfied by most known schemes.

1. Introduction

The streaming model of computation has emerged as an increasingly popular paradigm for the analysis of massive datasets, where the overwhelming size of the data places restrictions on the memory, computation time, and other re-

sources available to the algorithm. In the streaming model, the input data is implicitly defined through a stream of data elements that sequentially arrive one-by-one. One can also delete a previous occurrence of an item. An algorithm makes one pass over the stream and uses limited space to approximate a desired function of the input. In the formalization of Alon, Matias, and Szegedy (Alon et al., 1996), the data is represented as an underlying n -dimensional vector that is initialized to 0^n , and then undergoes a sequence of additive updates to its coordinates. The algorithm aims to optimize the space and time to process these updates, while being able to respond to queries about this underlying vector.

This model captures key resource requirements of algorithms for many practical settings such as Internet routers and traffic logs (Lall et al., 2006; Zhao et al., 2005; Sung et al., 2006; Venkataraman et al., 2005; Cormode & Muthukrishnan, 2005; Cormode & Garofalakis, 2008), databases (Selinger et al., 1979; Dasu et al., 2002; Flajolet & Nigel Martin, 1985; Lightstone, 2018; Shukla et al., 1996; Brown et al., 2005), sensor networks (Hu et al., 2022; Gaber, 2007), financial transaction data (Almeida et al., 2023), and scientific data streams (Chandrasekaran et al., 2003; Szalay et al., 2002). See (Muthukrishnan, 2005) for an overview of streaming algorithms and their applications.

A large body of work on streaming algorithms has been designed for *oblivious* streams, for which the sequence of updates may be chosen adversarially, but it is chosen independently of the randomness of the streaming algorithm. Recently, there is a growing body of work on robust streaming algorithms in the black-box adversarial model (Ben-Eliezer & Yogev, 2020; Ben-Eliezer et al., 2021; Hassidim et al., 2020; Woodruff & Zhou, 2021; Alon et al., 2021; Kaplan et al., 2021; Braverman et al., 2021; Menuhin & Naor, 2021; Attias et al., 2021; Ben-Eliezer et al., 2022; Chakrabarti et al., 2022), in which the adversary can monitor the outputs of the streaming algorithm and choose future stream updates based on these outputs. While useful in a number of applications, there are other settings where the adversary may also have access to the internal state of the algorithm, and this necessitates looking at a stronger adversarial model known as the white-box adversarial streaming model.

¹Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Correspondence to: Ying Feng <yingfeng@andrew.cmu.edu>, Aayush Jain <aayushja@andrew.cmu.edu>, David P. Woodruff <dwoodruf@cs.cmu.edu>.

1.1. White-box Adversarial Streaming Model

We consider the white-box adversarial streaming model introduced in (Ajtai et al., 2022), where a sequence of stream updates is chosen adaptively by an adversary who sees the full internal state of the algorithm at all times, including the parameters and the previous randomness used by the algorithm. The interaction between an adversary and the streaming algorithm is informally depicted as the following game.

Consider a multi-round, two-player game between StreamAlg, the streaming algorithm, and an adversary \mathcal{A} . Prior to the beginning of the game, fix a query function f , which asks for a function of an underlying dataset that will be implicitly defined by the stream. In each round:

1. \mathcal{A} computes an update x for the stream, which depends on all previous stream updates, all previous internal states, and the randomness used by StreamAlg.
2. StreamAlg acquires a fresh batch of random bits, uses x to update its data structures, and (if asked) outputs a response to the query function f .
3. \mathcal{A} observes the random bits, the internal state of StreamAlg, and the response.

The goal of \mathcal{A} is to make StreamAlg output an incorrect query at some round during the game.

The white-box adversarial model captures characteristics of many real-world attacks, where an adaptive adversary has access to the entirety of the internal states of the system. In comparison to the oblivious stream model or the black-box adversarial model (Ben-Eliezer et al., 2021), this model allows for richer adversarial scenarios.

For example, consider a distributed setting where a centralized server collects statistics of a database generated by remote users. The server may send components of its internal state S to the remote users in order to optimize the total communication over the network. The remote users may use S in some process that generates downstream data. Thus, future inputs depend on the internal state S of the algorithm run by the central coordinator. In such settings, the white-box robustness of the algorithms is crucial for optimal selection of query plans (Selinger et al., 1979), online analytical processing (Shukla et al., 1996), data integration (Brown et al., 2005), and data warehousing (Dasu et al., 2002).

Many persistent data structures provide the ability to quickly access previous versions of information stored in a repository shared across multiple collaborators. The internal persistent data structures used to provide version control may be accessible and thus visible to all users of the repository. These users may then update the persistent data structure in

a manner that is not independent of previous states (Driscoll et al., 1989; Fiat & Kaplan, 2003; Kaplan, 2004).

Dynamic algorithms often consider an adaptive adversary who generates the updates upon seeing the entire data structure maintained by the algorithm during the execution (Chan, 2010; Chan & He, 2021; Roghani et al., 2022). For example, (Wajc, 2020) assumes the entire state of the algorithm (including the set of randomness) is available to the adversary after each update, i.e., a white-box model.

Moreover, robust algorithms and adversarial attacks are important topics in machine learning (Szegedy et al., 2014; Goodfellow et al., 2014), with a large body of recent literature focusing on adversarial robustness of machine learning models against white-box attacks (Ilyas et al., 2018; Madry et al., 2018; Schmidt et al., 2018; Tramèr et al., 2018; Cubuk et al., 2018; Kurakin et al., 2017; Liu et al., 2017).

1.2. Prior Work in the White-Box Model

We aim to design *white-box adversarially robust* (WAR) streaming algorithms with provable correctness guarantees against such powerful adversaries. In general, it is difficult to design non-trivial WAR algorithms, due to the fact that many widely used space-saving techniques are subject to attacks when all parameters and generated randomness of the system are immediately revealed to the adversary. In the streaming community, both (Ajtai et al., 2022) and (Feng & Woodruff, 2023) explicitly study the white-box model for data streams and suggest the use of cryptography in designing WAR algorithms (see also (Ben-Eliezer et al., 2021) for the use of cryptography for black box robustness). A similar model to the white-box model has also been independently developed in the cryptography community in the study of property preserving hashing (Boyle et al., 2019).

In this work, the central problem we consider is k -sparse recovery as follows:

Sparse Recovery: given an input vector x , if x contains at most k non-zero entries, recover x . Otherwise, report *invalid input*.

A primitive for sparse recovery in the white-box model can be used for tasks such as distinct element estimation, low rank approximation of matrices and tensors, and finding a maximum matching (Ajtai et al., 2022; Feng & Woodruff, 2023). Sparse recovery is also itself a critical data acquisition and processing problem that arises in signal and image processing, machine learning, data networking, and medicine, see (Ajtai et al., 2022; Feng & Woodruff, 2023) and the references therein.

We note that a weaker definition of sparse recovery assumes that the input vector x is promised to be k -sparse, and only under this promise recovers x . If the input is not k -sparse,

the algorithm can output any vector. This problem can in fact be solved deterministically (and thus also in the white-box model) with low memory (see, e.g., (Berinde et al., 2008), for a version with fast update time, though other simpler schemes based on MDS codes exist). The drawback under this weaker definition though is that, without knowing the sparsity of the input ahead of time, a user of the algorithm cannot tell whether the output is a correct recovery or garbage. In the latter case, the client might ignorantly use the garbage output in subsequent applications and propagate such an error.

(Feng & Woodruff, 2023) solves this problem based on the assumed hardness of the Short-Integer-Solution (SIS) problem, which is a problem that is studied in lattice-based cryptography. The robust version of the property preserving hashing (Boyle et al., 2019) can also be applied to solve the same problem. However, a notable drawback of all prior solutions proposed is their reliance on either a *random oracle* or a prohibitively long random string. This is because a key component of these algorithms is a giant, random matrix \mathbf{A} for which it is hard to find a vector \mathbf{x} with small entries for which $\mathbf{A}\mathbf{x}$ is sparse. Long random strings are especially problematic in data streams, where the resource measure is the total space complexity, including the randomness, and there is a large body of work on derandomizing such algorithms, see, e.g., (Nisan, 1992; Indyk, 2006; Kane et al., 2010; Kacham et al., 2023). Since it was previously unknown how to generate \mathbf{A} pseudorandomly while maintaining this hardness property, previous work (Feng & Woodruff, 2023) relies on a random oracle to heuristically compress such a matrix \mathbf{A} . This assumes that the algorithm is given read access to a long string of random bits, which is often implemented with hash-based heuristic functions such as AES or SHA256. Unfortunately, this heuristic cannot be proven secure using a standard reduction-style security proof. This motivates a central question of this work: Do there exist efficient streaming algorithms in the white-box model without a random oracle?

1.3. Our Results

Our algorithms make use of the hardness assumptions of the Learning with Error (LWE) problem and its variant, Ring LWE. Both problems are popular in cryptography known for building homomorphic encryption schemes. In this section, we state the efficiency of our algorithms under the sub-exponential hardness assumptions of these problems. However, we stress that in contrast to (Feng & Woodruff, 2023), our algorithms are robust even under weaker polynomial hardness assumptions. We let n be the dimension of the underlying input vector, and for notational convenience, we assume that the entries of the vector are integers bounded in absolute value by $\text{poly}(n)$ at all times during the stream. We use $\tilde{O}(f)$ to denote $f \cdot \text{poly}(\log n)$.

Our main result is informally stated as follows:

Theorem 1.1 (Informal). *Assuming the sub-exponential hardness¹ of the Learning with Errors (LWE) problem, there is a WAR streaming algorithm for k -sparse recovery, which:*

- takes $\tilde{O}(k)$ bits of space,
- has $\tilde{O}(1)$ update time, and
- has $\tilde{O}(k^{1+c})$ report time for an arbitrarily small constant $c > 0$.

This matches (up to a poly-logarithmic factor) the optimal space and update time complexity for the weaker k -sparse recovery problem under the easier oblivious streaming setting. On the other hand, assuming polynomial security of the LWE problem, there is a multiplicative overhead of n^ϵ for arbitrary constant $\epsilon > 0$ on our space and time complexities.

Adversarial Model in the Distributed Setting. We introduce and formalize the white-box adversarial model for distributed data processing as a means of modeling the existence of one or more malicious parties in a distributed system. We also study the sparse recovery problem in a white-box distributed model. A WAR streaming algorithm can be naively converted to a WAR distributed algorithm under certain conditions, yet the resulting processing time of the servers is $\mathcal{O}(n)$ times that of the streaming update time. In contrast to this suboptimal time implied by the naive conversion, we show that we can achieve a near-optimal processing time in a white-box distributed model assuming the *polynomial* hardness of Ring-LWE, which is a variant of the LWE problem defined over a polynomial ring. This is because Ring-LWE supports SIMD (Single-Instruction Multiple-Data) style operations that have been exploited in the design of efficient homomorphic encryption schemes in cryptography. We informally state our result as follows:

Theorem 1.2 (Informal). *Assuming the hardness of Ring-LWE, there exists a WAR distributed protocol for k -sparse recovery, which:*

- uses $\tilde{O}(k)$ bits of communication between each server and the coordinator,
- has $\tilde{O}(n)$ processing time on each server, and
- the coordinator spends $\tilde{O}(\max(n, k^{1+c}))$ time to output the solution given the messages from the servers, for an arbitrarily small constant $c > 0$.

We summarize the complexities of our sparse recovery algorithms in Table 1, as compared to the best-known upper bounds for these problems in the white-box adversarial model.

¹Sub-exponential hardness means that we assume that for adversaries of size 2^{λ^β} , where λ is the security parameter and $\beta > 0$ is some constant, a sample from an LWE distribution is indistinguishable from a uniformly random sample.

Table 1. A summary of the bit complexities and runtime of our k -sparse recovery algorithm, as compared to the best-known upper bounds for these problems in the white-box adversarial model (Feng & Woodruff, 2023). The SIS assumption used in the previous work is proven to be equivalently hard to the LWE assumption we use, up to different parameters. The TIME columns refer to the update time in the streaming model and processing time in the distributed model. c denotes an arbitrarily small positive constant.

	MODEL	ASSUMPTION	SPACE	TIME	RAND. ORACLE?
PREV.	STREAMING	SUBEXP SIS	$\tilde{O}(k)$	$\tilde{O}(k)$	✓
	DISTRIBUTED	SUBEXP SIS	$\tilde{O}(k)$	$\tilde{O}(nk)$	✓
OURS	STREAMING	SUBEXP LWE	$\tilde{O}(k)$	$\tilde{O}(1)$	✗
	STREAMING	POLY LWE	$\tilde{O}(n^c + k)$	$\tilde{O}(n^c)$	✗
	DISTRIBUTED	SUBEXP LWE	$\tilde{O}(k)$	$\tilde{O}(n)$	✗
	DISTRIBUTED	POLY LWE	$\tilde{O}(n^c + k)$	$\tilde{O}(n^{1+c})$	✗
	DISTRIBUTED	POLY RING-LWE	$\tilde{O}(n^c + k)$	$\tilde{O}(n)$	✗

Extensions to matrix and tensor recoveries. Similar to (Feng & Woodruff, 2023), generalizing our construction produces low-rank matrix and tensor recovery algorithms, again with the crucial property that they do not assume a random oracle or store a large random seed. These results are given in Appendix G due to the page limit.

Consider a data stream updating the entries of an underlying matrix or tensor, which can be represented as its vectorization in a stream. Let $n \geq m$ be the dimensions of an underlying input matrix, and let n be the dimension parameter of an underlying data tensor in \mathbb{R}^{n^d} for some $d \in \mathcal{O}(1)$ (we assume all dimensions of the tensor are the same only for ease of notation). We have the following results in the streaming model:

Theorem 1.3 (Informal). *Assuming the sub-exponential hardness of the Learning with Errors (LWE) problem, there is a WAR streaming algorithm for rank- k matrix recovery, which takes $\tilde{O}(nk)$ bits of space. There is also a WAR streaming algorithm for rank- k tensor recovery, which takes $\tilde{O}(nk^{\lceil \log d \rceil})$ bits of space.*

We remark that the above matrix recovery algorithm performs $\tilde{O}(nk)$ measurements using n -sparse matrices, which can be converted into a distributed protocol with $\tilde{O}(n^2k)$ processing time. This improves upon the matrix recovery algorithm in the previous work (Feng & Woodruff, 2023) which, if implemented in a distributed setting, would require $\tilde{O}(n^3k)$ processing time at each server.

Additional Applications. Our k -sparse recovery algorithm can be used to construct an L_0 -norm estimation algorithm following the same recipe in (Feng & Woodruff, 2023). In short, by running our k -sparse recovery algorithm with the sparsity parameter k set to $n^{1-\alpha}$ for some constant α , we can get an n^α -approximation to the L_0 norm of a vector. In comparison to (Feng & Woodruff, 2023), our construction again removes the need for a random oracle or the need to store a lot of randomness.

1.4. Technical Overview

We leverage structural properties of lattice-based assumption to address all drawbacks listed above. In particular, departing from the previous construction of (Feng & Woodruff, 2023), we work with a suitable fully homomorphic encryption (FHE) scheme to achieve the desired property. We start with an intuitive explanation of the previous construction in (Feng & Woodruff, 2023).

Previous Construction One general framework for solving the sparse recovery problem is to run a weak recovery algorithm as a black-box subroutine, which potentially provides a garbage output, in addition to an additional “tester” scheme to verify whether the recovery output matches the original input (i.e., is a correct recovery) or is garbage. In the oblivious streaming setting, such a “tester” can be efficiently implemented using, e.g., a random polynomial “fingerprint” (Cormode & Jowhari, 2019). However, this approach is not WAR as the randomness of the sampler is revealed and thus not independent of the stream.

