

Encrypted Multi-Maps with Computationally-Secure Leakage

Seny Kamara*
Brown University

Tarik Moataz†
Brown University

Abstract

We initiate the study of structured encryption schemes with computationally-secure leakage. Specifically, we focus on the design of volume-hiding encrypted multi-maps; that is, of encrypted multi-maps that hide the response length to computationally-bounded adversaries. We describe the first volume-hiding STE schemes that do not rely on naive padding; that is, padding all tuples to the same length. Our first construction has efficient query complexity and storage but can be lossy. We show, however, that the information loss can be bounded with overwhelming probability for a large class of multi-maps (i.e., with lengths distributed according to a Zipf distribution). Our second construction is not lossy and can achieve storage overhead that is asymptotically better than naive padding for Zipf-distributed multi-maps. We also show how to further improve the storage when the multi-map is highly *concentrated* in the sense that it has a large number of tuples with a large intersection.

We achieve these results by leveraging computational assumptions. Not just for encryption but, more interestingly, to hide the volumes themselves. Our first construction achieves this using a pseudo-random function whereas our second construction achieves this by relying on the conjectured hardness of the planted densest subgraph problem which is a planted variant of the well-studied densest subgraph problem. This assumption was previously used to design public-key encryption schemes (Applebaum et al., *STOC '10*) and to study the computational complexity of financial products (Arora et al., *ICS '10*).

*seny@brown.edu.

†tarik_moataz@brown.edu.

Contents

1	Introduction	3
1.1	Naive Approaches	4
1.2	Our Techniques and Contributions	5
1.3	Related Work	7
2	Preliminaries	7
3	Definitions	7
4	The Pseudo-Random Transform	9
4.1	Analyzing the Number of Truncations	10
4.2	Analyzing the Storage Overhead	12
4.3	Concrete Parameters	13
5	A Volume-Hiding Multi-Map Encryption Scheme	15
6	DST: The Densest Subgraph Transform	17
6.1	Detailed Description	19
6.2	Analyzing the Load of a Bin	21
6.3	Query and Storage Efficiency	24
6.4	Concrete Parameters	25
7	AVLH: Advanced Volume Hiding Multi-Map Encryption Scheme	26
7.1	Security	27
8	Dynamic Volume Hiding Multi-Map Encryption Schemes	30
8.1	VLH ^d : a dynamic variant of VLH.	30
8.2	AVLH ^d : dynamic variant of AVLH.	31

1 Introduction

A structured encryption (STE) scheme encrypts a data structure in such a way that it can be privately queried. An STE scheme is secure if it reveals nothing about the structure and query beyond a well-specified and “reasonable” leakage profile [15, 12]. An important special case of STE is searchable symmetric encryption (SSE) which relies on encrypted multi-maps [15, 12, 27, 26, 33, 11, 10, 7, 8, 4, 18, 5, 16] to achieve optimal-time search. Another example is graph encryption which encrypts various kinds of graphs [12, 30]. STE has received a lot of attention due to its potential applications to cloud storage and database security. In recent years, much of the work on STE has focused on supporting more complex queries like Boolean [11, 34, 21, 23] and range queries [34, 21, 20, 35], more complex structures like relational databases [24] and improving security like achieving forward-privacy [37, 7, 8, 19, 1].

Leakage. One aspect of STE that is still poorly understood is its leakage. There are currently two directions to dealing with leakage. The first is cryptanalysis; that is, designing leakage attacks against various leakage profiles so that we can better understand their concrete security. This was initiated by Islam, Kuzu and Kantarcioglu in the context of SSE [22] and expanded to PPE by Naveed, Kamara and Wright [32] and to ORAM by Kellaris, Kollios, Nissim and O’Neill [28]. While there has been some progress on designing leakage attacks against STE [22, 9, 28, 29], these attacks remain mostly of theoretical interest due to the strong assumptions they rely on. Assumptions like knowledge of at least 80%–90% of client data in addition to knowledge of 5% of client queries [22, 9], or assuming clients make queries uniformly at random, often addition to assumptions about how client data is distributed [28, 29]. Nevertheless, these attacks do provide us with some guidance as to which leakage profiles to avoid when designing schemes. Another line of work related to leakage was initiated recently by Kamara, Moataz and Ohrimenko in [25] where they propose designing general-purpose techniques to suppress given leakage patterns. In [25], they show how to do this for the query equality (also known as the search pattern) without making use of ORAM simulation and, therefore, without incurring its poly-logarithmic multiplicative overhead.

Computationally-secure leakage. In this work, we consider a new approach to dealing with leakage. Our work starts from the observation that the presence of leakage does not necessarily imply that this leakage can be exploited. In fact, it could be that this leakage is not exploitable because it does not convey enough useful information to the adversary. Alternatively, it could be that the leakage does convey enough information but no computationally-bounded adversary can extract it. In other words, this leakage would be computationally-secure. The possibility of designing STE schemes with computationally-secure leakage patterns is interesting for several reasons. From a theoretical point of view, as far as we know, this question has never been considered before and it raises some intriguing foundational questions; one of which is what kind of computational assumptions would lend themselves to the design of secure leakage patterns? The traditional assumptions used in cryptography are usually algebraic or number-theoretic in nature and it is not clear how such assumptions could be used. From a more practical perspective, the ability to leverage “computationally-secure leakage” in the design of STE schemes could lead to a whole new set of techniques and, ultimately, to highly-efficient zero-leakage schemes—computationally speaking.

Volume-hiding EMMs. In this work, we initiate the study of computationally-secure leakage. In particular, we focus on the design of volume-hiding encrypted multi-maps or, more precisely, of encrypted multi-maps that hide the response length to computationally-bounded adversaries.¹ We focus on encrypted multi-maps because they are by far the most important encrypted structure; this is illustrated by the fact that they are central to the design of optimal-time single-keyword SSE [15, 12, 27, 10, 7, 31, 8, 19], sub-linear Boolean SSE [11, 23], graph encryption [12, 30], encrypted range structures [20, 17] and encrypted relational databases [24]. We consider the response length leakage pattern for several reasons. The first is that it is a very difficult leakage pattern to suppress. In fact, though encrypted search has been investigated since 2000, the first non-trivial construction to even partially hide the response length is the recent PBS scheme of [25].² In fact, response lengths are leaked even by ORAM-based solutions. The second reason we focus on response lengths is because of the recent volume attacks of Kellaris et al. [28]. Again, while these attacks are mostly of theoretical interest, they do suggest that the design of volume-hiding encrypted structures is well-motivated.

1.1 Naive Approaches

To better understand our techniques and the improvements they provide, we first describe two possible naive approaches to designing volume-hiding EMMs. Recall that a multi-map is a data structure that stores a set of pairs $\{(\ell, \mathbf{v})\}$, where ℓ is a label from a label space \mathbb{L} and \mathbf{v} is a tuple of values from some value space \mathbb{V} . Multi-maps support get and put operations. Get takes as input a label ℓ and returns its associated tuple \mathbf{v} whereas Put takes as input a label/value pair (ℓ, \mathbf{v}) and stores it. We denote the get operation by $\mathbf{v} := \text{MM}[\ell]$ and the put operation by $\text{MM}[\ell] := \mathbf{v}$.

Naive padding. The first approach to designing a volume-hiding multi-map encryption scheme is to pad the tuples of the plaintext multi-map MM to their maximum response length $t = \max_{\ell \in \mathbb{L}} \#\text{MM}[\ell]$ and encrypt the padded multi-map with any standard multi-map encryption scheme [12, 10, 7, 1]. It is easy to see that this hides the response lengths. Unfortunately, it also induce a non-trivial storage overhead.

Using ORAM. We now describe a volume-hiding construction based on ORAM. Note that, as far as we know, this construction has not appeared before and may be of independent interest. We first represent the multi-map MM as a dictionary by generating $N \stackrel{\text{def}}{=} \sum_{\ell \in \mathbb{L}} \#\text{MM}[\ell]$ pairs of the form $\{(\ell, v)_{\ell \in \mathbb{L}, v \in \text{MM}[\ell]}\}$ and storing them in a dictionary DX . We then add $t - 1$ dummy label/value pairs to DX , where t is the maximum response length of a label in MM . DX is then stored and managed using ORAM. To get the tuple associated with a label ℓ , we first obviously access DX . There are two cases: if $\#\text{MM}[\ell] = t$, then we retrieve all pairs associated with ℓ ; otherwise if $\#\text{MM}[\ell] < t$, we retrieve an additional $t - \#\text{MM}[\ell]$ dummies.

It is clear that this hides the response length since the ORAM simulation hides the query equality and, therefore, an adversary can't distinguish between a dummy label and a real label. From an efficiency perspective, if we use a state-of-the-art ORAM [38] then the storage overhead

¹Our constructions also reveal the query equality—even to a bounded adversary—but the latter can be suppressed using the cache-based transform from [25].

²The PBS construction has two variants. One can hide the response length on non-repeating sub-patterns but has a probability of failure in the sense that the client might not receive all its query responses. The second variant is always correct but reveals the sequence response length on non-repeating sub-patterns.

is $O(N)$. The communication complexity, however, is $O(t \cdot \log^2 N)$ which includes a multiplicative poly-logarithmic factor in addition to a logarithmic round complexity.

1.2 Our Techniques and Contributions

In this work, we describe two volume-hiding multi-map encryption schemes: VLH and AVLH. Both our constructions work by first transforming an input multi-map into a volume-hiding multi-map and encrypting the result with a custom multi-map encryption scheme (that itself makes black-box use of a standard multi-map encryption scheme). These constructions avoid the limitations of the naive approaches described above either by improving on the storage of naive padding or avoiding the multiplicative poly-logarithmic overhead of the ORAM-based solution.

A time-efficient construction. Our first construction relies on a simple transformation we call the *pseudo-random transform* which is parameterized by a public parameter λ and makes use of a small-domain pseudo-random function as follows. Each tuple \mathbf{v} in the multi-map is transformed into a new tuple \mathbf{v}' of size $n' = \lambda + F_K(n)$, where $n = \#\mathbf{v}$. If $n' > n$, then the elements of \mathbf{v} are stored in \mathbf{v}' and the latter is padded to have length n' . If $n' \leq n$ then only the first n' items of \mathbf{v} are stored in \mathbf{v}' which effectively truncates \mathbf{v} (we think of the case $n' = n$ as a padding). Note that the multi-map that results from this process is volume-hiding since each tuple has pseudo-random length. Perhaps surprisingly, we also show that if the lengths of the input multi-map are Zipf-distributed then the storage overhead and the number of truncations can be kept relatively small with overwhelming probability (in the number of labels). Our scheme VLH essentially consists of transforming a multi-map using the pseudo-random transform and encrypting it with a standard multi-map encryption scheme. The query complexity of VLH is $O(\lambda + \nu)$, where ν is the largest value in the domain of F . While the pseudo-random transform leads to an efficient construction, it is lossy since tuples can be truncated. In many practical settings, however, truncations are not necessarily an issue. For example, in the case of SSE where EMMs are used to store document identifiers clients can rank the document ids (say, by relevance) at setup time so that truncations only affect the low-ranked documents. Nevertheless, we also consider the problem of designing non-lossy volume-hiding EMMs.

A non-lossy transform. Our second construction relies on a different transformation we call the *dense subgraph transform*. Unlike the pseudo-random transform which introduces truncations, this approach is non-lossy. On the other hand, it is less efficient in terms of query complexity. Note, however, that it is hard to imagine any non-lossy construction being able to hide the response length of a query and having query complexity $o(t)$, where t is the maximum response length. Our goal, therefore, is to design a non-lossy scheme that improves on the *storage overhead* of the naive padding approach. At a high-level, our non-lossy transform works by re-arranging the data stored in the multi-map into bins according to a random bi-partite graph. Roughly speaking, we construct a random (regular) bi-partite graph with labels in one set and bins in the other. We then assign the values in a label's tuple to the bins that are incident to the label. The bins are then padded to hide their size. To ensure that this re-arrangement is still efficiently queryable, we show how to represent the structure encoded in the bi-partite graph and the data stored in the bins with a pair of standard data structures; specifically, a multi-map and a dictionary. We show that, with the right choice of parameters, this version of our transformation already yields a volume-hiding multi-map structure with better storage overhead than the naive padding approach. More precisely, we show that the

naive approach produces a volume-hiding multi-map of size $S_{\text{NV}} = \Omega(N)$, where N is the size of the original multi-map, whereas our approach yields a volume-hiding multi-map of size $O(N)$ with overwhelming probability in N . Interestingly, we also show that if the tuple-lengths of the input multi-map are Zipf-distributed then our transformation yields a multi-map of size $o(S_{\text{NV}})$ with overwhelming probability. We note that this version of the transformation already makes use of computational assumptions. In particular, it uses a pseudo-random function to generate the edges of the random bi-partite graph which allows us to “compress” the size of our data structures by storing random seeds as opposed to all the graph’s edges. To query our transformed multi-map on some label ℓ , it suffices to retrieve the bins incident to ℓ . Intuitively, this is volume-hiding because the bins are padded and the number of bins is fixed. Furthermore, it hides other leakage patterns because the tuple values are assigned to bins randomly.

Concentration and planted subgraphs. The version of the transformation described so far already improves over naive padding (with overwhelming probability) but we show that for a certain class of multi-maps we can do even better—though at the cost of increased query complexity. Specifically, we consider multi-maps that have a large number of tuples with a large intersection. We refer to this property as *concentration* and describe a version of the dense subgraph transform that leverages the multi-map’s concentration to improve storage efficiency even more. At a high-level, the idea is as follows. A concentrated multi-map has a number of redundant values which our transformation assigns to multiple bins. In our improved transform, we instead assign each these redundant values to a single bin and add edges between these bins and all the labels whose tuples they appear in. The rest of the bi-partite graph is generated (pseudo-)randomly as above. This has the benefit of inducing smaller bins and, therefore, of requiring less padding. The bi-partite graph, however, is not random anymore (even ignoring our use of a pseudo-random function to generate edges). We observe, however, that by adding the edges to the bins of the redundant values, we are effectively *planting* a small dense subgraph inside of a larger random graph. And while the resulting graph is clearly not random anymore, it can be shown to be computationally indistinguishable from a random graph. In fact, this reduces to the *planted densest subgraph problem* which has been used in the past by Applebaum et al. in the context of cryptography [2] and by Arora et al. in the context of computational complexity [3]. Based on this assumption, we can show that for multi-maps with concentration parameters within a certain range (in turn determined by the densest subgraph assumption) the transformed multi-map is of size $O(N - \sqrt{m} \cdot \text{polylog}(m))$ with overwhelming probability and where m is the number of labels in the original multi-map. If the input multi-map is Zipf-distributed, then the output multi-map has size $o(S_{\text{NV}})$.

