



How to Compress Encrypted Data

Nils Fleischhacker¹✉, Kasper Green Larsen², and Mark Simkin³

¹ Ruhr University Bochum, Bochum, Germany
mail@nilsfleischhacker.de

² Aarhus University, Aarhus, Denmark
larsen@cs.au.dk

³ Ethereum Foundation, Aarhus, Denmark
mark.simkin@ethereum.org

Abstract. We study the task of obviously compressing a vector comprised of n ciphertexts of size ξ bits each, where at most t of the corresponding plaintexts are non-zero. This problem commonly features in applications involving encrypted outsourced storages, such as searchable encryption or oblivious message retrieval. We present two new algorithms with provable worst-case guarantees, solving this problem by using only homomorphic additions and multiplications by constants. Both of our new constructions improve upon the state of the art asymptotically and concretely.

Our first construction, based on sparse polynomials, is perfectly correct and the first to achieve an asymptotically optimal compression rate by compressing the input vector into $\mathcal{O}(t\xi)$ bits. Compression can be performed homomorphically by performing $\mathcal{O}(n \log n)$ homomorphic additions and multiplications by constants. The main drawback of this construction is a decoding complexity of $\Omega(\sqrt{n})$.

Our second construction is based on a novel variant of invertible bloom lookup tables and is correct with probability $1 - 2^{-\kappa}$. It has a slightly worse compression rate compared to our first construction as it compresses the input vector into $\mathcal{O}(\xi\kappa t / \log t)$ bits, where $\kappa \geq \log t$. In exchange, both compression and decompression of this construction are highly efficient. The compression complexity is dominated by $\mathcal{O}(n\kappa / \log t)$ homomorphic additions and multiplications by constants. The decompression complexity is dominated by $\mathcal{O}(\kappa t / \log t)$ decryption operations and equally many inversions of a pseudorandom permutation.

1 Introduction

It is well known that in general encrypted data cannot be compressed. In this work, we study the task of compressing encrypted data, when a small amount

N. Fleischhacker—Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA - 390781972.

K. G. Larsen—Supported by Independent Research Fund Denmark (DFR) Sapere Aude Research Leader grant No. 9064-00068B.

© International Association for Cryptologic Research 2023

C. Hazay and M. Stam (Eds.): EUROCRYPT 2023, LNCS 14004, pp. 551–577, 2023.

https://doi.org/10.1007/978-3-031-30545-0_19

of knowledge about the structure of the underlying plaintexts is known. More concretely, we consider encryptions of vectors $\mathbf{m} = (m_1, \dots, m_n)$ where at most t distinct coordinates in \mathbf{m} are non-zero. In the context of outsourced storage applications the task of compressing such vectors appears naturally.

In searchable encryption [3, 26], we have a client Charlie, who holds a vector (m_1, \dots, m_n) of data elements and wants to store it remotely on server Sally. To hide the contents of the data elements, Charlie encrypts the data vector under a secret key only she knows before sending it to Sally. Later on, Charlie may want to search through the vector and retrieve all elements that match some secret keyword. A series of recent works [1, 6–8, 18, 27] have shown how to construct searchable encryption schemes from fully homomorphic encryption [12, 22] with reasonable concrete efficiency. Conceptually, these approaches are comprised of two major steps. First Charlie sends a short keyword-dependent hint to the server, who uses it to obliviously transform the vector of ciphertexts (c_1, \dots, c_n) into a new vector $\tilde{\mathbf{c}} = (\tilde{c}_1, \dots, \tilde{c}_n)$, where for $i \in \{1, \dots, n\}$ the ciphertext \tilde{c}_i is either an encryption of the original message m_i in c_i or zero, depending on whether m_i was matching the keyword of Charlie or not. In the second step, the server obliviously compresses the vector $\tilde{\mathbf{c}}$ under the assumption that no more than t ciphertexts were matching the keyword and sends it back to Charlie. If the assumption about the sparsity of the vector $\tilde{\mathbf{c}}$ was correct, then Charlie successfully decodes the vector and obtains the desired result. If the assumption was not correct, then Charlie may not be able to retrieve the output.

The best compression algorithm used in this context is due to Choi et al. [8], which compresses $\tilde{\mathbf{c}}$ into a bit string of length $\Omega(t\xi(\kappa + \log n))$, where ξ is the length of one ciphertext entry. Under the assumption that the plaintext messages are of a specific form¹, the authors show that Charlie can correctly decode the vector with probability $1 - 2^{-\kappa}$. Both compressing and decoding are computationally concretely efficient.

In the oblivious message retrieval setting, recently introduced by Liu and Tromer [19], we have a server Sally, who keeps a public bulletin board, and multiple clients Charlie, Chucky, and Chris. Each of the clients can post encrypted messages for any of the other clients on the bulletin board, but would like to hide who is the recipient of which message. At some point, for example, Chucky may want to retrieve all messages that are intended for him. Naively, he could simply download all contents from Sally's bulletin board, but this would incur a large bandwidth overhead that is linear in the total number of messages stored by Sally. Instead, the idea behind oblivious message retrieval is to let Chucky generate a short identity-dependent hint that can be used by Sally to obliviously generate a short message that contains all relevant encrypted messages for Chucky. Conceptually, the construction of Liu and Tromer follows the exact same blueprint as the searchable encryption scheme outlined above. First Sally obliviously filters her vector with the hint provided by Chucky and then she

¹ This assumption can be removed at the cost of doubling the size of the compressed vector and additionally assuming that one is not only given $\tilde{\mathbf{c}}$, but also some auxiliary vector $\hat{\mathbf{c}}$ as the output of the first step of their protocol.

obviously compresses the filtered vector under the assumption that not too many messages are addressed to Chucky.

From an efficiency perspective, the solution of Liu and Tromer is rather expensive. To compress, Sally performs $\Omega(tn + \kappa t \log t)$ homomorphic additions and $\Omega(tn)$ homomorphic multiplications by constants, where κ is the correctness error defined as in the searchable encryption example. Sally's message to Chucky is $\Omega(\xi t + \xi \kappa t \log t)$ bits long. To decode the result from Sally's message, Chucky needs to perform gaussian elimination on a matrix of size $\mathcal{O}(t) \times \mathcal{O}(t)$, which incurs a computational overhead of $\Omega(t^3)$. The authors provide heuristic optimizations of their constructions that improve their performance significantly, but unfortunately these come without asymptotic bounds or provable correctness guarantees.

Taking a step back and looking at the two applications described above from a more abstract point of view, one can recognize that both follow a very similar blueprint. In the first step, both apply some vastly different techniques to convert a vector of ciphertexts into a sparse vector containing only the desired entries. In the second step, both works solve the identical problem. They both need to compress a sparse homomorphically encrypted vector with nothing but the knowledge of how many entries are non-zero, and in particular without any knowledge about which entries are zero and which ones are not. How to compress such a sparse encrypted vector is the topic of this work.

1.1 Our Contribution

We present two new algorithms, one based on polynomials and one based on algorithmic hashing, for compressing sparse encrypted vectors, which both improve upon the prior state of the art in terms of compression rate. Our algorithms only rely on homomorphic additions and homomorphic multiplications by constants. Both of our constructions have provable worst-case bounds for all their parameters.

Compressing via Polynomials. Our first construction (Sect. 4) is perfectly correct and is based on the concept of sparse polynomial interpolation. Its compression rate has an asymptotically optimal dependence on t , as the compressed vector is merely $\mathcal{O}(\xi \cdot t)$ bits large. During compression one needs to perform $\mathcal{O}(n \log n)$ homomorphic additions and equally many homomorphic multiplications by constants. The main bottleneck of this solution is the decompression complexity of $\tilde{\mathcal{O}}(t \cdot \sqrt{n})$, which depends on the length n of the original vector. Although the compression rate is much better than that of previous works, such as those Choi et al. [8] and that of Liu and Tromer [19], this construction suffers from a slower decompression time.

Compressing via Hashing. Our second construction (Sect. 5) is a randomized hashing based solution, which is correct with probability $1 - 2^{-\kappa}$, where the probability is taken over the random coins of the compression algorithm. We develop a novel data structure that is heavily inspired by the invertible bloom

lookup tables of Goodrich and Mitzenmacher [14], but can be applied efficiently to encrypted data. Both compression and decompression are highly efficient. During compression one needs to perform $\mathcal{O}(n\kappa/\log t)$ homomorphic additions and multiplications by constants, where $\kappa \geq \log t$. During decompression the main costs come from $\mathcal{O}(\kappa t/\log t)$ many decryptions and equally many evaluations of a pseudorandom permutation. In contrast to our polynomial based solution, however, the compressed vector is $\mathcal{O}(\xi\kappa t/\log t)$ bits large. Nevertheless, this construction outperforms all prior works in terms of compression rate, while having either superior or comparable compression and decompression complexities.

1.2 Strawman Approach

When the sparse vector is encrypted using a fully homomorphic encryption scheme, conceptually simple solutions to the compression problem exist. For instance, Sally could just homomorphically sort all entries in the vector and then only send back the t largest entries. Such solutions, however, require her to perform multiplications of encrypted values. This is problematic for multiple reasons. Multiplications of encrypted values are much more computationally expensive than homomorphic additions or multiplications by constants. Since Sally may potentially store a very large database, we would like to minimize her computational overhead. Furthermore, if the data is encrypted using a somewhat homomorphic encryption scheme, then the multiplicative depth of the circuit that can be executed on the vector by Sally is bounded. Ideally, we would like the compression step to be concretely efficient and not require the use of any multiplications of encrypted values. For these reasons, we only focus on compression algorithms that require homomorphic additions and multiplications by constants in this work.

1.3 Additional Related Works

In addition to what has already been discussed above, there are several other works that are related to ours. Johnson, Wagner, and Ramchandran [16] showed that, assuming messages from a source with bounded entropy, it is possible to compress one-time pad encryptions without knowledge of the encryption key through a clever application of Slepian-Wolf coding [24]. Their result only applied to linear stream ciphers but was later extended to block ciphers using certain chaining modes by Klinc et al. [17]. These results do not apply to our setting, where we focus on compressing vectors encrypted using more complex homomorphic public-key encryption schemes.

In the context of fully homomorphic encryption, multiple works [4, 20, 25] have studied the question of how to optimize the encryption rate, i.e., the size of the ciphertext relative to the size of the plaintext, by packing multiple plaintexts into one ciphertext. These results are related, but do not allow for obviously “removing” irrelevant encryptions of zero.