Instead, (Feng & Woodruff, 2023) uses a random matrix \mathbf{A} to implement a WAR tester, which is assumed to be heuristically compressed by an idealized hash function. During the stream, (Feng & Woodruff, 2023) maintains a hash $h = \mathbf{A}\mathbf{x}$, where \mathbf{x} is the data vector implicitly defined by the stream. In the testing stage, it acknowledges an output \mathbf{x}' to be a correct recovery if and only if \mathbf{x}' is k -sparse and $\mathbf{A}\mathbf{x}'$ equals the hash h . The idea is that if there exists a k -sparse collision \mathbf{x}' satisfying $\mathbf{A}(\mathbf{x} - \mathbf{x}') = 0$ and $\mathbf{x}' \neq \mathbf{x}$, then an adversary can brute-force over all k -sparse vectors to find such $\mathbf{x} - \mathbf{x}'$, which is a solution to the SIS problem. To spawn n matrix columns and support the brute-force reduction, the security parameter (i.e., the number of rows of \mathbf{A}) needs to be in $\omega(k \log n)$, assuming the sub-exponential security of SIS.

Efficient Reduction. The number of matrix rows in the previous construction depends on the sparsity parameter k due to a brute-forcing step in the reduction. To grant enough time for iterating through all k -sparse vectors (with $\text{poly}(n)$ -bounded entries), the security parameter has to be greater than $k \log n$, which results in $\tilde{O}(k)$ time to process every single update. Our first idea is to construct a poly-time reduction in order to remove such a dependency on k . This allows us to achieve better space and time efficiency, in addition to basing the construction on a polynomial hardness assumption.

The observation is that the non-existence of a sparse collision is an overly strong condition. Instead, for our purpose, we only need to ensure that such a collision, if it exists, cannot be found by the poly-time weak recovery scheme that the algorithm runs. Thus, in the reduction, instead of brute-forcing over all k -sparse vectors, we let the adversary run the same poly-time weak recovery scheme that the algorithm does. If a collision \mathbf{x}' satisfying $\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}'$ is output by the scheme, then the adversary can still break the security assumption.

Reducing the Amount of Randomness. The rest of our effort will be attributed to reducing the size of the state of the algorithm. Rather than using a truly random giant matrix to hash the data as in the previous construction, we would like to design a short digest (ideally linear-sized in the security parameter) and generate the columns of the hash matrix on the fly. This naturally yields correlated columns as opposed to the truly random columns in an SIS matrix. We aim to show that in this setting, the chance of finding a collision is still negligible.

Streaming Friendly Collision-Resistant Hash Our main insight from cryptography to address the issue above is that we could construct a hash function family \mathcal{H} , with special properties to help solve our streaming problem:

- The hash key $h \leftarrow \mathcal{H}$ is small in size, ideally polynomial in the security parameter λ (or even polylogarithmic assuming subexponential time assumptions).
- The hash should be computable in a streaming fashion and it should be updatable where each update takes polynomial time in the security parameter λ (or even polylogarithmic time assuming subexponential time assumptions).
- The hash key should be pseudorandom. In other words, in our real-world implementation, the white-box adversarial streaming model does not allow us to use any structured randomness (such as with hidden planted secrets). In *the proof*, however, we could introduce these secrets which will be important for our construction.
- The hash should satisfy collision-resistance.

Constructing Streaming Friendly Hashing from FHE.

How could we build such a hash function? The SIS based hash function does have a fast update time. On the other hand, we could take an arbitrary collision resistant function that compresses the number of bits by a factor of $1/2$ and use it to perform a Merkle tree style hashing to hash strings of arbitrary length. Such a hash will have $\text{poly}(\lambda)$ -sized hashing keys. Unfortunately, the first idea suffers from a large hashing key size, which could be shortened in the random oracle (RO) model by using the RO to store the SIS matrix. Moreover, the second proposal will not support positive and negative stream updates to the coordinates of the underlying vector.

Our main idea is to leverage a suitably chosen Fully Homomorphic Encryption (FHE) scheme satisfying mild structural properties (that are satisfied currently by most schemes based on LWE and Ring-LWE) to build such a hash function. Our construction bears a resemblance to the construction of a collision resistant function from a private information retrieval scheme, where the private information retrieval scheme is implemented by a fully homomorphic encryption scheme.

We directly explain the idea in the context of our construction. In our hash function, the hash key consists of $L = O(\log n)$ ciphertexts $\text{ct}_1, \dots, \text{ct}_L$. We assume that our FHE scheme has pseudorandom ciphertexts and pseudorandom public keys. In our actual scheme, these ciphertexts will be generated as random strings. It is only in the proof that they will correspond to encryptions of carefully chosen values.

The idea is that from these L ciphertexts, we will derive n evaluated ciphertexts $\widehat{\text{ct}}_j$ for $j \in [n]$ by evaluating L ciphertexts on simple functions C_1, \dots, C_n (described later). These n ciphertexts will correspond to columns of the SIS type hashing. The hash of a vector \mathbf{x} is computed as the linear combination $\sum_{i \in [n]} \widehat{\text{ct}}_i \mathbf{x}_i$ where \mathbf{x}_i is the i^{th} coordinate. Here the sum is done over the field on which the FHE ciphertexts live.

The circuits C_1, \dots, C_n are chosen keeping the security proof in mind. The ciphertexts $\text{ct}_1, \dots, \text{ct}_L$ are “thought” to encrypt a value $m \in \{0, 1\}^L$ that will be chosen in the proof, and each $C_i(\cdot)$ on input m outputs 1 if $i = m$ and 0 otherwise. Observe that these circuits are polynomial in $L = O(\log n)$ sized and therefore can be evaluated on $\text{ct}_1, \dots, \text{ct}_L$ in polynomial in $(\lambda, \log n)$ time guaranteeing fast updates.

How do we prove security? The idea is that if an adversary breaks collision-resistance it will produce two vectors $\mathbf{x} \neq \mathbf{x}'$, so that the corresponding hashes $\sum_{i \in [n]} \widehat{\text{ct}}_i \mathbf{x}_i =$

$\sum_{i \in [n]} \widehat{\text{ct}}_i x'_i$. Suppose one could know exactly one coordinate $v \in [n]$ so that $x_v \neq x'_v$. Then we could set $m = v$. In this case $\widehat{\text{ct}}_v$ encrypts 1 and the other ciphertexts encrypt 0. In this case, $\sum_{i \in [n]} \widehat{\text{ct}}_i x_i = \sum_{i \in [n]} \widehat{\text{ct}}_i x'_i$ would lead to a contradiction if our FHE scheme has a special property (that is satisfied by common FHE schemes). The property is that if $\widehat{\text{ct}}_i$ encrypts bit $\mu_i \in \{0, 1\}$ and x_i are small norm then there is a decryption algorithm that uniquely binds $\sum_{i \in [n]} \widehat{\text{ct}}_i x_i$ to $\sum_{i \in [n]} \mu_i x_i$. If this holds, then the equality $\sum_{i \in [n]} \widehat{\text{ct}}_i x_i = \sum_{i \in [n]} \widehat{\text{ct}}_i x'_i$ would imply $x_v = x'_v$ which is a contradiction.

This leads us to two remaining problems:

1. Given that we do not know x and x' ahead of time, how do we decide which index m to encrypt such that it will help identify whether $x = x'$?
2. Moreover, a WAR algorithm cannot actually encrypt FHE ciphertexts, since all randomness and parameters used by the algorithm will be exposed to the adversary. Instead, we want our digest to be unstructured and truly random.

FHE with a Pseudorandom Property. We solve both issues above by considering a pseudorandom property of the FHE scheme and show that a random guess of the index m actually suffices. We assume that the distributions of public keys and ciphertexts of the FHE scheme are indistinguishable, respectively, from some truly random distributions from the perspective of the adversary.

We sample truly random digests in our construction, and perform a random guess of $m \in [n]$ in the proof. If an adversary finds a collision pair $x \neq x'$ with probability γ , then with probability $p \geq \gamma/n$, our guess pins down a coordinate of disagreement $x_m \neq x'_m$. In the proof, we then switch the ciphertexts to encrypt m . The chance of guessing the disagreement is still p , as otherwise one can distinguish the pseudorandom ciphertext. However, as argued earlier, $x_m \neq x'_m$ roughly implies $\sum_{i \in [n]} \widehat{\text{ct}}_i x_i \neq \sum_{i \in [n]} \widehat{\text{ct}}_i x'_i$, and thus two distinct vectors cannot form a hash collision that fools our “tester”.

In our complexity analysis, we instantiate the FHE scheme using GSW (Gentry et al., 2013), which satisfies the additional properties required by our construction. We remark that in addition to GSW, most current FHE schemes based on LWE and Ring-LWE also exhibit these properties, and thus they are not an artifact obtained from a particular FHE scheme.

Ring-LWE for the Distributed Setting. In the distributed model, each remote server receives as input a length- n vector. The goal is to design a protocol such that each server

can efficiently compute a short summary of their partition of data. Adapting the above streaming algorithm to this setting, each server needs to evaluate n ciphertexts in order to compute a linear combination $\sum_{i \in [n]} \widehat{\text{ct}}_i x_i$.

Using a polynomially secure LWE-based FHE scheme, it is unclear how to speed up such n -many individual evaluations as each computation takes $\text{poly}(\lambda)$ operations, resulting in a total complexity of $n \text{poly}(\lambda)$, where λ is the security parameter. Therefore, we appeal to the Ring-LWE assumption, which is an algebraic variant of LWE that works over the ring $\mathbb{Z}_p[x]/(x^N + 1)$. Over this ring, we can exploit the SIMD property of Ring-LWE by packing the data as ring elements, i.e., we represent the input vectors as partitioned segments of length proportional to the security parameter (rather than the vector length n), and efficiently operate on them segment-by-segment using the Fast Fourier Transform. This improves the processing time of each server to be almost linear in the dimension of the data. We refer the reader to Section F for more details.

Summary. Leveraging FHE or lattice-based assumptions in algorithm design are relatively new techniques. We think it would be good to explore their applications more broadly in robust algorithms, beyond the recovery problems or even the adversarial settings considered here.

2. Preliminaries

Basic Notations. For a finite set S , we write $x \stackrel{\$}{\leftarrow} S$ to denote a uniform sample x from S . Our logarithms are in base 2. For $n \in \mathbb{R}$, we use $\lceil n \rceil$ to denote rounding n up to the nearest integer. For a vector x , $\|x\|_0$ denotes the ℓ_0 norm of x , which is the number of its non-zero entries. We abbreviate PPT for probabilistic polynomial-time. A function $\text{negl} : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if for every constant $c > 0$ there exists an integer N_c such that $\text{negl}(x) < x^{-c}$ for all $x > N_c$.

2.1. Lattice Preliminaries

We define the Learning with Errors problem LWE.

Definition 2.1 (Learning with Errors). *Suppose we are given an integer dimension $g \in \mathbb{N}$, a modulus $q \in \mathbb{N}$, a sample complexity $h \geq \Omega(g \cdot \log q)$ and a parameter $\sigma \in \mathbb{R}^{\geq 0}$. Define an error distribution as the discrete Gaussian $\chi_{\sigma}^h := \mathcal{G}_{\sigma, \mathbb{Z}^h}$. We define $\text{LWE}_{g, h, q, \chi_{\sigma}^h}$ as follows: sample a matrix $\mathbf{A} \leftarrow \mathbb{Z}_q^{g \times h}$ and a secret $s \in \mathbb{Z}_q^{1 \times g}$. Sample an error vector $e \leftarrow \chi_{\sigma}^h$. The problem is to distinguish $(\mathbf{A}, s\mathbf{A} + e \pmod{q})$ from (\mathbf{A}, r) where r is random in $\mathbb{Z}_q^{1 \times h}$.*

Under quantum and classical reductions it can be shown that $\text{LWE}_{g, h, q, \chi_{\sigma}}$ is harder to solve than Gap-SVP_{γ} with $\gamma = g \frac{q}{\sigma}$, which is a worst-case problem over lattices. We discuss more about the lattice problems in Appendix A.

We also consider the Ring-LWE problem, which is an average-case lattice problems over ideal lattices.

Definition 2.2 (Ring-LWE). *Let $q \geq 1$ be a prime number, and let g be a power of 2. Let g, σ be positive integers. The Ring-LWE $_{g,h,q,\chi_\sigma}$ problem is defined as follows: Consider $R = \frac{\mathbb{Z}_q[x]}{x^g+1}$. Sample at random $\mathbf{a}_1, \dots, \mathbf{a}_h \leftarrow R$ and a secret polynomial $\mathbf{s} \leftarrow R$. Sample h error polynomials $\mathbf{e}_1, \dots, \mathbf{e}_h$ such that each \mathbf{e}_i is a g -dimensional vector $\mathbf{e}_i \leftarrow \chi_\sigma = \mathcal{G}_{\sigma, \mathbb{Z}^g}$. The problem is to distinguish the tuple $(\mathbf{a}_1, \dots, \mathbf{a}_h, \{\mathbf{s}\mathbf{a}_i + \mathbf{e}_i\}_{i \in [h]})$ from $(\mathbf{a}_1, \dots, \mathbf{a}_h, \{\mathbf{r}_i\}_{i \in [h]})$ where $\mathbf{r}_1, \dots, \mathbf{r}_h$ are random polynomials in R .*

The Ring-LWE problem is as hard as worst-case problems over ideal lattices (Lyubashevsky et al., 2010). Ring LWE is sometimes advantageous to use over LWE because of shorter parameters. One ring LWE element is of dimension g , and requires only one ring element as a coefficient, as opposed to g vectors in the case of LWE. Moreover, ring multiplication can be sped up using number-theoretic transforms (Smart & Vercauteren, 2011).

2.2. Fully Homomorphic Encryption

Notation. We denote the security parameter of a cryptographic scheme by λ . We write PPT_λ to denote probabilistic poly-time in λ . We say two ensembles of distributions $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\mathcal{D}'_\lambda\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable, denoted by \approx_c , if for any non-uniform PPT adversary \mathcal{A} there exists a negligible function negl such that \mathcal{A} can distinguish between the two distributions with probability at most $\text{negl}(\lambda)$.

Remark 2.1. We also define (\mathcal{T}, ϵ) -indistinguishability, where \mathcal{T} and ϵ are two functions in λ . We say that two distribution ensembles $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\mathcal{D}'_\lambda\}_{\lambda \in \mathbb{N}}$ are (\mathcal{T}, ϵ) -indistinguishable, denoted by $\approx_{(\mathcal{T}, \epsilon)}$, if there exists an integer λ_0 such that for all $\lambda > \lambda_0$, for every adversary \mathcal{A} , \mathcal{A} can distinguish between the two distributions with probability at most $\epsilon(\lambda)$.

Definition 2.3 (FHE). *A fully homomorphic encryption FHE scheme is a tuple of PPT_λ algorithms $\text{FHE} = (\text{FHE.Setup}, \text{FHE.Enc}, \text{FHE.Eval}, \text{FHE.Dec})$ with the following properties:*

- $\text{FHE.Setup}(1^\lambda) \rightarrow (\text{sk}, \text{pk})$: On input a security parameter λ , the setup algorithm outputs a key pair (pk, sk) .
- $\text{FHE.Enc}(\text{pk}, \mu) \rightarrow \text{ct}$: On input a public key pk and a message $\mu \in \{0, 1\}$, the encryption algorithm outputs a ciphertext ct .
- $\text{FHE.Eval}(\text{pk}, C, \text{ct}_1, \dots, \text{ct}_\ell) \rightarrow \widehat{\text{ct}}$: On input a public key pk , a circuit $C : \{0, 1\}^\ell \rightarrow \{0, 1\}$ of depth at

most $\text{poly}(\lambda)$, and a tuple of ciphertexts $\text{ct}_1, \dots, \text{ct}_\ell$, the evaluation algorithm outputs an evaluated ciphertext $\widehat{\text{ct}}$.

