Mitigating Leakage in Secure Cloud-Hosted Data Structures: Volume-Hiding for Multi-Maps via Hashing

Sarvar Patel
sarvar@google.com
Google LLC

Kevin Yeo
kwlyeo@google.com
Google LLC

Giuseppe Persiano
giuper@gmail.com
Università di Salerno

Moti Yung
moti@google.com
Google LLC
Columbia University

ABSTRACT

Volume leakage has recently been identified as a major threat to the security of cryptographic cloud-based data structures by Kellaris *et al.* [CCS'16] (see also the attacks in Grubbs *et al.* [CCS'18] and Lacharité *et al.* [S&P'18]). In this work, we focus on volume-hiding implementations of *encrypted multi-maps* as first considered by Kamara and Moataz [Eurocrypt'19]. Encrypted multi-maps consist of outsourcing the storage of a multi-map to an untrusted server, such as a cloud storage system, while maintaining the ability to perform private queries. Volume-hiding encrypted multi-maps ensure that the number of responses (volume) for any query remains hidden from the adversarial server. As a result, volume-hiding schemes can prevent leakage attacks that leverage the adversary's knowledge of the number of query responses to compromise privacy.

We present both conceptual and algorithmic contributions towards volume-hiding encrypted multi-maps. We introduce the first formal definition of *volume-hiding leakage functions*. In terms of design, we present the first *volume-hiding encrypted multi-map* *dprfMM* whose storage and query complexity are both asymptotically optimal. Furthermore, we experimentally show that our construction is practically efficient. Our server storage is smaller than the best previous construction while we improve query complexity by a factor of 10-16x.

In addition, we introduce the notion of *differentially private volume-hiding leakage functions* which strikes a better, tunable balance between privacy and efficiency. To accompany our new notion, we present a *differentially private volume-hiding encrypted multi-map* *dpMM* whose query complexity is the volume of the queried key plus an additional logarithmic factor. This is a significant improvement compared to all previous volume-hiding schemes whose query overhead was the maximum volume of any key. In natural settings, our construction improves the average query overhead by a factor of 150-240x over the previous best volume-hiding construction even when considering small privacy budget of $\epsilon = 0.2$.

CCS CONCEPTS

• Security and privacy → Management and querying of encrypted data;

KEYWORDS

cloud storage, privacy, encrypted search, volume-hiding

ACM Reference Format:

Sarvar Patel, Giuseppe Persiano, Kevin Yeo, and Moti Yung. 2019. Mitigating Leakage in Secure Cloud-Hosted Data Structures: Volume-Hiding for Multi-Maps via Hashing. In *2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19), November 11–15, 2019, London, United Kingdom*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3319535.3354213>

1 INTRODUCTION

In this paper, we study *structured encryption* (STE), first introduced by Chase and Kamara [16], which is a cryptographic primitive used to study the security of cloud-hosted data structures. Structured encryption schemes enable the owner of a data structure to encrypt the data structure and outsource the storage of encrypted data structure to a potentially untrusted third-party such as a cloud storage system. Additionally, STE schemes allow the data owner to perform data structure operations on the outsourced encrypted data structure without revealing any information to the server beyond some well-defined and “sensible” leakage function.

An important example of a STE scheme is the *encrypted multi-map* (EMM) [17, 32] primitive which enables the storage of keys associated to a sequence of (possibly) multiple values. Furthermore, multi-maps allows its owner to query for a key and receive all values associated with the key. EMM's form the basis of many important applications. Two such applications are searching over a corpus of encrypted documents and performing queries over encrypted relational databases. As a result, the construction of both efficient and private encrypted multi-maps is a very important problem. We now explain these applications in more detail.

Searchable encryption is a primitive first introduced by Song *et al.* [51] and has continued to be well-studied in the past decade. One can see some of the following as examples of works in searchable encryption [2, 4–6, 9–11, 13–15, 17–19, 21, 30, 34, 42, 44, 47, 52]. Searchable encryption enables a data owner with a corpus of documents to encrypt and store the corpus to an untrusted third-party while maintaining the ability to privately search for documents

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CCS '19, November 11–15, 2019, London, United Kingdom

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6747-9/19/11.

<https://doi.org/10.1145/3319535.3354213>

containing specified keywords. In many cases, searchable encryption uses EMMs as the main underlying primitive to map keywords to documents that contain keywords. Although, EMMs have also been used in many other ways.

Encrypted databases are another important application of EMMs. The goal is to encrypt databases whose storage will be outsourced to an untrusted third party while enabling the data owner to privately perform database operations. We note that earlier attempts [49] at encrypted databases used property-preserving encryption schemes such as deterministic [6] and order-preserving [7, 8] encryption schemes. Work by Naveed *et al.* [43] show that encrypted databases built from property-preserving encryption have security vulnerabilities. Recently, Kamara and Moataz [31] present an encrypted SQL database scheme that foregoes the use of property-preserving encryption. Instead, they rely on EMMs to perform SQL operations privately.

As a result of the above important applications, it is clear that constructing EMMs that are both efficient and private is important. While efficiency is clear to evaluate, assessing the privacy of an EMM (and STE schemes in general) is a challenging problem. So far, our only measure of privacy is a “sensible” or “reasonable” leakage function, which is both a vague and subjective qualifier. There has been a lot of work that attempt to understand various leakage profiles and determine whether certain leakage profiles may be abused to compromise privacy. The first leakage-abuse attacks was presented by Islam *et al.* [29]. Many follow up works [12, 26–28, 35, 40, 43, 50, 54] consider either different leakage profiles and/or weaker assumptions. These attacks significantly further our understanding of the dangers of various types of leakage profiles. Furthermore, these attacks present guidance on the necessary requirements of private STE schemes. Therefore, an important line of research is to construct STE schemes with smaller leakage that protect against these attacks.

In an attempt to mitigate the risk of leakage profiles, the seminal work of Kamara and Moataz [32] introduce the notion of volume-hiding EMM schemes. These schemes ensure that the number of values (volume) associated with a single key is never leaked to the adversary. Several of the above leakage-abuse attacks rely on the knowledge of volume to compromise privacy. As a result, volume-hiding schemes can foil such attacks. In this work, we continue the study of volume-hiding constructions by presenting schemes with better query and storage efficiency.

1.1 Our Contributions

In this work, we make both conceptual and algorithmic contributions in the area of volume-hiding EMMs. In particular, we present formal definitions for volume-hiding EMMs and introduce the notion of differentially private volume-hiding EMMs. Furthermore, we present efficient constructions for both types of EMMs.

Throughout this section, we will phrase the efficiency improvements of our constructions using multipliers. As an example, a 2x improvement in communications means that our new construction uses half the communication of the previous best construction.

Volume-Hiding EMMs. We start by describing our contributions to volume-hiding encryption schemes for multi-maps. To our knowledge, we present the first formal security definition for volume-hiding leakage functions. Our security definition is built on top of the typical simulation-based security definitions for STE schemes allowing the proof techniques from STE schemes to also be used for volume-hiding STE schemes. Accompanying our conceptual definitional contributions, we present the following efficient, volume-hiding STE scheme for general multi-maps:

THEOREM 1 (INFORMAL). *Consider any multi-map MM with n total values and let ℓ be the maximum number of values associated with a key of MM . Then there exists a volume-hiding encrypted multi-map with communication complexity of $O(\ell)$, server storage of $O(n)$ and client storage of size $f(\lambda)$, for every function $f(\lambda) = \omega(1)$. The leakage of the encrypted multi-map consists of only the query equality pattern and the values ℓ and n .*

Our construction is lossless (that is, it always returns all the values associated with a key) and is asymptotically optimal in terms of storage and communication complexity. Indeed, in a lossless construction each of the n values must be stored at least once and our construction results in server storage of about $2n$. More precisely, it can be instantiated to use $(2 + \alpha)n$ storage, for every constant $\alpha > 0$. Moreover, due to volume-hiding, it must be that each query returns at least ℓ records and our construction will always return exactly 2ℓ records. Finally, the client storage typically must include at least a private key. Our construction requires storing a private key in addition to just $\omega(1)$ values. We show experimentally that our volume-hiding scheme is concretely efficient and improves on previous best volume-hiding scheme. In particular, our scheme uses less server storage and improves query overhead by a factor of 10-16x over [32] when encrypting multi-maps of size 1-67 MB in the plaintext and consisting of 2^{16} - 2^{22} total values.