Another line of works [5,11,13] studies the compressed sensing problem, where the task is to design a matrix A such that it is possible to recover, possibly high-dimensional, but sparse vectors x from a vector of measurements Ax of small dimension. In general these works aim to recover an approximation of x even when given somewhat noisy measurements. In our case, we are interested in the simpler problem of exact recovery of a sparse vector. Our construction based on polynomials can be seen as a matrix-vector multiplication. Looking ahead, our matrix A will be a carefully chosen Vandermonde matrix that allows for very efficient matrix vector multiplication. The server will multiply this public matrix with the encrypted vector and send back the result that can be decoded by the client. Our second construction, based on hashing, does not fall into this category of algorithms.

2 Preliminaries

Notation. Given a possibly randomized function $f : X \rightarrow Y$, we will sometimes abuse notation and write $f(\mathbf{x}) := (f(x_1), \dots, f(x_n))$ for $\mathbf{x} \in X^n$. For a set X , we write $x \leftarrow X$ to denote the process of sampling a uniformly random element $x \in X$. For a vector $\mathbf{v} \in X^n$, we write v_i to denote its i -th component. For a matrix $M \in X^{n \times m}$, we write $M[i, j]$ to denote the cell in the i -th row and j -th column. We write $[n]$ to denote the set $\{1, \dots, n\}$. For a set X^n , we define the scissor operator $\mathfrak{S}(X^n) := \{(x_1, \dots, x_n) \in X^n \mid x_i \neq x_j \ \forall i, j \in [n]\}$ to denote the subset of X^n consisting only of those vectors with unique entries.

Definition 1 (Sparse Vector Representation). Let \mathbb{F}_q be a field and let $\mathbf{a} \in \mathbb{F}_q^n$ be a vector. The sparse representation of \mathbf{a} is the set $\text{sparse}(\mathbf{a}) := \{(i, a_i) \mid a_i \neq 0\}$.

2.1 Homomorphic Encryption

Informally, a homomorphic encryption scheme allows for computing an encryption of $f(\mathbf{m})$, when only given the description of f and an encryption of message vector \mathbf{m} . Throughout the paper, we assume that functions are represented as circuits composed of addition and multiplication gates.

Definition 2. A homomorphic encryption scheme \mathcal{E} is defined by a tuple of PPT algorithms (Gen, Enc, Eval, Dec) that work as follows:

Gen(1^λ): The key generation algorithm takes the security parameter 1^λ as input and returns a secret key sk and public key pk . The public key implicitly defines a message space \mathcal{M} and ciphertext space \mathcal{C} . We denote the set of all public keys as \mathcal{P} .

Enc(pk, m): The encryption algorithm takes the public key pk and message $m \in \mathcal{M}$ as input and returns a ciphertext $c \in \mathcal{C}$.

Eval(pk, f, \mathbf{c}): The evaluation algorithm takes the public key pk , a function $f : \mathcal{M}^n \rightarrow \mathcal{M}^m$, and a vector $\mathbf{c} \in \mathcal{C}^n$ of ciphertexts as input and returns a new vector of ciphertexts $\tilde{\mathbf{c}} \in \mathcal{C}^m$.

$\text{Dec}(\text{sk}, c)$: The deterministic decryption algorithm takes the secret key sk and ciphertext $c \in \mathcal{C}$ as input and returns a message $m \in \mathcal{M} \cup \{\perp\}$.

Throughout the paper we assume that the ciphertext size is fixed and does not increase through the use of the homomorphic evaluation algorithm. We extend the definition of Enc and Dec to vectors and matrices of messages and ciphertexts respectively, by applying them componentwise, i.e., for any matrix $\mathbf{M} \in \mathcal{M}^{n \times m}$, we have $\text{Enc}(\text{pk}, \mathbf{M}) = \mathbf{C}$ with $\mathbf{C} \in \mathcal{C}^{n \times m}$ and $C[i, j] = \text{Enc}(\text{pk}, M[i, j])$ and equivalently $\text{Dec}(\text{sk}, \mathbf{C}) = \mathbf{M}'$ with $\mathbf{M}' \in \mathcal{M}^{n \times m}$ and $M'[i, j] = \text{Dec}(\text{sk}, C[i, j])$. Let \mathcal{E} be an additively homomorphic encryption scheme with message space $\mathcal{M} = \mathbb{F}_q$ for some prime power q . And let $f : \mathbb{F}_q^2 \rightarrow \mathbb{F}_q$, $f(a, b) := a + b$ and let $g_\alpha : \mathbb{F}_q \rightarrow \mathbb{F}_q$, $g_\alpha(a) := \alpha \cdot a$ for any constant $\alpha \in \mathbb{F}_q$. For notational convenience we write $\text{Eval}(\text{pk}, f, (c_1, c_2)^\top)$ as $c_1 \boxplus c_2$ and $\text{Eval}(\text{pk}, g_\alpha, c)$ as $\alpha \boxtimes c$ with pk being inferrable from context. For the sake of simplicity we restrict ourselves to homomorphic encryption schemes with unique secret keys, i.e. for a given pk , there exists at most one sk , such that $(\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^\lambda)$. We write $\text{Gen}^{-1}(\text{pk})$ to denote the – not efficiently computable – unique secret key.

Later on in the paper, it will be convenient for us to talk about ciphertexts that may not be fresh encryptions, but still allow for some homomorphic operations to be performed on them.

Definition 3 (\mathcal{Z} -Validity). Let $(\text{Gen}, \text{Enc}, \text{Eval}, \text{Dec})$ be a homomorphic encryption scheme, let \mathcal{Z} be a class of circuits, and let pk be a public key. A vector \mathbf{c} of ciphertexts is \mathcal{Z} -valid for pk , iff for all functions $f \in \mathcal{Z}$ it holds that $\perp \notin \text{Dec}(\text{Gen}^{-1}(\text{pk}), \mathbf{c})$ and $\text{Dec}(\text{Gen}^{-1}(\text{pk}), \text{Eval}(\text{pk}, f, \mathbf{c})) = f(\text{Dec}(\text{sk}, \mathbf{c}))$. We denote by $\text{vld}(\mathcal{Z}, \text{pk})$ the set of ciphertext vectors \mathcal{Z} -valid for pk .

2.2 Polynomial Kung Fu

Let $f(x) = \sum_{i=0}^d a_i \cdot x^i \in \mathbb{F}_q[x]$ be a polynomial with coefficients from a finite field \mathbb{F}_q . The degree of f is defined as the largest exponent in any monomial with a non-zero coefficient. We say that f is s -sparse, if the number of non-zero monomials is at most s or more formally if $|\{a_i \mid a_i \neq 0 \wedge i \in [n]\}| \leq s$. It is well-known that any polynomial of degree at most d can be interpolated from $d+1$ evaluation points. In this work, we will make use of the less well-known fact that sparse polynomials can be interpolated from a number of evaluation points that is linear in the polynomial’s sparsity. The first algorithms for interpolating sparse univariate and multivariate polynomials were presented by Prony [21] and Ben-Or and Tiwari [2] respectively. We will make use of the following result by Huang and Gao [15] for sparse interpolation of univariate polynomials over finite fields:

Theorem 4 ([15]). Let $f \in \mathbb{F}_q[x]$ be an s -sparse univariate polynomial of degree at most d with coefficients from a finite field \mathbb{F}_q . Let $\omega \in \mathbb{F}_q$ be a primitive $2(s+1)$ -th root of unity. There exists an algorithm *Interpolate* that takes evaluations $f(\omega^0), \dots, f(\omega^{2s+1})$ as input and returns the coefficients of f in sparse representation in time $\tilde{O}(s \cdot \sqrt{d})$.

The algorithm of Huang and Gao relies on a subroutine for finding discrete logarithms. Using Shank’s algorithm [23] for this step, we obtain the computational complexity stated in the above theorem with deterministic performance guarantees.

Another tool we will use is the Fast Fourier Transform, (re-)discovered by Cooley and Tukey [9] which allows for evaluating a degree d polynomial given as a list of coefficients at $\ell \leq d$ evaluation points simultaneously in time $\mathcal{O}(d \log d)$. More precisely, for a fixed set of evaluation points $(\omega^0, \dots, \omega^\ell)$, one can represent the circuit taking the polynomial coefficients (a_0, \dots, a_d) as input and returning $(f(\omega^0), \dots, f(\omega^\ell))$ as a series of $\mathcal{O}(\log d)$ alternating layers of $\mathcal{O}(d)$ addition or multiplication by constants gates respectively.

Theorem 5 [Fast Fourier Transform]. *Let $d, \ell \in \mathbb{N}$ with $d \geq \ell$. Let $f = \sum_{i=0}^d a_i \cdot x^i \in \mathbb{F}_q[x]$ be a polynomial of degree at most d with coefficients from a finite field \mathbb{F}_q . Let $\omega \in \mathbb{F}_q$ be a primitive ℓ -th root of unity. There exists an arithmetic circuit FFT comprised of a series of $\mathcal{O}(\log d)$ alternating layers of $\mathcal{O}(d)$ addition or multiplication by constants gates respectively that takes (a_0, \dots, a_d) as well as $(\omega^0, \dots, \omega^\ell)$ as input and returns $(f(\omega^0), \dots, f(\omega^\ell))$.*

2.3 Invertible Bloom Lookup Tables

An invertible Bloom lookup table (IBLT) is a data structure first introduced by Goodrich and Mitzenmacher [14] that supports three operations called Insert, Peel, and List. The insertion operations adds elements to the data structure, the deletion operations removes them² and the list operation recovers all currently present elements with high probability, if not too many elements are present.

The data structure consists of two $\gamma \times 8t$ matrices C , the count matrix and V the valueSum matrix. It further requires t -wise independent hash functions $h_i : \{0, 1\}^* \rightarrow [8t]$ for $i \in [\gamma]$. Initially all values are set to 0. To insert an element x into the data structure, we locate the cells $C_{i, h_i(x)}$ and $V_{i, h_i(x)}$ for $i \in [\gamma]$ and add 1 to each counter and x to each valueSum. To remove an element, we perform the inverse operations. To list all elements currently present in (C, V) , we repeatedly perform a peeling operation until (C, V) is empty. The peeling operation finds a cell with counter 1, adds that corresponding valueSum value to the output list and deletes the element from (C, V) . The only way the list operation may fail is if (C, V) is not empty, but the peeling operation cannot find any cell with counter 1. It has been shown by Goodrich and Mitzenmacher that this probability decreases exponentially in $\gamma \log t$. We formally describe the algorithms in Fig. 1.