- $\text{FHE.Dec}(\text{pk}, \text{sk}, \widehat{\text{ct}}) \rightarrow \widehat{\mu}$: On input a public key pk , a secret key sk , and a ciphertext $\widehat{\text{ct}}$, the decryption algorithm outputs a message $\widehat{\mu} \in \{0, 1, \perp\}$.

We also require an FHE scheme to satisfy compactness, correctness, and security properties, which are elaborated in Appendix B.

Additional properties. Our construction of robust streaming algorithms will be based on an FHE scheme satisfying a few additional properties. We refer to this type of FHE scheme as a *pseudorandom FHE (PFHE) scheme*.

Definition 2.4 (Pseudorandom FHE (PFHE)). *Let FHE be a fully homomorphic encryption scheme (Definition 2.3). We call FHE a (\mathcal{T}, ϵ) -pseudorandom FHE scheme if for all security parameters λ , it satisfies the following additional properties:*

- FHE.Eval is deterministic.
- (Pseudorandom pk). *There exists an explicit distribution $\mathcal{D}_{\widehat{\text{pk}}(\lambda)}$, such that the following two distributions are (\mathcal{T}, ϵ) -indistinguishable:*
 $\{\text{pk} : (\text{pk}, \text{sk}) \leftarrow \text{FHE.Setup}(1^\lambda)\} \approx_{(\mathcal{T}, \epsilon)} \mathcal{D}_{\widehat{\text{pk}}(\lambda)}$.
- (Pseudorandom ct). *There exists an explicitly distribution $\mathcal{D}_{\widehat{\text{ct}}(\lambda)}$, such that the following two distributions are (\mathcal{T}, ϵ) -indistinguishable:*

$$\left\{ \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{FHE.Setup}(1^\lambda) \\ \text{ct} : \begin{array}{l} \mu \in \{0, 1\} \\ \text{ct} \leftarrow \text{FHE.Enc}(\text{pk}, \mu) \end{array} \end{array} \right\} \approx_{(\mathcal{T}, \epsilon)} \mathcal{D}_{\widehat{\text{ct}}(\lambda)}$$

- (\mathbb{Z} Linear Homomorphism). *Given $(\text{pk}, \text{sk}) \leftarrow \text{FHE.Setup}(1^\lambda)$, the set of all ciphertexts is a subset of a vector space over \mathbb{Z}_q for some $q = q(\lambda)$. Let $\mathcal{M} = \sum_{i \in [k]} x_i \cdot \text{ct}_i \pmod q$ be an arbitrary linear combination of ciphertexts. Denote $\mu_i \leftarrow \text{FHE.Dec}(\text{pk}, \text{sk}, \text{ct}_i)$. Suppose \mathcal{M} satisfies:*

1. $x_i \in \mathbb{Z}$ for all $i \in [k]$
2. $k \cdot \max_{i \in [k]} x_i \ll q$

then we have the following: $|\mathcal{M}| \leq \text{poly}(\lambda)$. Also, there exists a $\text{poly}(\lambda)$ -time decoding algorithm LinearDec , such that $\Pr[\text{LinearDec}(\text{pk}, \text{sk}, \mathcal{M}) = \sum_{i \in [k]} x_i \cdot \mu_i] = 1$.

As in a standard FHE scheme, we also require PFHE to satisfy compactness, correctness, and security.

There exist FHE schemes based on LWE and Ring-LWE (Gentry et al., 2013; Brakerski et al., 2012; Brakerski & Vaikuntanathan, 2011) assumptions. We will focus on the GSW scheme in our streaming construction and the BV scheme in the distributed setting. Both of them satisfy the above definition (or a slight variant) of PFHE.

3. Streaming K -Sparse Recovery

Sparse Recovery Problem We study the sparse recovery problem for vectors in white-box adversarial streaming settings. There exist deterministic (thus WAR) streaming algorithms that recover k -sparse vectors *assuming that the input vector is k -sparse*. However, their outputs can be arbitrary when the input violates such an assumption. In contrast, our algorithm has the crucial property that it can detect if the input violates the sparsity constraint. We describe the problem as follows:

Definition 3.1 (k -Sparse Recovery). *Given a sparsity parameter k and an input vector $\mathbf{x} \in \mathbb{Z}^n$, output \mathbf{x} if $\|\mathbf{x}\|_0 \leq k$ and report \perp otherwise.*

Here and throughout, for the sparsity to make sense we always assume $k \leq n$, though k can be a function in n . Before tackling this problem, we first consider a relaxation for it, which only guarantees correct recovery when the input vector \mathbf{x} is k -sparse. In other words, we have no constraint on the function output when $\|\mathbf{x}\|_0 > k$.

Definition 3.2 (Relaxed k -Sparse Recovery). *Given a sparsity parameter k and an input vector $\mathbf{x} \in \mathbb{Z}^n$, output $\mathbf{x} \in \mathbb{Z}^n$ if $\|\mathbf{x}\|_0 \leq k$.*

As mentioned earlier, there exists a space- and time-efficient, deterministic streaming algorithm that solves the relaxed k -sparse recovery problem, which we will refer to as StreamAlg_0 in our algorithm.

With StreamAlg_0 and a PFHE, we are now ready to construct a WAR streaming algorithm for the k -sparse recovery problem.

Construction 3.1. Let StreamAlg_0 be an algorithm for the Relaxed k -Sparse Recovery problem. Let PFHE be a pseudorandom FHE. Let $\lambda = \mathcal{S}(n)$ be the security parameter of PFHE, for some function \mathcal{S} in n that we will specify later. Choose $q \in \text{poly}(n)$ with $q \gg n \cdot N$ to be the modulus of PFHE ciphertexts, where $N \in \text{poly}(n)$ is a bound on the stream length. All computations will be modulo q . Our construction is as follows:

- $\text{StreamAlg.Setup}(n)$:
Run $\text{StreamAlg}_0.\text{Setup}(n)$. Define $\ell = \lceil \log n \rceil$. For each $i \in [n]$, define a circuit $C_i : \{0, 1\}^\ell \rightarrow \{0, 1\}$, such that: $C_i(\mu_1, \mu_2, \dots, \mu_\ell) =$

$$\begin{cases} 1 & \text{if } (\mu_1, \mu_2, \dots, \mu_\ell) \text{ is the bit representation of } i \\ 0 & \text{otherwise.} \end{cases}$$

Take $\mathcal{D}_{\widetilde{\text{pk}}(\lambda)}$ to be the distribution of pseudo-public key in Definition 2.4. Sample $\widetilde{\text{pk}} \leftarrow \mathcal{D}_{\widetilde{\text{pk}}(\lambda)}$. Similarly, take $\mathcal{D}_{\widetilde{\text{ct}}(\lambda)}$ to be the distribution of pseudo-ciphertext in Definition 2.4. Sample $\widetilde{\text{ct}}_1, \widetilde{\text{ct}}_2, \dots, \widetilde{\text{ct}}_\ell \leftarrow \mathcal{D}_{\widetilde{\text{ct}}(\lambda)}$. Store $(\widetilde{\text{pk}}, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_\ell)$ and $\text{sketch} := 0$.

- $\text{StreamAlg.Update}(i_t, \Delta_t)$:
Run $\text{StreamAlg}_0.\text{Update}(i_t, \Delta_t)$. Evaluate $\widehat{\text{ct}}_{i_t} \leftarrow \text{PFHE.Eval}(\widetilde{\text{pk}}, C_{i_t}, \widetilde{\text{ct}}_1, \widetilde{\text{ct}}_2, \dots, \widetilde{\text{ct}}_\ell)$. Then update $\text{sketch} \leftarrow \text{sketch} + \Delta_t \cdot \widehat{\text{ct}}_{i_t}$.
- $\text{StreamAlg.Report}()$:
Let $\mathbf{x}' \leftarrow \text{StreamAlg}_0.\text{Report}()$. If $\|\mathbf{x}'\|_0 > k$, output \perp . Otherwise, evaluate $\widehat{\text{ct}}_i \leftarrow \text{PFHE.Eval}(\widetilde{\text{pk}}, C_i, \widetilde{\text{ct}}_1, \widetilde{\text{ct}}_2, \dots, \widetilde{\text{ct}}_\ell)$ for all $i \in [n]$. Then compute $\text{hash} \leftarrow \sum_{i \in [n]} \widehat{\text{ct}}_i \cdot \mathbf{x}'_i$. Output \mathbf{x}' if $\text{hash} = \text{sketch}$, and output \perp otherwise.

We state the white-box robustness of our construction as follows. We delay the proof to Appendix D.

Theorem 3.1 (WAR of Construction 3.1). *Suppose StreamAlg_0 is a streaming algorithm for Relaxed k -Sparse Recovery, and PFHE is a pseudorandom FHE scheme with (\mathcal{T}, ϵ) -security. Let $\lambda = \mathcal{S}(n)$ be the PFHE security parameter for some function \mathcal{S} .*

If for all $n \in \mathbb{N}$, functions \mathcal{T}, ϵ , and \mathcal{S} satisfy:

- $\mathcal{T}(\mathcal{S}(n)) = \Omega(\text{poly } n)$
- $\epsilon(\mathcal{S}(n)) \leq \text{negl}(n)$ for some negligible function negl .

then the StreamAlg scheme in Construction 3.1 is a WAR streaming algorithm for the k -Sparse Recovery problem.

4. Efficiency

In this section, we discuss the space and time complexity of Construction 3.1. We start by instantiating StreamAlg_0 and PFHE used in our construction.

Deterministic Relaxed Sparse Recovery. There exist space- and time-efficient streaming algorithms for the relaxed k -sparse recovery problems. Given a sparse vector $\mathbf{x} \in \mathbb{Z}^n$, many of these algorithms perform some linear measurement $\Phi : \mathbb{Z}^n \rightarrow \mathbb{Z}^r$ on \mathbf{x} with $r \ll n$ to get a compressed representation $\Phi(\mathbf{x})$, which is often referred to as a *sketch* of \mathbf{x} . This sketch is later used to reconstruct \mathbf{x} . The linearity property is useful for incrementally maintaining the sketch during the stream. Many algorithms implement the above specification using a random measurement matrix

Φ . However, in our attempt to remove the random oracle assumption within sub-linear total space, we cannot afford to store an n -column random matrix during the stream. Instead, we use algorithms which are deterministic, such that the measurements themselves incur no space overhead.

Theorem 4.1 ((Gilbert et al., 2017)). *There exists a streaming algorithm StreamAlg_0 for the Relaxed k -Sparse Recovery problem, such that given an input parameter n , for an integer stream with entries bounded by $\text{poly}(n)$:*

- StreamAlg_0 .Setup, Update, and Report are all deterministic.
- StreamAlg_0 takes $\tilde{O}(k)$ bits of space.
- StreamAlg_0 .Update runs in $\tilde{O}(1)$.
- StreamAlg_0 .Report runs in $\tilde{O}(k^{1+c})$ for an arbitrarily small constant $c > 0$.

Pseudorandom FHE. We focus on GSW (Gentry et al., 2013), a standard LWE based fully homomorphic encryption scheme, for our instantiation. To start with, we briefly review the construction of GSW. We omit Dec since it is not involved in our scheme. Instead, we will describe a decoding algorithm LinearDec for linear combinations of ciphertexts (which largely resembles the original Dec of GSW) when showing the \mathbb{Z} Linear Homomorphism property of GSW.

- $\text{Setup}(1^\lambda)$: Choose a modulus q , a lattice dimension parameter g , and error distribution χ parametrized by λ . Also, choose some $h = \Theta(g \log q)$. Let $params = (g, h, q, \chi)$ be parameters for $\text{LWE}_{g,h,q,\chi}$.
- $\text{SecretKeyGen}(params)$: Sample $\bar{s} \leftarrow \mathbb{Z}_q^g$. Output $sk = s^\top = [-\bar{s}^\top \mid 1]$.
- $\text{PublicKeyGen}(params, sk)$: Sample a matrix $\bar{\mathbf{A}} \leftarrow \mathbb{Z}_q^{g \times h}$ uniformly and an error vector $e \leftarrow \chi^h$. Let $\alpha = \bar{s}^\top \cdot \bar{\mathbf{A}} + e^\top$. Set $pk = \mathbf{A}$ to be the $(g+1)$ -row matrix consisting of g rows of $\bar{\mathbf{A}}$ followed by α .
- $\text{Enc}(params, pk, \mu)$: Sample a uniform matrix $\mathbf{R} \leftarrow \{0, 1\}^{h \times h}$. For $pk = \mathbf{A}$, output $\mathbf{A} \cdot \mathbf{R} + \mu \cdot \mathbf{G}$ where \mathbf{G} is a gadget matrix.

We assume that our set of LWE parameters $params$ is (\mathcal{T}, ϵ) -secure, such that for any $\mathcal{T}(\lambda)$ time adversary, the advantage of distinguishing an LWE sample from a uniformly random sample is at most $\epsilon(\lambda)$. We verify in Appendix E that GSW satisfies the additional properties required by PFHE.

Efficiency Based on Subexponential LWE. We choose the lattice dimension parameter $g = \lambda$, the modulus q and the noise σ with subexponential modulus-to-noise ratio, and

the sample complexity $h \in \Theta(g \cdot \log q)$. The subexponential hardness of LWE requires the below indistinguishability to hold for adversaries of size $\mathcal{T}(\lambda) = 2^{\lambda^\beta}$ for some constant $\beta > 0$, with subexponentially small distinguishing advantage $\epsilon(\lambda) = 2^{-\lambda^c}$ for some constant $c > 0$ denoted by \approx_ϵ :

$$\left\{ \begin{array}{l} (\mathbf{A}, s \cdot \mathbf{A} + e) \text{ mod } q : \\ \mathbf{A} \leftarrow \mathbb{Z}_q^{g \times h} \\ s \leftarrow \mathbb{Z}_q^{1 \times g} \\ e \leftarrow \chi_\sigma^{1 \times h} \end{array} \right\} \approx_\epsilon \left\{ (\mathbf{A}, \mathbf{u}) : \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{g \times h} \\ \mathbf{u} \leftarrow \mathbb{Z}_q^{1 \times g} \end{array} \right\}$$

We set the security parameter be $\lambda = \log^r n$ for some constant $r > 1/\beta$ and $r > 1/c$. This satisfies the relationship between functions required by the Robustness Theorem 3.1: $\mathcal{T}(\log^r n) = 2^{(\log n)^{r\beta}} = \Omega(\text{poly } n)$ and $\epsilon(\log^r n) = 2^{-(\log n)^{rc}} \leq \text{negl}(n)$ for some negligible functions negl .

Given these parameters, using GSW to implement PFHE and Theorem 4.1 to implement StreamAlg_0 , we get the claimed result.

Corollary 4.1. *Assuming the subexponential hardness of LWE, the StreamAlg scheme in Construction 3.1 has the following complexity:*

- StreamAlg takes $\tilde{O}(k)$ bits of space.
- StreamAlg .Update runs in time $\tilde{O}(1)$.
- StreamAlg .Report runs in time $\tilde{O}(k^{1+c})$ for an arbitrarily small constant $c > 0$.

In Appendix E, we prove the above corollary and show the complexity of our algorithm under different implementations of StreamAlg_0 and hardness assumptions.

Impact statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgements

David P. Woodruff was supported by a Simons Investigator Award and NSF Grant No. CCF-2335412.