Differentially Private, Volume-Hiding EMMs. As another contribution, we directly address a statement appearing in [32]: “it is hard to imagine any non-lossy construction being able to hide response length of a query and have query complexity $o(t)$, where t is the maximum response length.” We introduce the notion of a *differentially private volume-hiding* leakage function as a weakening of volume-hiding. Kellaris *et al.* [36] studied the related notion of differentially private volume-hiding for ORAM.

We present the following differentially private, volume-hiding encryption scheme whose query complexity for many keys are much smaller than the maximum volume:

THEOREM 2 (INFORMAL). *Consider any multi-map MM with n total values and let, for every key, $\ell(\text{key})$ denote the number of values, the volume, associated with key. Then there exists an encrypted multi-map with communication complexity $O(\ell(\text{key}) + f(\lambda))$, for any function $f(\lambda) = \omega(\log \lambda)$, with server storage of $O(n)$, and client storage of size $g(\lambda)$, for any function $g(\lambda) = \omega(1)$. The scheme is lossless except with probability negligible in λ . The leakage of the encrypted multi-map consists of only query equality pattern, the number of total values n as well as differentially private leakage of volumes of the queried keys.*

We note that, unlike volume-hiding schemes, our differentially private scheme is able to achieve query complexity that is only

dependent on the volume of the queried key. In particular, the query complexity will be at most the volume of the query key plus an additional logarithmic factor. In typical scenarios, most keys have smaller volume than the maximum volume. Using experiments, we show that our schemes can reduce query overhead by a factor of 150-240x over the previous best volume-hiding scheme [32] when encrypting multi-maps consisting of 2^{16} - 2^{22} values and occupy 1-67 MB in the plaintext.

By defining our new security notion of differential privacy volume-hiding, we attempt to strike a tunable balance between efficiency and privacy. The previous notion of volume-hiding considered strong privacy that resulted in large query complexity. When defending future, feasible attacks, volume-hiding might be too strong resulting in unnecessary query overhead. We thus look at the differential privacy framework to obtain a more efficient protocol at the cost of weakening the security notion in a meaningful way. We observe that all previous volume attacks on STE schemes [25, 29, 35, 40, 43] are generic and do not exploit weaknesses of specific constructions. Rather, they are attacks to the ideal world functionality. Therefore, our generic framework will provide protection against all of these, as well as future, attacks to ideal world functionalities.

1.2 Comparison to Previous Works

In this section, we compare our new constructions with those that were presented in the work on volume-hiding EMMs by Kamara and Moataz [32] as well as some naive approaches. To do this, we briefly review their EMM constructions and compare them to our work. For a more detailed exposition on these constructions, we refer the reader to [32]. We summarize our comparisons in Table 1. All shown efficiencies are taken from the suggested concrete parameters sections of [32]. We denote the number of keys by m , the total number of values by n , and the maximum number of values associated with a key by ℓ .

Naive Padding. The simplest approach to volume-hiding encryption schemes is naive padding. Given any multi-map, a new multi-map can be constructed where each key has ℓ values. For keys with less than ℓ values, additional dummy values are padded until there are ℓ values. We note that the resulting storage is $m \cdot \ell$. In many cases, $(m \cdot \ell) \gg n$ where n is the number of total values. For example, when $m = n/2$ and $\ell = n/2$, the storage of naive padding is $m \cdot \ell = \Theta(n^2)$ which is quadratically larger than the number of total values. Additionally, naive padding results in larger overhead if the multi-map follows naturally appearing distributions such as Zipf's distribution. In practice, plaintext multi-maps typically use storage on the order of the number total values n . Therefore, naive padding uses too much storage in almost all cases and our storage goal for volume-hiding EMMs should be close to n .

ORAMs. Another naive approach is to use oblivious RAMs (ORAMs), which were introduced by Goldreich and Ostrovsky [23]. ORAMs are a powerful primitive that enable access to storage hosted by a potentially untrusted server such as a cloud storage provider without leaking any information beyond the number of accesses. As a result, there has been a lot of work in ORAM [3, 24, 39, 46, 53] leading to logarithmic overhead constructions. Furthermore, ORAM lower bounds [41, 48] have shown that these are the best possible constructions. A simple approach is to take any STE scheme

and replace each access using an ORAM access to suppress leakage. Furthermore, one can make an STE scheme volume-hiding by making fake ORAM accesses until ℓ records have been retrieved. The resulting scheme would have $\Theta(\ell \cdot \log n)$ overhead using the best theoretical constructions and $\Theta(\ell \cdot \log^2 n)$ using the more concretely efficient ORAM constructions. In either case, the overhead of ORAM is too large for practical use cases.

Pseudorandom Transform. From a high level, the pseudorandom transform [32] takes an input multi-map and generates a new multi-map such that the number of values associated with each key is generated using a pseudorandom function. As a result, there may exist keys whose volume in the new multi-map is smaller compared to the input multi-map. Therefore, the construction is lossy as truncation occurs and several values are removed from the new multi-map. The authors show that the number of truncated keys is small when the input multi-map is Zipf's distributed. Furthermore, the authors of [32] show that the storage overhead is $(m \cdot \ell)/2$ which is better than naive padding. Unfortunately, this transform is not practical as there are no guarantees for data loss (truncation) on general multi-maps and the storage overhead is closer to naive padding of $m \cdot \ell$ than our goal of n .

Densest Subgraph Transform for General Multi-Maps. This transform considers a bipartite graph where each of the m keys are in one part and there are b empty bins in the other part. Each of the m keys are assigned ℓ bins chosen uniformly at random. Each of the, at most ℓ , values associated with a key is placed into one of the key's assigned bin such that at most one value appears in each bin. After this is done for all keys, the bins are padded with dummy values to the size of the maximum bin. By balls-and-bins analysis, the size of the bins must be $\Omega(\log n)$ size. As a result, the query overhead becomes at least $\Theta(\ell \cdot \log n)$. On the other hand, the storage overhead becomes $\Theta(n)$. We note that this construction is strictly worse in query overhead than our construction as it requires an extra $\log n$ overhead while achieving the same asymptotic storage overhead of $\Theta(n)$. Furthermore, we experimentally show that our constructions achieve 10-16x query overhead improvements.

Densest Subgraph Transform for Concentrated Multi-Maps. The last construction of [32] considered concentrated multi-maps. That is, multi-maps where a large number of keys share many same values. In this case, they modify the previous construction such that this concentrated set of values will appear only once in a single set of bins. As a result, the storage overhead becomes $\Theta(n - \sqrt{m} \cdot \text{polylog}(m))$. Unfortunately, this construction requires assuming that the planted densest subgraph problem is hard, which has been not heavily studied. Due to the lack of cryptanalysis, it is very difficult to find concrete parameters that can be used in practice where security can be trusted. Furthermore, the resulting query complexity is very large. As a result, this construction does not seem to be usable in practice at the moment.