Theorem 6 ([14]). *Let h_1, \dots, h_γ be t -wise independent hash functions, then for any $X = \{x_1, \dots, x_t\}$ it holds that*

$$\Pr [B := \text{Insert}((0^2)^{\gamma \times 8t}, X) : \text{List}(B) \neq X] \leq \mathcal{O}(2^{-(\gamma-2) \log t}),$$

² For the present discussion, we assume that only previously inserted elements are deleted.




Fig. 1. An invertible Bloom lookup table

where the probability is taken over the random choices of h_1, \dots, h_γ .

Remark 1. The construction of an IBLT can be modified to store tuples of values, by maintaining multiple valueSum matrices, one for each component. As long as one of the components remains unique among all inserted values, it is sufficient to use this component as input to the has functions, without affecting Theorem 6. We will make use of this in our construction in Sect. 5.

3 Ciphertext Compression

In this section we formally define the concept of a ciphertext compression scheme (Compress, Decompress). Intuitively, the compression algorithm takes the public encryption key pk as well as a vector of ciphertexts \mathbf{c} from some family \mathcal{F}_{pk} of ciphertext vectors as input and returns some compressed representation thereof. The decompression algorithm gets the compressed representation as well as the secret decryption key as input and should return the decryption of \mathbf{c} .

Definition 7 (Ciphertext Compression Scheme). Let $\mathcal{E} = (\text{Gen}, \text{Enc}, \text{Eval}, \text{Dec})$ be a homomorphic public key encryption scheme with ciphertext size $\xi = \xi(\lambda)$. Let \mathcal{P} be the public key space of \mathcal{E} . For each $\text{pk} \in \mathcal{P}$ let \mathcal{F}_{pk} be a set of ciphertext vectors. A δ -compressing, $(1 - \epsilon)$ -correct ciphertext compression scheme for the family $\mathcal{F} := \{\mathcal{F}_{\text{pk}} \mid \text{pk} \in \mathcal{P}\}$ is a pair of PPT algorithms (Compress, Decompress), such that for any $(\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^\lambda)$ and any $\mathbf{c} \in \mathcal{F}_{\text{pk}}$ the output length of $\text{Compress}(\text{pk}, \mathbf{c})$ is at most $\delta\xi|\mathbf{c}|$ and it holds that

$$\Pr[\text{Decompress}(\text{sk}, \text{Compress}(\text{pk}, \mathbf{c})) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] = 1 - \epsilon(\lambda),$$

where the probability is taken over the random coins of the compression and decompression algorithms.

Remark 2. Note, that a ciphertext compression scheme gives no guarantee whatsoever in the case where $c \notin \mathcal{F}_{pk}$.

4 Compression via Sparse Polynomials

In this section we present our first construction, which is based on the idea of interpolating sparse polynomials. Given the right building blocks, the construction is conceptually very simple. We simply view the sparse encrypted vector (c_1, \dots, c_n) as the coefficient representation of sparse polynomial. Using the Fast Fourier Transform, we homomorphically evaluate this polynomial efficiently at some sufficient number of points. These encrypted evaluations will constitute the compression of the vector. To obtain the original vector during decompression, we simply decrypt the evaluation points and interpolate the corresponding sparse polynomial.

Definition 8 (Fast Fourier Functions). *The class of fast fourier functions is the set of functions $\mathcal{Z}_{FFT}^\ell = \{f_x^\ell \mid \mathbf{x} \in \mathbb{F}_q^\ell\}$ with*

$$f_x^\ell : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^\ell, \quad f_x(\mathbf{a}) := \text{FFT}(\mathbf{a}, \mathbf{x}).$$

Definition 9. ($\mathcal{Z}_{FFT}^{2(t+1)}$ -Valid Low Hamming Weight Ciphertext Vectors). *Let $\mathcal{E} = (\text{Gen}, \text{Enc}, \text{Eval}, \text{Dec})$ be a homomorphic public key encryption scheme. For any $pk \in \mathcal{P}$, let*

$$\mathcal{F}_{t,pk}^{FFT} := \{c \in \text{vld}(\mathcal{Z}_{FFT}^{2(t+1)}, pk) \mid \text{hw}(\text{Dec}(\text{Gen}^{-1}(pk), c)) < t\}.$$

We then define the family of $\mathcal{Z}_{FFT}^{2(t+1)}$ -valid ciphertext vectors with low hamming weight as $\mathcal{F}_t^{FFT} := \{\mathcal{F}_{t,pk}^{FFT} \mid pk \in \mathcal{P}\}$.




Fig. 2. A ciphertext compression scheme for \mathcal{F}_t^{FFT} based on sparse polynomials. Here $\text{FFT}(\cdot, (\omega^0, \dots, \omega^{2t+1}))$ refers to the circuit of the function $f_{(\omega^0, \dots, \omega^{2t+1})}$ from Definition 8, i.e., the FFT circuit with the hardcoded second input $(\omega^0, \dots, \omega^{2t+1})$.

Theorem 10. *Let $\mathcal{E} = (\text{Gen}, \text{Enc}, \text{Eval}, \text{Dec})$ be an additively homomorphic encryption scheme with message space $\mathcal{M} = \mathbb{F}_q$ with ciphertext size $\xi = \xi(\lambda)$. Let $n, t \in \mathbb{N}$ be integers such that $n < q$ and let $\omega \in \mathbb{F}_q$ be a $2(t+1)$ -th primitive root of unity. Then $(\text{Compress}, \text{Decompress})$ from Fig. 2 is a $2(t+1)/n$ -compressing perfectly correct ciphertext compression scheme for family $\mathcal{F}_t^{\text{FFT}}$.*

Proof. Let \mathbf{c} be an arbitrary, but fixed $\mathcal{Z}_{\text{FFT}}^{2(t+1)}$ -valid ciphertext vector and let S be an arbitrary vector in sparse representation. Due to the validity condition on \mathbf{c} we know that

$$\begin{aligned} & \Pr \left[\begin{array}{l} \tilde{\mathbf{c}} \leftarrow \text{Eval}(\text{pk}, \text{FFT}(\cdot, (\omega^0, \dots, \omega^{2t+1})), \mathbf{c}) \\ \mathbf{m} \leftarrow \text{Dec}(\text{sk}, \tilde{\mathbf{c}}) \\ S = \text{Interpolate}(\mathbf{m}) \end{array} \right] \\ &= \Pr \left[\begin{array}{l} \mathbf{m} \leftarrow \text{Dec}(\text{sk}, \mathbf{c}) \\ \mathbf{m}' \leftarrow \text{FFT}(\mathbf{m}, (\omega^0, \dots, \omega^{2t+1})) \\ S = \text{Interpolate}(\mathbf{m}') \end{array} \right]. \end{aligned}$$

Furthermore, the $\mathcal{Z}_{\text{FFT}}^{2(t+1)}$ -validity of \mathbf{c} tells us that $\text{Dec}(\text{sk}, \mathbf{c})$ is a vector of hamming weight at most t or, when viewed as a polynomial in coefficient representation, a t -sparse polynomial of degree at most n . From Theorem 4 it follows this sparse polynomial can be correctly interpolated from its $2(t+1)$ evaluations produced by $\text{FFT}(\text{Dec}(\text{sk}, \mathbf{c}), (\omega^0, \dots, \omega^{2t+1}))$ and therefore

$$\text{Interpolate}(\text{FFT}(\text{Dec}(\text{sk}, \mathbf{c}), (\omega^0, \dots, \omega^{2t+1}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c})).$$

The output of Compress is a vector of $2(t+1)$ ciphertexts of size ξ and thus the scheme is $2(t+1)/n$ compressing.

5 Compression via IBLTs

In this section we present our second construction, which is on a variant of invertible Bloom lookup tables. Given a vector of ciphertexts \mathbf{c} the idea is to homomorphically insert the corresponding non-zero plaintexts into an (encrypted) IBLT. The encrypted IBLT would then constitute the compression of the vector.

This approach encounters two problems: First, the insertion operation of an IBLT requires hashing the value to choose the cells to insert it in, which we cannot do because we do not have access to the plaintext. Second, since we do not *know* which of the ciphertexts correspond to a non-zero it is unclear how to only insert those.

The first problem can be solved using Remark 1 by actually storing pairs (d, m_d) . Since the index d is both publically known and unique, we can rely on only hashing d to derive the positions to insert the values.

The second problem is a bit trickier to solve. To build some intuition, we can first consider an easier compression problem where in addition to \mathbf{c} we are given

an additional vector of ciphertexts \mathbf{h} containing “zero hints”. I.e., h_d decrypts to 0 if c_i decrypts to 0 and h_d decrypts to 1, if c_d decrypts to anything else.

The IBLT then gets initialized as three matrices, a count matrix \mathbf{C} and two valueSum matrices \mathbf{M} and \mathbf{D} with each cell containing an encryption of zero. To insert the content of ciphertext c_d into the IBLT, we can then for $i \in [\gamma]$ compute $j := h_i(d)$ and insert the value by setting $\mathbf{C}[i, j] := \mathbf{C} \boxplus h_d$, $\mathbf{M}[i, j] := \mathbf{M} \boxplus c_d$, and $\mathbf{D}[i, j] := \mathbf{D} \boxplus (d \boxtimes h_d)$.

Note that for any ciphertext corresponding to zero, this results in zero being added to all entries, which is equivalent to not inserting the value at all. Decompressing then involves decrypting all three matrices and using the List algorithm to extract pairs (d, m_d) giving us a sparse representation of the plaintext vector corresponding to \mathbf{c} . By the correctness guarantee of an IBLT, this works as long as not too many ciphertexts decrypt to a non-zero value.

However, actually getting such “zero hint” ciphertexts may not be feasible in all scenarios, especially if the encryption scheme is *only* additively homomorphic. This means we need to somehow simulate having a count matrix without these zero hints.

The trick that we use is to choose a vector of random values \mathbf{k} that we will use to “recognize” cells that only contain a single message. We will still initialize two matrices \mathbf{M} and \mathbf{K} but inserting into the IBLT is now done by setting $\mathbf{M}[i, j] := \mathbf{M} \boxplus c_d$, and $\mathbf{K}[i, j] := \mathbf{K} \boxplus (k_d \boxtimes c_d)$. Note now, that after decryption, for any cell (i, j) that only contains a single value m_d , we have that $\mathbf{M}[i, j] = m_d$ and $\mathbf{K}[i, j] = k_d \cdot m_d$. By checking if $\mathbf{K}[i, j]/\mathbf{M}[i, j]$ corresponds to one of the values in \mathbf{k} , we can thus recognize which cells contain only a single value and which index it corresponds to, allowing us to peel the message from the IBLT. In section we prove in a helpful lemma in Sect. 5.2 we prove that we can bound the probability that this recognition procedure produces false positives.