References

- Aharonov, D. and Regev, O. Lattice problems in np comp. *J. ACM*, 52:749–765, 2005.
- Ajtai, M., Braverman, V., Jayram, T., Silwal, S., Sun,

- A., Woodruff, D. P., and Zhou, S. The white-box adversarial data stream model. In *Proceedings of the 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, PODS '22, pp. 15–27, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392600. doi: 10.1145/3517804.3526228. URL <https://doi.org/10.1145/3517804.3526228>.
- Almeida, A., Brás, S., and Sargento, S., e. a. Time series big data: a survey on data stream frameworks, analysis and algorithms. *J Big Data*, 2023. doi: <https://doi.org/10.1186/s40537-023-00760-1>.
- Alon, N., Matias, Y., and Szegedy, M. The space complexity of approximating the frequency moments. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pp. 20–29, New York, NY, USA, 1996. Association for Computing Machinery. ISBN 0897917855. doi: 10.1145/237814.237823. URL <https://doi.org/10.1145/237814.237823>.
- Alon, N., Ben-Eliezer, O., Dagan, Y., Moran, S., Naor, M., and Yogev, E. Adversarial laws of large numbers and optimal regret in online classification. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 447–455, 2021.
- Attias, I., Cohen, E., Shechner, M., and Stemmer, U. A framework for adversarial streaming via differential privacy and difference estimators. *CoRR*, abs/2107.14527, 2021.
- Ben-Eliezer, O. and Yogev, E. The adversarial robustness of sampling. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, PODS, pp. 49–62, 2020.
- Ben-Eliezer, O., Jayaram, R., Woodruff, D. P., and Yogev, E. A framework for adversarially robust streaming algorithms. *SIGMOD Rec.*, 50(1):6–13, 2021.
- Ben-Eliezer, O., Eden, T., and Onak, K. Adversarially robust streaming via dense-sparse trade-offs. In *5th Symposium on Simplicity in Algorithms*, SOSA, 2022. (to appear).
- Berinde, R., Indyk, P., and Ruzic, M. Practical near-optimal sparse recovery in the ℓ_1 norm. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pp. 198–205, 2008. doi: 10.1109/ALLERTON.2008.4797556.
- Boyle, E., LaVigne, R., and Vaikuntanathan, V. Adversarially robust property-preserving hash functions. In Blum, A. (ed.), *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPICs*, pp. 16:1–16:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- Brakerski, Z. and Vaikuntanathan, V. Fully homomorphic encryption from ring-lwe and security for key dependent messages. In Rogaway, P. (ed.), *Advances in Cryptology – CRYPTO 2011*, pp. 505–524, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. ISBN 978-3-642-22792-9.
- Brakerski, Z., Gentry, C., and Vaikuntanathan, V. (leveled) fully homomorphic encryption without bootstrapping. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pp. 309–325, 2012.
- Braverman, V., Hassidim, A., Matias, Y., Schain, M., Silwal, S., and Zhou, S. Adversarial robustness of streaming algorithms through importance sampling. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems, NeurIPS*, 2021. (to appear).
- Brown, P., Haas, P. J., Myllymaki, J., Pirahesh, H., Reinwald, B., and Sismanis, Y. Toward automated large-scale information integration and discovery. In *Data Management in a Connected World, Essays Dedicated to Hartmut Wedekind on the Occasion of His 70th Birthday*, pp. 161–180, 2005.
- Chakrabarti, A., Ghosh, P., and Stoeckl, M. Adversarially robust coloring for graph streams. In *13th Innovations in Theoretical Computer Science Conference, ITCS*, 2022. (to appear).
- Chan, T. M. A dynamic data structure for 3-d convex hulls and 2-d nearest neighbor queries. *J. ACM*, 57(3):16:1–16:15, 2010.
- Chan, T. M. and He, Q. More dynamic data structures for geometric set cover with sublinear update time. In *37th International Symposium on Computational Geometry, SoCG*, pp. 25:1–25:14, 2021.
- Chandrasekaran, S., Cooper, O., Deshpande, A., Franklin, M. J., Hellerstein, J. M., Hong, W., Krishnamurthy, S., Madden, S. R., Reiss, F., and Shah, M. A. Telegraphcq: Continuous dataflow processing. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data*, SIGMOD '03, pp. 668, New York, NY, USA, 2003. Association for Computing Machinery. ISBN 158113634X. doi: 10.1145/872757.872857. URL <https://doi.org/10.1145/872757.872857>.
- Cormode, G. and Garofalakis, M. Streaming in a connected world: Querying and tracking distributed data streams. In *Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology*, EDBT '08, pp. 745,

- New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781595939265. doi: 10.1145/1353343.1353442. URL <https://doi.org/10.1145/1353343.1353442>.
- Cormode, G. and Jowhari, H. Lp samplers and their applications: A survey. *ACM Comput. Surv.*, 52(1), feb 2019. ISSN 0360-0300. doi: 10.1145/3297715. URL <https://doi.org/10.1145/3297715>.
- Cormode, G. and Muthukrishnan, S. What’s new: finding significant differences in network data streams. *IEEE/ACM Transactions on Networking*, 13(6):1219–1232, 2005. doi: 10.1109/TNET.2005.860096.
- Cubuk, E. D., Zoph, B., Mané, D., Vasudevan, V., and Le, Q. V. Autoaugment: Learning augmentation policies from data. *CoRR*, abs/1805.09501, 2018.
- Dasu, T., Johnson, T., Muthukrishnan, S., and Shkapenyuk, V. Mining database structure; or, how to build a data quality browser. In *Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data*, pp. 240–251, 2002.
- Driscoll, J. R., Sarnak, N., Sleator, D. D., and Tarjan, R. E. Making data structures persistent. *J. Comput. Syst. Sci.*, 38(1):86–124, 1989.
- Feng, Y. and Woodruff, D. P. Improved algorithms for white-box adversarial streams. In *ICML*, 2023.
- Fiat, A. and Kaplan, H. Making data structures confluent persistent. *J. Algorithms*, 48(1):16–58, 2003.
- Flajolet, P. and Nigel Martin, G. Probabilistic counting algorithms for data base applications. *Journal of Computer and System Sciences*, 31(2):182–209, 1985. ISSN 0022-0000. doi: [https://doi.org/10.1016/0022-0000\(85\)90041-8](https://doi.org/10.1016/0022-0000(85)90041-8). URL <https://www.sciencedirect.com/science/article/pii/0022000085900418>.
- Forbes, M. A. and Shpilka, A. On identity testing of tensors, low-rank recovery and compressed sensing. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC ’12, pp. 163–172, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450312455. doi: 10.1145/2213977.2213995. URL <https://doi.org/10.1145/2213977.2213995>.
- Gaber, M. M. *Data Stream Processing in Sensor Networks*, pp. 41–48. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. ISBN 978-3-540-73679-0. doi: 10.1007/3-540-73679-4_4. URL https://doi.org/10.1007/3-540-73679-4_4.
- Gentry, C., Sahai, A., and Waters, B. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Canetti, R. and Garay, J. A. (eds.), *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pp. 75–92. Springer, Heidelberg, August 2013. doi: 10.1007/978-3-642-40041-4_5.
- Gilbert, A. C., Li, Y., Porat, E., and Strauss, M. J. For-all sparse recovery in near-optimal time. *ACM Trans. Algorithms*, 13(3), mar 2017. ISSN 1549-6325. doi: 10.1145/3039872. URL <https://doi.org/10.1145/3039872>.
- Goldreich, O. and Goldwasser, S. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000. URL <https://doi.org/10.1006/jcss.1999.1686>.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *CoRR*, abs/1412.6572, 2014. URL <http://arxiv.org/abs/1412.6572>.
- Hassidim, A., Kaplan, H., Mansour, Y., Matias, Y., and Stemmer, U. Adversarially robust streaming algorithms via differential privacy. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems, NeurIPS*, 2020.
- Haviv, I. and Regev, O. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In Johnson, D. S. and Feige, U. (eds.), *39th ACM STOC*, pp. 469–477. ACM Press, June 2007. doi: 10.1145/1250790.1250859.
- Hu, Y., Qu, A., Wang, Y., and Work, D. B. Streaming data preprocessing via online tensor recovery for large environmental sensor networks. *ACM Trans. Knowl. Discov. Data*, 16(6), jul 2022. ISSN 1556-4681. doi: 10.1145/3532189. URL <https://doi.org/10.1145/3532189>.
- Ilyas, A., Engstrom, L., and Madry, A. Prior convictions: Black-box adversarial attacks with bandits and priors. *CoRR*, abs/1807.07978, 2018.
- Indyk, P. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *J. ACM*, 53(3):307–323, 2006.
- Jafarpour, S. *Deterministic Compressed Sensing*. PhD thesis, Princeton University, 2011.
- Kacham, P., Pagh, R., Thorup, M., and Woodruff, D. P. Pseudorandom hashing for space-bounded computation with applications in streaming, 2023.

- Kane, D. M., Nelson, J., and Woodruff, D. P. On the exact space complexity of sketching and streaming small norms. In Charikar, M. (ed.), *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pp. 1161–1178. SIAM, 2010.
- Kaplan, H. Persistent data structures. In *Handbook of Data Structures and Applications*. Chapman and Hall/CRC, 2004.
- Kaplan, H., Mansour, Y., Nissim, K., and Stemmer, U. Separating adaptive streaming from oblivious streaming using the bounded storage model. In *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO, Proceedings, Part III*, pp. 94–121, 2021.
- Khot, S. Hardness of approximating the shortest vector problem in lattices. In *45th FOCS*, pp. 126–135. IEEE Computer Society Press, October 2004. doi: 10.1109/FOCS.2004.31.
- Kurakin, A., Goodfellow, I. J., and Bengio, S. Adversarial machine learning at scale. In *5th International Conference on Learning Representations, ICLR, Conference Track Proceedings*, 2017.
- Lall, A., Sekar, V., Ogihara, M., Xu, J., and Zhang, H. Data streaming algorithms for estimating entropy of network traffic. *SIGMETRICS Perform. Eval. Rev.*, 34(1):145–156, jun 2006. ISSN 0163-5999. doi: 10.1145/1140103.1140295. URL <https://doi.org/10.1145/1140103.1140295>.
- Lightstone, S. Physical database design for relational databases. In *Encyclopedia of Database Systems, Second Edition*. Springer, 2018.
- Liu, Y., Chen, X., Liu, C., and Song, D. Delving into transferable adversarial examples and black-box attacks. In *5th International Conference on Learning Representations, ICLR, Conference Track Proceedings*, 2017.
- Lyubashevsky, V., Peikert, C., and Regev, O. On ideal lattices and learning with errors over rings. In Gilbert, H. (ed.), *EUROCRYPT 2010*, volume 6110 of *LNCS*, pp. 1–23. Springer, Heidelberg, May / June 2010. doi: 10.1007/978-3-642-13190-5_1.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations, ICLR, Conference Track Proceedings*, 2018.
- Menuhin, B. and Naor, M. Keep that card in mind: Card guessing with limited memory. *CoRR*, abs/2107.03885, 2021.
- Muthukrishnan, S. Data streams: Algorithms and applications. *Foundations and Trends® in Theoretical Computer Science*, 1(2):117–236, 2005. ISSN 1551-305X. doi: 10.1561/04000000002. URL <http://dx.doi.org/10.1561/04000000002>.
- Nisan, N. Pseudorandom generators for space-bounded computation. *Comb.*, 12(4):449–461, 1992.
- Phillips, J. M., Verbin, E., and Zhang, Q. Lower bounds for number-in-hand multiparty communication complexity, made easy, 2015.
- Roghani, M., Saberi, A., and Wajc, D. Beating the folklore algorithm for dynamic matching. In *13th Innovations in Theoretical Computer Science Conference, ITCS, 2022*. (to appear).
- Schmidt, L., Santurkar, S., Tsipras, D., Talwar, K., and Madry, A. Adversarially robust generalization requires more data. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS.*, pp. 5019–5031, 2018.
- Selinger, P. G., Astrahan, M. M., Chamberlin, D. D., Lorie, R. A., and Price, T. G. Access path selection in a relational database management system. In *Proceedings of the 1979 ACM SIGMOD International Conference on Management of Data*, pp. 23–34. ACM, 1979.
- Shoup, V. Fast construction of irreducible polynomials over finite fields. *Journal of Symbolic Computation*, 17(5):371–391, 1994. ISSN 0747-7171. doi: <https://doi.org/10.1006/jsc.1994.1025>. URL <https://www.sciencedirect.com/science/article/pii/S074771718471025X>.
- Shukla, A., Deshpande, P., Naughton, J. F., and Ramasamy, K. Storage estimation for multidimensional aggregates in the presence of hierarchies. In *VLDB’96, Proceedings of 22th International Conference on Very Large Data Bases*, pp. 522–531, 1996.
- Smart, N. and Vercauteren, F. Fully homomorphic SIMD operations. *Cryptology ePrint Archive*, Report 2011/133, 2011. <https://eprint.iacr.org/2011/133>.
- Sung, M., Kumar, A., Li, L., Wang, J., and Xu, J. Scalable and efficient data streaming algorithms for detecting common content in internet traffic. In *22nd International Conference on Data Engineering Workshops (ICDEW’06)*, pp. 27–27, 2006. doi: 10.1109/ICDEW.2006.130.
- Szalay, A., Gray, J., and VandenBerg, J. Petabyte scale data mining: Dream or reality? *Proceedings of SPIE - The International Society for Optical Engineering*, 08 2002. doi: 10.1117/12.461427.

- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. Intriguing properties of neural networks. *International Conference on Learning Representations*, 2014.
- Tramèr, F., Kurakin, A., Papernot, N., Goodfellow, I. J., Boneh, D., and McDaniel, P. D. Ensemble adversarial training: Attacks and defenses. In *6th International Conference on Learning Representations, ICLR, Conference Track Proceedings*, 2018.
- Venkataraman, S., Song, D. X., Gibbons, P. B., and Blum, A. New streaming algorithms for fast detection of superspreaders. In *Network and Distributed System Security Symposium*, 2005. URL <https://api.semanticscholar.org/CorpusID:538509>.
- von zur Gathen, J. and Gerhard, J. *Fast polynomial evaluation and interpolation*. Cambridge University Press, 3 edition, 2013. doi: 10.1017/CBO9781139856065.013.
- Wajc, D. Rounding dynamic matchings against an adaptive adversary. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC*, pp. 194–207, 2020.
- Woodruff, D. P. and Zhang, Q. Tight bounds for distributed functional monitoring. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing, STOC '12*, pp. 941–960, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450312455. doi: 10.1145/2213977.2214063. URL <https://doi.org/10.1145/2213977.2214063>.
- Woodruff, D. P. and Zhou, S. Tight bounds for adversarially robust streams and sliding windows via difference estimators. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pp. 1183–1196, 2021.
- Zhao, Q. G., Kumar, A., Wang, J., and Xu, J. J. Data streaming algorithms for accurate and efficient measurement of traffic and flow matrices. *SIGMETRICS Perform. Eval. Rev.*, 33(1):350–361, jun 2005. ISSN 0163-5999. doi: 10.1145/1071690.1064258. URL <https://doi.org/10.1145/1071690.1064258>.

Appendix A contains additional preliminaries for lattices. Appendix B contains the compactness, correctness, and security properties of FHE. In Appendix C, we recall the streaming model, the structure of a streaming algorithm, and the formal definition of white-box adversarial robustness.

We prove the robustness of our streaming K-Sparse Recovery algorithm in Appendix D and show the efficiency of the algorithm in Appendix E. Finally, in Appendix F and G, we describe our results for distributed protocols and low-rank matrix and tensor recovery, respectively.

A. Additional Lattice Preliminaries

In this section we introduce more background on the lattice problems, which helps justify our hardness assumptions based on LWE and Ring-LWE.