Our non-lossy construction. As mentioned above, the dense subgraph transform produces multi-maps that we represent using a combination of a dictionary and a (standard) multi-map. To encrypt this particular representation, we design a new scheme called AVLH. The resulting construction has query complexity $O(t \cdot (\text{polylog}(m) - \sqrt{m}))$ for multi-maps with concentration parameters within a certain range.

Dynamism. Our VLH and AVLH constructions are for static multi-maps. While there are many important applications of static EMMs, we describe how to extend these constructions to handle updates. This results in two additional constructions, VLH^d and AVLH^d. The former handles three kinds of updates: tuple addition, tuple deletion and tuple edits; and the latter handles tuple edits.

1.3 Related Work

Structured encryption was introduced by Chase and Kamara [12] as a generalization of searchable symmetric encryption which was first considered by Song, Wagner and Perrig [36] and formalized by Curtmola, Garay, Kamara and Ostrovsky [15]. Multi-map encryption schemes are a special case of STE and have been used to achieve optimal-time single-keyword SSE [15, 27, 10, 7, 19, 8], sub-linear Boolean SSE [11, 23], encrypted range search [20, 17], encrypted relational databases [24] and graph encryption [12, 30]. The first leakage attack against volume leakage was described by Kollios, Kellaris, Nissim and O’Neill [28] under the assumption of uniform query distributions. In [25], Kamara, Moataz and Ohrimenko describe an STE scheme called PBS which partially hides the volume pattern. More precisely, the first variant of PBS reveals only the sequence response length (i.e., the sum of the response lengths of a given query sequence) on non-repeating query sequences. The second variant reveals nothing (beyond a public parameter independent of the volume) on non-repeating query sequences. The planted densest graph problem was first used as a computational assumption by Applebaum, Barak and Wigderson in [2] for the purpose of designing public-key encryption schemes under new assumptions. It was later used by Arora, Barak, Brunnermeier and Ge to study the computational complexity of financial products [3].

2 Preliminaries

Notation. The set of all binary strings of length n is denoted as $\{0,1\}^n$, and the set of all finite binary strings as $\{0,1\}^*$. We write $x \leftarrow \chi$ to represent an element x being sampled from a distribution χ , and $x \xleftarrow{\$} X$ to represent an element x being sampled uniformly at random from a set X . The output x of an algorithm \mathcal{A} is denoted by $x \leftarrow \mathcal{A}$. Given a sequence \mathbf{v} of n elements, we refer to its i th element as v_i or $\mathbf{v}[i]$. If S is a set then $\#S$ refers to its cardinality and 2^S to its powerset.

Basic cryptographic primitives. A private-key encryption scheme is a set of three polynomial-time algorithms $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ such that Gen is a probabilistic algorithm that takes a security parameter k and returns a secret key K ; Enc is a probabilistic algorithm that takes a key K and a message m and returns a ciphertext c ; Dec is a deterministic algorithm that takes a key K and a ciphertext c and returns m if K was the key under which c was produced. Informally, a private-key encryption scheme is secure against chosen-plaintext attacks (CPA) if the ciphertexts it outputs do not reveal any partial information about the plaintext even to an adversary that can adaptively query an encryption oracle. We say a scheme is random-ciphertext-secure against chosen-plaintext attacks (RCPA) if the ciphertexts it outputs are computationally indistinguishable from random even to an adversary that can adaptively query an encryption oracle. In addition to encryption schemes, we also make use of pseudo-random functions (PRF), which are polynomial-time computable functions that cannot be distinguished from random functions by any probabilistic polynomial-time adversary.

3 Definitions

Structured encryption schemes encrypt data structures in such a way that they can be privately queried. There are several natural forms of structured encryption. The original definition of

[12] considered schemes that encrypt both a structure and a set of associated data items (e.g., documents, emails, user profiles etc.). In [13], the authors also describe *structure-only* schemes which only encrypt structures. Another distinction can be made between *interactive* and *non-interactive* schemes. Interactive schemes produce encrypted structures that are queried through an interactive two-party protocol, whereas non-interactive schemes produce structures that can be queried by sending a single message, i.e, the token. One can also distinguish between *response-hiding* and *response-revealing* schemes: the former reveal the response to queries whereas the latter do not. We recall here the syntax of an interactive response-hiding structured encryption scheme.

Definition 3.1 (Structured encryption). *An interactive response-hiding structured encryption scheme $\Sigma_{\text{DS}} = (\text{Setup}, \text{Query})$ for data type DS consists of the following polynomial-time algorithms and protocols:*

- $(K, \text{EDS}) \leftarrow \text{Setup}_{\mathbf{C}}(1^k, \text{DS})$: *is a probabilistic algorithm that takes as input a security parameter 1^k and a structure DS of type DS and outputs a secret key K and an encrypted structure EDS.*
- $(r, \perp) \leftarrow \text{Query}_{\mathbf{C}, \mathbf{S}}(\text{tk}; \text{EDS})$: *is an interactive protocol executed between a client \mathbf{C} and a server \mathbf{S} . The client inputs a token tk and the server inputs an encrypted structure EDS. The client receives a response r and the server receives \perp .*

We refer the reader to, for example [1], for syntax definitions of dynamic STE.

Security. The standard notion of security for STE guarantees that: (1) an encrypted structure reveals no information about its underlying structure beyond the setup leakage $\mathcal{L}_{\mathbf{S}}$; (2) that the query protocol reveals no information about the structure and the queries beyond the query leakage $\mathcal{L}_{\mathbf{Q}}$. If this holds for non-adaptively chosen operations then the scheme is said to be non-adaptively secure. If, on the other hand, the operations can be chosen adaptively, the scheme is said to be adaptively-secure.

Definition 3.2 (Adaptive security of interactive STE). *Let $\Sigma = (\text{Setup}, \text{Query})$ be an interactive STE scheme and consider the following probabilistic experiments where \mathcal{A} is a stateful semi-honest adversary, \mathcal{S} is a stateful simulator, $\mathcal{L}_{\mathbf{S}}$ and $\mathcal{L}_{\mathbf{Q}}$ are leakage profiles and $z \in \{0, 1\}^*$:*

Real $_{\Sigma, \mathcal{A}}(k)$: *given z the adversary \mathcal{A} outputs a structure DS and receives EDS from the challenger, where $(K, \text{EDS}) \leftarrow \text{Setup}(1^k, \text{DS})$. The adversary then adaptively chooses a polynomial number of queries and, for each, executes the Query protocol with the challenger, where the adversary plays the server and the challenger plays the client. Finally, \mathcal{A} outputs a bit b that is output by the experiment.*

Ideal $_{\Sigma, \mathcal{A}, \mathcal{S}}(k)$: *given z the adversary \mathcal{A} generates a structure DS which it sends to the challenger. Given z and leakage $\mathcal{L}_{\mathbf{S}}(\text{DS})$ from the challenger, the simulator \mathcal{S} returns an encrypted structure EDS to \mathcal{A} . The adversary then adaptively chooses a polynomial number of queries and, for each one, executes the Query protocol with the simulator, where the adversary plays the server and the simulator plays the client (note that here, the simulator is allowed to deviate from Query). Finally, \mathcal{A} outputs a bit b that is output by the experiment.*

We say that Σ is adaptively $(\mathcal{L}_{\mathbf{S}}, \mathcal{L}_{\mathbf{Q}})$ -secure if there exists a PPT simulator \mathcal{S} such that for all PPT adversaries \mathcal{A} , for all $z \in \{0, 1\}^*$,

$$|\Pr[\text{Real}_{\Sigma, \mathcal{A}}(k) = 1] - \Pr[\text{Ideal}_{\Sigma, \mathcal{A}, \mathcal{S}}(k) = 1]| \leq \text{negl}(k).$$

Let $F : \{0, 1\}^k \times \{0, 1\}^* \rightarrow \{0, 1\}^s$ be a pseudo-random function, $\text{rank} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be ranking function and $\lambda \in \mathbb{N}$ be a public parameter. Consider the transform PRT defined as follows:

- $\text{PRT}(1^k, \lambda, \text{MM})$:
 1. sample a key $K \xleftarrow{\$} \{0, 1\}^k$;
 2. instantiate an empty multi-map MM' ;
 3. for all $\ell \in \mathbb{L}_{\text{MM}}$,
 - (a) let $\mathbf{r} := \text{MM}[\ell]$ and $n_\ell = \#\mathbf{r}$;
 - (b) compute $\mathbf{r}' := \text{rank}(\mathbf{r})$;
 - (c) let $n'_\ell = \lambda + F_K(\ell \| n_\ell)$;
 - (d) if $n'_\ell > n_\ell$, set $\text{MM}'[\ell] := (\mathbf{r}', \perp_1, \dots, \perp_{n'_\ell - n_\ell})$;
 - (e) otherwise, set $\text{MM}'[\ell] := (r'_1, \dots, r'_{n'_\ell})$;
 4. output MM' .
- $\text{Get}(\ell, \text{MM})$: output $\text{MM}[\ell]$.

Figure 1: The pseudo-random transform.

Modeling leakage. Every STE scheme is associated with leakage which itself can be composed of multiple *leakage patterns*. The collection of all these leakage patterns forms the scheme’s *leakage profile*. Leakage patterns are (families of) functions over the various spaces associated with the underlying data structure. For concreteness, we borrow the nomenclature introduced in [25] and recall some well-known leakage patterns that we make use of in this work:

- the *query equality pattern* is the function family $\text{qeq} = \{\text{qeq}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{qeq}_{k,t} : \mathbb{D}_k \times \mathbb{Q}_k^t \rightarrow \{0, 1\}^{t \times t}$ such that $\text{qeq}_{k,t}(\text{DS}, q_1, \dots, q_t) = M$, where M is a binary $t \times t$ matrix such that $M[i, j] = 1$ if $q_i = q_j$ and $M[i, j] = 0$ if $q_i \neq q_j$. The query equality pattern is referred to as the search pattern in the SSE literature;
- the *response identity pattern* is the function family $\text{rid} = \{\text{rid}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{rid}_{k,t} : \mathbb{D}_k \times \mathbb{Q}_k^t \rightarrow [2^{[n]}]^t$ such that $\text{rid}_{k,t}(\text{DS}, q_1, \dots, q_t) = (\text{DS}[q_1], \dots, \text{DS}[q_t])$. The response identity pattern is referred to as the access pattern in the SSE literature;
- the *response length pattern* is the function family $\text{rlen} = \{\text{rlen}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{rlen}_{k,t} : \mathbb{D}_k \times \mathbb{Q}_k^t \rightarrow \mathbb{N}^t$ such that $\text{rlen}_{k,t}(\text{DS}, q_1, \dots, q_t) = (|\text{DS}[q_1]|, \dots, |\text{DS}[q_t]|)$;
- the *domain size pattern* is the function family $\text{dsize} = \{\text{dsize}_{k,t}\}_{k,t \in \mathbb{N}}$ with $\text{dsize}_{k,t} : \mathbb{D}_k \rightarrow \mathbb{N}$ such that $\text{dsize}_{k,t}(\text{DS}) = \#\mathbb{Q}$.

4 The Pseudo-Random Transform

We describe the pseudo-random transform (PRT) in Figure 1 and provide a high level description below.

Overview. PRT is a data structure transformation that takes as input a multi-map MM , a security parameter k and a public parameter λ . It first generates a random key K and initializes

an empty multi-map MM' . For each label ℓ in the multi-map, it ranks the tuple $\mathbf{r} := \text{MM}[\ell]$; resulting in a ranked tuple \mathbf{r}' . It then evaluates a PRF on the label ℓ concatenated to the length n_ℓ of \mathbf{r} . The output of this PRF evaluation is then added to λ in order to compute a new length n'_ℓ . There are two possible cases that can occur at this point: (1) if n'_ℓ is larger than n_ℓ , the ranked response is padded with dummies and inserted in $\text{MM}'[\ell]$; (2) if n'_ℓ is at most n_ℓ , the ranked response is truncated to its first n'_ℓ elements and inserted in $\text{MM}'[\ell]$. Note that for ease of exposition and without loss of generality, we consider the case where $n'_\ell = n_\ell$ a padding. Finally, the transform outputs the multi-map MM' . The get algorithm simply outputs the tuple corresponding to the label ℓ .

A note on probabilistic analysis. Throughout this work, we model pseudo-random functions as random functions for the purposes of probabilistic analysis. It should be understood that all our bounds will have an additional negligible value in the security parameter.

4.1 Analyzing the Number of Truncations

For any label ℓ of the multi-map, the transform can pad or truncate its ranked response depending on the output of the PRF. In this Section, we will analyze the number of truncations induced by our transformation. The number of truncations is defined as

$$\#\{\ell \in \mathbb{L}_{\text{MM}} : \#\text{MM}'[\ell] < \#\text{MM}[\ell]\}.$$

In the worst-case, the number of truncations can be $\#\mathbb{L}_{\text{MM}}$ which occurs when every label in MM is truncated. We will show, however, that in practice this is very unlikely to occur. In particular, we will show that for real-world distributions of response lengths, the number of truncations is small with high probability. Note that if we set $\lambda \geq \max_{\ell \in \mathbb{L}} \#\text{MM}[\ell]$, then truncations can never occur since $\#\text{MM}'[\ell] \geq \max_{\ell \in \mathbb{L}} \#\text{MM}[\ell] \geq \#\text{MM}[\ell]$. We therefore only consider settings in which $\lambda < \max_{\ell \in \mathbb{L}} \#\text{MM}[\ell]$.