There still remains the problem that simply using a random vector \mathbf{k} and storing it, which would require $\mathcal{O}(n)$ storage and $\mathcal{O}(n)$ computation to recognize the entries. To solve this issue we introduce the concept of wunderbar pseudorandom vectors in Sect. 5.1, which allows us to store a compact $\mathcal{O}(\lambda)$ representation of a pseudorandom vector \mathbf{k} and recognition of vector entries in time $\mathcal{O}(\text{polylog}(n))$.

5.1 Wunderbar Pseudorandom Vectors

The concept of a pseudorandom vector is conceptually similar to that of pseudorandom sets introduced in [10], except that we do not require puncturability. The idea is that it allows us to sample a short description of a long vector, which is indistinguishable from a random vector with unique entries. Importantly, we require that there exists an efficient algorithm that can recover the position of a given entry in the vector in time independent of the vector length. Naively one can always find the position in linear time in the vector length. This is, however, not good enough for our application, which is why we require the pseudorandom vector to be “wunderbar”. In particular, we want the description length of the vector to be in $\mathcal{O}(\lambda)$ and getting individual entries as well as index recovery should be possible in $\mathcal{O}(\text{polylog}(n))$.

Definition 11. A pseudorandom vector with index recovery for an efficiently sampleable universe $K = K(\lambda)$ consists of a triple of ppt algorithms (Sample, Entry, Index) such that

Sample($1^\lambda, n$): The sampling algorithm takes as input the security parameter λ and the vector length n in unary and outputs the description of a pseudorandom vector s .

Entry(s, i): The deterministic retrieving algorithm takes as input a description s and an index $i \in [n]$ and outputs a value $k_i \in K$.

Index(s, k): The deterministic index recovery algorithm takes as input a description s and a value k and outputs either an index $i \in [n]$ or \perp .

A pseudorandom vector with index recovery is correct, if for all vector lengths $n = \text{poly}(\lambda)$ and all seeds $s \leftarrow \text{Gen}(1^\lambda, 1^n)$ it holds that:

1. For all indices $i \in [n]$ it holds that $\text{Index}(s, \text{Entry}(s, i)) = i$.
2. For all $k^* \notin \{\text{Entry}(s, i) \mid i \in [n]\}$ it holds that $\text{Index}(s, k^*) = \perp$.

The pseudorandom vector is wunderbar if the description of a vector has length $\mathcal{O}(\lambda)$ and the runtime of Entry and Index is $\mathcal{O}(\text{polylog}(n))$. A pseudorandom vector is secure, if for all $n = \text{poly}(\lambda)$ and all ppt algorithms \mathcal{A}

$$\left| \Pr \left[\begin{array}{l} s \leftarrow \text{Sample}(1^\lambda, 1^n), \\ \mathbf{k} := \begin{pmatrix} \text{Entry}(s, 1) \\ \vdots \\ \text{Entry}(s, n) \end{pmatrix} : \mathcal{A}(\mathbf{k}) \end{array} \right] - \Pr[\mathbf{k} \leftarrow \mathfrak{K}(K^n) : \mathcal{A}(\mathbf{k})] \right| \leq \text{negl}(\lambda)$$

Remark 3. For ease of notation we define two algorithms DummySample and DummyIndex that represent a dummy version of a pseudorandom vector with index recovery. I.e., DummySample($1^\lambda, n$) simply samples $\mathbf{k} \leftarrow \mathfrak{K}(K^n)$ and DummyIndex(\mathbf{k}, k) performs an exhaustive search and returns i iff $k_i = k$ and \perp if none of them match. Using this notation, the above security definition can be rewritten as

$$\left| \Pr \left[\begin{array}{l} s \leftarrow \text{Sample}(1^\lambda, 1^n), \\ \mathbf{k} := \begin{pmatrix} \text{Entry}(s, 1) \\ \vdots \\ \text{Entry}(s, n) \end{pmatrix} : \mathcal{A}(\mathbf{k}) \end{array} \right] - \Pr[\mathbf{k} \leftarrow \text{DummySample}(1^\lambda, n) : \mathcal{A}(\mathbf{k})] \right| \leq \text{negl}(\lambda)$$

Wunderbar Pseudorandom Vectors from Pseudorandom Permutations.

Let \mathbb{F}_{p^m} be a field such that $m \cdot \lceil \log p \rceil \geq \lambda$. We construct wunderbar pseudorandom vectors over a subset $K \subseteq \mathbb{F}_{p^m}$ from an arbitrary family of pseudorandom permutations over \mathbb{F}_2^λ . To do so we need an efficiently computable and

efficiently invertible injective function `binToField` mapping from \mathbb{F}_2^λ to \mathbb{F}_{p^m} . The exact function is irrelevant, but for concreteness, we specify it in the following. Let $\{0, 1\} : [q] \rightarrow \mathbb{F}_2^{\lceil \log q \rceil}$ denote the function that maps an integer to its canonical binary representation and let `proj` : $\mathbb{F}_2^{\lceil \log q \rceil} \rightarrow [q]$ be its inverse. Then we specify

$$\text{binToField} : \mathbb{F}_2^\lambda \rightarrow \mathbb{F}_{p^m} \quad \text{binToField}((b_1, \dots, b_\lambda)) := \sum_{i=0}^{m-1} c_i x^i$$

where

$$c_i := \sum_{j=1}^{\min\{\lceil \frac{\lambda}{m} \rceil, \lambda - i \lceil \frac{\lambda}{m} \rceil\}} 2^{j-1} b_{i \lceil \frac{\lambda}{m} \rceil + j}.$$

We further specify the inverse function as

$$\text{fieldToBin} : \mathbb{F}_{p^m} \rightarrow \mathbb{F}_2^\lambda \cup \{\perp\}$$

$$\text{fieldToBin}\left(\sum_{i=0}^{m-1} c_i x^i\right) := \begin{cases} \perp & \text{if } \exists c_i. c_i \geq 2^{\min\{\lceil \frac{\lambda}{m} \rceil, \lambda - i \lceil \frac{\lambda}{m} \rceil\}} \\ (b_1, \dots, b_\lambda) & \text{otherwise} \end{cases}$$

where

$$b_i := \text{bin}(c_{\lfloor i/\lceil \lambda/m \rceil \rfloor})_{i - \lfloor i/\lceil \lambda/m \rceil \rfloor \cdot \lceil \lambda/m \rceil}.$$

<code>Sample</code> ($1^\lambda, 1^n$)	<code>Entry</code> ((s, n), i)	<code>Index</code> ((s, n), k)
$s \leftarrow \mathbb{F}_2^\lambda$	$\mathbf{b} := \text{bin}(i)$	$\mathbf{b}' := \text{fieldToBin}(k)$
return (s, n)	$\mathbf{b}' := \text{PRP}(s, \mathbf{b})$	if $\mathbf{b}' \neq \perp$
	return <code>binToField</code> (\mathbf{b}')	$\mathbf{b} := \text{PRP}^{-1}(s, \mathbf{b}')$
		if <code>proj</code> (\mathbf{b}) $\in [n]$
		return <code>proj</code> (\mathbf{b})
		return \perp

Fig. 3. A wunderbar pseudorandom vector for $K \subseteq \mathbb{F}_{p^m}$ constructed from a family of pseudorandom permutations over \mathbb{F}_2^λ .

Theorem 12. *Let PRP be a secure family of pseudorandom permutations over some \mathbb{F}_2^λ . Then (Sample, Entry, Index) as described in Fig. 3 is a secure wunderbar pseudorandom vector with index recovery for universe $K = \{\text{binToField}(\mathbf{b}) \mid \mathbf{b} \in \mathbb{F}_2^\lambda\}$.*

Proof. We need to establish that the construction is correct, wunderbar, and secure. It is simple to see that the construction is correct:

$$\begin{aligned}
 & \text{Index}((s, n), \text{Entry}((s, n), i)) \\
 &= \text{Index}((s, n), \text{binToField}(\text{PRP}(s, \text{bin}(i)))) && \text{(Def. of Entry)} \\
 &= \text{proj}(\text{PRP}^{-1}(s, \text{fieldToBin}(\text{binToField}(\text{PRP}(s, \text{bin}(i))))) && \text{(Def. of Index)} \\
 &= \text{proj}(\text{PRP}^{-1}(s, \text{PRP}(s, \text{bin}(i)))) \\
 &= \text{proj}(\text{bin}(i)) \\
 &= i.
 \end{aligned}$$

Similarly, it is easy to see that the construction is wunderbar: the description consists of $s \in \mathbb{F}_2^\lambda$ and $n = \text{poly}(\lambda) \leq \lambda^c$ for some constant c . Therefore, it has size at most $\lambda + c \cdot \log \lambda \in \mathcal{O}(\lambda)$. The runtime of **Entry** is in fact independent of n and thus trivially in $\mathcal{O}(\text{polylog}(n))$ and the only computation in **Index** that depends on n , is the membership check $\text{proj}(\mathbf{b}) \in [n]$ which can be performed in time $\mathcal{O}(\log n) \subset \mathcal{O}(\text{polylog}(n))$.

It remains to show that the construction is secure. Let $n = n(\lambda) = \text{poly}(\lambda)$ and let \mathcal{A} be an arbitrary PPT algorithm, such that We construct an adversary \mathcal{B} against the pseudorandomness of as follows. \mathcal{B} takes as input the security parameter λ and is given access to an oracle. For each $i \in [n]$, query $\text{bin}(i)$ to the oracle, receiving back \mathbf{b}'_i and compute $k_i := \text{binToField}(\mathbf{b}'_i)$. Invoke $\mathcal{A}(\mathbf{k})$ and output whatever \mathcal{A} outputs. Clearly, \mathcal{B} is also PPT, needing a runtime overhead of just n oracle queries over simply running \mathcal{A} . We now consider two cases: if, on the one hand, the oracle is $\text{PRP}(s, \cdot)$, then for all $i \in [n]$ $k_i = \text{binToField}(\text{PRP}(s, \text{bin}(i))) = \text{Entry}((s, n), i)$. I.e., we have

$$\begin{aligned}
 & \Pr[s \leftarrow \mathbb{F}_2^\lambda : \mathcal{B}^{\text{PRP}(s, \cdot)} = 1] \\
 &= \Pr[s \leftarrow \text{Sample}(1^\lambda, 1^n), \mathbf{k} := (\text{Entry}(s, 1), \dots, \text{Entry}(s, n))^\top : \mathcal{A}(\mathbf{k})]. \tag{1}
 \end{aligned}$$

If, on the other hand, the oracle is a truly random permutation g , then for all $i \in [n]$ it holds that $k_i = \text{binToField}(g(\text{bin}(i)))$ and therefore

$$\begin{aligned}
 & \Pr[g \leftarrow \Pi(\mathbb{F}_2^\lambda) : \mathcal{B}^{g(\cdot)} = 1] \\
 &= \Pr[g \leftarrow \Pi(\mathbb{F}_2^\lambda); \forall i \in [n]. k_i = \text{binToField}(g(\text{bin}(i))) : \mathcal{A}(\mathbf{k})] \tag{2} \\
 &= \Pr[(\mathbf{b}'_1, \dots, \mathbf{b}'_n) \leftarrow \mathfrak{R}((\mathbb{F}_2^\lambda)^n); \mathbf{k} = (\text{binToField}(\mathbf{b}'_i))_{i \in [n]} : \mathcal{A}(\mathbf{k})] \tag{3} \\
 &= \Pr[\mathbf{k} \leftarrow \mathfrak{R}(K^n) : \mathcal{A}(\mathbf{k})]. \tag{4}
 \end{aligned}$$

Here, Eq. 3 holds because g is a uniformly chosen random permutation and therefore the values $g(\text{bin}(i))$ are uniformly distributed conditioned on not being duplicates and Eq. 4 holds because binToField is an injective function into K .