General definitions. A *lattice* \mathcal{L} is a discrete subgroup of \mathbb{R}^g , or equivalently the set $\mathcal{L}(\mathbf{b}_1, \dots, \mathbf{b}_m) = \{\sum_{i=1}^g x_i \mathbf{b}_i : x_i \in \mathbb{Z}\}$ of all integer combinations of g vectors $\mathbf{b}_1, \dots, \mathbf{b}_m \in \mathbb{R}^g$ that are linearly independent. Such \mathbf{b}_i 's form a *basis* of \mathcal{L} . The lattice \mathcal{L} is said to be *full-rank* if $g = m$. We denote by $\lambda_1(\mathcal{L})$ the so-called first minimum of \mathcal{L} , defined to be the length of a shortest non-zero vector of \mathcal{L} .

Discrete Gaussian and Related Distributions For any $s > 0$, define $\rho_s(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/s^2)$ for all $\mathbf{x} \in \mathbb{R}^g$. We write ρ for ρ_1 . For a discrete set S , we extend ρ to sets by $\rho_s(S) = \sum_{\mathbf{x} \in S} \rho_s(\mathbf{x})$. Given a lattice \mathcal{L} , the *discrete Gaussian* $\mathcal{G}_{\mathcal{L},s}$ is the distribution over \mathcal{L} such that the probability of a vector $\mathbf{y} \in \mathcal{L}$ is proportional to $\rho_s(\mathbf{y})$: $\Pr_{X \leftarrow \mathcal{G}_{\mathcal{L},s}}[X = \mathbf{y}] = \frac{\rho_s(\mathbf{y})}{\rho_s(\mathcal{L})}$.

The discrete Gaussian is not bounded but with overwhelming probability a sample from $\mathcal{G}_{\mathcal{L},s}$ is bounded by $s\sqrt{g}$ in ℓ_2 norm.

We now define a worst-case lattice problem, the gap version of the shortest vector problem, and discuss its relation to the average-case problems (LWE and Ring-LWE) that we used.

Definition A.1. For any function $\gamma = \gamma(g) \geq 1$, the decision problem GapSVP_γ (*Gap Shortest Vector Problem*) is defined as follows: the input is a basis \mathbf{B} for a lattice $\mathcal{L} \subset \mathbb{R}^g$ and a real number $d > 0$ that satisfy the promise that $\lambda_1(\mathcal{L}(\mathbf{B})) \leq d$, or $\lambda_1(\mathcal{L}(\mathbf{B})) \geq \gamma d$. The goal is to decide whether $\lambda_1(\mathcal{L}(\mathbf{B})) \leq d$, or $\lambda_1(\mathcal{L}(\mathbf{B})) \geq \gamma d$.

This problem is known to be NP hard for $\gamma = g^{o(1)}$ (Khot, 2004; Haviv & Regev, 2007), and is known to be inside $\text{NP} \cap \text{coNP}$ for $\gamma = O(\sqrt{g})$ (Aharonov & Regev, 2005; Goldreich & Goldwasser, 2000). The problem is believed to be intractable by polynomial time adversaries even for $\gamma = 2^{g^\epsilon}$ for $0 < \epsilon < 1$.

Under quantum and classical reductions it can be shown that $\text{LWE}_{g,h,q,\chi_\sigma}$ is harder to solve than Gap-SVP_γ with $\gamma = g^{\frac{q}{\sigma}}$. On the other hand, The Ring-LWE problem is as hard as worst-case problems over ideal lattices (Lyubashevsky et al., 2010).

B. Properties of Fully Homomorphic Encryption

We elaborate the compactness, correctness, and security properties that we requires for an FHE scheme.

Definition B.1 (FHE Compactness). We say that a FHE scheme is compact if there exists a polynomial $\text{poly}(\cdot)$ such that for all security parameters λ , circuits $C : \{0, 1\}^\ell \rightarrow \{0, 1\}$ of depth at most $\text{poly}(\lambda)$, and $\mu_i \in \{0, 1\}$ for $i \in [\ell]$, the following holds: Given $(\text{pk}, \text{sk}) \leftarrow \text{FHE.Setup}(1^\lambda)$, $\text{ct}_i \leftarrow \text{FHE.Enc}(\text{pk}, \mu_i)$ for $i \in [\ell]$, $\hat{\text{ct}} \leftarrow \text{FHE.Eval}(\text{pk}, C, \text{ct}_1, \dots, \text{ct}_\ell)$, we always have $|\hat{\text{ct}}| \leq \text{poly}(\lambda)$.

Definition B.2 (FHE Correctness). We say that a FHE scheme is correct if for all security parameters λ , circuit $C : \{0, 1\}^\ell \rightarrow \{0, 1\}$ of depth at most $\text{poly}(\lambda)$, and $\mu_i \in \{0, 1\}$ for $i \in [\ell]$, the following holds: Given $(\text{pk}, \text{sk}) \leftarrow \text{FHE.Setup}(1^\lambda)$, $\text{ct}_i \leftarrow \text{FHE.Enc}(\text{pk}, \mu_i)$ for $i \in [\ell]$, $\hat{\text{ct}} \leftarrow \text{FHE.Eval}(\text{pk}, C, \text{ct}_1, \dots, \text{ct}_\ell)$,

$$\Pr[\text{FHE.Dec}(\text{pk}, \text{sk}, \hat{\text{ct}}) = f(\mu_1, \dots, \mu_\ell)] = 1.$$

Definition B.3 (FHE Security). We say that a FHE scheme satisfies security if for all security parameter λ , the following holds: For any PPT_λ adversary \mathcal{A} , there exists a negligible function negl such that

$$\text{adv}_{\mathcal{A}, \text{FHE}} := 2 \cdot |1/2 - \Pr[\text{Expt}_{\mathcal{A}, \text{FHE}}(1^\lambda) = 1]| < \text{negl}(\lambda),$$

where the experiment $\text{Expt}_{\mathcal{A}, \text{FHE}}$ is defined as follows:

$\text{Expt}_{\mathcal{A}, \text{FHE}}(1^\lambda) :$

1. On input the security parameter 1^λ , the challenger runs $(\text{pk}, \text{sk}) \leftarrow \text{FHE.Setup}(1^\lambda)$, and $\text{ct} \leftarrow \text{FHE.Enc}(\text{pk}, b)$ for $b \xleftarrow{\$} \{0, 1\}$. It provides (pk, ct) to \mathcal{A} .
2. \mathcal{A} outputs a guess b' . The experiment outputs 1 if and only if $b = b'$.

Remark B.1. The security definition above is in terms of PPT_λ adversaries and negl advantage, but it can be lifted to a (\mathcal{T}, ϵ) setting similar to Remark 2.1 as follows: We say that a FHE scheme is (\mathcal{T}, ϵ) -secure if there exists an integer λ_0 , such that for all $\lambda > \lambda_0$, for any non-uniform probabilistic $\mathcal{T}(\lambda)$ time adversary, the advantage $\text{adv}_{\mathcal{A}, \text{FHE}}$ of the above experiment is at most $\epsilon(\lambda)$.

C. Streaming Model

Notation. Let $\mathbb{1}(p)$ for some proposition p be an indicator variable, which equals 1 if p is true and 0 otherwise. We denote the input dimension parameter of an algorithm by n . We write PPT_λ to denote probabilistic poly-time in λ .

Streaming Model. The streaming model captures key resource requirements of algorithms for database, machine learning, and network tasks, where the size of the data is significantly larger than the available storage. This model was formalized by Alon, Matias, and Szegedy (Alon et al., 1996) as a vector undergoing additive *updates* to its coordinates. In this model, the input data takes the form of a sequence of updates. These updates assume the existence of a fixed-length underlying vector which represents the current dataset, and dynamically modify this vector one coordinate at a time.

Given a dimension parameter $n \in \mathbb{N}$, the model behaves as follows:

- At the beginning of the stream, an underlying vector \mathbf{x} is initialized to 0^n .
- The input stream is a sequence of $T \leq \text{poly}(n)$ additive updates (x_1, \dots, x_T) . Each update x_t for $t \in [T]$ is a tuple $x_t := (i_t, \Delta_t)$, where $i_t \in [n]$ denotes a coordinate of \mathbf{x} and Δ_t is an arbitrary value.
 $x_t := (i_t, \Delta_t)$ is interpreted as updating $\mathbf{x}_{i_t} \leftarrow \mathbf{x}_{i_t} + \Delta_t$. Thus at any time $t' \in [T]$ during the stream, \mathbf{x} is the accumulation of all past updates, with its i -th coordinate for $i \in [n]$: $\mathbf{x}_i = \sum_{t \leq t'} \mathbb{1}(i_t = i_{t_0}) \cdot \Delta_t$.

A streaming algorithm receives the stream of updates and uses limited memory to compute a function of \mathbf{x} . We stress that during the stream, the underlying vector \mathbf{x} is not explicitly available to the algorithm, but is instead implicitly defined by the sequence of past updates.

In this work, we consider a bounded integer stream, which assumes that throughout the stream, \mathbf{x} is promised to be in the range $\{-N, -N+1, \dots, N-1, N\}^n$ for some integer $N \in \text{poly}(n)$.

Streaming Algorithm. A streaming algorithm maintains some internal state during the stream, which evolves upon receiving new updates. It takes one pass over all updates and then uses its maintained state to compute a function of the underlying dataset \mathbf{x} . In the streaming model, the dimension parameter n (the length of \mathbf{x}) is typically large. Therefore, a streaming algorithm usually optimizes the size of the state so that it takes a sublinear number of bits of space in n to represent.

We define the general framework for streaming algorithms as follows:

Definition C.1 (Streaming Algorithm). *Given a query function ensemble $\mathcal{F} = \{f_n : \mathbb{Z}^n \rightarrow X\}_{n \in \mathbb{N}}$ for some domain X , a streaming algorithm StreamAlg that computes \mathcal{F} contains a tuple of operations $\text{StreamAlg} = (\text{StreamAlg.Setup}, \text{StreamAlg.Update}, \text{StreamAlg.Report})$, with the following properties:*

- $\text{StreamAlg.Setup}(n) :$ On input a dimension parameter n , the setup operation initializes an internal data structure \mathcal{DS} within the memory of StreamAlg .
- $\text{StreamAlg.Update}(x_t) :$ On input an additive update $x_t := (i_t, \Delta_t)$, the update operation modifies \mathcal{DS} .
- $\text{StreamAlg.Report}() \rightarrow r :$ The report operation uses \mathcal{DS} to compute a query response $r \in X$.

A response r is said to be correct if $r = f_n(\mathbf{x})$, where $\mathbf{x} \in \mathbb{Z}^n$ is the accumulation of all past updates such that $\mathbf{x}_i = \sum_t \mathbb{1}(i_t = i) \cdot \Delta_t$ for $i \in [n]$.

In this work, we only consider exact computations of query functions ($r = f_n(\mathbf{x})$), though the above definition can also be generalized to approximation problems ($r \approx f_n(\mathbf{x})$).

White-box Adversarial Stream. In the conventional *oblivious* stream model, the sequence of updates may be chosen adversarially, but is assumed to be chosen independently of the randomness of a streaming algorithm `StreamAlg`. In contrast, in the white-box adversarial streaming model, a sequence of stream updates (x_1, \dots, x_T) is adaptively generated by an adversary who sees the *full internal state* of `StreamAlg` at all times, including \mathcal{DS} and any other parameters and randomness used by `StreamAlg`. We say that a streaming algorithm is *white-box adversarially robust* if it guarantees to output correct responses even against such powerful adversaries. This is described as the following game:

Definition C.2 (White-Box Adversarially Robustness). *We say a streaming algorithm `StreamAlg` is white-box adversarially robust (WAR) if for all dimension parameters $n \in \mathbb{N}$, the following holds: For any PPT adversary \mathcal{A} , the following experiment $\text{Expt}_{\mathcal{A}, \text{StreamAlg}}(n)$ outputs 1 with probability at most $\text{negl}(n)$:*

$\text{Expt}_{\mathcal{A}, \text{StreamAlg}}(n)$:

1. The challenger and \mathcal{A} agree on a query function ensemble $\mathcal{F} = \{f_n : \mathbb{Z}^n \rightarrow X\}_{n \in \mathbb{N}}$ for some domain X . On input a dimension parameter n , let $f_n \in \mathcal{F}$ be the query function of the experiment.
2. The challenger runs `StreamAlg.Setup`(n). It then provides all internal states of `StreamAlg` to \mathcal{A} , including \mathcal{DS} and any other parameters and randomness previously used by `StreamAlg`.
3. For some $T \in \text{poly}(n)$, at each time $t \in [T]$:
 - \mathcal{A} outputs an adaptive update $x_t := (i_t, \Delta_t)$ for $t \in [T]$. Also, \mathcal{A} may issue a query.
 - The challenger runs `StreamAlg.Update`(x_t), and outputs a response $r_t \leftarrow \text{StreamAlg.Report}()$ if it is queried. The challenger then provides all internal states of `StreamAlg` to \mathcal{A} , including \mathcal{DS} and any other parameters and randomness used by `StreamAlg`.
4. The experiment outputs 1 if and only if at some time $t \in [T]$ in the stream, `StreamAlg` output an incorrect response $r_t \neq f_n(\mathbf{x})$, where \mathbf{x} is the underlying data vector with $x_i = \sum_{j \in [t]} \mathbb{1}(i_j = i) \cdot \Delta_j$ for $i \in [n]$.

D. Proof of White-box Robustness of Construction 3.1

For the concreteness of our proof, we assume that the relaxed k -sparse recovery algorithm `StreamAlg0` that is used in our construction is deterministic and runs in polynomial time. Both conditions are satisfied by many existing sparse recovery algorithms.

Also, recall that in the FHE preliminaries section, we define the (\mathcal{T}, ϵ) -security of a FHE scheme, where \mathcal{T} and ϵ are functions in λ . Also, in Construction 3.1, we set the PFHE security parameter $\lambda = \mathcal{S}(n)$ for some function \mathcal{S} .

In this section, we will show that if for all $n \in \mathbb{N}$, functions \mathcal{T} , ϵ , and \mathcal{S} satisfy:

- $\mathcal{T}(\mathcal{S}(n)) = \Omega(\text{poly } n)$
- $\epsilon(\mathcal{S}(n)) \leq \text{negl}(n)$ for some negligible functions negl .

then the `StreamAlg` scheme in Construction 3.1 is a WAR streaming algorithm that solves the k -Sparse Recovery problem.

Remark D.1. Roughly, we need $\text{poly}(n) \subseteq \mathcal{T}(\mathcal{S}(n))$ to convert between the runtime of a PFHE adversary and a `StreamAlg` adversary. We also need $\epsilon(\mathcal{S}(n)) \leq \text{negl}(n)$ to convert between the advantage of a PFHE adversary and a `StreamAlg` adversary.

Proof (of Theorem 3.1). We first show that at any time during the stream, if the underlying vector \mathbf{x} is k -sparse, then the query response of `StreamAlg` is correct.

Lemma D.1. *Suppose PFHE satisfies Definition 2.4 and `StreamAlg0` solves the Relaxed k -Sparse Recovery problem. Then the following holds: at any time $t \in [T]$, if the current underlying vector \mathbf{x} satisfies $\|\mathbf{x}\|_0 \leq k$, then `StreamAlg.Report` correctly outputs \mathbf{x} .*

Proof. By the correctness of the StreamAlg_0 scheme, when $\|\mathbf{x}\|_0 \leq k$, StreamAlg_0 .Report outputs $\mathbf{x}' = \mathbf{x}$. Since Definition 2.4 restricts PFHE.Eval to be deterministic, we essentially have $\text{hash} = \sum_{i \in [n]} \widehat{ct}_i \cdot \mathbf{x}'_i$ and $\text{sketch} = \sum_{i \in [n]} \widehat{ct}_i \cdot \mathbf{x}_i$ for the same set of ciphertexts \widehat{ct}_i . Thus $\text{hash} = \text{sketch}$ follows from $\mathbf{x}' = \mathbf{x}$. In this case, StreamAlg.Report correctly outputs \mathbf{x} as desired. \square

It remains to show the correctness of responses in the case when the underlying data vector is denser than k . Our proof proceeds via a sequence of hybrid experiments between an adversary \mathcal{A} and a challenger.