Zipf-distributed multi-maps. To get a concrete bound on the number of truncations, we have to make an assumption on how the response lengths of the multi-map are distributed. Here, we will assume that they are distributed according to the Zipf distribution which is a standard assumption in information retrieval [14, 39]. More precisely, we say that a multi-map MM is $\mathcal{Z}_{a,b}$ -distributed if its r^{th} response has length

$$\frac{r^{-b}}{H_{a,b}} \cdot N$$

where $N \stackrel{\text{def}}{=} \sum_{\ell \in \mathbb{L}} \#\text{MM}[\ell]$ is the volume of MM and $H_{a,b}$ is the harmonic number $\sum_{i=1}^a i^{-b}$. Throughout, we will consider multi-maps that are $\mathcal{Z}_{m,1}$ -distributed where $m = \#\mathbb{L}_{\text{MM}}$. From this assumption, it follows that the set of all response lengths is

$$L = (L_1, \dots, L_m) = \left(\frac{N}{1 \cdot H_{m,1}}, \dots, \frac{N}{m \cdot H_{m,1}} \right),$$

Note that we consider the case where $b = 1$ for ease of exposition but our analyses generalize to any b .

Theorem 4.1. *If MM is $\mathcal{Z}_{m,1}$ -distributed, then with probability at least $1 - \varepsilon$ the number of truncations is at most*

$$m \cdot \left(\frac{N}{\nu \cdot H_{m,1}^2} \cdot \left(H_{\rho,2} - \frac{\lambda \cdot H_{m,1}}{N} \cdot H_{\rho,1} \right) + \sqrt{\frac{\ln(1/\varepsilon)}{2m}} \right),$$

where $\rho = \lfloor N/(\lambda \cdot H_{m,1}) \rfloor$.

Proof. Let ℓ_1, \dots, ℓ_m denote the labels of MM sorted in decreasing order by their original response lengths. Let X_i be the random variable that outputs $F_K(\ell_i | n_i)$, where n_i is the response length of ℓ_i . Modeling F as a random function, note that X_i is the uniform distribution over $[0, \nu - 1]$ so, for all $j \in [0, \nu - 1]$, $\Pr[X_i = j] = 1/\nu$. Now let Y_i be the random variable that outputs the i th response length n_i . Since MM is $\mathcal{Z}_{m,1}$ -distributed, we have that $\Pr[Y_i = y] = 1/(x \cdot H_{m,1})$, where $y = N/(x \cdot H_{m,1})$ for all $x \in [m]$.

First notice that no truncation occurs if and only if $X_i + \lambda - Y_i \geq 0$. In the following, we compute the probability that this event occurs:

$$\begin{aligned} \Pr[X_i \geq Y_i - \lambda] &= \sum_{y \in L} \Pr[X_i \geq Y_i - \lambda | Y_i = y] \cdot \Pr[Y_i = y] \\ &= \sum_{\substack{y \in L \\ y < \lambda}} \Pr[Y_i = y] + \sum_{\substack{y \in L \\ y \geq \lambda}} \Pr[X_i \geq y - \lambda] \cdot \Pr[Y_i = y] \tag{1} \\ &= \sum_{\substack{y \in L \\ y < \lambda}} \Pr[Y_i = y] + \sum_{\substack{y \in L \\ y \geq \lambda}} \left(\sum_{j=y-\lambda}^{\nu-1} \Pr[X_i = j] \right) \cdot \Pr[Y_i = y] \\ &= \sum_{\substack{y \in L \\ y < \lambda}} \Pr[Y_i = y] + \sum_{\substack{y \in L \\ y \geq \lambda}} \frac{\nu - y + \lambda}{\nu} \cdot \Pr[Y_i = y] \\ &= 1 + \sum_{\substack{y \in L \\ y \geq \lambda}} \frac{\lambda - y}{\nu} \cdot \Pr[Y_i = y]. \tag{2} \end{aligned}$$

Equation (1) follows from that fact that for all $y \in L$ such that $y - \lambda \leq 0$, the probability that X_i is larger than a negative integer is always equal to 1. Given Equation (2), we compute the non-negative integer $\rho \in [m]$ such that

$$\frac{N}{(\rho + 1) \cdot H_{m,1}} < \lambda \leq \frac{N}{\rho \cdot H_{m,1}},$$

and reformulate the summation in the Equation (2) such that

$$\begin{aligned}
\Pr [X_i \geq Y_i - \lambda] &= 1 + \frac{1}{\nu} \sum_{j=1}^{\rho} \left(\lambda - \frac{N}{j \cdot H_{m,1}} \right) \cdot \Pr \left[Y_i = \frac{N}{j \cdot H_{m,1}} \right] \\
&= 1 + \frac{1}{\nu} \sum_{j=1}^{\rho} \left(\lambda - \frac{N}{j \cdot H_{m,1}} \right) \cdot \frac{1}{j \cdot H_{m,1}} \\
&= 1 - \frac{1}{\nu \cdot H_{m,1}} \left(\sum_{j=1}^{\rho} \frac{N}{j^2 \cdot H_{m,1}} - \lambda \sum_{j=1}^{\rho} \frac{1}{j} \right) \\
&= 1 - \frac{N}{\nu \cdot H_{m,1}^2} \left(H_{\rho,2} - \frac{\lambda \cdot H_{m,1}}{N} \cdot H_{\rho,1} \right).
\end{aligned}$$

Computing the negation of the event above, we obtain the probability that the i th output of the PRF induces a truncation:

$$\Pr [X_i < Y_i - \lambda] = \frac{N}{\nu \cdot H_{m,1}^2} \left(H_{\rho,2} - \frac{\lambda \cdot H_{m,1}}{N} \cdot H_{\rho,1} \right). \quad (3)$$

Now let Z_i be an indicator random variable that equals 1 when $X_i < Y_i - \lambda$, and 0 otherwise. Let $Z = \sum_{j=1}^m Z_i$ and note that the Z_i 's are independent and identically distributed Bernoulli distributions with parameter p equal to Equation (3). Z denotes the number of truncations and has expectation,

$$\mathbb{E}[Z] = m \cdot \frac{N}{\nu \cdot H_{m,1}^2} \left(H_{\rho,2} - \frac{\lambda \cdot H_{m,1}}{N} \cdot H_{\rho,1} \right).$$

Using Hoeffding's first inequality, we have that for all $\delta \geq 0$,

$$\Pr [Z \geq \mathbb{E}[Z] + m \cdot \delta] \leq \exp(-2 \cdot m \cdot \delta^2).$$

Setting $\varepsilon = \exp(-2 \cdot m \cdot \delta^2)$, we have

$$\Pr \left[Z < \mathbb{E}[Z] + m \cdot \sqrt{\frac{\ln(1/\varepsilon)}{2 \cdot m}} \right] \geq 1 - \varepsilon,$$

from which the Theorem follows. ■

4.2 Analyzing the Storage Overhead

As detailed above, PRT can truncate or pad the responses in the multi-map. This has a direct impact on the storage overhead of the transformed multi-map since padding increases the storage overhead while truncations decrease it. In the following, we show that the size of the transformed multi-map MM' can be upper bounded with high probability without any assumptions on the distribution of response lengths.

Theorem 4.2. *With probability at least $1 - \varepsilon$, the size of the transformed multi-map is at most*

$$m \cdot \left(\frac{\nu - 1}{2} + (\nu - 1) \cdot \sqrt{\frac{\ln(1/\varepsilon)}{2m}} + \lambda \right),$$

where $\lambda \geq 0$.

Proof. Let ℓ_1, \dots, ℓ_m denote the labels of MM sorted in decreasing order by their original response lengths. Let X_i be the random variable that outputs $F_K(\ell_i | n_i)$, where n_i is the response length of ℓ_i . Note that X_i is the uniform distribution over $[0, \nu - 1]$ so, for all $j \in [0, \nu - 1]$, $\Pr[X_i = j] = 1/\nu$. Let $X = \sum_{i=1}^m X_i$ from which it follows that $\mathbb{E}[X] = m \cdot (\nu - 1)/2$.

Since the X_i 's are bounded by $[0, \nu - 1]$, using Hoeffding's second inequality we have that, for all $\delta \geq 0$,

$$\Pr \left[X \geq m \cdot \left(\frac{\nu - 1}{2} + \delta \right) \right] \leq \exp \left(\frac{-2 \cdot m \cdot \delta^2}{(\nu - 1)^2} \right).$$

Given that the transform adds a constant λ to every new response length, the total storage overhead is $Z = X + \lambda \cdot m$. It follows then that

$$\Pr \left[Z \geq m \cdot \left(\frac{\nu - 1}{2} + \delta + \lambda \right) \right] \leq \exp \left(\frac{-2 \cdot m \cdot \delta^2}{(\nu - 1)^2} \right).$$

Setting

$$\delta = (\nu - 1) \cdot \sqrt{\frac{\ln(1/\varepsilon)}{2m}}$$

we have

$$\Pr \left[Z < m \cdot \left(\frac{\nu - 1}{2} + (\nu - 1) \cdot \sqrt{\frac{\ln(1/\varepsilon)}{2m}} + \lambda \right) \right] \geq 1 - \varepsilon.$$

■

4.3 Concrete Parameters

In this Section, we will provide concrete parameters for PRT. Our goal is to find parameters that will provide a good balance between a small number of truncations and a small storage overhead. To study this, we first introduce two naive transformations that achieve extreme tradeoffs between truncations and storage:

- the *naive padding transform* is a transformation that pads the response of every label with dummies \perp so that the length of the new responses are all set to the maximum response's length $\max_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$. Note that there are no truncations in this case and the size of the transformed multi-map is

$$S_{\text{NV}} \stackrel{\text{def}}{=} m \cdot \max_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell].$$

- the *naive truncating transform* truncates the responses of every label to the minimum response length $\min_{\ell \in \mathbb{L}} \#\text{MM}[\ell]$. Note that the number of truncations in this case is $T_{\text{NV}} \stackrel{\text{def}}{=} \#\mathbb{L}_{\text{MM}} = m$ and the storage overhead is $m \cdot \min_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$, which is optimal.

In the following Corollary, we set concrete values for λ and s so that we can achieve the best of both worlds. Specifically, we show that if the input multi-map is $\mathcal{Z}_{m,1}$ -distributed, then by setting the output length of the PRF to $s = \log(L_1 + 1)$, where L_1 is the maximum response length, and setting $\lambda = O(\nu \cdot \alpha)$, where $1/2 < \alpha < 1$, then we can achieve storage overhead $\alpha \cdot \mathcal{S}_{\text{NV}}$ with $\beta \cdot T_{\text{NV}}$ truncations with high probability, where β is a function of α and m .

Corollary 4.3. *Let $1/2 < \alpha < 1$. If MM is $\mathcal{Z}_{m,1}$ -distributed and if*

$$\log \nu = \log(L_1 + 1) \quad \text{and} \quad \lambda = (\nu - 1) \cdot (2\alpha - 1)/4$$

then with probability at least:

- $1 - \exp(-m \cdot (2\alpha - 1)^2/8)$, the total volume of the transformed multi-map is at most $\alpha \cdot \mathcal{S}_{\text{NV}}$.
- $1 - \exp(-2m/\log^2 m)$, the number of truncations is at most

$$\frac{1}{\log m} \cdot H_{\lfloor \frac{4}{2\alpha-1} \rfloor, 2} \cdot T_{\text{NV}}.$$

Proof. Based on Theorem 4.2, by setting

$$\lambda = (\nu - 1) \cdot \frac{2\alpha - 1}{4} \quad \text{and} \quad \varepsilon_1 = \exp\left(-\frac{m \cdot (2\alpha - 1)^2}{8}\right)$$

we get a storage overhead of at most

$$m \cdot \left(\frac{\nu - 1}{2} + (\nu - 1) \cdot \frac{2\alpha - 1}{4} + (\nu - 1) \cdot \frac{2\alpha - 1}{4} \right) = \alpha \cdot m \cdot (\nu - 1),$$

with probability at least $1 - \varepsilon_1$. It is clear that we have to set $\alpha < 1$ to achieve better storage overhead than the naive padding approach. On the other hand, we have to set $\alpha > 1/2$ since λ must be a strictly nonnegative integer. More precisely, $\lambda > 0$ implies

$$\begin{aligned} (\nu - 1) \cdot \frac{2\alpha - 1}{4} &> 0 \\ 2\alpha - 1 &> 0. \end{aligned}$$

Similarly, based on Theorem 4.1, if we set $\lambda = (\nu - 1) \cdot (2\alpha - 1)/4$ and $\varepsilon_2 = \exp(-2m/\log^2 m)$, the number of truncations is at most

$$\theta \stackrel{\text{def}}{=} m \cdot \left(\frac{N}{(\nu - 1) \cdot H_{m,1}^2} \left(H_{\rho,2} - \frac{\lambda \cdot H_{m,1}}{N} \cdot H_{\rho,1} \right) + \frac{1}{\log m} \right)$$

where $\rho = \lfloor N/(\lambda \cdot H_{m,1}) \rfloor$. If we set $\log \nu = \log(L_1 + 1) = \log(N/H_{m,1} + 1)$ we get,

$$\begin{aligned} \theta &\leq m \cdot \left(\frac{1}{H_{m,1}} \left(H_{\rho,2} - \frac{2\alpha - 1}{4} \cdot H_{\rho,1} \right) + \frac{1}{\log m} \right) \\ &\leq m \cdot \left(\frac{1}{H_{m,1}} \cdot H_{\rho,2} + \frac{1}{\log m} \right) \\ &< \frac{m}{\log m} \cdot H_{\rho,2} \end{aligned} \tag{4}$$

$$= \frac{m}{\log m} \cdot H_{\lfloor \frac{4}{2\alpha-1} \rfloor, 2} \tag{5}$$

Let $\text{STE}_{\text{EMM}} = (\text{Setup}, \text{Get})$ be a static multi-map encryption scheme and PRT be the pseudo-random transform. Consider the scheme VLH = (Setup, Get) defined as follows:

- $\text{Setup}(1^k, \lambda, \text{MM})$:

1. generate a PRT-transform of MM by computing $\text{MM}' = \text{PRT}(1^k, \lambda, \text{MM})$;
2. encrypt the transform by computing

$$(K, st, \text{EMM}) \leftarrow \text{STE}_{\text{EMM}}.\text{Setup}(1^k, \text{MM}');$$

3. output (K, st, EMM) .

- $\text{Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM})$: \mathbf{C} and \mathbf{S} execute

$$(r, \perp) \leftarrow \text{STE}_{\text{EMM}}.\text{Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM}).$$

Figure 2: VLH: A volume-hiding multi-map encryption scheme.