Final Comparison. To make a final comparison, it is clear that our volume-hiding construction is more concretely efficient and practical compared to all previous works. In particular, it achieves the best storage overhead of just $(2 + \alpha)$ times the number of total values. Additionally, the query overhead is only 2ℓ . In contrast, the schemes from [32] either might lose a lot of information for general

	Query Complexity	Server Storage for General MMs	Lossy with Non-Negligible Probability	Server Storage for Concentrated MMs	Computational Assumption
Naive Padding	$\Theta(\ell)$	$\Theta(m \cdot \ell)$		$\Theta(m \cdot \ell)$	One-Way Functions
Pseudorandom Transform [32]	$\Theta(\ell)$	$\Theta(m \cdot \ell)$	✓	$\Theta(m \cdot \ell)$	One-Way Functions
Dense Subgraph Transform [32]	$\Theta(\ell \cdot \log n)$	$\Theta(n)$		$\Theta(n)$	One-Way Functions
Dense Subgraph and Planting Transform [32]	$\Theta\left(\ell \cdot \frac{n}{\text{polylog}(m)}\right)$	$\Theta(n)$		$\Theta(n - \sqrt{m} \cdot \text{polylog}(m))$	Planted Densest Subgraph Problem
Ours (dprfMM)	2ℓ	$(2 + \alpha)n$		$(2 + \alpha)n$	One-Way Functions
Ours (dpMM)	$2\ell(k) + \omega(\log \lambda)$	$(2 + \alpha)(m + n)$		$(2 + \alpha)(m + n)$	One-Way Functions

Table 1: This table compares previous volume-hiding encryption schemes for multi-maps from [32] with the constructions of this paper. For notation, ℓ represent the volume of the largest key of the input multi-map, m denotes the number of unique keys and n denotes the total number of values over all keys. For any queried key k , we denote by $\ell(k)$ as the number of values associated with the queried key k . The value α may be any positive constant. Finally, we refer to $2^{-\lambda}$ as the probability of losing data in our differentially private solution.

multi-maps and have large storage overhead (pseudorandom transform), require computational assumptions that have not been well studied (densest subgraph for concentrated multi-maps) or require a much larger query overhead of $\Theta(\ell \cdot \log n)$ (densest subgraph for general multi-maps) than our construction. Furthermore, our differentially private volume-hiding construction uses slightly more server storage while enabling significantly smaller query overhead.

1.3 Our Techniques

In this section, we outline our techniques that enable us to construct better volume-hiding, encryption schemes for multi-maps. We start by considering a simple construction. Consider any key k and its associated values $v_1, \dots, v_{\ell(k)}$ where $\ell(k)$ is the number of values associated with key k . The server stores a dictionary in the following manner. The value v_i will appear at location $F_K(k \parallel i)$ in the dictionary where F is a pseudorandom function (PRF), K is the key of the PRF and $k \parallel i$ is the concatenation of the key k to the index i . Assuming the output range of the PRF F is large (like $\Omega(n^2)$ where n is the total number of values), there would be no collisions in the dictionary. To query for a key k , the client would simply send the values $F_K(k \parallel 1), \dots, F_K(k \parallel \ell(k))$ and the server would return the associated encrypted values from the dictionary.

Going towards the problem of volume-hiding, it is not clear how to modify this EMM construction to hide the volumes of all the keys. The first naive approach would be to send ℓ PRF values for any key where ℓ is the maximum value of any key. Unfortunately, this does not work. Consider a key k where $\ell(k) < \ell$. If the client sends the values $F_K(k \parallel 1), \dots, F_K(k \parallel \ell)$, the server will see that the values $F_K(k \parallel \ell(k)+1), \dots, F_K(k \parallel \ell)$ do not exist in the server-held dictionary. As a result, the server can quickly determine the volume of key k is exactly $\ell(k)$. One way to mitigate this leakage is to simply populate the dictionary with the missing values for all keys k . However, this is similar to the naive padding approach which would result in $m \cdot \ell$ storage which is much larger than our goal of n (the number of total values) in most cases.

Therefore, we seem to have contrasting problems now. We wish to make sure that all PRF values for those with no value, $F_K(k \parallel \ell(k) + 1), \dots, F_K(k \parallel \ell)$, point to some non-empty entry in the dictionary. On the other hand, we do not want to simply insert dummy values in for each of the PRF values and increase server storage. The main idea to overcome these problems is to reduce the output space of the PRF F . For example, we could attempt to reduce the output range of the PRF F to be $\Theta(n)$. Now, the problem is to place each of the n values into a location that is determined by F . The remaining empty locations can be simply padded with dummy values without, hopefully, incurring a huge storage overhead like before. This is the exact problem that is considered in *hashing* where the goal is to place n items into a space of $\Theta(n)$ locations such that each item can be easily retrieved by looking into only a small number of locations specified by a PRF F . One approach to this problem would be to simply place a value into its location specified by the PRF F . Each location will be padded to the size of the largest bin. The classical balls-and-bins analysis shows the size of the largest will contain $\Omega(\log n)$ items. To query, a key will be download its ℓ associated bins which would incur an $\Omega(\ell \cdot \log n)$ query complexity. We note an approach similar to this idea is used by the densest subgraph transform in [32].

To avoid the additional logarithmic overhead, we will use *cuckoo hashing* introduced by Pagh and Rodler [45] for our constructions. In particular, we will use a variant presented by Kirsch *et al.* [38] where an additional stash exists to exponentially decrease the probability of failure. From a high level, cuckoo hashing consists of two tables of size $(1 + \alpha)n$ for a small constant $0 < \alpha < 1$. For value v_i associated with key k , the value v_i will be placed into one of the three locations: $F_K(k \parallel i \parallel 0)$ in table one or $F_K(k \parallel i \parallel 1)$ in table two or in the stash. The stash will be stored by the client. The remaining empty locations in each table will be filled by dummy values and all values in both tables will be encrypted. Let us now reconsider a query for a key k where $\ell(k) < k$. The client can send the values $\{F_K(k \parallel i \parallel 0), F_K(k \parallel i \parallel 1)\}_{i \in [\ell]}$ and the server will

return all 2ℓ encrypted values in the tables. We prove that this construction is volume-hiding as the server will be unaware whether a location is filled with a real or dummy value.

Additionally, we present a concrete optimization that further improves the communication costs of our volume-hiding encryption scheme for multi-maps. In the above scheme, the client must send the output of 2ℓ PRF evaluations. However, the PRF evaluations may be easily arranged to correspond to a consecutive set of inputs. As a result, we are able to use delegatable PRFs [22, 37] where the 2ℓ PRF evaluations can be sent using smaller communication. In our specific application, the client will always be sending the prefix of exactly 2ℓ evaluations. We are able to convey these 2ℓ PRF evaluations using exactly one PRF evaluation in such a way that the server can securely expand the single PRF evaluation to the required 2ℓ PRF evaluations. This optimization reduces the bandwidth from the client to the server from 2ℓ PRF outputs to just one PRF output.

Finally, we describe our modifications to construct a differentially private volume-hiding EMM. A first attempt at differentially private volume-hiding would be to send $\ell(k) + X$ PRF evaluations for any key k with $\ell(k)$ values and X is drawn from the Laplacian distribution. The problem is that $\text{Lap}(1/\epsilon)$ may be negative resulting in a lossy construction. Instead, we will pick some public parameter $f(\lambda)$ such that the probability that X drawn from the Laplacian is smaller than $f(\lambda)$ occurs with probability at most $2^{-\lambda}$. By picking large enough λ and $f(\lambda) = \omega(\log \lambda)$, we can guarantee that our construction is non-lossy except with small probability that is not observable in practice. Note, the query complexity of our differentially private volume-hiding EMM is $\ell(k) + \omega(\log \lambda)$ which is, on average, significantly smaller than the largest volume ℓ . The only caveat is that we now need to know the volume of any queried key. To do this, we will store a count table which, for each key k with positive volume, stores the volume of the key $\ell(k)$. The count table will be stored on the server using cuckoo hashing.