- **Hybrid₀**: This is a real WAR experiment $\text{Expt}_{\mathcal{A}, \text{StreamAlg}}(n)$ from Definition C.2, where the challenger runs the StreamAlg scheme from Construction 3.1. On input a dimension parameter n , the challenger samples $\widetilde{\text{pk}} \leftarrow \mathcal{D}_{\widetilde{\text{pk}}(\lambda)}$ and $\widetilde{\text{ct}}_1, \widetilde{\text{ct}}_2, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil} \leftarrow \mathcal{D}_{\widetilde{\text{ct}}(\lambda)}$ from the distributions of pseudo-public key and pseudo-ciphertext (Definition 2.4) and provides them to \mathcal{A} .

\mathcal{A} outputs $T \in \text{poly}(n)$ -many adaptive updates $x_t := (i_t, \Delta_t)$ upon seeing past internal states of StreamAlg . At any time $t \in [T]$, these updates implicitly define an underlying vector \mathbf{x} , whose i -th coordinate $\mathbf{x}_i = \sum_t \mathbb{1}(i_t = i) \cdot \Delta_t$ for $i \in [n]$.

- **Hybrid₁**: This is the same experiment as **Hybrid₀**, except that \mathcal{A} now runs some additional extraction on itself:
 - At each time $t \in [T]$, \mathcal{A} computes the underlying vector \mathbf{x} defined by all of its past updates.
 - \mathcal{A} runs an additional instance of the StreamAlg_0 scheme from Definition 3.2 (which solves the Relaxed k -Sparse Recovery problem):
 1. Before \mathcal{A} generates its first update, it initializes its StreamAlg_0 instance to be in the same state as the StreamAlg_0 scheme run by the challenger. (i.e., \mathcal{A} runs $\text{StreamAlg}_0.\text{Setup}(n)$ using the same set of randomness used by the challenger.)
 2. At each time $t \in [T]$, after \mathcal{A} generates an update x_t , it calls $\text{StreamAlg}_0.\text{Update}(x_t)$ and then $\text{StreamAlg}_0.\text{Report}()$ to get a query response \mathbf{x}' .

We note that if \mathcal{A} runs in $\text{poly}(n)$ time in **Hybrid₀**, then it still runs in $\text{poly}(n)$ time in **Hybrid₁** since computing \mathbf{x} is trivial and StreamAlg_0 runs in poly -time.

For an adversary \mathcal{A} , we write $\text{Hyb}_i(\mathcal{A})$ to denote the output of **Hybrid_i**.

Lemma D.2. For any adversaries \mathcal{A} , $|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = 0$.

Proof. This just follows from the fact that the additional extractions run by \mathcal{A} do not affect the output of the experiments. \square

- **Hybrid₂**: Same as **Hybrid₁**, except that at the beginning of the experiment, the challenger samples a random coordinate $m \xleftarrow{\$} [n]$ which will be hidden from \mathcal{A} . The experiment outputs 1 if at some time $t \in [T]$ both of the following holds:
 1. StreamAlg outputs an incorrect response $r_t \neq f_n(\mathbf{x})$ (i.e., the same condition for $\text{Hyb}_1(\mathcal{A}) = 1$),
 2. The m -th coordinates of \mathbf{x} and \mathbf{x}' extracted by \mathcal{A} satisfy $\mathbf{x}_m \neq \mathbf{x}'_m$,

We show that with the additional success condition, the success probability of **Hybrid₂** drops by at most a $1/n$ multiplicative factor compared to **Hybrid₁**.

Lemma D.3. For any adversaries \mathcal{A} , $\Pr[\text{Hyb}_2(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_1(\mathcal{A}) = 1]/n$.

Before proving Lemma D.3, we make the following observation:

Claim D.1. At any time $t \in [T]$ during the **Hybrid₂** experiment, if StreamAlg outputs an incorrect response $r_t \neq f_n(\mathbf{x})$, and \mathbf{x} and \mathbf{x}' are the vectors extracted by \mathcal{A} at time t , then $\mathbf{x}' \neq \mathbf{x}$.

Proof. By Lemma D.1, $r_t \neq f_n(\mathbf{x})$ can occur only when $\|\mathbf{x}\|_0 > k$ and $r_t \neq \perp$. In this case, by the construction of StreamAlg.Report , r_t should be a k -sparse vector, and thus different from \mathbf{x} . Also, using deterministic $\text{StreamAlg}_0.\text{Update}$ and $\text{StreamAlg}_0.\text{Report}$, \mathcal{A} exactly extracts $\mathbf{x}' = r_t$, so $\mathbf{x}' \neq \mathbf{x}$. \square

Lemma D.3 then follows from Claim D.1:

Proof (of Lemma D.3). Claim D.1 states that $\mathbf{x}' \neq \mathbf{x}$, and in particular, \mathbf{x}' and \mathbf{x} disagree on at least one coordinate. Recall that the random coordinate m is hidden from \mathcal{A} at all times, so with probability at least $1/n$, m collides with a coordinate where \mathbf{x}' and \mathbf{x} disagree. In this case, both conditions in Hybrid_2 are satisfied, so $\Pr[\text{Hyb}_2(\mathcal{A}) = 1 \mid \text{Hyb}_1(\mathcal{A}) = 1] \geq 1/n$. \square

- **Hybrid₃**: Same as **Hybrid₂**, except that the challenger now generates an actual public key $\text{pk} : (\text{pk}, \text{sk}) \leftarrow \text{PFHE.Setup}(1^\lambda)$ to replace $\widetilde{\text{pk}}$.
- **Hybrid₄**: Same as **Hybrid₃**, except that the challenger now uses the random coordinate $m \stackrel{\$}{\leftarrow} [n]$ (whose bits are denoted as $m_1 m_2 \cdots m_{\lceil \log n \rceil}$) that was sampled privately as the message in the ciphertexts: It generates $\lceil \log n \rceil$ actual ciphertexts $\text{ct}_i \leftarrow \text{PFHE.Enc}(\text{pk}, m_i)$ for $i \in [\lceil \log n \rceil]$ to replace $(\widetilde{\text{ct}}_1, \widetilde{\text{ct}}_2, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$. i.e., $(\text{ct}_1, \text{ct}_2, \dots, \text{ct}_{\lceil \log n \rceil})$ encrypts the bit representation of m .

Lemma D.4. *Given that PFHE satisfies Definition 2.4, for any PPT_n adversaries \mathcal{A} , $|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| \leq \text{negl}(n)$ for some negligible function negl .*

Proof. The only difference in \mathcal{A} 's views in **Hybrid₂** and **Hybrid₃** is the way the challenger generates $\widetilde{\text{pk}}$ or pk . In **Hybrid₂**, the challenger samples $\widetilde{\text{pk}} \leftarrow \mathcal{D}_{\widetilde{\text{pk}}(\lambda)}$; and in **Hybrid₃**, the challenger runs $\text{pk} : (\text{pk}, \text{sk}) \leftarrow \text{PFHE.Setup}(1^\lambda)$. Thus \mathcal{A} distinguishing these two hybrids will distinguish the two distributions $\{\text{pk} : (\text{pk}, \text{sk}) \leftarrow \text{FHE.Setup}(1^\lambda)\}$ and $\mathcal{D}_{\widetilde{\text{pk}}(\lambda)}$ with the same advantage. Also by the (\mathcal{T}, ϵ) -indistinguishability in Definition 2.4, for any $\text{poly}(n) \leq \mathcal{T}(\lambda)$ -sized adversaries, the latter two distributions can be distinguished with probability at most $\epsilon(\lambda) \leq \text{negl}(n)$. \square

Lemma D.5. *Given that PFHE satisfies Definition 2.4, for any PPT_n adversaries \mathcal{A} , $|\Pr[\text{Hyb}_3(\mathcal{A}) = 1] - \Pr[\text{Hyb}_4(\mathcal{A}) = 1]| \leq \text{negl}(n)$ for some negligible function negl .*

Proof. Similar to Lemma D.4, the only difference in \mathcal{A} 's views in **Hybrid₃** and **Hybrid₄** is the way the challenger generates $(\widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ or $(\text{ct}_1, \dots, \text{ct}_{\lceil \log n \rceil})$. Recall that we have: $\{\text{ct}\} \approx_{(\mathcal{T}, \epsilon)} \mathcal{D}_{\widetilde{\text{ct}}(\lambda)}$ from Definition 2.4. Therefore for any PPT_n adversary \mathcal{A} , the advantage of distinguishing $(\widetilde{\text{ct}}_1, \widetilde{\text{ct}}_2, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ in **Hybrid₃** from a ciphertext sequence $(\text{ct}_1, \text{ct}_2, \dots, \text{ct}_{\lceil \log n \rceil})$ where each $\text{ct}_i \leftarrow \text{PFHE.Enc}(\text{pk}, m_i)$ is at most $\lceil \log n \rceil \cdot \epsilon(\lambda) = \text{negl}(n)$. Thus \mathcal{A} cannot distinguish between **Hybrid₃** and **Hybrid₄** with non-negligible advantage, as otherwise either the pseudorandom-ct property in Definition 2.4 would be broken. \square

Now that we have related the success probability between the above consecutive hybrid experiments, we will show next that $\Pr[\text{Hyb}_4(\mathcal{A}) = 1] = 0$.

Lemma D.6. *Given that PFHE satisfies Definition 2.4, for any PPT_n adversaries \mathcal{A} , $\Pr[\text{Hyb}_4(\mathcal{A}) = 1] = 0$.*

Proof. As in the proof for Claim D.1, if $\text{Hyb}_4(\mathcal{A}) = 1$, at some time t we must have $r_t = \mathbf{x}' \neq \mathbf{x}$. Additionally, we have:

- $\mathbf{x}'_m \neq \mathbf{x}_m$, which is among the condition of $\text{Hyb}_4(\mathcal{A}) = 1$, and
- $\text{hash} = \text{sketch}$, which is verified in StreamAlg.Report .

Recall the StreamAlg scheme computes $\text{hash} := \sum_{i \in [n]} \widehat{\text{ct}}_i \cdot \mathbf{x}'_i$ and $\text{sketch} := \sum_{i \in [n]} \widehat{\text{ct}}_i \cdot \mathbf{x}_i$, where $\widehat{\text{ct}}_i \leftarrow \text{PFHE.Eval}(\text{pk}, C_i, \text{ct}_1, \dots, \text{ct}_{\lceil \log n \rceil})$, for $C_i(\mu) = 1$ iff $\mu = i$. Therefore, given that $(\text{ct}_1, \dots, \text{ct}_{\lceil \log n \rceil})$ encrypts m in **Hybrid₄**, $\widehat{\text{ct}}_m$ should encode 1, and $\widehat{\text{ct}}_i$ encodes 0 for all $i \neq m$. Thus by the linear homomorphism property of PFHE,

there exists a deterministic decoding algorithm such that $hash$ decodes to x'_m and $sketch$ decodes to x_m . However, since $x'_m \neq x_m$, this essentially implies $hash \neq sketch$, contradicting the condition in StreamAlg.Report. □

Finally, summarizing Lemma D.2 to Lemma D.6:

1. $|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = 0$
2. $\Pr[\text{Hyb}_2(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_1(\mathcal{A}) = 1]/n$
3. $|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| = \text{negl}(n)$
4. $|\Pr[\text{Hyb}_3(\mathcal{A}) = 1] - \Pr[\text{Hyb}_4(\mathcal{A}) = 1]| = \text{negl}(n)$
5. $\Pr[\text{Hyb}_4(\mathcal{A}) = 1] = 0$

We have that for any adversaries \mathcal{A} , $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] \leq \text{negl}(n) \cdot n = \text{negl}(n)$. i.e., The real WAR game between \mathcal{A} and a challenger running the StreamAlg scheme from Construction 3.1 outputs 1 with probability at most $\text{negl}(n)$. □

E. Efficiency

Notation. We use b_i to denote the standard basis vector, whose i -th entry equals 1 and 0 anywhere else.

E.1. (Alternative) Deterministic Relaxed Sparse Recovery

The deterministic Relaxed k -Sparse Recovery algorithm in Theorem 4.1 achieves nearly optimal space and time. However, the measurement is “constructed” non-explicitly using the probabilistic method. For explicit construction, we consider the following result:

Theorem E.1 ((Jafarpour, 2011)). *There exists a streaming algorithm StreamAlg₀ for k -Sparse Recovery without Detection, such that given an input parameter n , for an integer stream with entries bounded by $\text{poly}(n)$:*

- StreamAlg₀.Setup, Update, and Report are all deterministic.
- StreamAlg₀ takes $\tilde{\mathcal{O}}(k)$ bits of space.
- StreamAlg₀.Update runs in (amortized) $\tilde{\mathcal{O}}(1)$.
- StreamAlg₀.Report runs in $\tilde{\mathcal{O}}(k^2)$.

The above theorem except the amortized runtime has been shown in (Jafarpour, 2011).

Proof (of the runtime in Theorem E.1). The measurement Φ in Theorem E.1 is a $2k \times n$ Vandermonde matrix explicitly defined² as $\Phi_{i,j} = j^{i-1}$ over \mathbb{Z}_p , where $p > 2M$ is a large enough prime for $M \in \text{poly}(n)$ being the bound on the magnitude of the stream entries. The reconstruction algorithm uses the same idea as the algebraic algorithm for decoding Reed-Solomon codes. It uses the sketch to construct an error-locator polynomial; the roots of this polynomial identify the signals appearing in the sparse superposition.

We note that the Vandermonde matrix is dense but structured. Thus to achieve the claimed $\tilde{\mathcal{O}}(1)$ update time, we buffer every $2k$ stream updates and then efficiently process them in one batch. We need the following result for fast multi-point evaluation of polynomials and the Transposition Principle:

Theorem E.2 ((von zur Gathen & Gerhard, 2013)). *Let R be a ring, and let $q \in R[x]$ be a degree- d polynomial. Then given at most d distinct points $x_1, \dots, x_d \in R$, all values $q(x_1), \dots, q(x_d)$ can be computed using $\mathcal{O}(d \log^2 d \log \log d)$ total operations over R .*

²For our discussion, we choose a Vandermonde matrix that is easy to describe. Though many other Vandermonde matrices also satisfy the specifications.

Theorem E.3 (Transposition Principle, (Shoup, 1994)). *Let K be a field and \mathbf{A} be an $r \times w$ matrix over K . Suppose there exists an arithmetic circuit of size s that can compute $\mathbf{A}\mathbf{y}$ for arbitrary length- w vector \mathbf{y} over K . Then there exists an $\mathcal{O}(s)$ -time algorithm that transforms this circuit to compute $\mathbf{A}^\top \mathbf{x}$ for arbitrary length- r vector \mathbf{x} over K , with the size of the new circuit in $\mathcal{O}(s)$.*

Theorem E.2 and Theorem E.3 together show the following update time for Vandermonde measurements:

Proposition E.1. *Suppose Φ is a $2k \times n$ Vandermonde matrix as described above. Let x_1, x_2, \dots, x_{2k} be a sequence of updates with $x_t := (i_t, \Delta_t)$ for each $t \in [2k]$. Define $\mathbf{v} = \sum_{t \in [2k]} \mathbf{b}_{i_t} \cdot \Delta_t$, then $\Phi(\mathbf{v})$ can be computed using $\tilde{\mathcal{O}}(k)$ time.*

Proof (of Proposition E.1). Without loss of generality, we assume all updated coordinates i_1, i_2, \dots, i_{2k} are distinct, as one can always compress multiple updates to the same coordinate as a single update. Then i_1, i_2, \dots, i_{2k} index $2k$ columns of Φ and $2k$ coordinates of \mathbf{v} (which are the only coordinates that can possibly be non-zero). Let \mathbf{A} be the indexed $2k \times 2k$ column sub-matrix and \mathbf{v}' be the $2k$ -length vector containing only i_1, i_2, \dots, i_{2k} -th entries of \mathbf{v} in their original order. We have $\Phi\mathbf{v} = \mathbf{A}\mathbf{v}'$.