Where Inequality (4) follows from the fact that

$$H_{m,1} < \log m + \zeta + \frac{1}{2m},$$

where ζ is the Euler-Mascheroni constant and Equality (5) follows from

$$\rho = \left\lfloor \frac{N}{\lambda \cdot H_{m,1}} \right\rfloor = \left\lfloor \frac{\nu - 1}{\lambda} \right\rfloor = \left\lfloor \frac{4}{2\alpha - 1} \right\rfloor.$$

Given the inequality above, we obtain

$$\Pr \left[Z < \frac{m}{\log m} \cdot H_{\lfloor \frac{4}{2\alpha - 1} \rfloor, 2} \right] \geq 1 - \varepsilon_2,$$

from which the Theorem follows. ■

5 A Volume-Hiding Multi-Map Encryption Scheme

In this Section, we use the PRT to construct a volume-hiding multi-map encryption scheme. Our construction is described in detail in Figure 2 and works as follows.

Overview. The construction, VLH = (Setup, Get), makes black-box use of an underlying multi-map encryption scheme $\text{STE}_{\text{EMM}} = (\text{Setup}, \text{Get})$. VLH.Setup takes as input a security parameter k , a public parameter λ and a multi-map MM. It applies the PRT transform on MM which results in a new multi-map MM' . It then encrypts MM with STE_{EMM} , resulting in an encrypted multi-map EMM, a state st and a key K which it returns as its own output. To execute a Get query ℓ on EMM, the client and the server execute $\text{STE}_{\text{EMM}}.\text{Get}$ on ℓ and EMM.

Efficiency. Assuming that STE_{MM} is an optimal-time multi-map encryption scheme [15, 10, 7, 1], the get complexity of VLH is $O(\lambda + n'_\ell)$, where $n'_\ell \in \{0, \dots, \nu - 1\}$. Therefore, the worst-case complexity is $O(\lambda + \nu)$ while the best-case complexity is $O(\lambda)$. The expected complexity is $O(\lambda + 2^{s-1})$.

The storage overhead of VLH is

$$O(N) = O\left(\sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}'[\ell]\right) = O\left(\lambda \cdot m + \sum_{\ell \in \mathbb{L}_{\text{MM}}} n'_\ell\right),$$

where, again, $n'_\ell \in \{0, \dots, \nu - 1\}$. So based on Corollary 4.3, when $\lambda = (\nu - 1) \cdot (2\alpha - 1)/4$ and $1/2 < \alpha < 1$, the storage overhead of VLH is

$$O(\alpha \cdot (\nu - 1) \cdot m)$$

with high probability.

Correctness. The correctness of VLH is affected by the number of truncations induced by PRT. Based on Corollary 4.3, we can show that the number of truncations performed by VLH is at most $O(m/\log m)$ under the same assumptions stated in the corollary.

Security. We now describe the leakage profile of VLH assuming STE_{MM} is instantiated with one of the standard optimal-time multi-map encryption schemes [15, 10, 7, 1] all of which have leakage profile

$$\Lambda_{\text{MM}} = (\mathcal{L}_S, \mathcal{L}_Q) = (\text{trlen}, (\text{qeq}, \text{rlen})).$$

Theorem 5.1. *If STE_{EMM} is a $(\text{trlen}, (\text{qeq}, \text{rlen}))$ -secure multi-map encryption scheme and F in PRT is a pseudo-random function, then VLH is a $(\text{dsize}, \text{qeq})$ -secure multi-map encryption scheme.*

Proof sketch: Let \mathcal{S}_{MM} be the simulator guaranteed to exist from the adaptive security of Σ_{MM} . Consider the simulator \mathcal{S}_{VLH} , that works as follows:

Simulating EMM: Given $\text{dsize}(\text{MM}) = m$, it computes $\theta = \sum_{i=1}^m r_i$, where $r_1, \dots, r_m \stackrel{\$}{\leftarrow} \{0, 1\}^k$, and outputs $\text{EMM} \leftarrow \mathcal{S}_{\text{MM}}(\theta)$.

Simulating Get: Given $\gamma = \text{qeq}(\text{MM}, \mathbf{q})$, it uses $\mathcal{S}_{\text{MM}}(\gamma, \rho)$, to simulate a get on EMM, where $\rho = \lambda + r$ and $r \stackrel{\$}{\leftarrow} \{0, 1\}^k$ (recall that λ is public). $\Sigma_{\text{MM}}.\text{Get}$ is an interactive protocol so here \mathcal{S}_{VLH} uses \mathcal{S}_{MM} to play the role of the client \mathbf{C} .

It remains to show that for all probabilistic polynomial-time adversaries \mathcal{A} , the probability that $\mathbf{Real}(k)$ outputs 1 is negligibly-close to the probability that $\mathbf{Ideal}(k)$ outputs 1. This can be shown with a standard sequence of games:

- **Game₀:** is the same as $\mathbf{Real}_{\Sigma, \mathcal{A}}(k)$.
- **Game₁:** is the same as **Game₀** except that the pseudo-random function F is replaced with a random function.
- **Game₂:** is the same as **Game₁** except that the values generated from the random function are replaced with values chosen uniformly at random.

- **Game₃**: is the same as **Game₁** except that **EMM** is replaced by the output of $\mathcal{S}_{\text{MM}}(\theta)$, the execution of $\Sigma_{\text{MM}}.\text{Get}((K, st, \ell), \text{EMM})$ is replaced with a simulated execution between $\mathcal{S}_{\text{MM}}(\gamma, \theta)$ and the adversary.

Note that **Game₃** is the **Ideal**(k) experiment. The probabilities that **Game₀** and **Game₁** output 1 are negligibly-close, otherwise the pseudo-randomness of F would be violated. The probabilities that **Game₁** and **Game₂** output 1 are equal by the properties of a random function. The probabilities that **Game₂** and **Game₃** output 1 are negligibly-close, otherwise the $(\text{trlen}, (\text{req}, \text{rlen}))$ -security of the multi-map encryption scheme Σ_{MM} would be violated. ■

6 DST: The Densest Subgraph Transform

In this section, we introduce a new data structure transformation called the *dense subgraph transform* (DST). Unlike the PRT which achieves efficient storage by increasing truncations (and therefore losing information), this new transform improves on the storage complexity of PRT without losing any information. The transformation is randomized and, surprisingly, we can show that, with high probability, it incurs no asymptotic storage overhead. Furthermore, we can also show that, for the case of Zipf-distributed multi-maps, it produces multi-maps that are asymptotically-smaller than the naive padding and truncating transforms described in Section 4.3.

The DST takes a multi-map MM as input and creates a new multi-map MM' that is volume-hiding. This new structure results from re-arranging the data in the input multi-map according to a random bi-partite graph. To ensure this re-arrangement is still efficiently queryable, we represent it using a pair of standard data structures which include a multi-map MM_G and a dictionary DX . As we will show, the storage complexity of the final representation depends on certain properties of the bi-partite graph which are, in turn, inherited from the original multi-map.

Below, we provide a high-level overview of our transformation. A more detailed description is given in Section 6.1. The overview is divided in two parts: (1) a variant for general multi-maps; and (2) a variant for what we refer to as *concentrated* multi-maps. Note that the transformation handles both cases but achieves better results for the later. We then provide a more detailed description in Figure 6.

General multi-maps. Given a multi-map MM we begin creating a bi-partite graph with \mathbb{L}_{MM} as the top vertices and a set of n empty bins as the bottom vertices. For each label/vertex ℓ in V_{top} , we randomly select t bins and insert in each bin a single value of the tuple $\text{MM}[\ell]$. Here, t is the maximum tuple size in the multi-map. If $\#\text{MM}[\ell] < t$, then some of the selected bins won't receive a value. At the end of this process, we pad all bins so that they all have the same size. Note that this process creates a bi-partite graph where the edges incident to some top vertex/label ℓ correspond to the bins selected for that label/vertex. We now create two data structures to represent and efficiently process this bi-partite graph. The first is a dictionary that maps bin identifiers to the bin's contents. The second is a multi-map that maps a label to the identifiers of the bins associated to it. To retrieve the values associated to a given label ℓ , we query the multi-map on ℓ to retrieve its t bin identifiers and then query the dictionary on each of the t bin identifier to retrieve the contents of the bins.

It is already clear from this high-level description that all labels will have exactly the same response length: $t \cdot \alpha$, where α is the maximum size of a bin. It can be shown that with the right choice of parameters, this transformation results in a small amount of padding compared to the naive approach.

Concentrated multi-maps. The storage overhead of our approach can be greatly improved when the multi-map satisfies a certain property we refer to as *concentration*. At a high level, a multi-map is concentrated if there exists a large number of values that appear in the tuples of a large number of labels. More formally, we define this property as follows.

Definition 6.1 (Concentrated multi-maps). *Let $\mu, \tau > 0$. We say that a multi-map MM is (μ, τ) -concentrated if there exists a set of μ labels $\ell_1, \dots, \ell_\mu \in \mathbb{L}_{\text{MM}}$ such that,*

$$\# \bigcap_{i=1}^{\mu} \text{MM}[\ell_i] = \tau.$$

*We refer to this set of labels as MM 's concentrated component and denote it $\widehat{\mathbb{L}}_{\text{MM}}$. Throughout, we will assume the existence of an efficient algorithm **FindComp** that takes as input a multi-map and outputs the multi-map's concentrated component $\widehat{\mathbb{L}}_{\text{MM}}$. If no such component exists, the algorithm outputs \perp . We sometimes refer to the value $(\mu - 1) \cdot \tau$ as the multi-map's concentration.*

A (μ, τ) -concentrated multi-map has μ labels with an intersection of size τ which means that there is some redundancy in the structure. Unfortunately, the previous approach does not take advantage of this since it stores all the values in the multi-map independently. To exploit this redundancy, we will proceed as follows. We will dedicate a random subset of the bins to store the tuple values of the multi-map's concentrated component. Because the component's tuples have a large intersection, we will avoid storing the same values over and over again. At a high-level, we modify the process as follows. We first choose a random subset of τ bins and store, in each one, one value from the intersection $\bigcap_{\ell \in \widehat{\mathbb{L}}_{\text{MM}}} \text{MM}[\ell]$. We then add an edge between each of these bins and the labels/vertices in the concentrated component. This will result in μ labels/vertices that share the same bins. Notice the improved storage overhead as we don't store the values in the intersection in multiple bins. For the remaining labels, we follow a similar process to the one presented in the generic case.

Finding the concentration component. Our DST transform relies on an efficient algorithm **FindComp** to find the concentration component of a multi-map. We now describe such an algorithm. The algorithm first determines the set of labels $\widehat{\mathbb{L}}_{\text{MM}}$ with tuples of size larger than $\bar{\tau}$. For $i \in [\lambda]$, it selects $\bar{\mu}$ labels uniformly at random with replacement from $\widehat{\mathbb{L}}_{\text{MM}}$. We refer to this set as $\bar{\mathbb{L}}_{\text{MM}}^i$. The algorithm then computes $\tau_i = \# \left(\bigcap_{\ell \in \bar{\mathbb{L}}_{\text{MM}}^i} \text{MM}[\ell] \right)$ where λ is the number of times the random selection is computed. The algorithm finally determines $\rho = \text{argmax}_{i \in [\lambda]} \{ \tau_i : \tau_i \in O(\bar{\tau}) \}$ and outputs $\widehat{\mathbb{L}}_{\text{MM}} = \bar{\mathbb{L}}_{\text{MM}}^\rho$ if such ρ exists and $\widehat{\mathbb{L}}_{\text{MM}} = \perp$ otherwise. Notice that the algorithm runs in $O(\lambda \cdot \bar{\mu})$ time. So it is sufficient to choose λ and $\bar{\mu}$ to be polynomial in m . In our setting, we need to set the parameters to align with the densest subgraph assumption (described in Definition 7.2) so we need $\bar{\mu} = \Omega(m^{0.5+\delta})$ and $\bar{\tau} = \tilde{O} \left(\sqrt{\max_{\ell \in \mathbb{L}_{\text{MM}}} \# \text{MM}[\ell]} \right)$, for some positive δ . We note that this algorithm is only an example and we believe that more efficient algorithms can be designed.

A storage optimization. Notice that the auxiliary multi-map MM_G associates to every label a randomly selected set of t bins. In particular this means that the size of MM_G is $O(m \cdot t)$ which could be rather large. Fortunately, storing the identifiers of each bin is not necessary. Instead, we can choose the bins to assign to a label using a pseudo-random function and store the key in MM_G . This will reduce the size of MM_G to $O(m)$.

6.1 Detailed Description

We now provide a detailed description of the DST. The pseudo-code is in Figure 6. The transform makes black-box use of two pseudo-random functions F and H .

Setup. The Setup algorithm takes as input a security parameter 1^k , an integer n and a multi-map MM . It instantiates a bi-partite graph $G = (V_{\text{top}}, V_{\text{bot}}, E)$ where the top vertices $V_{\text{top}} = \mathbb{L}_{\text{MM}}$ are the labels in MM , the bottom vertices V_{bot} are n empty bins denoted $\mathbf{B} = \{B_1, \dots, B_n\}$ and the set of edges E is empty.

The set of edges are generated as follows. Setup first computes the concentrated component of the multi-map $\widehat{\mathbb{L}}_{\text{MM}} := \text{FindComp}(\text{MM})$. If no concentrated component exists, FindComp outputs \perp . If $\#\widehat{\mathbb{L}}_{\text{MM}} \neq \perp$, it then pseudo-randomly chooses τ bins $\mathbf{B}' = \{B'_1, \dots, B'_\tau\}$, where $\tau = \#(\bigcap_{\ell \in \widehat{\mathbb{L}}_{\text{MM}}} \text{MM}[\ell])$. More precisely, it samples a k -bit value rand^* uniformly at random and chooses the bins indexed by the set

$$\left\{ F_{K_1}(\text{rand}^* \| 1), \dots, F_{K_1}(\text{rand}^* \| \tau) \right\}.$$

Note that all these τ positions have to be *distinct*. If not, then it keeps resampling a new k -bit value rand^* uniformly at random until no collisions are found. Note however that the probability p that no collision occurs, modeling F as a random function, is equal to

$$p = \prod_{i=0}^{\tau-1} \left(1 - \frac{i}{n}\right) \geq \left(1 - \frac{\tau}{n}\right)^\tau \approx e^{-\tau^2/n},$$

which tends to 1 when $\tau = o(n)$ — which aligns with our concrete parameters that we are going to detail in Section 6.4.