2 DEFINITIONS

Structured encryption schemes considers the problem of encrypting a data structure. Additionally, the encrypted data structure should enable the original data owner to be able to perform queries in a private manner when the encrypted data structure is held by a potentially untrusted third party server. The first definitions of structured encryption were presented by Chase and Kamara [16]. While we can consider generic definitions for encrypting any data structure, we focus our definitions on encrypting multi-maps exclusively.

2.1 Encrypted Multi-Maps

We start by considering the multi-map (MM) primitive, which maintains a set of m key to value vector pairs $\text{MM} = \{(\text{key}_i, \vec{v}_i)\}_{i \in [m]}$ where each key_i is from the *key universe* $\mathcal{K}$ and $\vec{v}_i$ is a tuples (or vector) of values from the *value universe* $\mathcal{V}$. For convenience, we suppose that all keys are unique. That is, $\text{key}_i \neq \text{key}_j$ when $i \neq j$. Note, this assumption is without loss of generality as if there are two equal keys $\text{key}_i = \text{key}_j$, then the two tuples of values can be combined into a single tuple of values. For any key_i , we denote by the number of values associated with key_i by $\ell(\text{key}_i)$ (that is,

$\ell(\text{key}_i) := |\vec{v}_i|$). Note, the number of values associated with different keys can be different. We denote the maximum number of values associated by ℓ (that is, $\ell := \max_{i \in [m]} \ell(\text{key}_i) = \max_{i \in [m]} |\vec{v}_i|$). Note, we will refer to ℓ as the *width* of the multi-map. We denote the total number of values by $n := \sum_{i \in [m]} \ell(\text{key}_i) = \sum_{i \in [m]} |\vec{v}_i|$.

In terms of functionality, a multi-map provides a query operation which takes as input key from the key universe $\mathcal{K}$. The output of the query operation with input key for a multi-map MM will be the tuple of values associated with the key stored in MM. For example, if $\text{key} = \text{key}_i$ for some $i \in [m]$, then the query operation will return the tuples $\vec{v}_i$. On the other hand, if $\text{key} \neq \text{key}_i$ for all $i \in [m]$, then the output of the query operation will be $\perp$.

Moving towards encrypted multi-maps, we now present a formal definition of structured encryption for multi-maps. Our STE definition will consider query algorithms with r rounds of interaction where the data structure must correctly answer to queries using r rounds of communication between the client and the server. However, we will only consider non-interactive ($r = 1$) or two-rounds of interaction ($r = 2$) in this paper. We focus on schemes with few rounds of interaction due to their practicality and efficiency as they only use a small number of roundtrips of communication between the client and server per query.

Definition 3 (r-Interactive, Structured Encryption for MM's). A r -interactive structured encryption scheme for multi-maps $\Sigma_{\text{MM}} = (\text{Setup}, \text{Query}, \text{Reply}, \text{Result})$ where $r \geq 1$ consists of the $r + 2$ following polynomial time algorithms:

- (1) $(K, \text{EMM}) \leftarrow \text{Setup}(1^\lambda, \text{MM})$ is an algorithm executed by the client $\mathbb{C}$ that takes as input the security parameter λ as well as the input multi-map MM. The output of Setup is a private key K as well as the encrypted multi-map EMM. The private key K will be held by the client $\mathbb{C}$ while the encrypted multi-map EMM is sent to the server $\mathbb{S}$.
- (2) For $i \in \{0, \dots, r-1\}$, we define the query algorithm as $\text{Request}_i \leftarrow \text{Query}^i(K, \text{key}, \text{Response}_{i-1})$ that is executed by the client $\mathbb{C}$ that takes as input the private key K , the key from the key universe $\mathcal{K}$ as well as the server's response, Response_{i-1} from the previous query. For the first query ($i = 0$), Response_{-1} will be $\perp$. The output is a search request Request_i which will be sent to the server $\mathbb{S}$.
- (3) For $i \in \{0, \dots, r-1\}$, we define the reply algorithm as $\text{Response}_i \leftarrow \text{Reply}^i(\text{Request}_i, \text{EMM})$ that is executed by the server $\mathbb{S}$ that takes as input Request_i and the encrypted multi-map EMM. The output of Reply will be Response_i which is the server's response for the i -th query.
- (4) $\vec{v} \leftarrow \text{Result}(K, \text{Response}_{r-1})$ is an algorithm executed by the client $\mathbb{C}$ that takes as input Response_{r-1} as well as the private key K . The output will be the tuple of values $\vec{v}$ associated with key used to generate $\text{Request}_0, \dots, \text{Request}_{r-1}$ by the client $\mathbb{C}$.

For non-interactive encryption schemes with $r = 1$, we will simply drop all subscripts and superscripts related to Query and Reply as there will only be one algorithm for each.

2.2 Adaptive Security and Leakage Functions

The security notion of structured encryption for multi-maps is parameterized by two leakage functions, $\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Query}}$. Leakage

function $\mathcal{L}_{\text{Setup}}$ gives an upper bound on the information leaked by the encrypted multi-map EMM generated by Setup. Leakage function $\mathcal{L}_{\text{Query}}$ instead bounds the information leaked by the request generated by the client using Query. Note that of the algorithms used to perform a query (that is, Query, Reply and Result) Query is the only one that can potentially leak information to the server $\mathbb{S}$ through its output Request. We next formally present the security notion using the simulation-based game definitions. In particular, we will consider the adaptive variant where the query operations may be chosen adaptively by the adversary after seeing the leakage from previous queries.

Let $\Sigma = (\text{Setup}, \{\text{Query}^i, \text{Reply}^i\}_{i=0, \dots, r-1}, \text{Result})$ be a structured encryption scheme for multi-maps with r rounds of interaction. To define adaptive security, we consider the following real and ideal experiments where $\mathcal{A}$ is a stateful, honest-but-curious PPT adversary, $\mathcal{S} = (\text{SimSetup}, \text{SimQuery})$ is a stateful, PPT simulator and $\mathcal{L}_{\text{Setup}}$ and $\mathcal{L}_{\text{Query}}$ are the leakage functions for the setup and query process respectively.

$\text{Real}_{\Sigma, \mathcal{A}}(1^\lambda) :$

- (1) The adversary $\mathcal{A}$ selects an input multi-map MM and gives it to the challenger.
- (2) The challenger $\mathcal{C}$ executes $(K, \text{EMM}) \leftarrow \text{Setup}(1^\lambda, \text{MM})$ and sends EMM to the adversary $\mathcal{A}$.
- (3) The adversary $\mathcal{A}$ will adaptively pick a polynomial number of queries $\text{key}_1, \dots, \text{key}_{\text{poly}(\lambda)}$. For each query key_j , the challenger $\mathcal{C}$ executes the $\{\text{Query}^i, \text{Reply}^i\}_{i=0, \dots, r-1}$ with $\mathcal{A}$. $\mathcal{A}$ receives $\text{Request}_i \leftarrow \text{Query}^i(K, \text{key}_j, \text{Response}_{i-1})$ for all i and key_j .
- (4) Finally, the adversary $\mathcal{A}$ outputs a bit $b \in \{0, 1\}$.