Combining Eq. 1 and Eq. 4 we get

$$\begin{aligned} & \left| \Pr \left[\begin{array}{l} s \leftarrow \text{Sample}(1^\lambda, 1^n), \\ \mathbf{k} := \begin{pmatrix} \text{Entry}(s, 1) \\ \vdots \\ \text{Entry}(s, n) \end{pmatrix} : \mathcal{A}(\mathbf{k}) \end{array} \right] - \Pr[\mathbf{k} \leftarrow \mathfrak{K}(K^n) : \mathcal{A}(\mathbf{k})] \right| \\ &= \left| \Pr[s \leftarrow \mathbb{F}_2^\lambda : \mathcal{B}^{\text{PRP}(s, \cdot)} = 1] - \Pr[g \leftarrow \{f : \mathbb{F}_2^\lambda \rightarrow \mathbb{F}_2^\lambda\} : \mathcal{B}^{g(\cdot)} = 1] \right| \\ &\leq \text{negl}(\lambda) \end{aligned}$$

where the last inequality follows from the fact that PRP is pseudorandom. \square

5.2 A Helpful Lemma

We prove a helpful lemma which allows to bound the probability of false positives when attempting to detect cells with only a single entry in the IBLT. Recall, that we have two matrices \mathbf{M} and \mathbf{K} , where the cells of \mathbf{M} contain sums of messages m_d and the cells of \mathbf{K} contain sums of $k_d \cdot m_d$ for a random vector \mathbf{k} . We check for cells containing only a single message, i.e. cells that can be peeled, by checking whether $\mathbf{K}[i, j]/\mathbf{M}[i, j]$ corresponds to one of the values in \mathbf{k} . A false positive could occur, if for some set of at least two non-zero messages corresponding to indices $I \subseteq [n]$ it happens to hold that

$$k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i}$$

for some $j \in [n]$. The lemma states that we can bound the probability of this occurring by choosing the entries of \mathbf{k} from a large enough space.

Lemma 13. *Let $K \subseteq \mathbb{F}_q$, $(m_1, \dots, m_n) \in \mathbb{F}_q^n$ and $I \subseteq [n]$ be arbitrary such that $\sum_{i \in I} m_i \neq 0$ and there exist $i, i' \in I$ with $0 \notin \{m_i, m_{i'}\}$. It holds that*

$$\Pr \left[\mathbf{k} \leftarrow K^n : \exists j \in [n]. k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i} \right] \leq \frac{n}{|K|}$$

Proof. Using a union bound we have

$$\begin{aligned} & \Pr \left[\mathbf{k} \leftarrow K^n : \exists j \in [n]. k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i} \right] \\ &\leq \sum_{j \in [n]} \Pr \left[\mathbf{k} \leftarrow K^n : k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i} \right]. \end{aligned} \tag{5}$$

It thus remains to bound the above probability for individual j . Let $j \in [n]$ be arbitrary but fixed and let $\xi = 1$ if $j \in I$ and $\xi = 0$ otherwise. It then holds that

$$\begin{aligned} & \Pr \left[\mathbf{k} \leftarrow K^n : k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i} \right] \\ &= \Pr \left[\mathbf{k} \leftarrow K^n : k_j \cdot \sum_{i \in I} m_i = \sum_{i \in I} k_i m_i \right] \\ &= \Pr \left[\mathbf{k} \leftarrow K^n : k_j \cdot \left(\sum_{i \in I} m_i - \xi m_j \right) = \sum_{i \in I} k_i m_i - k_j \xi m_j \right] \\ &= \Pr \left[\mathbf{k} \leftarrow K^n : k_j \cdot \sum_{i \in I \setminus \{j\}} m_i = \sum_{i \in I \setminus \{j\}} k_i m_i \right] \end{aligned}$$

We now consider two cases. If $\sum_{i \in I \setminus \{j\}} m_i = 0$, let $j' \in I \setminus \{j\}$ be an index, such that $m_{j'} \neq 0$. Note that such an index always exists by the condition on I . We then have

$$\begin{aligned} & \Pr \left[\mathbf{k} \leftarrow K^n : k_j \cdot \overbrace{\sum_{i \in I \setminus \{j\}} m_i}^{=0} = \sum_{i \in I \setminus \{j\}} k_i m_i \right] \\ &= \Pr \left[\mathbf{k} \leftarrow K^n : 0 = \sum_{i \in I \setminus \{j\}} k_i m_i \right] \\ &= \Pr \left[\mathbf{k} \leftarrow K^n : k_{j'} = \frac{\sum_{i \in I \setminus \{j, j'\}} k_i m_i}{-m_{j'}} \right] \end{aligned}$$

where $(-m_{j'})^{-1}$ is always defined by the condition that $m_{j'} \neq 0$. Since the right hand side of the equality is independent of $k_{j'}$, the probability that the equality holds is at most $1/|K|$ for any choice of $k_i, i \neq j'$. Thus, in this case

$$\Pr \left[\mathbf{k} \leftarrow K^n : k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i} \right] \leq \frac{1}{|K|}. \tag{6}$$

In the other case, i.e., if $\sum_{i \in I \setminus \{j\}} m_i \neq 0$, $(\sum_{i \in I \setminus \{j\}} m_i)^{-1}$ is well defined and we have

$$\Pr \left[\mathbf{k} \leftarrow K^n : k_j \cdot \sum_{i \in I \setminus \{j\}} m_i = \sum_{i \in I \setminus \{j\}} k_i m_i \right] = \Pr \left[\mathbf{k} \leftarrow K^n : k_j = \frac{\sum_{i \in I \setminus \{j\}} k_i m_i}{\sum_{i \in I \setminus \{j\}} m_i} \right].$$

Here again, the right hand side of the equality is independent of k_j . Thus, the probability that the equality holds is at most $1/|K|$ for any choice of $k_i, i \neq j$ and also in this case it holds that

$$\Pr \left[\mathbf{k} \leftarrow K^n : k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i} \right] \leq \frac{1}{|K|} \tag{7}$$

Finally, combining Eq. 5 with Eq. 6 and Eq. 7, we get

$$\Pr \left[\mathbf{k} \leftarrow K^n : \exists j \in [n]. k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i} \right] \leq \frac{n}{|K|} \quad \square$$

By observing that the statistical distance between K^n and $\mathfrak{K}(K^n)$ is at most $n^2/|K|$ due to the birthday bound, we obtain the following corollary.

Corollary 14. *Let $K \subseteq \mathbb{F}_q$, $(m_1, \dots, m_n) \in \mathbb{F}_q^n$ and $I \subseteq [n]$ be arbitrary such that $\sum_{i \in I} m_i \neq 0$ and there exist $i, i' \in I$ with $0 \notin \{m_i, m_{i'}\}$. It holds that*

$$\Pr \left[\mathbf{k} \leftarrow \mathfrak{K}(K^n) : \exists j \in [n]. k_j = \frac{\sum_{i \in I} k_i m_i}{\sum_{i \in I} m_i} \right] \leq \frac{n^2 + n}{|K|}$$

5.3 Construction of Ciphertext-Compression from IBLTs

Populating the IBLT involves homomorphically evaluating an inner product between the encrypted vector and a plain vector. Therefore, the compression scheme can only work for ciphertext vectors that allow the evaluation of inner product functions defined in the following.

Definition 15 (Inner Product Functions). *The class of inner product functions is the set of functions $\mathcal{Z}_{\text{ip}} = \{f_{\mathbf{a}} \mid \mathbf{a} \in \mathbb{F}_q^n\}$ with*

$$f_{\mathbf{a}} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q, \quad f_{\mathbf{a}}(\mathbf{x}) := \langle \mathbf{a}, \mathbf{x} \rangle.$$

The family of ciphertext vectors the construction is applicable to is then exactly those ciphertext vectors with low hamming weight and allow the evaluation of inner product functions. We define this family as follows.

Definition 16. (\mathcal{Z}_{ip} -Valid Low Hamming Weight Ciphertext Vectors). *Let $\mathcal{E} = (\text{Gen}, \text{Enc}, \text{Eval}, \text{Dec})$ be a homomorphic public key encryption scheme. For any $\text{pk} \in \mathcal{P}$, let*

$$\mathcal{F}_{t, \text{pk}}^{\text{ip}} := \{ \mathbf{c} \in \text{vld}(\mathcal{Z}_{\text{ip}}, \text{pk}) \mid \text{hw}(\text{Dec}(\text{Gen}^{-1}(\text{pk}), \mathbf{c})) < t \}.$$

We then define the family of \mathcal{Z}_{ip} -valid ciphertext vectors with low hamming weight as $\mathcal{F}_t^{\text{ip}} := \{ \mathcal{F}_{t, \text{pk}}^{\text{ip}} \mid \text{pk} \in \mathcal{P} \}$.