For all $\ell \in \widehat{\mathbb{L}}_{\text{MM}}$, it: (1) adds an edge between ℓ and every bin in \mathbf{B}' ; and (2) adds an edge between ℓ and $t - \tau$ bins outside of \mathbf{B}' . To do the latter, it indexes the bins in $\mathbf{B} \setminus \mathbf{B}'$ from 1 to $n - \tau$, samples a k -bit value rand_ℓ uniformly at random, and chooses the bins indexed by the set

$$\left\{ H_{K_2}(\text{rand}_\ell \| 1) + \text{slide}_1, \dots, H_{K_2}(\text{rand}_\ell \| t - \tau) + \text{slide}_{t-\tau} \right\},$$

where slide_i , for $i \in \{1, \dots, t - \tau\}$, is an integer used to deterministically map back the smaller output of H in $[n - \tau]$ to the corresponding bin identifier in $[n]$ and is computed as follows. First, it orders the set of bins in \mathbf{B}' in a numerical order such that

$$\mathbf{B}' = \left(B_{\text{pos}_1}, \dots, B_{\text{pos}_\tau} \right),$$

where $\text{pos}_i < \text{pos}_j$, for $i, j \in [\tau]$. Then it defines the following quantities based on which the slide value is determined—refer to Figure 3 for an illustration of the computation,

$$\text{gap}_i = \begin{cases} [1, \text{pos}_i - 1] & \text{if } i = 1 \\]\text{pos}_{i-1} - (i - 1), \text{pos}_i - (i - 1)[& \text{if } i \in \{2, \dots, \tau\} \\]\text{pos}_{i-1} - (i - 1), n - \tau] & \text{if } i = \tau + 1 \end{cases}$$

Then, for $i \in \{1, \dots, n - \tau\}$, identify $j \in \{1, \dots, \tau + 1\}$ such that $H_{K_2}(\text{rand}_\ell \| i) \in \text{gap}_j$, then set $\text{slide}_i = j - 1$.

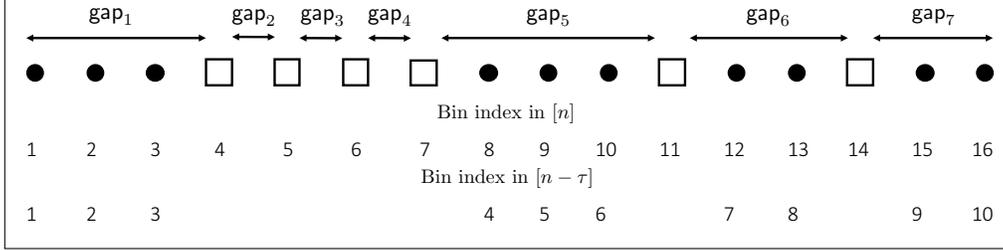


Figure 3: Gaps computation for $n = 16$ and $\tau = 6$. \square denotes bins being part of $\mathbf{B}' = \{B_4, B_5, B_6, B_7, B_{11}, B_{14}\}$ while \bullet denotes bins in $\mathbf{B} \setminus \mathbf{B}'$.

Note that rand has also to be chosen in such a way that the selected $t - \tau$ positions are distinct. If not, similarly to above, it resamples a new k -bit value uniformly at random until no collision occurs. The probability that no collision occurs is approximately equal to $e^{-(t-\tau)^2/n}$ which tends to 1 when $t = o(n)$ —which aligns with our concrete parameterization as we are going to detail in Section 6.4.

For each $\ell \notin \widehat{\mathbb{L}}_{\text{MM}}$, it samples a k -bit value rand_ℓ uniformly at random and adds an edge between ℓ and all bins indexed by the set

$$\left\{ F_{K_1}(\text{rand}_\ell \| 1), \dots, F_{K_1}(\text{rand}_\ell \| t) \right\}.$$

Again, if a collision is found, it keeps resampling a new k -bit value uniformly at random until all positions are distinct. The probability that no collision occurs is approximately equal to $e^{-t^2/n}$. Notice that at the end of this process, each vertex has degree exactly t .

Now **Setup** will use the graph to load the bins in V_{bot} as follows. For each $\ell \notin \widehat{\mathbb{L}}_{\text{MM}}$, it stores one value from the tuple $\text{MM}[\ell]$ in one of the bins that are incident to ℓ . When inserting into a bin, the algorithm concatenates each value with ℓ (this will be helpful at query time). If $\#\text{MM}[\ell] < t$, then some of the incident bins will not receive any value. For all $\ell \in \widehat{\mathbb{L}}_{\text{MM}}$, it stores one element from $\text{MM}[\ell] \setminus \bigcap_{\ell \in \widehat{\mathbb{L}}_{\text{MM}}} \text{MM}[\ell]$ in the bins from $\mathbf{B} \setminus \mathbf{B}'$ that are incident to ℓ —again concatenating each value with ℓ . Also if $\#\text{MM}[\ell] < t$, then some of the incident bins will not receive any value. Finally, it stores each value from the set

$$\mathbf{r}' := \bigcap_{\ell \in \widehat{\mathbb{L}}_{\text{MM}}} \text{MM}[\ell].$$

in a distinct bin in \mathbf{B}' in such a way that every bin in \mathbf{B}' will contain one value in \mathbf{r}' . Here, the algorithm concatenates the values with \star . The algorithm then pads all the bins to have the same size.

Finally, it creates a dictionary DX and a multi-map MM_G . The dictionary maps bin identifiers to bin contents. The multi-map MM_G maps labels $\ell \notin \widehat{\mathbb{L}}_{\text{MM}}$ to rand_ℓ and labels $\ell \in \widehat{\mathbb{L}}_{\text{MM}}$ to $(\text{rand}_\ell, \text{rand}^*)$. It outputs $\text{MM}' = (\text{MM}_G, \text{DX})$.

The storage complexity of MM' is $O(m + n \cdot \lambda)$, where λ is the maximum load of a bin.

Get. Get operations on $\text{MM}' = (\text{MM}_G, \text{DX})$ work as follows. Given a label ℓ , we first query MM_G on ℓ . If $\ell \notin \widehat{\mathbb{L}}_{\text{MM}}$, then MM_G returns rand_ℓ from which we compute the bin identifiers $\{F_{K_1}(\text{rand}_\ell \| i)\}_{i=1}^t$. We can then query DX on the bin identifiers to recover the bins and output the elements concatenated with ℓ . If $\ell \in \widehat{\mathbb{L}}_{\text{MM}}$, MM_G returns a pair $(\text{rand}_\ell, \text{rand}^*)$ from which we can compute the sets

$$\left\{ F_{K_1}(\text{rand}^* \| i) \right\}_{i=1}^\tau \quad \text{and} \quad \left\{ H_{K_2}(\text{rand}_\ell \| i) + \text{slide}_i \right\}_{i=1}^{t-\tau}$$

which we, in turn, use to query DX and recover the bins. From these bins the algorithm recovers the elements concatenated with ℓ and \star . The complexity of gets is $O(t \cdot \lambda)$ where, again, λ is the maximum load of a bin.

6.2 Analyzing the Load of a Bin

As seen in the previous Section, an important quantity to evaluate the query and storage efficiency of our transformation is the maximum load of a bin. In this Section, we will show that, with high probability, the maximum load can be upper bounded by $(N - N_{\text{ds}})/n$ where N_{ds} is the size of the concentrated component and n is the number of bins. Before stating our result, we recall a generalization of Chernoff's inequality for the binomial distribution.

Lemma 6.2. *Let X_1, \dots, X_m be independent random variables over $\{0, 1\}$ such that $\Pr[X_i = 1] = p_i$ and $\Pr[X_i = 0] = 1 - p_i$. If $X = X_1 + \dots + X_m$, then*

$$\Pr[X \geq \mathbb{E}[X] + \theta] \leq \exp\left(-\frac{\theta^2}{2(\mathbb{E}[X] + \theta/3)}\right).$$

Theorem 6.3. *With probability at least $1 - \varepsilon$, the maximum load of a bin is at most*

$$\frac{N - N_{\text{ds}}}{n} + \frac{\ln(1/\varepsilon)}{3} \left(1 + \sqrt{1 + \frac{18(N - N_{\text{ds}})}{n \cdot \ln(1/\varepsilon)}}\right),$$

where $N_{\text{ds}} = (\mu - 1) \cdot \tau$.

Proof. Without loss of generality, the label space \mathbb{L}_{MM} can be rewritten as

$$\mathbb{L}_{\text{MM}} = \left(\ell_1, \dots, \ell_{m-\mu}, \ell_{m-\mu+1}, \dots, \ell_m\right),$$

where the concentrated component $\widehat{\mathbb{L}}_{\text{MM}} = (\ell_{m-\mu+1}, \dots, \ell_m)$. For all $i \in [n]$ and $j \in [m]$ let $X_{i,j}$ be the random variable that outputs 1 if the i th bin is assigned to the j th label. Notice that for all $i \in [n]$ and $j \in [m - \mu]$, $X_{i,j}$ is a Bernoulli distribution with parameter n_{ℓ_j}/n , where $n_{\ell_j} = \#\text{MM}[\ell_j]$. This follows from the fact that for all ℓ outside the concentrated component **Setup**

Let $n \in \mathbb{N}$ be a public parameter, $F : \{0, 1\}^k \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ and $H : \{0, 1\}^k \times \{0, 1\}^* \rightarrow \{0, 1\}^{n'}$ be two pseudo-random functions with $n' < n$. Consider the transform DST defined as follows:

- $\text{DST}(1^k, n, \text{MM})$:

1. instantiate an empty dictionary DX, an empty multi-map MM_G , and a bi-partite graph $G = ((\mathbb{L}_{\text{MM}}, \mathbf{B}), \text{E})$ where $\mathbf{B} = (B_1, \dots, B_n)$ and $\text{E} = \emptyset$;
2. compute $\widehat{\mathbb{L}}_{\text{MM}} \leftarrow \text{FindComp}(\text{MM})$, set $\tau = \#(\bigcap_{\ell \in \widehat{\mathbb{L}}_{\text{MM}}} \text{MM}[\ell])$ and $t := \max_{\ell \in \mathbb{L}_{\text{MM}}} \text{MM}[\ell]$;
3. sample two keys $K_1 \xleftarrow{\$} \{0, 1\}^k$ and $K_2 \xleftarrow{\$} \{0, 1\}^k$;
4. for all $\ell \in \mathbb{L}_{\text{MM}} \setminus \widehat{\mathbb{L}}_{\text{MM}}$,
 - (a) sample $\text{rand}_\ell \xleftarrow{\$} \{0, 1\}^k$ and output

$$(i_1, \dots, i_t) := \left\{ F_{K_1}(\text{rand}_\ell \| 1), \dots, F_{K_1}(\text{rand}_\ell \| t) \right\},$$

if there exist distinct $i, j \in [t]$ for which $i_i = i_j$ redo the sampling. Add to E

$$\left\{ (\ell, i_j) : j \in [t] \right\};$$

- (b) parse $\text{MM}[\ell]$ as (r_1, \dots, r_{n_ℓ}) and put $r_j \| \ell$ in B_{i_j} for all $j \in [n_\ell]$;
5. if $\widehat{\mathbb{L}}_{\text{MM}} \neq \perp$, sample $\text{rand}^* \xleftarrow{\$} \{0, 1\}^k$ and set $\mathbf{B}' = (B_{i_1}, \dots, B_{i_\tau})$ where

$$(i_1, \dots, i_\tau) := \left\{ F_{K_1}(\text{rand}^* \| 1), \dots, F_{K_1}(\text{rand}^* \| \tau) \right\},$$

if there exist distinct $i, j \in [\tau]$ for which $i_i = i_j$ redo the sampling. Otherwise set $\mathbf{B}' = \perp$;

6. compute

$$\mathbf{r}' := \bigcap_{\ell \in \widehat{\mathbb{L}}_{\text{MM}}} \text{MM}[\ell] = (r'_1, \dots, r'_\tau);$$

7. put $r'_j \| \star$ in B_j for all $j \in [\tau]$ and $B_j \in \mathbf{B}'$;
8. for all $\ell \in \widehat{\mathbb{L}}_{\text{MM}}$, add to E

$$\left\{ (\ell, j) : B_j \in \mathbf{B}' \right\};$$

Figure 4: DST: The Dense Subgraph Transform (Part 1).

• $\text{DST}(1^k, n, \text{MM})$:

9. for all $\ell \in \widehat{\mathbb{L}}_{\text{MM}}$,

(a) sample $\text{rand}_\ell \xleftarrow{\$} \{0, 1\}^k$ and output

$$(i_1, \dots, i_{t-\tau}) := \left\{ H_{K_2}(\text{rand}_\ell \| 1), \dots, H_{K_2}(\text{rand}_\ell \| t - \tau) \right\},$$

if there exist distinct $i, j \in [t - \tau]$ for which $i_i = i_j$, redo the sampling. Add to E

$$\left\{ (\ell, i_j + \text{slide}_j) : j \in [t - \tau] \right\};$$

where slide_j is computed as follows

i. order \mathbf{B}' in a numerical order such that

$$\mathbf{B}' := (B_{\text{pos}_1}, \dots, B_{\text{pos}_\tau});$$

ii. if $i_j \in [1, \text{pos}_1]$, set $\text{slide}_j = 0$;

iii. if $i_j \in]\text{pos}_{i-1} - (i-1), \text{pos}_i - (i-1)[$, set $\text{slide}_j = i - 1$, for any $i \in \{2, \dots, \tau\}$;

iv. if $i_j \in]\text{pos}_\tau - \tau, n - \tau[$, set $\text{slide}_j = \tau$;

(b) parse $\text{MM}[\ell]$ as (r_1, \dots, r_{n_ℓ}) ;

(c) if $n_\ell > \tau$, then for all $r_j \in \text{MM}[\ell] \setminus \mathbf{r}'$, then put $r_j \| \ell$ in $B_{i_j + \text{slide}_j}$;

10. set $\theta = \max_{i \in [n]} \#B_i$ and set for all $i \in [n]$

$$B_i = (B_i, \perp_1, \dots, \perp_{\theta - \#B_i});$$

11. for all $i \in [n]$, set $\text{DX}[i] = B_i$;

12. for all $\ell \in \mathbb{L}_{\text{MM}}$, if $\ell \in \widehat{\mathbb{L}}_{\text{MM}}$, set $\text{MM}[\ell] := (\text{rand}_\ell, \text{rand}^*)$, otherwise set $\text{MM}[\ell] := \text{rand}_\ell$;

13. output the key $K = (K_1, K_2)$ and $\text{MM}' = (\text{DX}, \text{MM}_G)$.