$\text{Ideal}_{\Sigma, \mathcal{A}, \mathcal{S}}(1^\lambda) :$

- (1) The adversary $\mathcal{A}$ generates an input multi-map MM .
- (2) The simulator receives $\mathcal{L}_{\text{Setup}}(\text{MM})$ and returns an encrypted multi-map EMM.
- (3) The adversary $\mathcal{A}$ will adaptively pick a polynomial number of queries $\text{key}_1, \dots, \text{key}_{\text{poly}(\lambda)}$. For each query key_j , the simulator $\mathcal{S}$ computes Request_i using only $\mathcal{L}_{\text{Query}}(\text{key}_1, \dots, \text{key}_i, \text{MM})$ as input, for all $i \in \{0, \dots, r-1\}$. Request_i is forwarded to the adversary $\mathcal{A}$.
- (4) Finally, the adversary $\mathcal{A}$ outputs a bit $b \in \{0, 1\}$.

Definition 4 (Adaptive Security for r -Interactive, Structured Encryption for MMs). The r -interactive, structured encryption scheme for multi-maps Σ is *adaptively* $(\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Query}})$ -secure if there exists a stateful, PPT simulator $\mathcal{S}$ such that for all PPT adversaries $\mathcal{A}$:

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

2.3 Typical Leakage Functions

In this section, we describe typical leakage function that have been considered for structured encryption for multi-maps. We will follow the terminology introduced by Kamara *et al.* [33].

- **Query Equality:** This leakage reports the equality pattern leaking whether two queries are to the same key or not. Formally, for a sequence of queried keys $\text{key}_1, \dots, \text{key}_t$,

$\text{req}(\text{key}_1, \dots, \text{key}_t) = M$ consists of a $t \times t$ matrix such that $M[i][j] = 1$ if and only if $\text{key}_i = \text{key}_j$.

- **Response Length:** This leakage consists of the number of values (volume) associated with queried keys. Formally, for a sequence of queried keys $\text{key}_1, \dots, \text{key}_t$ and a multi-map MM , That is, $\text{rlen}(\text{MM}, \text{key}_1, \dots, \text{key}_t) = (|\text{MM}[\text{key}_i]|)_{i \in [t]}$ where $\text{MM}[\text{key}_i]$ refers to the tuple of values associated with key_i .
- **Maximum Response Length:** This leakage pattern consists of the maximum number of values associated with any key in the multi-map. Formally, for any multi-map MM , $\text{mrLen}(\text{MM}) = \ell = \max_{\text{key} \in \mathcal{K}} |\text{MM}[\text{key}]|$.
- **Domain Size:** This leakage pattern refers to the total number of values in the multi-map. Formally, for any multi-map MM , $\text{dsize}(\text{MM}) = n = \sum_{\text{key} \in \mathcal{K}} |\text{MM}[\text{key}]|$.

From a high level, the goal of volume-hiding STE schemes is to suppress response length leakage. We will formally define volume-hiding leakage functions in the next section.

2.4 Volume-Hiding Leakage Functions

Roughly speaking, a *volume-hiding* leakage function ensures that the number of values associated with any single key, the *volume* of the key, is not revealed and that only the maximum volume of a key is leaked.

We start with the following definition.

Definition 5. The *signature* of a multi-map MM is the sequence of pairs $((\text{key}, \ell(\text{key})))_{\text{key} \in \mathcal{K}}$ where $\ell(\text{key})$ is the length of tuple of values associated with key in the multi-map.

We next define the game $\text{VHGame}_{\eta}^{\mathcal{A}, \mathcal{L}}(n, \ell)$ for leakage functions $\mathcal{L} = (\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Query}})$, adversary $\mathcal{A}$, and $\eta \in \{0, 1\}$.

$\text{VHGame}_{\eta}^{\mathcal{A}, \mathcal{L}}(n, \ell) :$

- (1) $\mathcal{A}$ generates two signatures $S_0 = \{(\text{key}, \ell_0(\text{key}))\}_{\text{key} \in \mathcal{K}}$ and $S_1 = \{(\text{key}, \ell_1(\text{key}))\}_{\text{key} \in \mathcal{K}}$ with n total values and maximum volume ℓ . That is,
 - $\sum_{\text{key} \in \mathcal{K}} \ell_0(\text{key}) = \sum_{\text{key} \in \mathcal{K}} \ell_1(\text{key}) = n$;
 - $\max_{\text{key} \in \mathcal{K}} \ell_0(\text{key}) = \max_{\text{key} \in \mathcal{K}} \ell_1(\text{key}) = \ell$.
- (2) The challenger $\mathcal{C}$ receives signatures S_0 and S_1 from the adversary $\mathcal{A}$ and generates a multi-map MM with the signature S_η . Specifically, the challenger $\mathcal{C}$ generates $\ell_\eta(\text{key})$ arbitrary values for each $\text{key} \in \mathcal{K}$. The challenger $\mathcal{C}$ then sends $\mathcal{L}_{\text{Setup}}(\text{MM})$ to the adversary $\mathcal{A}$.
- (3) The adversary adaptively picks keys $\text{key}_1, \dots, \text{key}_t$ for query operations. For each key_i , the challenger $\mathcal{C}$ will compute $\mathcal{L}_{\text{Query}}(\text{MM}, \text{key}_1, \dots, \text{key}_i)$ which is sent to the adversary.
- (4) Finally, the adversary $\mathcal{A}$ outputs a bit $b \in \{0, 1\}$.

We denote by $p_{\eta}^{\mathcal{A}, \mathcal{L}}(n, \ell)$ the probability that $\mathcal{A}$ outputs 1 when playing game $\text{VHGame}_{\eta}^{\mathcal{A}, \mathcal{L}}(n, \ell)$.

Definition 6 (Volume-Hiding Leakage Functions). A leakage function $\mathcal{L} = (\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Query}})$ is *volume-hiding* if and only if for all adversaries $\mathcal{A}$ and for all values $n \geq \ell \geq 1$,

$$p_0^{\mathcal{A}, \mathcal{L}}(n, \ell) = p_1^{\mathcal{A}, \mathcal{L}}(n, \ell).$$

Given the above definition of volume-hiding leakage functions, we can now define volume-hiding, STE schemes for multi-maps.

Definition 7 (Volume-Hiding, r -Interactive, Structured Encryption for MMs). A r -interactive, structured encryption scheme for multi-maps, Σ , is *volume-hiding* if there exists a leakage function $\mathcal{L} = (\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Query}})$ such that

- (1) Σ is adaptively $\mathcal{L}$ -secure according to Definition 4;
- (2) $\mathcal{L}$ is a volume-hiding leakage function according to Definition 6.

We stress that our definition of a volume-hiding leakage function is formalized through a game in which the adversary only receives the leakage associated with the setup and queries of the STE scheme. Specifically, the adversary does not see any encrypted multi-map from the setup phase or requests and responses from queries. As a result, we are able to formalize “volume-hiding” as a property of the leakage function as opposed to the specific implementation of any STE scheme. As a consequence, our definition applies to any adversary $\mathcal{A}$ regardless of their computational power. However, this does not imply that a volume-hiding STE scheme is secure with respect to unbounded adversaries as the STE scheme is proven to be only $\mathcal{L}$ -secure against computational adversaries.

2.5 Differentially Private Volume-Hiding Leakage Functions

To properly define the notion of volume-hiding within the framework of differential privacy, we consider a *sanitizer* San . A sanitizer San is a randomized algorithm that takes a signature S and samples a *sanitized signature*. Roughly speaking, we will say that a pair $(\mathcal{L}, \text{San})$, consisting of a leakage function $\mathcal{L}$ and a sanitizer San , is (ϵ, δ) -differentially private if the probabilities that an adversary $\mathcal{A}$ outputs 1 in games VHGame_0 and VHGame_1 , respectively played on the sanitized versions $\text{San}(S_0)$ and $\text{San}(S_1)$ of two *neighboring signatures* S_0 and S_1 , are related through ϵ and δ . Let us now proceed more formally.