Theorem 17. *Let $\mathcal{E} = (\text{Gen}, \text{Enc}, \text{Eval}, \text{Dec})$ be an additively homomorphic encryption scheme with message space $\mathcal{M} = \mathbb{F}_q$ for some prime power q with ciphertext size $\xi = \xi(\lambda)$. Let $\lambda, \kappa, t, n \in \mathbb{N}$ be integers and let $\gamma := \lceil \frac{\kappa}{\log t} \rceil + 2$. Let PRF be a family of pseudorandom functions $\text{PRF} : [\gamma] \times [2^\lambda] \rightarrow [8t]$ and let $(\text{Sample}, \text{Entry}, \text{Index})$ be a wunderbar pseudorandom vector with index recovery for a universe $K = K(\lambda) \subseteq \mathbb{F}_q$ with $|K| \geq 2^\kappa (8t\gamma)(n^3 + n^2)$. Then $(\text{Compress}, \text{Decompress})$ from Fig. 4 is a $(\mathcal{O}(\lambda) + 16t\gamma\xi)/(n\xi)$ -compressing $(1 - \mathcal{O}(2^{-\kappa}) - \text{negl}(\lambda))$ -correct ciphertext compression scheme for family $\mathcal{F}_t^{\text{ip}}$.*

Proof. The output of the compression algorithm consists of a s_1 , s_2 and $16t\gamma$ ciphertexts. Since the pseudorandom vector is wunderbar, it holds that $|s_1| = \mathcal{O}(\lambda)$ and s_2 is chosen as a λ -bit string. Therefore, it is easy to see that the scheme is $(\mathcal{O}(\lambda) + 16t\gamma\xi)/(n\xi)$ -compressing. It remains to prove that it is correct. To do so we define a series of six hybrid schemes in Figs. 5 through 9.

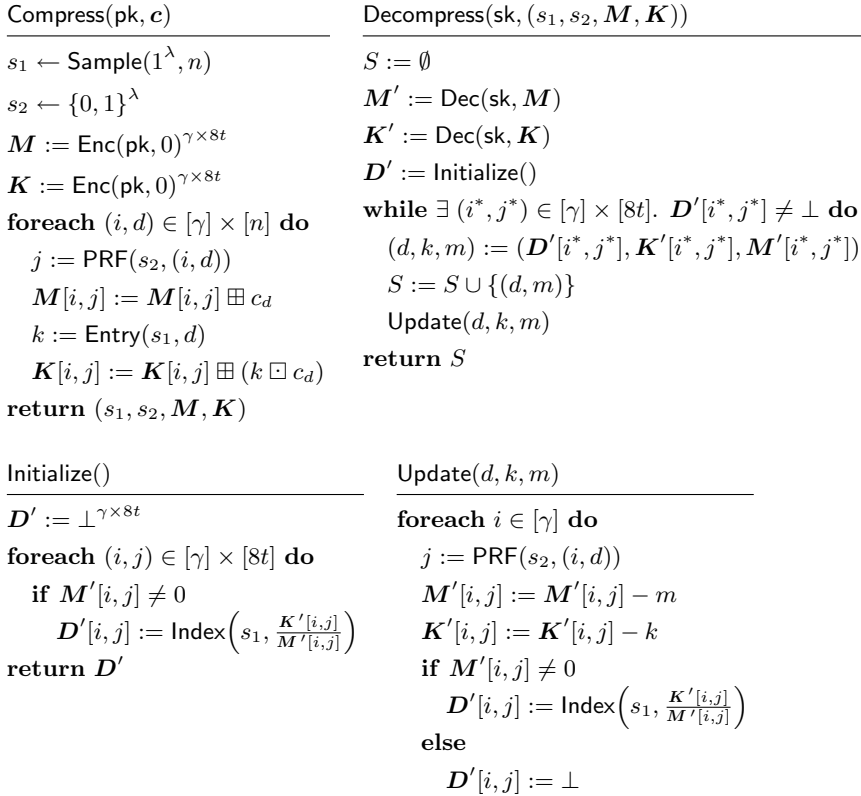


Fig. 4. A ciphertext compression scheme based on invertible bloom lookup tables and wunder pseudorandom vectors.

Claim 18. For any \mathcal{Z}_{ip} -valid vector of ciphertexts it holds that

$$\Pr[\text{Decompress}(\text{sk}, \text{Compress}(\text{pk}, \mathbf{c})) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \\ = \Pr[\text{Decompress}_1(\text{Compress}_1(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))]$$

Proof. Following Definition 15 we denote by f_a the inner product function $f_a : \mathbb{F}_q^n \rightarrow \mathbb{F}_q, f_a(\mathbf{x}) := \langle \mathbf{a}, \mathbf{x} \rangle$. Further, we denote

$$\mathbf{v}_{i,j} := \begin{pmatrix} \delta_{j, \text{PRF}(s_2, (i, 1))} \\ \vdots \\ \delta_{j, \text{PRF}(s_2, (i, n))} \end{pmatrix} \quad \mathbf{w}_{i,j} := \begin{pmatrix} \text{Entry}(s_1, 1) \cdot \delta_{j, \text{PRF}(s_2, (i, 1))} \\ \vdots \\ \text{Entry}(s_1, n) \cdot \delta_{j, \text{PRF}(s_2, (i, n))} \end{pmatrix}$$

Now, let M_0, K_0, M'_0, K'_0 and M_1, K_1, M'_1, K'_1 denote the relevant matrices in the actual scheme and hybrid 1 respectively. We note, that since \mathbf{c} is \mathcal{Z}_{ip} -valid,

Compress ₁ (m)	Decompress ₁ ((s_1, s_2, M, K))
$s_1 \leftarrow \text{Sample}(1^\lambda, n)$	$S := \emptyset$
$s_2 \leftarrow \{0, 1\}^\lambda$	$M' := M$
$M := 0^{\gamma \times 8t}$	$K' := K$
$K := 0^{\gamma \times 8t}$	$D' := \text{Initialize}()$
foreach $(i, d) \in [\gamma] \times [n]$ do	while $\exists (i^*, j^*) \in [\gamma] \times [8t]. D'[i^*, j^*] \neq \perp$ do
$j := \text{PRF}(s_2, (i, d))$	$(d, k, m) := (D'[i^*, j^*], K'[i^*, j^*], M'[i^*, j^*])$
$M[i, j] := M[i, j] + m_d$	$S := S \cup \{(d, m)\}$
$k := \text{Entry}(s_1, d)$	Update (d, k, m)
$K[i, j] := K[i, j] + (k \cdot m_d)$	return S
return (s_1, s_2, M, K)	

Fig. 5. The first hybrid scheme works exactly as the actual ciphertext compression scheme, except that it operates on *plaintext* messages instead of encrypted messages. I.e., ciphertexts are now decrypted before compression instead of between compression and decompression.

it holds for all $(i, j) \in [\gamma] \times [8t]$ that

$$\begin{aligned}
 M'_0[i, j] &= \text{Dec}(\text{sk}, M_0[i, j]) && \text{(Definition 3)} \\
 &= \text{Dec}(\text{sk}, \bigoplus_{d \in \{[n] \mid \text{PRF}(s_2, (i, d))=j\}} c_d) \\
 &= \text{Dec}(\text{Eval}(\text{pk}, f_{v_{i,j}}, \mathbf{c})) \\
 &= f_{v_{i,j}}(\text{Dec}(\text{sk}, \mathbf{c})) \\
 &= \sum_{d \in \{[n] \mid \text{PRF}(s_2, (i, d))=j\}} \text{Dec}(\text{sk}, c_d) = M'_1[i, j]
 \end{aligned}$$

as well as

$$\begin{aligned}
 K'_0[i, j] &= \text{Dec}(\text{sk}, K_0[i, j]) && \text{(Definition 3)} \\
 &= \text{Dec}(\text{sk}, \bigoplus_{d \in \{[n] \mid \text{PRF}(s_2, (i, d))=j\}} c_d \cdot \text{Entry}(s_1, d)) \\
 &= \text{Dec}(\text{Eval}(\text{pk}, f_{w_{i,j}}, \mathbf{c})) \\
 &= f_{w_{i,j}}(\text{Dec}(\text{sk}, \mathbf{c})) \\
 &= \sum_{d \in \{[n] \mid \text{PRF}(s_2, (i, d))=j\}} \text{Dec}(\text{sk}, c_d) \cdot \text{Entry}(s_1, d) = K'_1[i, j]
 \end{aligned}$$

Since the computation on M', K' is otherwise identical between the two hybrids, the claim immediately follows. □

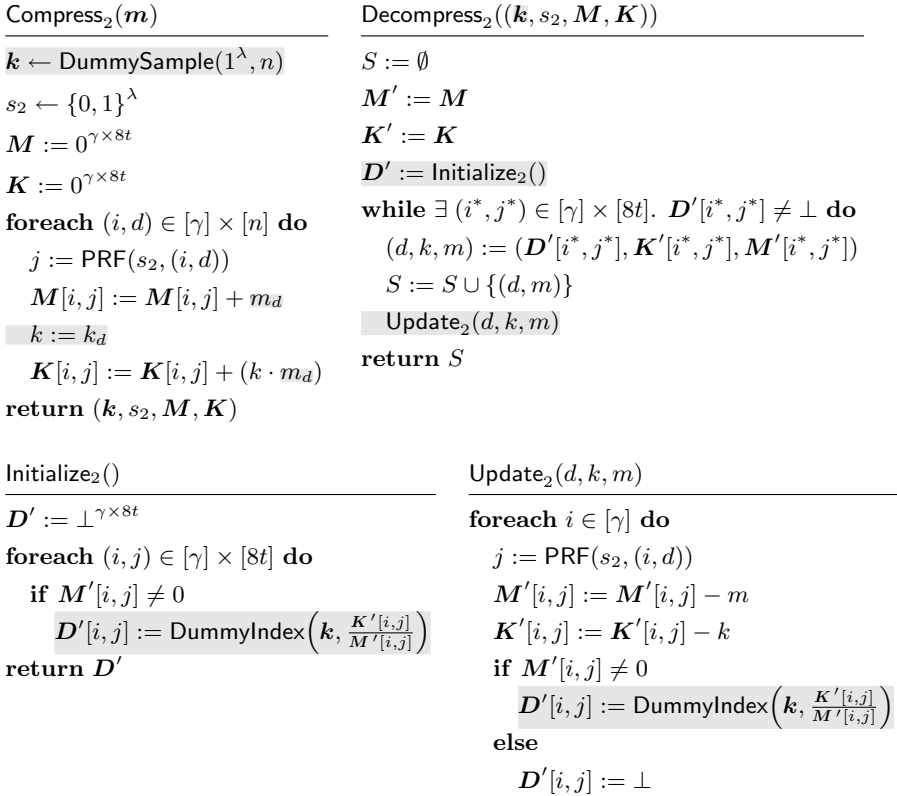


Fig. 6. The second hybrid scheme works exactly as the first hybrid scheme, except that instead of using the wunder pseudorandom vector it uses the dummy sampler and the dummy index recovery to work with a uniformly random vector $\mathbf{k} \in \mathfrak{K}^n$.