• $\text{Get}(K, \ell, \text{MM})$:

1. parse K as (K_1, K_2) and MM as (DX, MM_G) and instantiate an empty set **Result**;

2. if $\text{MM}_G[\ell] = \text{rand}$, then

(a) add $\text{DX}[\ell_i]$ to **Result**, where for all $i \in [t]$,

$$\ell_i := F_{K_1}(\text{rand} \| i);$$

(b) keep all values of the form $\cdot \| \ell$;

3. if $\text{MM}_G[\ell] = (\text{rand}, \text{rand}^*)$, then

(a) add $\text{DX}[\ell_i]$ to **Result**, where for all $i \in [t - \tau]$,

$$\ell_i := H_{K_2}(\text{rand} \| i) + \text{slide}_i,$$

and for all $i \in [\tau]$,

$$\ell_i := F_{K_1}(\text{rand}^* \| i);$$

(b) keep all values of the form $\cdot \| \ell$ or $\cdot \| \star$;

4. output **Result**.

Figure 5: DST: The Dense Subgraph Transform (Part 2).

assigns bins uniformly at random. This is not the case, however, for labels outside the concentrated component. For all $i \in [n]$ and $j \in \{m - \mu + 1, \dots, m\}$, $X_{i,j}$ is a Bernoulli distribution with parameter n'_{ℓ_j}/n where

$$n'_{\ell_j} = \begin{cases} 0 & \text{if } n_{\ell_j} = \tau \\ n_{\ell_j} - \tau & \text{otherwise.} \end{cases}$$

This follows from the fact that **Setup** assigns ℓ to $n'_{\ell_j} - \tau$ bins chosen uniformly at random. Finally, **Setup** also select τ bins uniformly at random to form \mathbf{B}' . Let Y_i be the random variable that outputs 1 if the i th bin is selected and 0 otherwise. It is clear that Y_i is a Bernoulli distribution with parameter τ/n .

Let $X_i = Y_i + \sum_{j=1}^m X_{i,j}$ and note that X_i is the load of the i th bin. Its expectation is

$$\begin{aligned} \mathbb{E}[X_i] &= \sum_{j=1}^{m-\mu} \frac{n_{\ell_j}}{n} + \sum_{j=m-\mu+1}^m \frac{n'_{\ell_j}}{n} + \frac{\tau}{n} \\ &= \sum_{j=1}^{m-\mu} \frac{n_{\ell_j}}{n} + \sum_{j=m-\mu+1}^m \frac{n_{\ell_j} + \tau}{n} - \frac{\mu \cdot \tau}{n} + \frac{\tau}{n} \\ &= \sum_{j=1}^{m-\mu} \frac{n_{\ell_j}}{n} + \sum_{j=m-\mu+1}^m \frac{n_{\ell_j}}{n} - \frac{\mu \cdot \tau}{n} + \frac{\tau}{n} \\ &= \frac{N}{n} - \frac{N_{\text{ds}}}{n} \end{aligned}$$

where $N_{\text{ds}} = (\mu - 1) \cdot \tau$ and $N = \sum_{\ell \in \mathbb{L}_{\text{MM}}} n_{\ell}$. The third equality follows from the definition of n'_{ℓ} above. By Lemma 6.2 and by solving the quadratic equation

$$\frac{\theta^2}{2(\mathbb{E}[X] + \theta/3)} = \ln(1/\varepsilon)$$

we have

$$\Pr \left[X < \frac{N - N_{\text{ds}}}{n} + \frac{\ln(1/\varepsilon)}{3} \left(1 + \sqrt{1 + \frac{18(N - N_{\text{ds}})}{n \cdot \ln(1/\varepsilon)}} \right) \right] \geq 1 - \varepsilon.$$

■

6.3 Query and Storage Efficiency

We now give the storage and query efficiency of the DST transform.

Storage efficiency. The output of DST consists of a multi-map MM_G and a dictionary DX . The multi-map MM_G has tuples of size 1 or 2 depending on the label. That is, the size of the multi-map is upper bounded by $2m$. The dictionary DX stores the content of the padded bins. From Theorem 6.3 and the union bound, we have that the size of the dictionary is at most

$$N - N_{\text{ds}} + \frac{n \cdot \ln(1/\varepsilon)}{3} \cdot \left(1 + \sqrt{1 + \frac{18(N - N_{\text{ds}})}{n \cdot \ln(1/\varepsilon)}} \right)$$

with probability $1 - n \cdot \varepsilon$.

Get efficiency. The `Get` algorithm first retrieves either a random value or a pair of random values from MM_G . In the former case, t PRF evaluations are computed and t bins are retrieved. In the later case, $2t$ PRF evaluations are computed (using F and H) and t bins are retrieved.³ Assuming that both MM_G and DX are data structures with optimal query complexity, the `Get` query complexity is at most

$$t \cdot \frac{N - N_{\text{ds}}}{n} + \frac{t \cdot \ln(1/\varepsilon)}{3} \cdot \left(1 + \sqrt{1 + \frac{18(N - N_{\text{ds}})}{n \cdot \ln(1/\varepsilon)}}\right)$$

with probability $1 - n \cdot \varepsilon$.

6.4 Concrete Parameters

In this Section, we propose concrete parameters for the `DST`. In particular, we will be interested in parameters that guarantee better storage overhead than the naive padding transform. Note that we do not compare to the naive truncating approach since the `DST` does not lose any information.

General multi-maps. Recall that the naive padding transform has a storage overhead

$$S_{\text{NV}} \stackrel{\text{def}}{=} m \cdot \max_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell] = \Omega(N),$$

where $N \stackrel{\text{def}}{=} \sum_{\ell \in \mathbb{L}} \#\text{MM}[\ell]$. From Theorem 6.3, we have the following corollary.

Corollary 6.4. *Let $n \geq 1$ and $m \geq 0$. If $N > n \log n$, then with probability at least $1 - 1/e^{N/5n}$, the size of the resulting multi-map is at most $O(N)$.*

Notice that if the original multi-map is $\mathcal{Z}_{m,1}$ -distributed (but not necessarily concentrated), then $S_{\text{NV}} = N \cdot m / H_{m,1}$ where $H_{m,1} = \Theta(\log m)$ is the harmonic number (please refer to Section 4). It follows that, in this case, $N = o(S_{\text{NV}})$ so the storage overhead of `DST` is *small-o* of the overhead of the naive padding transform.

Concentrated multi-maps. We now consider a multi-map MM with a concentrated component of size $(\mu - 1) \cdot \tau$. We show below that in this case, the storage overhead induced by `DST` can be considerably smaller than the storage overhead of the naive padding transform. The following Corollary is a consequence of Theorem 6.3.

Corollary 6.5. *Let $n \geq 1$ and $m \geq 0$. If $N > n \log n$,*

$$\mu = O\left(\sqrt{m} \cdot \text{polylog}(m)\right) \quad \text{and} \quad \tau = O\left(\text{polylog}(m)\right),$$

then with probability at least $1 - 1/e^{N/5n}$, the size of the resulting multi-map is at most

$$O\left(N - \sqrt{m} \cdot \text{polylog}(m)\right).$$

As above, if the original multi-map MM is $\mathcal{Z}_{m,1}$ -distributed, then the storage overhead of `DST` is *small-o* of the overhead of the naive padding transform.

³Note that the computation of the `slidei`'s is $O(t)$ which is dominated by the total query complexity.

A remark on security. As we will see in Section 7, the parameters μ and τ have to satisfy certain constraints for our multi-map encryption scheme to be secure. In particular, the parameters have to be chosen in such a way that they verify the densest subgraph assumption which we detail in Definition 7.2. We note here that to satisfy both this assumption and the constraints of Corollary 6.5, it is sufficient that

$$\mu = O\left(\sqrt{m} \cdot \text{polylog}(m)\right), \quad \tau = O(t) = O\left(\text{polylog}(m)\right) \quad \text{and} \quad n = \Theta\left(\text{polylog}(m)\right).$$

Note that this is only an example and is not the only choice of parameters that can be used.

7 AVLH: Advanced Volume Hiding Multi-Map Encryption Scheme

In this Section, we use the DST to construct a volume-hiding multi-map encryption scheme. Our construction is described in detail in Figure 6 and works as follows.

Overview. The construction, $\text{AVLH} = (\text{Setup}, \text{Get})$, makes black-box use of an underlying response-hiding dictionary encryption scheme $\text{STE}_{\text{DX}}^{\text{RH}} = (\text{Setup}, \text{Get})$. AVLH.Setup takes as input a security parameter k , a public parameter n , and a multi-map MM . It first applies the DST transform on MM which results in a key $K_1 = (K_{1,1}, K_{1,2})$ and two structures: a multi-map MM_G and a dictionary DX . It then encrypts the dictionary DX , resulting in an encrypted dictionary EDX , a state st_{DX} and a key K_2 . It finally outputs a key $K = (K_1, K_2)$, a state $st = (\text{MM}_G, st_{\text{DX}})$ and an encrypted multi-map $\text{EMM} = \text{EDX}$. To execute AVLH.Get , the client differentiates two cases: if $\text{MM}_G[\ell]$ is a tuple composed of a single value rand , then the client and server execute the $\text{STE}_{\text{DX}}^{\text{RH}}.\text{Get}$ on ℓ_i where ℓ_i is a new label equal to $F_{K_{1,1}}(\text{rand}||i)$, for all $i \in [t]$, and $t = \max_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$. In this case the client \mathbf{C} only outputs values of the form $\cdot||\ell$. Otherwise, if $\text{MM}_G[\ell]$ is a tuple composed of a pair $(\text{rand}, \text{rand}^*)$, then the client and server execute $\text{STE}_{\text{DX}}^{\text{RH}}.\text{Get}$ on ℓ_i where now ℓ_i is equal to $F_{K_{1,1}}(\text{rand}^*||i)$ for $i \in [\tau]$, and $H_{K_{1,2}}(\text{rand}||i) + \text{slide}_i$ for all $i \in \{1, \dots, t - \tau\}$. Note that slide_i , for which the computation was detailed in Section 6.1, is used to deterministically map the smaller output of H in $[n - \tau]$ into a value in $[n]$. In this case, the client \mathbf{C} only outputs values of the form $\cdot||\ell$ or $\cdot||\star$.

Efficiency. Assuming that $\text{STE}_{\text{DX}}^{\text{RH}}$ is an optimal-time dictionary encryption scheme [15, 10, 7, 1], the get complexity of AVLH is $O(t \cdot \lambda)$ where $t = \max_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$ and λ is the load of a bin. Given the parameters detailed in the previous section and if $N > n \cdot \log n$ then the get complexity is

$$O\left(\frac{N - \sqrt{m} \cdot \text{polylog}(m)}{\text{polylog}(m)}\right).$$

when

$$\mu = O\left(\sqrt{m} \cdot \text{polylog}(m)\right), \quad \tau = O(t) = O\left(\text{polylog}(m)\right) \quad \text{and} \quad n = \Theta\left(\text{polylog}(m)\right).$$

The storage overhead of AVLH is, with high probability,

$$O\left(\sum_{i=1}^n \#\text{DX}[\ell_i]\right) = O(n \cdot \lambda) = O(N - \sqrt{m} \cdot \text{polylog}(m)).$$

7.1 Security

We will now study the security of our construction. More precisely, we will show that it is volume-hiding in the sense that its query leakage does not include the response length. The proof relies on a computational assumption known as the densest subgraph assumption. We first recall this assumption and then proceed to stating our security theorem.

The densest subgraph problem. The hardness of the (decisional) *densest subgraph problem* was first used by Applebaum, Barak, and Wigderson in [2] to design public-key encryption schemes based on new assumptions. It was later used by Arora et al. [3] to study the hardness of financial products. Informally, the DSP asks whether it is possible to distinguish between a random regular bi-partite graph and a random regular bi-partite graph with a planted random subgraph.

Definition 7.1 (The (decision) densest subgraph problem.). *Let $m, n, t, \mu, \nu, \tau > 0$. The decisional unbalanced expansion problem is to distinguish between the two following distributions:*

- \mathcal{R} samples an (m, n, t) -bi-partite graph uniformly at random. In other words, for each vertex in V_{top} it samples t neighbors from V_{bot} uniformly at random.
- \mathcal{P} is obtained as follows. First, two sets $T \subset V_{\text{top}}$ and $B \subset V_{\text{bot}}$, such that $\#T = \mu$ and $\#B = 2\nu$, are sampled uniformly at random. Then, for each vertex in T , we choose $t - \tau$ random neighbors in V_{bot} and τ random neighbors in B . For each vertex in $V_{\text{top}} \setminus T$, we choose t random neighbors in V_{bot} .

The following hardness assumption, used in [2, 3], is based on state-of-the-art algorithms of Bhaskara, Charikar, Chlamtac, Feige, and Vijayaraghavan in [6].

Definition 7.2 (The DSP assumption). *There is no $\varepsilon > 0$ and PPT adversary \mathcal{A} that can distinguish between \mathcal{R} and \mathcal{P} with advantage ε when*

$$n = o(m \cdot t), \quad \left(\frac{\mu \cdot \tau^2}{\nu}\right)^2 = o\left(\frac{m \cdot t^2}{n}\right), \quad \nu = \Omega(n^{0.5+\delta}),$$

$$\mu = \Omega(m^{0.5+\delta}) \quad \text{and} \quad \tau = \tilde{O}(\sqrt{t})$$

for some positive δ .