Observe that the column submatrix \mathbf{A} is still a Vandermonde matrix, and each row i of \mathbf{A}^\top takes the form $[1, \alpha_i, \alpha_i^2, \dots, \alpha_i^{2k-1}]$ for some $\alpha_i \in [n]$. For an arbitrary vector $u \in \mathbb{Z}_p^{2k}$, we interpret it as the coefficient vector of a degree- $2k$ polynomial in \mathbb{Z}_p . Then $\mathbf{A}^\top \mathbf{u}$ can be computed as evaluating the polynomial at $2k$ points $\alpha_1, \alpha_2, \dots, \alpha_{2k}$, which takes $\tilde{\mathcal{O}}(k)$ time using the fast multi-point evaluation in Theorem 4.1. By the Transposition Principle, this implies asymptotically the same runtime for evaluating $\mathbf{A}\mathbf{v}'$ thus $\Phi\mathbf{v}$. \square

Therefore, once we buffer the updates into $2k$ -sized batches and amortize the cost of processing each batch to the future $2k$ updates, we get $\tilde{\mathcal{O}}(1)$ amortized runtime for $\text{StreamAlg}_0.\text{Update}$. \square

E.2. GSW Satisfies PFHE Properties

We check that GSW satisfies the additional properties required by PFHE:

- Eval is deterministic by the construction of GSW.
- (Pseudorandom pk, ct.) This is just the built-in security of GSW shown via the LWE hardness and the Left-over Hash lemma. Assuming the (\mathcal{T}, ϵ) -hardness of $\text{LWE}_{g,h,q,\chi}$, we have:

$$\left\{ \begin{array}{l} \mathbf{A} \leftarrow \text{PublicKeyGen} \\ (\mathbf{A}, \mathbf{A} \cdot \mathbf{R} + \mu \cdot \mathbf{G}) : \mathbf{R} \leftarrow \{0, 1\}^{h \times h} \\ \mu \in \{0, 1\} \end{array} \right\} \approx_{(\mathcal{T}, \epsilon)} \left\{ (\mathbf{U}_1, \mathbf{U}_2) : \begin{array}{l} \mathbf{U}_1 \leftarrow \mathbb{Z}_q^{(g+1) \times h} \\ \mathbf{U}_2 \leftarrow \mathbb{Z}_q^{(g+1) \times h} \end{array} \right\}.$$

- (\mathbb{Z} Linear Homomorphism.) Consider the following arbitrary linear combination, where each ciphertext $\text{ct}_i = \mathbf{A} \cdot \mathbf{R}_i + \mu_i \cdot \mathbf{G}$ for $i \in [k]$:

$$\mathcal{M} = \sum_{i \in [k]} x_i \cdot \text{ct}_i = \mathbf{A} \cdot \left(\sum_{i \in [k]} x_i \mathbf{R}_i \right) + \left(\sum_{i \in [k]} x_i \mu_i \right) \cdot \mathbf{G}$$

For $k \cdot \max_{i \in [k]} x_i \ll q$, it can be deterministically decrypted using the normal GSW decryption for integer messages. For completeness, we briefly describe LinearDec as follows:

- LinearDec(pk, sk, \mathcal{M}): Compute

$$\begin{aligned} \text{sk} \cdot \mathcal{M} &= [-\bar{\mathbf{s}}^\top \mid \mathbf{1}] \cdot \mathbf{A} \cdot \left(\sum_{i \in [k]} x_i \mathbf{R}_i \right) + \left(\sum_{i \in [k]} x_i \mu_i \right) \cdot [-\bar{\mathbf{s}}^\top \mid \mathbf{1}] \cdot \mathbf{G} \\ &= \mathbf{e}^\top \cdot \left(\sum_{i \in [k]} x_i \mathbf{R}_i \right) + \left(\sum_{i \in [k]} x_i \mu_i \right) \cdot [-\bar{\mathbf{s}}^\top \mid \mathbf{1}] \cdot \mathbf{G} \\ &\approx \left(\sum_{i \in [k]} x_i \mu_i \right) \cdot [-\bar{\mathbf{s}}^\top \mid \mathbf{1}] \cdot \mathbf{G} \end{aligned}$$

Since it is easy to solve LWE with respect to \mathbf{G} , we can recover $(\sum_{i \in [k]} x_i \mu_i) \cdot [-\bar{s}^\top \mid 1]$ and hence $\sum_{i \in [k]} x_i \mu_i$ (modulo q).

E.3. Efficiency of PFHE Based on Subexponential LWE

We use GSW to implement PFHE and set the security parameter to be $\lambda = \log^r n$ for some constant $r > 1/\beta$ and $r > 1/c$. We get the following proposition.

Proposition E.2. *Let $\mathcal{R}(\mathcal{A})$ denote the runtime of an algorithm \mathcal{A} . Assuming the above subexponential hardness of LWE, we have the following complexity regarding StreamAlg in Construction 3.1:*

1. $(\widetilde{\text{pk}}, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ and sketch can be stored using $\widetilde{\mathcal{O}}(1)$ total bits of space.
2. $\mathcal{R}(\text{StreamAlg.Update}) = \widetilde{\mathcal{O}}(1) + \mathcal{R}(\text{StreamAlg}_0.\text{Update})$.
3. $\mathcal{R}(\text{StreamAlg.Report}) = \widetilde{\mathcal{O}}(k) + \mathcal{R}(\text{StreamAlg}_0.\text{Report})$.

Proof. Recall that computations in StreamAlg are modulo q with $q \in \text{poly}(n)$ and $q \gg n \cdot N$, where $N \in \text{poly}(n)$ is the bound of the stream. Thus all entries have bit complexity in $\mathcal{O}(\log n)$.

1. $\widetilde{\text{pk}}$, sketch , and $\widetilde{\text{ct}}_i$ for $i \in [\log n]$ are all $\mathcal{O}(g) \times \mathcal{O}(h)$ matrices over \mathbb{Z}_q . For $g = \lambda \in \text{poly}(\log n)$ and $h \in \Theta(g \cdot \log q)$, they can be stored using $\mathcal{O}(\text{poly}(\log n)) = \widetilde{\mathcal{O}}(1)$ total bits of space.
2. The second bullet point considers the following operation:
 - Evaluate $\widehat{\text{ct}}_{i_t} \leftarrow \text{PFHE.Eval}(\widetilde{\text{pk}}, C_{i_t}, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_\ell)$. Then update sketch as $\text{sketch} \leftarrow \text{sketch} + \Delta_t \cdot \widehat{\text{ct}}_{i_t}$.

Updating sketch takes $\widetilde{\mathcal{O}}(1)$ time. For evaluation, recall that C_{i_t} is a point function checking whether the input bits equal i_t . This can be na

vely implemented using $\mathcal{O}(\log n)$ arithmetic operations, each of which runs in $\text{poly} \log n$ time as described in Eval_+ and Eval_\times above. Thus the total runtime is $\widetilde{\mathcal{O}}(1)$.

3. The third bullet point considers the following operation:

- (If $\|\mathbf{x}'\|_0 > k$, output \perp . Otherwise:) Evaluate $\widehat{\text{ct}}_i \leftarrow \text{PFHE.Eval}(\widetilde{\text{pk}}, C_i, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_\ell)$ for $i \in [n]$. Then compute $\text{hash} \leftarrow \sum_{i \in [n]} \widehat{\text{ct}}_i \cdot \mathbf{x}'_i$.

Observe that although the construction is described as evaluating n -many $\widehat{\text{ct}}_i$, one for each coordinate of \mathbf{x}' , at this line of the algorithm \mathbf{x}' is guaranteed to have at most k non-zero entries. Thus k evaluations (corresponding to non-zero coordinates of \mathbf{x}') suffice for computing the linear combination hash . Each evaluation is $\widetilde{\mathcal{O}}(1)$, so the total runtime is $\widetilde{\mathcal{O}}(k)$. □

E.4. Efficiency of PFHE Based on Polynomial LWE

Under the polynomial LWE assumption, we set $\lambda = n^w$ for an arbitrarily small constant $w > 0$. The (\mathcal{T}, ϵ) -security then reduces to the standard definition in terms of PPT adversaries and negligible advantage. We again choose $g = \lambda$ and $h \in \Theta(g \cdot \log q)$. The complexity is:

Proposition E.3. *Let $\mathcal{R}(\mathcal{A})$ denote the runtime of an algorithm \mathcal{A} . Assuming the polynomial hardness of LWE, we have the following complexity regarding StreamAlg in Construction 3.1:*

1. $(\widetilde{\text{pk}}, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ and sketch can be stored in $\widetilde{\mathcal{O}}(n^{w_1})$ bits of space.
2. $\mathcal{R}(\text{StreamAlg.Update}) = \widetilde{\mathcal{O}}(n^{w_2}) + \mathcal{R}(\text{StreamAlg}_0.\text{Update})$.

$$3. \mathcal{R}(\text{StreamAlg.Report}) = \tilde{\mathcal{O}}(k \cdot n^{w_3}) + \mathcal{R}(\text{StreamAlg}_0.\text{Report}).$$

where w_1, w_2 , and w_3 are arbitrarily small positive constants.

Proof. The proof is exactly the same as the proof for Proposition E.2, except that now the dimension of the ciphertext matrices blows up from $\text{poly log } n$ to n^w for an arbitrarily small $w > 0$. □

E.5. Complexity of Construction 3.1

Corollary 4.1 that we include in the main body follows from Theorem 4.1 and Proposition E.2. We summarize all combinations of instantiations in Table 2 below.

Table 2. A summary of the time and space complexities of the StreamAlg scheme in Construction 3.1, given different instantiations of StreamAlg₀ and PFHE. c denotes an arbitrarily small positive constant.

RELAXED REC. SCHEME	LWE ASSUMPTION	SPACE	UPDATE TIME	REPORT TIME
THEOREM 4.1	SUBEXP	$\tilde{\mathcal{O}}(k)$	$\tilde{\mathcal{O}}(1)$	$\tilde{\mathcal{O}}(k^{1+c})$
THEOREM E.1	SUBEXP	$\tilde{\mathcal{O}}(k)$	$\tilde{\mathcal{O}}(1)$	$\tilde{\mathcal{O}}(k^2)$
THEOREM 4.1	POLY	$\tilde{\mathcal{O}}(k + n^c)$	$\tilde{\mathcal{O}}(n^c)$	$\tilde{\mathcal{O}}(k \cdot (k^c + n^c))$
THEOREM E.1	POLY	$\tilde{\mathcal{O}}(k + n^c)$	$\tilde{\mathcal{O}}(n^c)$	$\tilde{\mathcal{O}}(k \cdot (k + n^c))$

F. Distributed K -Sparse Recovery from Ring LWE

F.1. Distributed Computation

In the distributed computation model, the dataset is split over multiple servers. The goal is to design a communication protocol for servers to collectively compute a function on the aggregated dataset. In this work, we consider the message-passing coordinator model (see, e.g., (Phillips et al., 2015)), where there are T servers and a single central coordinator. Each server $s_t : t \in [T]$ receives a partition of data in the form of an n -dimensional vector \mathbf{x}_t , and communicates with the central coordinator through a two-way, single-round communication channel. At the end of the communication, the coordinator needs to compute or approximate a function $f_n(\mathbf{x})$ of the aggregated vector $\mathbf{x} = \sum_{t \in [T]} \mathbf{x}_t$. We aim to design protocols that minimize the total amount of communication through the channel and speed up the computation at each server and the coordinator. See also (Woodruff & Zhang, 2012) and the references therein on the distributed model.

Similar to the streaming model, we focus on designing a robust protocol in a white-box adversarial distributed setting, where the dataset is generated and partitioned across each server by a white-box adversary who monitors the internal state of the coordinator and the servers. We define the model as follows:

Definition F.1 (Distributed Protocol). *Given a query function ensemble $\mathcal{F} = \{f_n : \mathbb{Z}^n \rightarrow X\}_{n \in \mathbb{N}}$ for some domain X , a distributed protocol DistProt that computes \mathcal{F} contains a tuple of algorithms DistProt = (DistProt.Setup, DistProt.Server, DistProt.Coordinator), with the following properties:*

- DistProt.Setup(n, ρ_0) $\rightarrow \zeta$: On input a parameter n , the coordinator samples randomness ρ_0 and uses ρ_0 to compute some state ζ . It sends ζ to all servers.
- DistProt.Server($\zeta, \mathbf{x}_t, \rho_t$) $\rightarrow p_t$: On input a state ζ and a data partition $\mathbf{x}_t \in \mathbb{Z}^n$, the t -th server samples randomness ρ_t and uses it to compute a packet p_t . It sends p_t back to the coordinator.
- DistProt.Coordinator(p_1, \dots, p_T, ρ_c) $\rightarrow r$: On input T packets p_1, \dots, p_T , the coordinator samples randomness ρ_c and computes a query response $r \in X$.

A response r is said to be correct if $r = f_n(\mathbf{x})$ where $\mathbf{x} = \sum_{t \in [T]} \mathbf{x}_t$.

As in the streaming setting, we again assume bounded integer inputs, such that an arbitrary subset sum of $\{\mathbf{x}_t : t \in [T]\}$ is in $[-N, N]^n$ for some $N \in \text{poly}(n)$.

Definition F.2. We say that a distributed protocol DistProt is white-box adversarially robust if for all dimension parameters $n \in \mathbb{N}$, the following holds: For any PPT adversary \mathcal{A} , the following experiment $\text{Expt}_{\mathcal{A}, \text{DistProt}}(n)$ outputs 1 with probability at most $\text{negl}(n)$:

$\text{Expt}_{\mathcal{A}, \text{DistProt}}(n)$:

1. The challenger and \mathcal{A} agree on a query function ensemble $\mathcal{F} = \{f_n : \mathbb{Z}^n \rightarrow X\}_{n \in \mathbb{N}}$ for some domain X . On input a dimension parameter n , let $f_n \in \mathcal{F}$ be the query function of the experiment.
2. The challenger samples all randomness $\rho_0, \rho_t : t \in [T]$, and ρ_c for the servers and the coordinator. It then provides $\rho_0, \rho_1, \dots, \rho_T$ and ρ_c to \mathcal{A} .
3. \mathcal{A} generates T partitions of data $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathbb{Z}^n$ and provides them to the challenger.
4. The challenger executes the protocol to sequentially compute the following:
 - (i) $\zeta \leftarrow \text{DistProt.Setup}(n, \rho_0)$.
 - (ii) $p_t \leftarrow \text{DistProt.Server}(\zeta, \mathbf{x}_t, \rho_t)$ for all $t \in [T]$.
 - (iii) $r \leftarrow \text{DistProt.Coordinator}(p_1, \dots, p_T, \rho_c)$.
5. The experiment outputs 1 if and only if at some time $r \neq f_n(\mathbf{x})$.

Remark F.1. We remark on the connection between the streaming model and the distributed model. There exists a well-known reduction from a streaming algorithm to a distributed protocol under the oblivious setting: Each server runs StreamAlg.Update on the same initial state $\zeta \leftarrow \text{StreamAlg.Setup}$, using all entries of \mathbf{x}_t (its data partition) as updates. It then sends the resulting sketch to the coordinator. As long as the sketches are linear, the coordinator can run StreamAlg.Report on the sum of all sketches to compute a query response.