Claim 19. *If (Sample, Entry, Index) is a secure pseudorandom vector, it holds for any key pair (sk, pk) and any vector c that*

$$\left| \begin{aligned} & \Pr[\text{Decompress}_1(\text{Compress}_1(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \\ & - \Pr[\text{Decompress}_2(\text{Compress}_2(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \end{aligned} \right| \leq \text{negl}(\lambda).$$

Proof. We construct an attacker \mathcal{A} against security of the pseudorandom vector as follows. On input \mathbf{k} , \mathcal{A} executes $\text{Decompress}_2(\text{Compress}_2(\text{Dec}(\text{sk}, \mathbf{c})))$, except that it uses its input \mathbf{k} instead of sampling a fresh one. If \mathbf{k} was chosen using $\mathbf{k} \leftarrow \text{DummySample}(1^\lambda, n)$, this is identical to a regular execution of $\text{Decompress}_2(\text{Compress}_2(\text{Dec}(\text{sk}, \mathbf{c})))$. If on the other hand \mathbf{k} was chosen by sampling $s_2 \leftarrow \text{Sample}(1^\lambda, n)$ and setting $\mathbf{k} := (\text{Entry}(s, 1), \dots, \text{Entry}(s, n))^T$, this is identical to a regular execution of $\text{Decompress}_1(\text{Compress}_1(\text{Dec}(\text{sk}, \mathbf{c})))$. Therefore, by the security of the pseudorandom vector

$\text{Compress}_3(m)$	$\text{Decompress}_3((k, s_2, M, K, C))$
$k \leftarrow \text{DummySample}(1^\lambda, n)$ $s_2 \leftarrow \{0, 1\}^\lambda$ $M := 0^{\gamma \times 8t}$ $K := 0^{\gamma \times 8t}$ $C := 0^{\gamma \times 8t}$ foreach $(i, d) \in [\gamma] \times [n]$ do $j := \text{PRF}(s_2, (i, d))$ $M[i, j] := M[i, j] + m_d$ $k := k_d$ $K[i, j] := K[i, j] + (k \cdot m_d)$ if $m_d \neq 0$ do $C[i, j] := C[i, j] + 1$ return (k, s_2, M, K, C)	$S := \emptyset$ $M' := M$ $K' := K$ $D' := \text{Initialize}_2()$ while $\exists (i^*, j^*) \in [\gamma] \times [8t]. C[i^*, j^*] = 1$ do $(d, k, m) := (D'[i^*, j^*], K'[i^*, j^*], M'[i^*, j^*])$ $S := S \cup \{(d, m)\}$ $\text{Update}_3(d, k, m)$ return S <hr style="border: 0.5px solid black;"/> $\text{Update}_3(d, k, m)$ foreach $i \in [\gamma]$ do $j := \text{PRF}(s_2, (i, d))$ $M'[i, j] := M'[i, j] - m$ $K'[i, j] := K'[i, j] - k$ $C[i, j] := C[i, j] - 1$ if $M'[i, j] \neq 0$ $D'[i, j] := \text{DummyIndex}\left(k, \frac{K'[i, j]}{M'[i, j]}\right)$ else $D'[i, j] := \perp$

Fig. 7. The third hybrid scheme works exactly as the second hybrid scheme, except that it maintains a matrix counting how many non-zero messages are mapped to each individual cell and deciding which messages to peel based on these exact counts instead of relying on the matrix K .

$$\begin{aligned}
 \text{negl}(\lambda) &\geq \left| \Pr \left[s \leftarrow \text{Sample}(1^\lambda, n), \mathbf{k} := \begin{pmatrix} \text{Entry}(s, 1) \\ \vdots \\ \text{Entry}(s, n) \end{pmatrix} : \mathcal{A}(\mathbf{k}) \right] \right. \\
 &\quad \left. - \Pr[\mathbf{k} \leftarrow \text{DummySample}(1^\lambda, n) : \mathcal{A}(\mathbf{k})] \right| \\
 &= \left| \Pr[\text{Decompress}_1(\text{sk}, \text{Compress}_1(\text{pk}, \mathbf{c})) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \right. \\
 &\quad \left. - \Pr[\text{Decompress}_2(\text{Compress}_2(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \right|
 \end{aligned}$$

□

Claim 20. *It holds that*

$$\left| \begin{aligned} & \Pr[\text{Decompress}_2(\text{Compress}_2(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \\ & - \Pr[\text{Decompress}_3(\text{Compress}_3(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \end{aligned} \right| \leq 2^{-\kappa}.$$

Proof. We first note two things

1. Whenever a correct element (m, d) is peeled, the resulting matrices $(\mathbf{C}', \mathbf{M}', \mathbf{K}', \mathbf{D})$ are identical to the scenario where $m_d = \text{Dec}(\text{sk}, c_d) = 0$ and all other messages are unchanged.
2. In the third hybrid scheme only correct elements are peeled.

The first observation follows because a correct peeling removes a message from the relevant cells by subtracting the corresponding values, which is equivalent to not adding them in the first place, which is exactly what happens if the message is zero. The second observation follows because we correctly keep track of the number of non-zero elements in each cell and only peel those, where a single non-zero element remains. By these observations, at any point during the execution of the decompression loop, there exists a vector $\mathbf{m}' \in \mathbb{F}_q^n$, such that for all (i, j)

$$\mathbf{K}'[i, j] := \begin{cases} d & \text{if } \frac{\sum_{\iota \in I_i} k_\iota m_\iota}{\sum_{\iota \in I_i} m_\iota} = k_d \\ \perp & \text{otherwise} \end{cases}$$

where $I_i = \{\iota \in [n] \mid \text{PRF}(s_1, (i, \iota)) = j\}$.

We denote by $E_{r,i,j}$ the event that before the r -th iteration of the main loop of Decompress_4 , it holds that $C[i, j] > 1$ but $K[i, j] \neq \perp$. Note that Decompress_4 and Decompress_3 behave identically unless at least one of $E_{r,i,j}$ for $(r, i, j) \in [n] \times [\gamma] \times [8t]$ occurs.

Therefore by a union bound and Corollary 14

$$\begin{aligned} & \left| \begin{aligned} & \Pr[\text{Decompress}_3(\text{sk}, \text{Compress}_3(\text{pk}, \mathbf{c})) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \\ & - \Pr[\text{Decompress}_4(\text{Compress}_4(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \end{aligned} \right| \\ & \leq \sum_{(r,i,j) \in [n] \times [\gamma] \times [8t]} \Pr[E_{r,i,j}] \\ & \leq \sum_{(r,i,j) \in [n] \times [\gamma] \times [8t]} \frac{n^2 + n}{|K|} \\ & = \frac{(8t\gamma)(n^3 + n^2)}{|K|} \leq \frac{(8t\gamma)(n^3 + n^2)}{2^\kappa \cdot (8t\gamma)(n^3 + n^2)} \leq 2^{-\kappa} \end{aligned}$$

□

Claim 21. *It holds that*

$$\left| \begin{aligned} & \Pr[\text{Decompress}_3(\text{Compress}_3(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \\ & = \Pr[\text{Decompress}_4(\text{Compress}_4(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \end{aligned} \right|.$$

Compress ₄ (m)	Decompress ₄ ($(\mathbf{k}, s_2, M, \mathbf{K}, C, D)$)
$\mathbf{k} \leftarrow \text{DummySample}(1^\lambda, n)$	$S := \emptyset$
$s_2 \leftarrow \{0, 1\}^\lambda$	$M' := M$
$M := 0^{\gamma \times 8t}$	$\mathbf{K}' := \mathbf{K}$
$\mathbf{K} := 0^{\gamma \times 8t}$	$D' := D$
$C := 0^{\gamma \times 8t}$	while $\exists (i^*, j^*) \in [\gamma] \times [8t]. C[i^*, j^*] = 1$ do
$D := 0^{\gamma \times 8t}$	$(d, k, m) := (D'[i^*, j^*], \mathbf{K}'[i^*, j^*], M'[i^*, j^*])$
foreach $(i, d) \in [\gamma] \times [n]$ do	$S := S \cup \{(d, m)\}$
$j := \text{PRF}(s_2, (i, d))$	Update ₄ (d, k, m)
$M[i, j] := M[i, j] + m_d$	return S
$k := k_d$	Update ₄ (d, k, m)
$\mathbf{K}[i, j] := \mathbf{K}[i, j] + (k \cdot m_d)$	foreach $i \in [\gamma]$ do
if $m_d \neq 0$ do	$j := \text{PRF}(s_2, (i, d))$
$C[i, j] := C[i, j] + 1$	$M'[i, j] := M'[i, j] - m$
$D[i, j] := D[i, j] + d$	$\mathbf{K}'[i, j] := \mathbf{K}'[i, j] - k$
return $(\mathbf{k}, s_2, M, \mathbf{K}, C, D)$	$C[i, j] := C[i, j] - 1$
	$D'[i, j] := D'[i, j] - d$

Fig. 8. The fourth hybrid scheme works exactly as the third hybrid scheme, except that the matrix D which before contained the indices of the messages *if* it could be inferred from the matrix K is now maintained with a sum of the indices of all messages mapped to the cell. This means that whenever $C[i, j] = 1$, $D[i, j]$ contains the index of the single non-zero message mapped to cell (i, j) .

Proof. The only difference between the two hybrids could occur, if when peeling a message from cell (i, j) , the content of $D[i, j]$ would differ between the two hybrids. However, this is not possible, since in hybrid three we have

$$\begin{aligned} D[i, j] &= \text{DummyIndex}(\mathbf{k}, \frac{\mathbf{K}[i, j]}{M[i, j]}) \\ &= \text{DummyIndex}(\mathbf{k}, \frac{k_d \cdot m_d}{m_d}) = \text{DummyIndex}(\mathbf{k}, k_d) = d \end{aligned}$$

just as in hybrid four. □

The fifth hybrid is identical to the fourth hybrid except that the now unnecessary matrix K is removed. The following claim trivially follows from the fact that K is not used in either hybrids.