Leakage profile. We now describe the leakage profile of AVLH assuming $\text{STE}_{\text{DX}}^{\text{RH}}$ is instantiated with one of the standard optimal-time dictionary encryption schemes [15, 10, 7, 1] all of which have leakage profile

$$\Lambda_{\text{DX}} = (\mathcal{L}_{\text{S}}, \mathcal{L}_{\text{Q}}) = (\text{trlen}, \text{qeq}).$$

Theorem 7.3. *If $\text{STE}_{\text{DX}}^{\text{RH}}$ is a $(\text{trlen}, \text{qeq})$ -secure dictionary encryption scheme, F and H are pseudo-random functions, and the DSP assumptions holds, then AVLH is a $((\text{trlen}, \text{conc}), \text{qeq})$ -secure multi-map encryption scheme; where conc is the leakage pattern that outputs a multi-map's concentration.*

Proof sketch: Let \mathcal{S}_{DX} be the simulator guaranteed to exist from the adaptive security of Σ_{DX} . Instantiate an empty graph $G = ((V_{\text{top}}, V_{\text{bot}}), E)$. Consider the simulator $\mathcal{S}_{\text{AVLH}}$, that works as follows:

Simulating EMM: Given $\theta = \text{trlen}(\text{MM})$ and $N_{\text{ds}} = \text{conc}(\text{MM})$, it computes and outputs $\text{EDX} \leftarrow \mathcal{S}_{\text{DX}}(\theta - N_{\text{DS}})$.

Simulating Get: Given $\gamma = \text{qeq}(\text{MM}, \mathbf{q})$, $\mathcal{S}_{\text{AVLH}}$ parses γ as a binary matrix M . If there exists no $j \in [i - 1]$ such that $M[q_i] = M[q_j]$ then it computes $(i_1, \dots, i_t) \xleftarrow{\$} [n]$ and adds the edges (i, i_l) to \mathbf{E} , the vertex i to \mathbf{V}_{top} and the vertex i_l to \mathbf{V}_{bot} , for all $l \in [t]$. Otherwise if there exists $j \in [i - 1]$ such that $M[q_j] = M[q_i]$, then it identifies (i_1, \dots, i_t) such that $(j, i_l) \in \mathbf{E}$, for $l \in [t]$. It then computes, for all $l \in [t]$, $\mathcal{S}_{\text{DX}}(\gamma_l)$, where $\gamma_j = \text{qeq}(\text{DX}, (i_1, \dots, i_j))$. Recall that $\Sigma_{\text{MM}}.\text{Get}$ is an interactive protocol so $\mathcal{S}_{\text{AVLH}}$ is using \mathcal{S}_{DX} to play the role of the client \mathbf{C} .

It remains to show that for all probabilistic polynomial-time adversaries \mathcal{A} , the probability that $\mathbf{Real}(k)$ outputs 1 is negligibly-close to the probability that $\mathbf{Ideal}(k)$ outputs 1. This can be shown with a standard sequence of games:

- **Game₀:** is the same as $\mathbf{Real}_{\Sigma, \mathcal{A}}(k)$.
- **Game₁:** is the same as **Game₀** except that the pseudo-random function F is replaced with a random function f .
- **Game₂:** is the same as **Game₁** except that the values of the random function f are sampled uniformly at random (and recalled when needed).
- **Game₃:** is the same as **Game₂** except that the pseudo-random function H is replaced with a random function h .
- **Game₄:** is the same as **Game₃** except that the values of the random function h are sampled uniformly at random (and recalled when needed).
- **Game₅:** is the same as **Game₂** except that the values previously generated from $[n]$ or from $[n - \tau]$, are now solely generated from $[n]$.
- **Game₆:** is the same as **Game₃** except that **EMM** is replaced by the output of $\mathcal{S}_{\text{DX}}(\text{trlen}^*)$, the execution of $\Sigma_{\text{MM}}.\text{Get}((K, st, \ell), \text{EMM})$ is replaced with a simulated execution between $\mathcal{S}_{\text{DX}}(\gamma_l)$, and the adversary, for $l \in [t]$ and where i_l are computed as above.

Note that **Game₆** is the $\mathbf{Ideal}(k)$ experiment. The probabilities that **Game₀** and **Game₁** output 1 are negligibly-close, otherwise the pseudo-randomness of F would be violated. The probabilities that **Game₁** and **Game₂** output 1 are equal by the definition of a random function. The probabilities that **Game₂** and **Game₃** output 1 are negligibly-close, otherwise the pseudo-randomness of H would be violated. The probabilities of **Game₃** and **Game₄** are equal, again, by the definition of a random function. The probabilities that **Game₄** and **Game₅** output 1 are negligibly-close, otherwise the DSP assumption would be violated. To see why, notice that that the bi-partite graph \mathbf{G} in **Game₄** is distributed similarly to the distribution \mathcal{P} whereas the distribution of the bi-partite graph \mathbf{G} in **Game₅** is distributed similarly to the distribution \mathcal{R} (see Definition 7.1). The probabilities that **Game₅** and **Game₆** output 1 are negligibly-close, otherwise the $(\text{trlen}, \text{qeq})$ -security of the dictionary encryption scheme Σ_{DX} would be violated. ■

Let $\text{STE}_{\text{DX}}^{\text{RH}} = (\text{Setup}, \text{Get})$ be a response-hiding dictionary encryption scheme and DST the densest subgraph transform. Consider the scheme $\text{AVLH} = (\text{Setup}, \text{Get})$ defined as follows:

- $\text{Setup}(1^k, n, \text{MM})$:

1. generate a DST -transform of MM by computing

$$(K_1, \text{MM}_G, \text{DX}) \leftarrow \text{DST}(1^k, n, \text{MM});$$

2. encrypt DX by computing

$$(K_2, st_{\text{DX}}, \text{EDX}) \leftarrow \text{STE}_{\text{DX}}^{\text{RH}}.\text{Setup}(1^k, \text{DX});$$

3. set $K = (K_1, K_2)$, $st = (\text{MM}_G, st_{\text{DX}})$, and $\text{EMM} = \text{EDX}$ and output (K, st, EMM) .

- $\text{Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM})$:

1. \mathbf{C} parses K as $((K_{1,1}, K_{1,2}), K_2)$, st as $(\text{MM}_G, st_{\text{DX}})$ and \mathbf{S} parses EMM as EDX ;

2. if $\text{MM}_G[\ell] = \text{rand}$, then

- (a) \mathbf{C} and \mathbf{S} execute $\text{STE}_{\text{DX}}^{\text{RH}}.\text{Get}_{\mathbf{C}, \mathbf{S}}((K_2, st_{\text{DX}}, \ell_i), \text{EDX})$, for all $i \in [t]$, where

$$\ell_i := F_{K_{1,1}}(\text{rand} \parallel i);$$

- (b) \mathbf{C} outputs values of the form $\cdot \parallel \ell$;

3. if $\text{MM}_G[\ell] = (\text{rand}, \text{rand}^*)$, then

- (a) \mathbf{C} and \mathbf{S} execute $\text{STE}_{\text{DX}}^{\text{RH}}.\text{Get}_{\mathbf{C}, \mathbf{S}}((K_2, st_{\text{DX}}, \ell_i), \text{EDX})$, where for all $i \in [\tau]$,

$$\ell_i := F_{K_{1,1}}(\text{rand}^* \parallel i),$$

and for all $i \in \{1, \dots, t - \tau\}$,

$$\ell_i := H_{K_{1,2}}(\text{rand} \parallel i) + \text{slide}_i,$$

where slide_j is computed as follows

- i. order

$$\left\{ F_{K_{1,1}}(\text{rand}^* \parallel i) \right\}_{i \in [\tau]}$$

as $(\text{pos}_1, \dots, \text{pos}_\tau)$;

- ii. if $H_{K_{1,2}}(\text{rand} \parallel i) \in [1, \text{pos}_1]$, set $\text{slide}_i = 0$;

- iii. if $H_{K_{1,2}}(\text{rand} \parallel i) \in]\text{pos}_{j-1} - (j-1), \text{pos}_j - (j-1)[$, set $\text{slide}_i = j - 1$, for any $j \in \{2, \dots, \tau\}$;

- iv. if $H_{K_{1,2}}(\text{rand} \parallel i) \in]\text{pos}_\tau - \tau, n - \tau]$, set $\text{slide}_i = \tau$;

- (b) \mathbf{C} outputs all values of the form $\cdot \parallel \ell$ or $\cdot \parallel \star$.

Figure 6: AVLH: An Advanced Volume Hiding Multi-Map Encryption Scheme.

Improving communication complexity. The communication (query) complexity of AVLH is equal to $O(t \cdot \lambda)$ where λ is the size of the bin and t the maximum response length. In the following we introduce a simple modification of AVLH such that the communication complexity becomes sub-linear in λ .

At a high level, the idea consists of replacing the retrieval of the entire bin’s content by an oblivious retrieval that only fetches the value of interest (note that a bin will always contain at most one value associated to any label). Therefore this technique would reduce the overhead from λ to the overhead of a single oblivious access into an array of size λ . The (informal) modified AVLH works as follows. At setup time, we parse the content of each bin as an array (a RAM) and encrypt it using a state-of-the-art ORAM algorithm. Note that now, instead of using a response-hiding dictionary, we use a response-revealing one. The get algorithm works similarly to the one in AVLH except that the dictionary’s get algorithm outputs an ORAM that we access separately. In terms of efficiency, the communication complexity becomes $O(t \cdot \log^2 \lambda)$ assuming that we use Path ORAM [38] as the underlying ORAM. Note that we can achieve better communication complexity by leveraging techniques from [25]. The storage complexity however remains the same since Path ORAM does not asymptotically increase the load of the bin.

8 Dynamic Volume Hiding Multi-Map Encryption Schemes

In this section, we show how to extend both VLH and AVLH to be dynamic. In particular, we will be interested in the following class of updates:

- *tuple addition*: this update operation adds a new tuple (ℓ, \mathbf{v}) to the multi-map where ℓ is a label that was not part of the original label space \mathbb{L}_{MM} .
- *tuple deletion*: this update operation removes an entire label/tuple pair (ℓ, \mathbf{v}) from the multi-map.
- *editing*: this update operation modifies the content of a specific tuple \mathbf{v} associated to ℓ by replacing an old value $v_{\text{old}} \in \mathbf{v}$ by a new one v_{new} .

In particular, we do not consider updates that add or remove a value to/from an existing tuple in the multi-map. In the following, we detail how to extend VLH to handle these three update operations and AVLH to handle the third update operation.

8.1 VLH^d: a dynamic variant of VLH.

We describe VLH^d in detail in Figure 7. It works as follows.

Overview. VLH^d = (Setup, Get, Put) makes black-box use of a dynamic response-hiding multi-map encryption scheme $\text{STE}_{\text{EMM}}^{\text{RH}} = (\text{Setup}, \text{Get}, \text{Put}, \text{Remove})$ and of the volume-hiding multi-map encryption scheme VLH = (Setup, Get).⁴ Both the Setup algorithm and the Get protocol are exactly the same as of those of VLH. The Put algorithm takes as input an update u and processes it as follows. If $u = (\text{add}, (\ell, \mathbf{v}))$, then the client first computes the PRT transform on a single-pair

⁴Note that the same multi-map encryption scheme $\text{STE}_{\text{EMM}}^{\text{RH}} = (\text{Setup}, \text{Get}, \text{Put}, \text{Remove})$ has to be used as the underlying multi-map encryption scheme for VLH.

multi-map defined as $\{(\ell, \mathbf{v})\}$ and outputs a new single-pair multi-map $\{(\ell, \mathbf{v}')\}$. The client and server then execute $\text{STE}_{\text{EMM}}^{\text{RH}}.\text{Put}$ on the label/tuple pair (ℓ, \mathbf{v}') . If $u = (\text{rm}, \ell)$, then the client and server execute $\text{STE}_{\text{EMM}}^{\text{RH}}.\text{Remove}$ on the label ℓ . If $u = (\text{edit}, (\ell, v_{\text{old}}, v_{\text{new}}))$, then the client and server first execute $\text{VLH}.\text{Get}$, the client receives the tuple \mathbf{v} associated to the label ℓ . The client and server then execute $\text{STE}_{\text{EMM}}^{\text{RH}}.\text{Remove}$ on the label ℓ . The client locally replaces the value v_{old} by v_{new} in the tuple \mathbf{v} and then executes $\text{STE}_{\text{EMM}}^{\text{RH}}.\text{Put}$ with the server on the modified label/tuple pair.

Efficiency analysis. In our analysis, assume $\text{STE}_{\text{EMM}}^{\text{RH}}$ is an optimal-time dynamic multi-map encryption scheme [27, 10, 7]. It is clear that the get and storage complexity of VLH^{d} are exactly the same as VLH . The Put complexity varies depending on the type of the update operation. If u is a tuple addition or a tuple edit, then the Put complexity is $O(\lambda + n'_\ell)$ where $n'_\ell \in \{0, \dots, 2^s - 1\}$. The worst-case is $O(\lambda + 2^s)$ while the best case is $O(\lambda)$. The expected complexity is $O(\lambda + 2^{s-1})$. If u is a *tuple deletion*, then the put complexity has constant time.

Security analysis. We now describe the leakage of VLH^{d} assuming that $\text{STE}_{\text{EMM}}^{\text{RH}}$ is instantiated with one of the standard optimal-time forward-private multi-map encryption schemes [7, 8, 1] all of which have leakage profile

$$\Lambda_{\text{MM}} = (\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U) = (\text{trlen}, (\text{qeq}, \text{rlen}), (\text{op}, \text{rlen}))$$

Theorem 8.1. *If $\text{STE}_{\text{EMM}}^{\text{RH}}$ is a $(\text{trlen}, (\text{qeq}, \text{rlen}), (\text{op}, \text{rlen}))$ -secure multi-map encryption scheme, F in PRT is a pseudo-random function and VLH is a (m, qeq) -secure multi-map encryption scheme, then VLH^{d} is a $(m, \text{qeq}, (\text{op}, \text{ueq}))$ -secure multi-map encryption scheme.*

The update equality pattern ueq leaks if and when a label edit has occurred. The proof of this theorem is similar to Theorem 5.1 and deferred to the full version of this work.

8.2 AVLH^d: dynamic variant of AVLH.

We describe AVLH^{d} in details in Figure 8 and it works as follows.