One may use the same recipe to construct a WAR distributed protocol from a WAR streaming algorithm. In general, the robustness is retained by such a reduction if StreamAlg.Update and StreamAlg.Report are deterministic. (Otherwise, the adversaries in the experiment $\text{Expt}_{\mathcal{A}, \text{DistProt}}$ would have additional views to all “future” randomness, while adversaries in the streaming experiment $\text{Expt}_{\mathcal{A}, \text{StreamAlg}}$ only monitor the past randomness already used by StreamAlg .) Our Construction 3.1 satisfies this condition.

F.2. Distributed Protocol Construction

For completeness, we briefly describe the distributed protocol for k -Sparse Recovery as follows, which essentially follows the same route as our streaming Construction 3.1.

Construction F.1. Similar to Construction 3.1, let PFHE be a pseudorandom FHE scheme. And let StreamAlg_0 be a streaming algorithm for the Relaxed k -Sparse Recovery problem. We assume that StreamAlg_0 performs a deterministic linear measurement Φ_n on an arbitrary length- n input vector, as described in the previous instantiation Section E.1.

- $\text{DistProt.Setup}(n)$: Let $\lambda = \mathcal{S}(n)$ be the security parameter of PFHE. Sample $\widetilde{\text{pk}} \leftarrow \mathcal{D}_{\widetilde{\text{pk}}(\lambda)}$ and $\widetilde{\text{ct}}_1, \widetilde{\text{ct}}_2, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil} \leftarrow \mathcal{D}_{\widetilde{\text{ct}}(\lambda)}$. Send $\zeta := (\widetilde{\text{pk}}, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ to all servers
- $\text{DistProt.Server}(\zeta, \mathbf{x}_t)$: Given $\zeta := (\widetilde{\text{pk}}, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$, evaluate $\widehat{\text{ct}}_i \leftarrow \text{PFHE.Eval}(\widetilde{\text{pk}}, C_i, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ for $i \in [n]$, where C_i is the same point function as in Construction 3.1.
Compute $\Phi_n(\mathbf{x}_t)$ and $\text{sketch}_t \leftarrow \sum_{i \in [n]} \widehat{\text{ct}}_i \cdot (\mathbf{x}_t)_i$, where $(\mathbf{x}_t)_i$ denotes the i -th coordinate of \mathbf{x}_t . Send $p_t := (\text{sketch}_t, \Phi_n(\mathbf{x}_t))$ to the coordinator.
- $\text{DistProt.Coordinator}(p_1, \dots, p_T)$: Let $\mathbf{x}' \leftarrow \text{StreamAlg}_0.\text{Report}(\sum_{t \in [T]} \Phi_n(\mathbf{x}_t))$, i.e., the output of running the reconstruction algorithm of StreamAlg_0 on the sum of all measurement results. If $\|\mathbf{x}'\|_0 > k$, output \perp .
Otherwise, evaluate $\widehat{\text{ct}}_i \leftarrow \text{PFHE.Eval}(\widetilde{\text{pk}}, C_i, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ for $i \in [n]$. Then compute $\text{hash} \leftarrow \sum_{i \in [n]} \widehat{\text{ct}}_i \cdot \mathbf{x}'_i$. Output \mathbf{x}' if $\text{hash} = \sum_{t \in [T]} \text{sketch}_t$, and output \perp otherwise.

F.3. Efficiency

In this section, we focus on showing that we can use BV (Brakerski & Vaikuntanathan, 2011), a Ring-LWE based FHE scheme, to achieve near-linear processing time on DistProt.Server under the *polynomial* ring-LWE assumption.

As a warm-up, we state the complexity of Construction F.1 using GSW to implement PFHE:

Proposition F.1. *Let $\mathcal{R}(\mathcal{A})$ denote the runtime of an algorithm or an evaluation \mathcal{A} , we have the following complexity regarding DistProt in Construction F.1:*

- *Assuming the subexponential hardness of LWE:*
 1. *The total bits of communication between an arbitrary server and the coordinator is $\tilde{O}(1) + \text{size}(\Phi_n(\mathbf{x}))$.*
 2. $\mathcal{R}(\text{DistProt.Coordinator}) = \tilde{O}(k) + \mathcal{R}(\text{StreamAlg}_0.\text{Report})$.
 3. $\mathcal{R}(\text{DistProt.Server}) = \tilde{O}(k) + \mathcal{R}(\Phi_n(\mathbf{x}))$.
- *Assuming the polynomial hardness of LWE: :*
 1. *The total bits of communication between an arbitrary server and the coordinator is $\tilde{O}(n^{w_1}) + \text{size}(\Phi_n(\mathbf{x}))$.*
 2. $\mathcal{R}(\text{DistProt.Coordinator}) = \tilde{O}(k \cdot n^{w_2}) + \mathcal{R}(\text{StreamAlg}_0.\text{Report})$.
 3. $\mathcal{R}(\text{DistProt.Server}) = \tilde{O}(k \cdot n^{w_3}) + \mathcal{R}(\Phi_n(\mathbf{x}))$.

where $\mathbf{x} \in [-N, N]^n$ is arbitrary, and w_1, w_2 , and w_3 are arbitrarily small positive constants.

Proof. The communication complexity and the runtime of DistProt.Coordinator are the same as the bit complexity and reporting time of StreamAlg in Construction 3.1 (except for some negligible addition operations). The runtime of DistProt.Server is roughly $n \cdot \mathcal{R}(\text{StreamAlg.Update})$, because every server has to perform n evaluations, one for each coordinate of its data partition. □

Efficiency Based on Ring-LWE. Before discussing the efficiency advantage, we briefly review some important features of BV to show that it satisfies a variant of the PFHE definition. See (Brakerski & Vaikuntanathan, 2011) for details of the BV construction.

Consider a ring $R := \frac{\mathbb{Z}[x]}{x^g+1}$ where g is the degree parameter. Let R_q denote R/qR for a prime modulus q , and let R_t be the message space for some prime $t \in \mathbb{Z}_q^*$ and $t \ll q/g$:

- The public keys of BV are two-dimensional vectors in R_q^2 , which are computationally indistinguishable from uniformly random samples from R_q^2 under the Ring-LWE assumption.

The ciphertexts are also vectors over R_q that are indistinguishable from uniform random.

- BV satisfies an analogy of \mathbb{Z} Linear Homomorphism, with coefficients of linear combinations $\mathcal{M} = \sum_{i \in [k]} x_i \cdot \text{ct}_i$ being small norm ring elements $x_i \in R_t$ as opposed to integer scalars. As long as the linear combination of corresponding messages ($\mathcal{M}_{msg} = \sum_{i \in [k]} x_i \cdot m_i$, where $m_i \in R_t$ is the message encrypted by ct_i) still lies in R_t , running the standard BV decoding algorithm on \mathcal{M} output the correct message.
- Eval and Dec are deterministic by the construction of BV.

We choose parameters for our Ring-LWE based construction as follows: Let $\lambda = n^w$ for an arbitrarily small constant $w > 0$. Let the dimension $g = \lambda$, the modulus of the message space $t > n \cdot N$, where $N \in \text{poly}(n)$ is the bound of the input data. The modulus q is chosen with subexponential modulus-to-noise ratio and $q \gg g \cdot t$ the Ring-LWE assumption states that for any $h \in \text{poly}(\lambda)$, it holds that

$$\{(a_i, a_i \cdot s + e_i) : e_i \leftarrow \chi\}_{i \in [h]} \approx_c \{(a_i, u_i) : u_i \leftarrow R_q\}_{i \in [h]}$$

where s is sampled from the noise distribution χ and a_i are uniform in R_q .

We have the following complexity:

Proposition F.2. Let $\mathcal{R}(\mathcal{A})$ denote the runtime of an algorithm or an evaluation \mathcal{A} . Assuming the polynomial hardness of Ring-LWE, a variant of DistProt in Construction F.1 has the following complexity:

1. The total bits of communication between an arbitrary server and the coordinator is $\tilde{\mathcal{O}}(n^{w_1}) + \text{size}(\Phi_n(\mathbf{x}))$.
2. $\mathcal{R}(\text{DistProt.Coordinator}) = \tilde{\mathcal{O}}(n) + \mathcal{R}(\text{StreamAlg}_0.\text{Report})$.
3. $\mathcal{R}(\text{DistProt.Server}) = \tilde{\mathcal{O}}(k \cdot n_2^w) + \mathcal{R}(\Phi_n(\mathbf{x}))$.

where $\mathbf{x} \in [-N, N]^n$ is arbitrary, and w_1, w_2 are arbitrarily small positive constants.

Proof. In comparison to the complexity under polytime standard LWE, we save a $n^{\Theta(1)}$ factor in the runtime of DistProt.Server. This is due to speeding up the following two steps:

1. Evaluate $\widehat{\text{ct}}_i \leftarrow \text{PFHE.Eval}(\widetilde{\text{pk}}, C_i, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ for all $i \in [n]$
2. Compute $\sum_{i \in [n]} \widehat{\text{ct}}_i \cdot \mathbf{x}_i$

The idea is to represent the messages as ring elements over R_q as opposed to individual scalars, and apply fast Fourier transforms to efficiently multiply them.

In a variant of Construction F.1, instead of performing n evaluations one for each coordinate $i \in [n]$, we partition the dimension n into n^{1-w} chunks. Each chunk is of length n^w , where $g = n^w$ is the degree parameter of the ring for an arbitrarily small constant $w > 0$. Such that a length- n vector \mathbf{x} is partitioned into n^{1-w} consecutive segments, each can be represented as a (small-normed in comparison to q) ring element $x_j \in R_q$ for $j \in [n^{1-w}]$.

We also define a new set of functions C'_j as follows:

$$C'_j(\mu_1, \dots, \mu_{\lceil \log n \rceil}) = \begin{cases} 1 & \text{if } (\mu_1, \dots, \mu_{\lceil \log n \rceil}) \text{ represents a value in } [jn^w, (j+1)n^w) \\ 0 & \text{otherwise.} \end{cases}$$

We evaluate $\widehat{\text{ct}}'_j \leftarrow \text{PFHE.Eval}(\widetilde{\text{pk}}, C'_j, \widetilde{\text{ct}}_1, \dots, \widetilde{\text{ct}}_{\lceil \log n \rceil})$ for all $j \in [n^{1-w}]$ to replace the above step 1. The total evaluation time is $\mathcal{O}(n^{1-w} \cdot n^w \text{poly}(\log(n^w))) = \tilde{\mathcal{O}}(n)$, since two ring elements can be multiplied in $\mathcal{O}(n^w \text{poly}(\log(n^w)))$ time using the FFT. Then the linear combination in the above step 2 is replaced by $\sum_{j \in [n^{1-w}]} \widehat{\text{ct}}'_j \cdot x_j$, which similarly takes $\tilde{\mathcal{O}}(n)$ time to compute.

Lastly, we check that given two vectors \mathbf{x} and \mathbf{x}' differ at a coordinate m , when $\text{ct}_1, \dots, \text{ct}_{\lceil \log n \rceil}$ encrypt the bits of m , we still have

$$\text{sketch} = \sum_{j \in [n^{1-w}]} \widehat{\text{ct}}'_j \cdot x_j \neq \sum_{j \in [n^{1-w}]} \widehat{\text{ct}}'_j \cdot x'_j = \text{hash}.$$

because they can be decrypted into distinct ring elements. This concludes that the above variant of Construction F.1 remains correct while achieving near-linear runtime of DistProt.Coordinator and DistProt.Server. \square

G. Low-Rank Matrix and Tensor Recovery

The low-rank matrix and tensor recovery problems are defined as follows:

Definition G.1 (Rank- k Matrix Recovery). Given a rank parameter k and an input matrix $\mathbf{X} \in \mathbb{Z}^{n \times m}$, output \mathbf{X} if $\text{rank}(\mathbf{X}) \leq k$ and report \perp otherwise.

Let \otimes denote the outer product of two vectors. We use the following definition for tensor rank:

Definition G.2 (CANDECOMP/PARAFAC (CP) rank). For a tensor $X \in \mathbb{R}^{n^d}$, consider it to be the sum of k rank-1 tensors: $X = \sum_{i=1}^k (x_{i1} \otimes x_{i2} \otimes \cdots \otimes x_{id})$ where $x_{ij} \in \mathbb{R}^n$. The smallest number of rank-1 tensors that can be used to express a tensor X is then defined to be the rank of the tensor.

Definition G.3 (Rank- k Tensor Recovery). Given a rank parameter k and an input tensor $\mathbf{X} \in \mathbb{Z}^{n^d}$, output \mathbf{X} if $CP - \text{rank}(\mathbf{X}) \leq k$ and report \perp otherwise.

The low-rank matrix and tensor recovery problems are natural generalizations of the k -sparse vector recovery problems. Using the same framework as the vector case, they can be solved by maintaining two hashes: one recovers the matrix or tensor but potentially produces garbage output, and the other uses a cryptographic tool to verify whether the output matches the input data. These are natural generalizations of the k -sparse vector recovery problems.

In (Feng & Woodruff, 2023), the recovery hash is measured by a random Gaussian matrix, and the verification hash is measured by an SIS matrix. Both measurements are assumed to be heuristically compressed by an idealized hash function. We instead replace the former with a deterministic measurement, which can be explicitly constructed and decoded in polynomial time (Forbes & Shpilka, 2012). The latter is replaced by the same FHE-based hashing as the one we used for the vector case in Construction 3.1, which operates on the vectorization of the data matrix or tensor. Similar to the vector case, our construction does not require using an idealized hash function or storing a separate large state.

In the streaming model, the complexities of our matrix and tensor recovery algorithms are as follows:

Proposition G.1. Let $\mathcal{R}(\mathcal{A})$ denote the runtime of an algorithm or an evaluation \mathcal{A} . Assuming the polynomial hardness of *LWE*,

- There exists a WAR streaming algorithm *StreamAlg* for the rank- k matrix recovery problem using $\tilde{\mathcal{O}}(nk)$ bits of space.
- There exists a WAR streaming algorithm *StreamAlg* for the rank- k tensor recovery problem using $\tilde{\mathcal{O}}(nk^{\lceil \log d \rceil})$ bits of space.

Proof. The recovery scheme for matrices and tensors takes $\tilde{\mathcal{O}}(nk)$ and $\tilde{\mathcal{O}}(nk^{\lceil \log n \rceil})$ bits of space, respectively (Forbes & Shpilka, 2012). Under the polynomial hardness assumption of *LWE*, we set $\lambda = n^c$ for an arbitrarily small positive constant c . Taking the sum of the two terms produces the claimed complexities. \square

In addition, the above streaming algorithms can be converted into WAR distributed protocols for matrix and vector recovery, following the same roadmap as in Section F.

Proposition G.2. Let $\mathcal{R}(\mathcal{A})$ denote the runtime of an algorithm or an evaluation \mathcal{A} . Assuming the polynomial hardness of *LWE*:

- There exists a WAR distributed protocol *DistProt* for the rank- k matrix recovery problem, using $\tilde{\mathcal{O}}(nk)$ bits of communication between each server and the coordinator, and $\tilde{\mathcal{O}}(n^2k)$ processing time at each server.
- There exists a WAR streaming algorithm *StreamAlg* for the rank- k tensor recovery problem, using $\tilde{\mathcal{O}}(nk^{\lceil \log d \rceil})$ bits of communication between each server and the coordinator.

Proof. The total number of bits of communication between each server and the coordinator equals the size of the state in the above streaming algorithms. For low-rank matrix recovery, the deterministic scheme performs $\tilde{\mathcal{O}}(nk)$ measurements using n -sparse matrices, and the runtime for computing the FHE-based hash is comparably small. This gives $\tilde{\mathcal{O}}(n^2k)$ processing time at each server. \square

We remark that the rank- k matrix recovery algorithm in the previous work (Feng & Woodruff, 2023) requires $\tilde{\mathcal{O}}(n^3k)$ processing time in the distributed setting, which our construction improves upon.