Definition 8. Signatures $S_0 = (\text{key}, \ell_0(\text{key}))_{\text{key} \in \mathcal{K}}$ and $S_1 = (\text{key}, \ell_1(\text{key}))_{\text{key} \in \mathcal{K}}$ are *neighbors* if there exist $\text{key}_0, \text{key}_1 \in \mathcal{K}$ such that

- (1) for all $\text{key} \notin \{\text{key}_0, \text{key}_1\}$, $\ell_0(\text{key}) = \ell_1(\text{key})$;
- (2) $\ell_0(\text{key}_0) = \ell_1(\text{key}_0) + 1$;
- (3) $\ell_0(\text{key}_1) = \ell_1(\text{key}_1) - 1$.

We define game dpVH_η , for $\eta = 0, 1$, as in the previous section with the only difference that the multi-set $\mathcal{D}$ is constructed by applying the sanitizer San to signature S_η .

$\text{dpVH}_\eta^{\mathcal{A}, \mathcal{L}, \text{San}}(n)$:

- (1) $\mathcal{A}$ outputs neighboring signatures $S_0 = (\text{key}, \ell_0(\text{key}))_{\text{key} \in \mathcal{K}}$ and $S_1 = (\text{key}, \ell_1(\text{key}))_{\text{key} \in \mathcal{K}}$ with n total values. That is, $\sum_{\text{key} \in \mathcal{K}} \ell_0(\text{key}) = \sum_{\text{key} \in \mathcal{K}} \ell_1(\text{key}) = n$.
- (2) The challenger C receives signatures S_0 and S_1 from $\mathcal{A}$ and samples a sanitized signature $S_{\text{San}} = \{(\text{key}, \ell_{\text{San}}(\text{key}))_{\text{key} \in \mathcal{K}}\}$ by running San on S_η .
- (3) The challenger C constructs a multi-map MM with signature S_{San} by picking $\ell_{\text{San}}(\text{key})$ arbitrary values for each $\text{key} \in \mathcal{K}$. C then sends $\mathcal{L}_{\text{Setup}}(\text{MM})$ to $\mathcal{A}$.
- (4) The adversary $\mathcal{A}$ adaptively picks keys $\text{key}_1, \dots, \text{key}_t$ for query operations. For each key_i , the challenger C computes and sends $\mathcal{L}_{\text{Query}}(\text{MM}, \text{key}_1, \dots, \text{key}_i)$ to the adversary.

- (5) Finally, the adversary $\mathcal{A}$ outputs a bit $b \in \{0, 1\}$.

We denote by $p_{\eta, \mathcal{A}}^{\mathcal{L}, S, \text{San}}$ the probability that adversary $\mathcal{A}$ outputs 1 when playing game $\text{dpVH}_\eta^{\mathcal{L}, S, \text{San}}$

Definition 9 (Differentially Private Volume-Hiding Leakage Functions). We say that $(\mathcal{L}, S)$ is (ϵ, δ) -differentially private volume-hiding if for all n and for all adversaries $\mathcal{A}$ that output two neighboring signatures with the same total number of values n , it holds that

$$p_{0, \mathcal{A}}^{\mathcal{L}, S, \text{San}}(n) \leq e^\epsilon \cdot p_{1, \mathcal{A}}^{\mathcal{L}, S, \text{San}}(n) + \delta.$$

If $\delta = 0$, we will simply say ϵ -differential private instead of (ϵ, δ) -differential private. Note also that a volume-hiding leakage function is $(0, \text{negl}(n))$ -differentially private with respect to the identity sanitizer Id that, for every signatures S , returns the signature itself.

Our definition follows the differentially private literature (see [20] as an example) where the sanitizer is used to preprocesses the data-base (the multi-map in our setting) before being stored. In our construction of Section 4, the multi-map will be implicitly sanitized at query time. The effect on the mitigating leakage will be identical to sanitization at setup time but it will result in a smaller server storage.

3 VOLUME-HIDING EMM

In this section, we describe our volume-hiding structured encryption scheme for multi-maps. To our knowledge, our construction is the first to achieve both asymptotically optimal query and storage complexity. Furthermore, we show that our construction is concretely efficient using experiments. In Section 5, we show that our constructions use less server storage and improve query overhead by a factor of 10-16x over the previous best constructions [32].

One of the major tools that will be used in our construction is cuckoo hashing, which we now describe.

Cuckoo Hashing. Cuckoo hashing was introduced by Pagh and Rodler [45] and consists of two algorithms Build and Search. Given n key-value pairs $(\text{key}_1, \text{value}_1), \dots, (\text{key}_n, \text{value}_n)$, Build constructs Table consisting two arrays T_1 and T_2 each with the capacity to hold $t = (1 + \alpha)n$ pairs for any constant $\alpha > 0$. The Build algorithm inserts the pairs one at the time using two hash functions $h_1, h_2 : \mathcal{K} \rightarrow [t]$. To insert pair $X = (\text{key}_i, \text{value}_i)$, Build places X in location $T_1[h_1(\text{key}_i)]$. If the location is empty, we are done. If the location is currently occupied by another pair, $Y = (\text{key}_j, \text{value}_j)$, Y is evicted and the algorithm attempts to inserting Y in location $T_2[h_2(\text{key}_j)]$. Again if the location is empty, we are done. Otherwise we evict the pair Z found in $T_2[h_2(\text{key}_j)]$ and try to insert Z in the location of T_1 specified by h_1 , and so on. Note that cuckoo hashing guarantees that a pair $(\text{key}_i, \text{value}_i)$ is found either at $T_1[h_1(\text{key}_i)]$ or $T_2[h_2(\text{key}_i)]$. Thus, Search only retrieves two table locations. We say that cuckoo hashing fails if a key-value pair is not inserted after $\Theta(\log n)$ evictions. It is possible to show that all n key-value pairs will be successfully inserted with large constant probability. Unfortunately, this failure probability is too large for privacy reasons and this has been used to compromise privacy of oblivious RAMs (see, e.g., [39]).

Cuckoo Hashing with a Stash. As a result, we resort to a modification of cuckoo hashing presented by Kirsch, Mitzenmacher and

Wieder [38] that introduces a stash S of some fixed capacity s . After $\Theta(\log n)$ evictions, if a key-value pair has not been inserted yet, the pair is inserted into the stash. Cuckoo hashing with a stash fails if strictly more than s items are attempted to be inserted into the stash. The introduction of the stash reduces the failure probability exponentially in the stash size:

THEOREM 10. *For constants $d, \alpha > 0$, there exists algorithm Build that stores n pairs into two arrays T_1 and T_2 of size $(1 + \alpha)n$ and a stash S . Then, the probability that, after all n pairs have been inserted with at most $d \log n$ evictions, S has size greater than s is $O(n^{-s})$.*

For a proof, see Theorem 2.1 of [38] and the extension to non-constant sized stashes in [24].

High-Level Description of vhMM . We start by describing our volume-hiding, multi-map encryption scheme from an intuitive level. Consider any multi-map MM storing $\{\text{key}_i, \tilde{v}_i\}_{i \in [m]}$ with n total values and a maximum volume of ℓ . The Setup algorithm will store the n values using cuckoo hashing with a stash with two tables, T_1 and T_2 , of size $t = (1 + \alpha)n$ for any constant $\alpha > 0$. We will use a PRF F that output values in the range $[t]$. Consider the j -th value, $\tilde{v}_i[j]$ associated with key_i . Then, $(\text{key}_i, \tilde{v}_i[j])$ will be assigned locations $T_1[F_K(\text{key}_i \parallel j \parallel 0)]$, $T_2[F_K(\text{key}_i \parallel j \parallel 1)]$ or the stash using the cuckoo hashing insertion algorithm. If we set the stash to be $f(n) = \omega(1)$, then the insertion will fail with probability at most $O(n^{f(n)}) = \text{negl}(n)$. All empty table locations will be filled with a dummy value. All table locations are then encrypted. The resulting encrypted table will be the encrypted multi-map EMM outsourced to the server. The small stash containing at most $f(n) = \omega(1)$ values will be stored by the client.