Claim 22. *It holds that*

$$\left| \Pr[\text{Decompress}_4(\text{Compress}_4(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] - \Pr[\text{Decompress}_5(\text{Compress}_5(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \right|$$

$\text{Compress}_6^{h(\cdot)}(m)$	$\text{Decompress}_6^{h(\cdot)}((M, C, D))$
$M := 0^{\gamma \times 8t}$	$S := \emptyset$
$C := 0^{\gamma \times 8t}$	$M' := M$
$D := 0^{\gamma \times 8t}$	$D' := D$
foreach $(i, d) \in [\gamma] \times [n]$ do	while $\exists (i^*, j^*) \in [\gamma] \times [8t]. C[i^*, j^*] = 1$ do
$j := h(i, d)$	$(d, m) := (D'[i^*, j^*], M'[i^*, j^*])$
$M[i, j] := M[i, j] + m_d$	$S := S \cup \{(d, m)\}$
if $m_d \neq 0$ do	$\text{Update}_5(d, m)$
$C[i, j] := C[i, j] + 1$	return S
$D[i, j] := D[i, j] + d$	$\text{Update}_6(d, m)$
return (M, C, D)	foreach $i \in [\gamma]$ do
	$j := h(i, d)$
	$M'[i, j] := M'[i, j] - m$
	$C[i, j] := C[i, j] - 1$
	$D'[i, j] := D'[i, j] - d$

Fig. 9. The sixth hybrid scheme works exactly as the fifth hybrid scheme, except that instead of using a pseudorandom function to derive j from (i, d) , it uses a truly random function h given as an oracle.

Claim 23. *If PRF is a secure pseudorandom function then it holds that*

$$\left| \Pr[\text{Decompress}_5(\text{Compress}_5(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] - \Pr[\text{Decompress}_6^{h(\cdot, \cdot)}(\text{Compress}_6^{h(\cdot, \cdot)}(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \right| \leq \text{negl}(\lambda).$$

Proof. The only difference between the two hybrids is the use of the function $\text{PRF}(s_2, \cdot, \cdot)$ in the fifth hybrid and $h(\cdot, \cdot)$ in the sixth hybrid. Thus, the claim follows from a straightforward reduction that, given access to an oracle o , executes $\text{Decompress}_6^{o(\cdot, \cdot)}(\text{Compress}_6^{o(\cdot, \cdot)}(\text{Dec}(\text{sk}, \mathbf{c})))$ and outputs 0 if the result equals $\text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))$ and 1 otherwise. \square

Claim 24. *It holds for any vector of ciphertexts with hamming weight at most t that*

$$\Pr[\text{Decompress}_6(\text{Compress}_6(\text{Dec}(\text{sk}, \mathbf{c}))) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \leq \mathcal{O}(2^{-\kappa})$$

Proof. Denote $\mathbf{m} := \text{Dec}(\text{sk}, \mathbf{c})$, $S := \{(d, m') \mid m' = m_d\}$ and $S_{\neq 0} := \{(d, m') \in S \mid m' \neq 0\} = \text{sparse}(\mathbf{m})$. Since \mathbf{c} has Hamming weight at most t , we have that $|S_{\neq 0}| \leq t$.

Comparing hybrid six with the definition of an IBLT in Fig. 1, and keeping in mind Remark 1 we can observe, that what Compress_6 actually outputs is simply an IBLT for pairs containing all elements of $S_{\neq 0}$ using hash functions $h(i, \cdot)$. Further, Decompress_6 is in fact the same as List with Update_6 being identical to Peel . And since $|S_{\neq 0}| \leq t$ and random functions are t -wise independent, it thus holds by Theorem 6 that

$$\begin{aligned} & \Pr[\text{Decompress}_6(\text{Compress}_6(\mathbf{m})) = \text{sparse}(\mathbf{m})] \\ &= \Pr[\text{List}(\text{Insert}(B_0, S_{\neq 0})) = S_{\neq 0}] \geq 1 - \mathcal{O}(2^{-(\gamma-2)\log t}) \geq 1 - \mathcal{O}(2^{-\kappa}) \end{aligned}$$

□

By combining all of the above claims and using the triangle inequality it follows that

$$\begin{aligned} & \Pr[\text{Decompress}(\text{sk}, \text{Compress}(\text{pk}, \mathbf{c})) = \text{sparse}(\text{Dec}(\text{sk}, \mathbf{c}))] \\ & \geq 1 - \mathcal{O}(2^{-\kappa}) - \text{negl}(\lambda) \end{aligned}$$

as claimed. □

References

1. Akavia, A., Feldman, D., Shaul, H.: Secure search on encrypted data via multi-ring sketch. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018: 25th Conference on Computer and Communications Security, pp. 985–1001. ACM Press, Toronto (2018). <https://doi.org/10.1145/3243734.3243810>
2. Ben-Or, M., Tiwari, P.: A deterministic algorithm for sparse multivariate polynomial interpolation (extended abstract). In: 20th Annual ACM Symposium on Theory of Computing, pp. 301–309. ACM Press, Chicago (1988). <https://doi.org/10.1145/62212.62241>
3. Boneh, D., Di Crescenzo, G., Ostrovsky, R., Persiano, G.: Public key encryption with keyword search. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 506–522. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24676-3_30
4. Brakerski, Z., Gentry, C., Halevi, S.: Packed ciphertexts in LWE-based homomorphic encryption. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 1–13. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36362-7_1
5. Candès, E.J., Romberg, J., Tao, T.: Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. IEEE Trans. Inf. Theory **52**(2), 489–509 (2006). <https://doi.org/10.1109/TIT.2005.862083>

6. Cheon, J.H., Kim, M., Kim, M.: Optimized search-and-compute circuits and their application to query evaluation on encrypted data. *IEEE Trans. Inf. Forens. Secur.* **11**(1), 188–199 (2016). <https://doi.org/10.1109/TIFS.2015.2483486>
7. Cheon, J.H., Kim, M., Lauter, K.: Homomorphic computation of edit distance. In: Brenner, M., Christin, N., Johnson, B., Rohloff, K. (eds.) *FC 2015*. LNCS, vol. 8976, pp. 194–212. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48051-9_15
8. Choi, S.G., Dachman-Soled, D., Gordon, S.D., Liu, L., Yerukhimovich, A.: Compressed oblivious encoding for homomorphically encrypted search. In: Vigna, G., Shi, E. (eds.) *ACM CCS 2021: 28th Conference on Computer and Communications Security*. pp. 2277–2291. ACM Press, Virtual Event (2021). <https://doi.org/10.1145/3460120.3484792>
9. Cooley, J.W., Tukey, J.W.: An algorithm for the machine calculation of complex Fourier series. *Math. Comput.* **19**(90), 297–301 (1965)
10. Corrigan-Gibbs, H., Kogan, D.: Private information retrieval with sublinear online time. In: Canteaut, A., Ishai, Y. (eds.) *EUROCRYPT 2020*. LNCS, vol. 12105, pp. 44–75. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_3
11. Donoho, D.L.: Compressed sensing. *IEEE Trans. Inf. Theory* **52**(4), 1289–1306 (2006). <https://doi.org/10.1109/TIT.2006.871582>
12. Gentry, C.: A Fully Homomorphic Encryption Scheme. Ph.D. thesis, Stanford, CA, USA (2009)
13. Gilbert, A., Indyk, P.: Sparse recovery using sparse matrices. *Proc. IEEE* **98**(6), 937–947 (2010). <https://doi.org/10.1109/JPROC.2010.2045092>
14. Goodrich, M.T., Mitzenmacher, M.: Invertible bloom lookup tables. In: 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 792–799. IEEE Computer Society Press (2011). <https://doi.org/10.1109/Allerton.2011.6120248>
15. Huang, Q.L., Gao, X.S.: Revisit sparse polynomial interpolation based on randomized kronecker substitution. In: England, M., Koepf, W., Sadykov, T.M., Seiler, W.M., Vorozhtsov, E.V. (eds.) *CASC 2019: 21st International Workshop on Computer Algebra in Scientific Computing*, vol. 11661, pp. 215–235. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-030-26831-2_15
16. Johnson, M., Wagner, D., Ramchandran, K.: On compressing encrypted data without the encryption key. In: Naor, M. (ed.) *TCC 2004*. LNCS, vol. 2951, pp. 491–504. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24638-1_27
17. Klinc, D., Hazay, C., Jagmohan, A., Krawczyk, H., Rabin, T.: On compression of data encrypted with block ciphers. In: Storer, J.A., Marcellin, M.W. (eds.) *DCC 2009: 19th Data Compression Conference*. pp. 213–222. IEEE Computer Society Press, Snowbird (2009). <https://doi.org/10.1109/DCC.2009.71>
18. Lauter, K., López-Alt, A., Naehrig, M.: Private computation on encrypted genomic data. In: Aranha, D.F., Menezes, A. (eds.) *LATINCRYPT 2014*. LNCS, vol. 8895, pp. 3–27. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-16295-9_1
19. Liu, Z., Tromer, E.: Oblivious message retrieval. In: Dodis, Y., Shrimpton, T. (eds.) *Advances in Cryptology - CRYPTO 2022, Part I*. LNCS, vol. 13507, pp. 753–783. Springer, Heidelberg (2022). https://doi.org/10.1007/978-3-031-15802-5_26
20. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Wagner, D. (ed.) *CRYPTO 2008*. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_31

21. de Prony, G.: Essai expérimental et analytique sur les lois de la dilatabilité des fluides élastiques et sur celles de la force expansive de la vapeur de l'eau et de la vapeur de l'alcool à différentes températures. *J. l'École Polytech.* **1**(22), 24–76 (1795)
22. Rivest, R.L., Adleman, L., Dertouzos, M.L.: On data banks and privacy homomorphisms. In: DeMillo, R.A., Lipton, R.J., Dobkin, D.P., Jones, A.K. (eds.) *Foundations of Secure Computation*, pp. 169–179. Academic Press (1978)
23. Shanks, D.: Class number, a theory of factorization, and genera. In: Lewis, D.J. (ed.) *1969 Number Theory Institute. Proceedings of Symposia in Pure Mathematics*, vol. 20, pp. 415–440. American Mathematical Society (1971)
24. Slepian, D., Wolf, J.: Noiseless coding of correlated information sources. *IEEE Trans. Inf. Theory* **19**(4), 471–480 (1973). <https://doi.org/10.1109/TIT.1973.1055037>
25. Smart, N.P., Vercauteren, F.: Fully homomorphic SIMD operations. *Designs Codes Cryptogr.* **71**(1), 57–81 (2012). <https://doi.org/10.1007/s10623-012-9720-4>
26. Song, D.X., Wagner, D., Perrig, A.: Practical techniques for searches on encrypted data. In: *2000 IEEE Symposium on Security and Privacy*, pp. 44–55. IEEE Computer Society Press, Oakland (2000). <https://doi.org/10.1109/SECPRI.2000.848445>
27. Yasuda, M., Shimoyama, T., Kogure, J., Yokoyama, K., Koshihara, T.: Secure pattern matching using somewhat homomorphic encryption. In: Juel, A., Parno, B. (eds.) *CCSW 2013: The ACM Cloud Computing Security Workshop*, pp. 65–76. ACM Press, Berlin (2013). <https://doi.org/10.1145/2517488.2517497>