Overview. The construction, AVLH^{d} , makes black box use of a dynamic response-hiding dictionary $\text{STE}_{\text{EDX}}^{\text{RH}} = (\text{Setup}, \text{Get}, \text{Put}, \text{Remove})$ and of the volume hiding multi-map encryption scheme $\text{AVLH} = (\text{Setup}, \text{Get})$.⁵ The Setup algorithm and the Get protocol are exactly the same as of those of AVLH . The Put algorithm takes as input an update u and processes it as follows. Parse u as $(\text{edit}, (\ell, \mathbf{v}))$, the client and server execute $(r, \perp) \leftarrow \text{AVLH}.\text{Get}_{\text{C,S}}((K, st, \ell), \text{EMM})$ where $r = (B_{i_1}, \dots, B_{i_t})$ and the client here does not dismiss any value from the retrieved bins. The client and server then execute $\text{STE}_{\text{EDX}}^{\text{RH}}.\text{Remove}$ on all retrieved bins. The client then identifies the bin that contains the value $\mathbf{v}_{\text{old}} \parallel \ell$ (or $\mathbf{v}_{\text{old}} \parallel \star$) that it replaces with $v_{\text{new}} \parallel \ell$ (or by $\mathbf{v}_{\text{new}} \parallel \star$ if concentrated). The client and server then execute $\text{STE}_{\text{EDX}}.\text{Put}$ on the pairs (i_j, B_{i_j}) , for all $j \in [t]$.

⁵Note that the same dictionary encryption scheme $\text{STE}_{\text{EDX}}^{\text{RH}} = (\text{Setup}, \text{Get}, \text{Put}, \text{Remove})$ has to be used as the underlying dictionary encryption scheme for AVLH .

Let $\text{STE}_{\text{EMM}}^{\text{RH}} = (\text{Setup}, \text{Get}, \text{Put}, \text{Remove})$ be a dynamic response-hiding multi-map encryption scheme, PRT be the pseudo-random transform and VLH = (Setup, Get) be the volume hiding multi-map encryption scheme. Consider the scheme $\text{VLH}^{\text{d}} = (\text{Setup}, \text{Get}, \text{Put})$ defined as follows:

- $\text{Setup}(1^k, \lambda, \text{MM})$: **C** and **S** execute $(K, st, \text{EMM}) \leftarrow \text{VLH.Setup}(1^k, \lambda, \text{MM})$.
- $\text{Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM})$: **C** and **S** execute $(r, \perp) \leftarrow \text{VLH.Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM})$.
- $\text{Put}_{\mathbf{C}, \mathbf{S}}((K, st, \lambda, u), \text{EMM})$:
 1. if $u = (\text{add}, (\ell, \mathbf{v}))$, then
 - (a) compute $\{(\ell, \mathbf{v}')\} \leftarrow \text{PRT}(1^k, \lambda, \{(\ell, \mathbf{v})\})$;
 - (b) **C** and **S** execute $(st', \text{EMM}') \leftarrow \text{STE}_{\text{EMM}}^{\text{RH}}.\text{Put}((K, st, (\ell, \mathbf{v}')), \text{EMM})$;
 2. if $u = (\text{rm}, \ell)$, then **C** and **S** execute $(st', \text{EMM}') \leftarrow \text{STE}_{\text{EMM}}^{\text{RH}}.\text{Remove}((K, st, \ell), \text{EMM})$;
 3. if $u = (\text{add}, v_{\text{old}}, v_{\text{new}})$, then
 - (a) **C** and **S** execute $(\mathbf{v}, \perp) \leftarrow \text{VLH.Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM})$;
 - (b) **C** and **S** execute $(st', \text{EMM}') \leftarrow \text{STE}_{\text{EMM}}^{\text{RH}}.\text{Remove}((K, st, \ell), \text{EMM})$;
 - (c) replace v_{old} with v_{new} in \mathbf{v} ;
 - (d) **C** and **S** execute $(st'', \text{EMM}'') \leftarrow \text{STE}_{\text{EMM}}^{\text{RH}}.\text{Put}((K, st', (\ell, \mathbf{v})), \text{EMM}')$.

Figure 7: VLH^{d} : A dynamic variant of VLH.

Let $\text{STE}_{\text{EDX}}^{\text{RH}} = (\text{Setup}, \text{Get}, \text{Put}, \text{Remove})$ be a dynamic response-hiding dictionary encryption scheme and AVLH = (Setup, Get) be the advanced volume hiding multi-map encryption scheme. Consider the scheme $\text{AVLH}^{\text{d}} = (\text{Setup}, \text{Get}, \text{Put})$ defined as follows:^a

- $\text{Setup}(1^k, n, \text{MM})$: **C** and **S** execute $(K, st, \text{EMM}) \leftarrow \text{AVLH.Setup}(1^k, n, \text{MM})$.
- $\text{Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM})$: **C** and **S** execute $(r, \perp) \leftarrow \text{AVLH.Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM})$.
- $\text{Put}_{\mathbf{C}, \mathbf{S}}((K, st, u), \text{EMM})$:
 1. parse u as $(\text{add}, v_{\text{old}}, v_{\text{new}})$;
 2. **C** and **S** execute $(r, \perp) \leftarrow \text{AVLH.Get}_{\mathbf{C}, \mathbf{S}}((K, st, \ell), \text{EMM})$;
 3. parse $r = (B_{i_1}, \dots, B_{i_t})$;
 4. **S** parses EMM as EDX ;
 5. **C** and **S** execute $(st', \text{EDX}') \leftarrow \text{STE}_{\text{EDX}}^{\text{RH}}.\text{Remove}((K, st, i_j), \text{EDX})$, for all $j \in [t]$;
 6. identify i_s such that $v_{\text{old}} \parallel \ell \in B_{i_s}$ and replace it with $v_{\text{new}} \parallel \ell$;
 7. **C** and **S** execute $(st', \text{EDX}') \leftarrow \text{STE}_{\text{EDX}}^{\text{RH}}.\text{Put}((K, st, (i_j, B_{i_j})), \text{EDX})$, for all $j \in [t]$.

^aThis pseudo-code only details the case where the edited label does not belong to the concentrated component.

Figure 8: AVLH^{d} : A dynamic variant of AVLH.

Efficiency analysis. We assume $\text{STE}_{\text{EDX}}^{\text{RH}}$ is an optimal-time dynamic dictionary encryption scheme [27, 10, 7]. Clearly, the get and the storage complexity of AVLH^{d} are exactly the same as AVLH. The Put complexity is equal to $O(t \cdot \lambda)$, where $t = \max_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$ is the maximum response length and λ is the size of the bin— which is the same as the get complexity. Refer to Section 7 for a more detailed and concrete analysis of the bin size λ .

Security analysis. We now describe the leakage of AVLH^{d} assuming that $\text{STE}_{\text{EDX}}^{\text{RH}}$ is instantiated with one of the standard optimal-time forward-private dictionary encryption scheme [7, 8, 1] all of which have a leakage profile at most

$$\Lambda_{\text{DX}} = (\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U) = (\text{trlen}, \text{qeq}, \text{op})$$

Theorem 8.2. *If $\text{STE}_{\text{EDX}}^{\text{RH}}$ is a $(\text{trlen}, \text{qeq}, \text{op})$ -secure dictionary encryption scheme and AVLH is a $(\text{trlen}, \text{qeq})$ -secure multi-map encryption scheme, then VLH^{d} is a $(\text{trlen}, \text{qeq}, (\text{op}, \text{ueq}))$ -secure multi-map encryption scheme.*

The proof of this theorem is similar to Theorem 7.3 and deferred to the full version of this work.

References

- [1] G. Amjad, S. Kamara, and T. Moataz. Breach-resistant structured encryption. *IACR Cryptology ePrint Archive*, 2018:195, 2018.
- [2] B. Applebaum, B. Barak, and A. Wigderson. Public-key cryptography from different assumptions. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 171–180. ACM, 2010.
- [3] S. Arora, B. Barak, M. Brunnermeier, and R. Ge. Computational complexity and information asymmetry in financial products. *Communications of the ACM*, 54(5):101–107, 2011.
- [4] G. Asharov, M. Naor, G. Segev, and I. Shahaf. Searchable symmetric encryption: Optimal locality in linear space via two-dimensional balanced allocations. In *(STOC '16)*, pages 1101–1114, New York, NY, USA, 2016. ACM.
- [5] G. Asharov, G. Segev, and I. Shahaf. Tight tradeoffs in searchable symmetric encryption. In *Annual International Cryptology Conference*, pages 407–436. Springer, 2018.
- [6] A. Bhaskara, M. Charikar, E. Chlamtac, U. Feige, and A. Vijayaraghavan. Detecting high log-densities: an $o(n^{1/4})$ approximation for densest k -subgraph. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 201–210. ACM, 2010.
- [7] R. Bost. Sophos - forward secure searchable encryption. In *ACM (CCS '16)*, 20016.
- [8] R. Bost, B. Minaud, and O. Ohrimenko. Forward and backward private searchable encryption from constrained cryptographic primitives. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1465–1482. ACM, 2017.

- [9] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. Leakage-abuse attacks against searchable encryption. In *ACM (CCS '15)*, pages 668–679. ACM, 2015.
- [10] D. Cash, J. Jaeger, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Dynamic searchable encryption in very-large databases: Data structures and implementation. In *(NDSS '14)*, 2014.
- [11] D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *Advances in Cryptology - CRYPTO '13*. Springer, 2013.
- [12] M. Chase and S. Kamara. Structured encryption and controlled disclosure. In *ASIACRYPT '10*, pages 577–594. Springer, 2010.
- [13] M. Chase and S. Kamara. Structured encryption and controlled disclosure. Technical Report 2011/010.pdf, IACR Cryptology ePrint Archive, 2010.
- [14] S. Chaudhuri, K. W. Church, A. C. König, and L. Sui. Heavy-tailed distributions and multi-keyword queries. In *ACM SIGIR 2007*.
- [15] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky. Searchable symmetric encryption: Improved definitions and efficient constructions. In *(CCS '06)*, 2006.
- [16] I. Demertzis, D. Papadopoulos, and C. Papamanthou. Searchable encryption with optimal locality: Achieving sublogarithmic read efficiency. In *Annual International Cryptology Conference*, pages 371–406. Springer, 2018.
- [17] I. Demertzis, S. Papadopoulos, O. Papapetrou, A. Deligiannakis, and M. Garofalakis. Practical private range search revisited. In *Proceedings of the 2016 International Conference on Management of Data*, pages 185–198. ACM, 2016.
- [18] I. Demertzis and C. Papamanthou. Fast searchable encryption with tunable locality. In *SIGMOD'17*, 2017.
- [19] M. Etemad, A. Küpçü, C. Papamanthou, and D. Evans. Efficient dynamic searchable encryption with forward privacy. *Proceedings on Privacy Enhancing Technologies*, 2018(1):5–20, 2018.
- [20] S. Faber, S. Jarecki, H. Krawczyk, Q. Nguyen, M. Rosu, and M. Steiner. Rich queries on encrypted data: Beyond exact matches. In *European Symposium on Research in Computer Security (ESORICS '15). Lecture Notes in Computer Science*, volume 9327, pages 123–145, 2015.
- [21] B. A. Fisch, B. Vo, F. Krell, A. Kumarasubramanian, V. Kolesnikov, T. Malkin, and S. M. Bellovin. Malicious-client security in blind seer: a scalable private dbms. In *IEEE Symposium on Security and Privacy*, pages 395–410. IEEE, 2015.
- [22] M. S. Islam, M. Kuzu, and M. Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In *(NDSS '12)*, 2012.

- [23] S. Kamara and T. Moataz. Boolean searchable symmetric encryption with worst-case sub-linear complexity. In *Advances in Cryptology - EUROCRYPT '17*, 2017.
- [24] S. Kamara and T. Moataz. SQL on structurally-encrypted databases. In *ASIACRYPT*, 2018.
- [25] S. Kamara, T. Moataz, and O. Ohrimenko. Structured encryption and leakage suppression. In *Annual International Cryptology Conference*, pages 339–370. Springer, 2018.
- [26] S. Kamara and C. Papamanthou. Parallel and dynamic searchable symmetric encryption. In *Financial Cryptography and Data Security (FC '13)*, 2013.
- [27] S. Kamara, C. Papamanthou, and T. Roeder. Dynamic searchable symmetric encryption. In *ACM (CCS '12)*, 2012.
- [28] G. Kellaris, G. Kollios, K. Nissim, and A. O. Neill. Generic attacks on secure outsourced databases. In *ACM Conference on Computer and Communications Security (CCS '16)*, 2016.
- [29] M.-S. Lacharité, B. Minaud, and K. G. Paterson. Improved reconstruction attacks on encrypted data using range query leakage. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 297–314. IEEE, 2018.
- [30] X. Meng, S. Kamara, K. Nissim, and G. Kollios. Grecs: Graph encryption for approximate shortest distance queries. In *(CCS 15)*, 2015.
- [31] I. Miers and P. Mohassel. Io-dsse: Scaling dynamic searchable encryption to millions of indexes by improving locality. Cryptology ePrint Archive, Report 2016/830, 2016. <http://eprint.iacr.org/2016/830>.
- [32] M. Naveed, S. Kamara, and C. V. Wright. Inference attacks on property-preserving encrypted databases. In *ACM Conference on Computer and Communications Security (CCS)*, CCS '15, pages 644–655. ACM, 2015.
- [33] M. Naveed, M. Prabhakaran, and C. Gunter. Dynamic searchable encryption via blind storage. In *IEEE Symposium on Security and Privacy (S&P '14)*, 2014.
- [34] V. Pappas, F. Krell, B. Vo, V. Kolesnikov, T. Malkin, S.-G. Choi, W. George, A. Keromytis, and S. Bellovin. Blind seer: A scalable private dbms. In *Security and Privacy (SP), 2014 IEEE Symposium on*, pages 359–374. IEEE, 2014.
- [35] R. Poddar, T. Boelter, and R. A. Popa. Arx: A strongly encrypted database system. Technical Report 2016/591.
- [36] D. Song, D. Wagner, and A. Perrig. Practical techniques for searching on encrypted data. In *IEEE S&P*, pages 44–55. IEEE Computer Society, 2000.
- [37] E. Stefanov, C. Papamanthou, and E. Shi. Practical dynamic searchable encryption with small leakage. In *(NDSS'14)*, 2014.
- [38] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas. Path oram: An extremely simple oblivious ram protocol. In *(CCS)*, 2013.
- [39] G. K. Zipf. The psycho-biology of language. 1935.