We now move onto the query operation for our STE scheme. The Query algorithm executed by the client will simply send the 2ℓ values $\{F_K(\text{key}_i \parallel j \parallel 0), F_K(\text{key}_i \parallel j \parallel 1)\}_{j \in [\ell]}$ to the server. The server executes Reply by returning the encrypted values located at $\{T_1[F_K(\text{key}_i \parallel j \parallel 0)], T_2[F_K(\text{key}_i \parallel j \parallel 1)]\}_{j \in [\ell]}$. The client retrieves the tuple of associated values by decrypting all table locations in the server's response as well as checking the stash.

Formal Description of vhMM . We now formally present our first volume-hiding STE scheme for multi-maps vhMM . In particular, our construction will use pseudorandom family of functions $\mathcal{F} = \{F_s\}_{s \in \{0,1\}^*}$ and an IND-CPA secure encryption scheme $\mathcal{E} = (\text{Enc}, \text{Dec})$. Furthermore, our construction is parameterized by a function $f(n) = \omega(1)$ and a constant $\alpha > 0$.

$(K, \text{EMM}) \leftarrow \text{vhMM.Setup}(1^\lambda, \text{MM} = \{\text{key}_i, \tilde{v}_i\}_{i \in [m]}):$

- (1) Randomly select PRF seed $K_{\text{PRF}} \leftarrow \{0,1\}^\lambda$.
- (2) Randomly select encryption key $K_{\text{Enc}} \leftarrow \{0,1\}^\lambda$.
- (3) Create two empty arrays, T_1, T_2 of length $t = (1 + \alpha)n$ where $n = \sum_{i \in [m]} |\tilde{v}_i|$.
- (4) Initialize $\text{Stash} \leftarrow \emptyset$.
- (5) For each $i \in [m]$ and each $j \in [|\tilde{v}_i|]$:
 - (a) Insert $(\text{key}_i, \tilde{v}_i[j])$ using the cuckoo hashing with a stash insertion algorithm where $(\text{key}_i, \tilde{v}_i[j])$ is assigned to one of $T_1[F_{K_{\text{PRF}}}(\text{key}_i \parallel j \parallel 0)], T_2[F_{K_{\text{PRF}}}(\text{key}_i \parallel j \parallel 1)]$ or Stash .
- (6) If Stash contains more than $f(n)$ items, abort.
- (7) For each location in T_1 or T_2 that is empty, insert $(\perp, \perp)$ into the location.
- (8) For each $i \in [t]$:

- (a) Set $T_1[i] \leftarrow \text{Enc}(K_{\text{Enc}}, T_1[i])$.
- (b) Set $T_2[i] \leftarrow \text{Enc}(K_{\text{Enc}}, T_2[i])$.
- (9) Set the private key as $K \leftarrow (K_{\text{PRF}}, K_{\text{Enc}}, \text{Stash})$.
- (10) Set $\text{EMM} \leftarrow (T_1, T_2)$.
- (11) Return (K, EMM) .

$\text{Request} \leftarrow \text{vhMM.Query}(K, \text{key}):$

- (1) Parse K as $(K_{\text{PRF}}, K_{\text{Enc}}, \text{Stash})$.
- (2) Return $(F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 0), F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 1))_{i \in [\ell]}$.

$\text{Response} \leftarrow \text{vhMM.Reply}(\text{Request}, \text{EMM}):$

- (1) Parse Request as $\{\text{ind}_{0,i}, \text{ind}_{1,i}\}_{i \in [\ell]}$.
- (2) Return $(T_1[\text{ind}_{0,i}], T_2[\text{ind}_{1,i}])_{i \in [\ell]}$.

$\tilde{v} \leftarrow \text{vhMM.Result}(K, \text{key}, \text{Response}):$

- (1) Parse K as $(K_{\text{PRF}}, K_{\text{Enc}}, \text{Stash})$.
- (2) Parse Response as $(\text{ct}_{0,i}, \text{ct}_{1,i})_{i \in [\ell]}$.
- (3) Set $\tilde{v} \leftarrow \emptyset$.
- (4) For each $i \in [\ell]$:
 - (a) Set $(\text{key}', \text{value}') \leftarrow \text{Dec}(K_{\text{Enc}}, \text{ct}_{0,i})$.
 - (b) If $\text{key}' = \text{key}$, $\tilde{v} \leftarrow \tilde{v} \cup \{\text{value}'\}$.
 - (c) Set $(\text{key}', \text{value}') \leftarrow \text{Dec}(K_{\text{Enc}}, \text{ct}_{1,i})$.
 - (d) If $\text{key}' = \text{key}$, $\tilde{v} \leftarrow \tilde{v} \cup \{\text{value}'\}$.
- (5) For each $(\text{key}', \text{value}') \in \text{Stash}$:
 - (a) If $\text{key}' = \text{key}$, $\tilde{v} \leftarrow \tilde{v} \cup \{\text{value}'\}$.
- (6) Return $\tilde{v}$.

Security. We now describe the leakage of vhMM due to algorithms Setup and Query. From a high level, the EMM reveals to the adversarial server $(2 + 2\alpha)n$ encrypted key-value pairs. As a result, $\mathcal{L}_{\text{Setup}} = \text{dsize}$ which is the leakage functions that reveals the number n of total values in MM . For each query, the adversarial server learns two things. First, the adversarial server learns query equality by observing identical 2ℓ PRF values. Secondly, the adversary learns the maximum volume of a key in MM . So, $\mathcal{L}_{\text{Query}} = (\text{req}, \text{mrlen})$. Note, that $\mathcal{L}_{\text{Query}}$ and $\mathcal{L}_{\text{Setup}}$ do not include the response length leakage rlen and thus, intuitively, vhMM is volume-hiding. We prove the following two lemmata in Appendix A.

LEMMA 1. *Let $\mathcal{L} = (\text{dsize}, (\text{req}, \text{mrlen}))$. For every constant $\alpha > 0$ and $f(n) = \omega(1)$, vhMM is an adaptive $\mathcal{L}$ -secure non-interactive STE scheme for multi-maps.*

LEMMA 2. *Leakage function $\mathcal{L} = (\text{dsize}, (\text{req}, \text{mrlen}))$ is volume-hiding.*

By combining the two lemmas above we obtain the main theorem of this section

THEOREM 11. *Non-interactive STE for multi-maps vhMM is volume-hiding.*

Efficiency. By Theorem 10, the probability that the Stash grows larger than $f(n)$ is at most $n^{-f(n)} = \text{negl}(n)$. The outsourced EMM contains exactly $(2 + 2\alpha)n$ encryptions of key-value pairs. The client generates 2ℓ PRF values during vhMM.Query and the server responds with 2ℓ key-value encryptions during vhMM.Reply . The client stores two private keys and the Stash consisting of at most $f(n)$ key-value pairs.

THEOREM 12. *Non-interactive STE for multi-maps vhMM aborts with negligible probability. It requires $(2 + 2\alpha)n$ storage on the server.*

Each query requires 4ℓ communication between the client and the server. The client stores $O(1)$ private keys as well as at most $f(n) = \omega(1)$ key-value pairs.

3.1 Improving Communication

In this section, we present an improvement to our construction vhMM from the previous section. First, we take a careful look at vhMM to look for any wasteful parts. In particular, consider the way that each of the n key-value pairs of the input multi-map are assigned to table locations. For vhMM , any value $\tilde{v}_i[j]$ associated with key_i will be assigned to locations $F_{K_{\text{PRF}}}(\text{key}_i \parallel j \parallel 0)$ and $F_{K_{\text{PRF}}}(\text{key}_i \parallel j \parallel 1)$. As a result, vhMM.Query requires the client to send 2ℓ PRF values. This seems quite wasteful as the 2ℓ table locations associated with key_i are very structured. We show that we can modify vhMM such that the Request generated by Query can be compressed to only contain a single PRF value as opposed to 2ℓ values. To do this, we will use *delegatable PRFs*.

Delegatable PRFs. A family of *delegatable PRFs* (dPRF), first discussed in [37], enables the owner of the secret seed K to *delegate* an untrusted party to compute $F_K(x)$ for all values $x \in S$, where S is a subset taken from a specified family of subsets. The delegation is obtained by computing a *token* tok_S that allows the computation of the PRF for all $x \in S$, without any further intervention of the owner without accessing the secret seed. The security requirement is that all values $F_K(x)$ for $x \notin S$ remain indistinguishable from truly random values even for an adversary that has access to tok_S . Note that this can be trivially achieved by having the owner release the value of $F_K(x)$ for all $x \in S$. Thus, the goal of dPRFs is to construct tokens of size $o(|S|)$. To construct efficient dPRFs, we will use the famous GGM PRF construction which we describe next.

The GGM PRF construction. The GGM construction [22] builds a family of PRFs $\mathcal{F} = \{F_K(\cdot)\}$ from a length-doubling pseudorandom generator G as follows. For any input v of length λ bits, $G(v)$ will result in 2λ bits. For convenience, we denote the first λ bits of the $G(v)$ as $G_0(v)$ and the last λ bits of $G(v)$ as $G_1(v)$. Consider a function $F_{\{0,1\}^\lambda} : \{0,1\}^t \rightarrow \{0,1\}^\lambda$ with a private key and output of the same length λ . The GGM construction uses a binary tree of height t where each node of the tree is labeled with a binary string encoding the path from the root to the node itself using the convention that left is encoded by a “0” and right is encoded by a “1” and the root at level 0 is labeled with the empty string $\perp$. Every node with label $x_1, \dots, x_i \in \{0,1\}^i$ is associated with value $F_K(x_1, \dots, x_i)$ computed in the following recursive manner. The root, labeled with the empty string $\perp$, is assigned the value $F_K(\perp) := K$. If a node is associated with value v , then its left child is associated with $G_0(v)$, the first λ bits of $G(v)$, and its right child is associated with $G_1(v)$, the last λ bits of $G(v)$. This recursive rule is equivalent to defining $F_K(x_1, \dots, x_i) := G_{x_i}(G_{x_{i-1}}(\dots G_{x_1}(K) \dots))$. The PRF output $F_K(X)$ for any t -bit string $X \in \{0,1\}^t$ is the value associated with the leaf with label X . The authors of [22] show that values associated with leaf nodes are indistinguishable from random values for any computational adversary.

Delegating Prefixes. In [37], the above GGM construction of a PRF is used to implement a delegatable PRF for subsets of strings with matching prefixes. More formally, the construction of [37] consists of two efficient algorithms: algorithm dPRF.GenTok that takes as

input the private dPRF key K and an s -bit prefix $X = x_1, \dots, x_s$ with $s \leq t$ and outputs a token tok_X ; and algorithm dPRF.Eval that takes the token for the s -bit prefix X , tok_X and string $Y \in \{0,1\}^{t-s}$ and computes the value $F_K(X \parallel Y)$. In other words, dPRF.Eval enables computation of the PRF evaluation of any string whose prefix is X .

dPRF.GenTok for a prefix X is implemented by returning $\text{tok}_X := F_K(X)$ where F is the GGM PRF construction. For any string Y with prefix X , dPRF.Eval is implemented by simply computing the values the node labeled Y from its ancestor node labeled by X in the GGM PRF construction. This can be done since the pseudorandom generator G is public. For security of dPRF, we only consider adversaries that perform prefix-free string queries. Prefix-free string queries refer to the fact that each query must not be a prefix of any query. In [37], it is shown that the resulting queried tokens are computationally indistinguishable from random values.

High-Level Description of dprfMM . We now describe the modifications to vhMM using prefix delegatable PRFs. The main idea is to replace all instances of a PRF in vhMM with a prefix dPRF. For cuckoo hashing, a value $\tilde{v}_i[j]$ associated with key_i will be assigned to the locations $F_K(\text{key}_i \parallel j \parallel 0)$ in table 1 and $F_K(\text{key}_i \parallel j \parallel 1)$ in table 2 where F is the GGM PRF construction. Furthermore, each $j \in [\ell]$ will be represented using a $\lceil \log_2 \ell \rceil$ -bit string.

We modify Query such that to query any key, the client will construct a token tok_{key} that will delegate the computation for the PRF for any value whose prefix matches key. In particular, this enables the adversarial server to compute the set of PRF values $\{F_K(\text{key} \parallel i \parallel 0), F_K(\text{key} \parallel i \parallel 1)\}_{i \in [\ell]}$ as they all have the shared prefix key. Note that tok_{key} consists of the single PRF value $F_K(\text{key})$. As a result, dprfMM reduces the client-to-server communication from 2ℓ PRF values to just a single PRF value.

Detailed Description of dprfMM . The functions Setup and Reply remain the same as vhMM except replacing the PRF F with the prefix dPRF F construction above.

Request $\leftarrow \text{dprfMM.Query}(K, \text{key})$:

- (1) Parse K as $(K_{\text{PRF}}, K_{\text{Enc}}, \text{Stash})$.
- (2) Compute $\text{tok}_{\text{key}} \leftarrow \text{dPRF.GenTok}(K_{\text{PRF}}, \text{key})$.
- (3) Return tok_{key} .

Response $\leftarrow \text{dprfMM.Reply}(\text{Request}, \text{EMM})$:

- (1) Parse Request as tok_{key} .
- (2) For each $i \in [\ell]$:
 - (a) Compute $F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 0) \leftarrow \text{dPRF.Eval}(\text{tok}_{\text{key}}, i \parallel 0)$.
 - (b) Compute $F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 1) \leftarrow \text{dPRF.Eval}(\text{tok}_{\text{key}}, i \parallel 1)$.
- (3) Return $\{T_1[F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 0)], T_2[F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 1)]\}_{i \in [\ell]}$.

Security. We note that the leakage of dprfMM is identical to vhMM where $\mathcal{L}_{\text{Setup}} = \text{dsize}$ and $\mathcal{L}_{\text{Query}} = (\text{req}, \text{mrlen})$. The proof of security can be found in Appendix B.

THEOREM 13. dprfMM is an adaptive $\mathcal{L} = (\text{dsize}, (\text{req}, \text{mrlen}))$ -secure STE scheme for multi-maps when $f(n) = \omega(1)$ and any constant $\alpha > 0$.

THEOREM 14. dprfMM is volume-hiding.

Efficiency. Note that dprfMM reduces the client-to-server communication from 2ℓ PRF values to one PRF seed of length equal to the

security parameter. Server-to-client communication stays 2ℓ which is asymptotically optimal.

THEOREM 15. *Construction dprfMM requires $(2 + 2\alpha)n$ storage on the server. Each query requires $\ell + 1$ communication between the client and the server. The client stores $O(1)$ private keys as well as at most $f(n) = \omega(1)$ key-value pairs. Both the client and server perform $O(\ell)$ computation.*

4 DIFFERENTIALLY PRIVATE VOLUME-HIDING EMM

In this section, we describe our construction of a STE scheme for multi-maps which is differentially private volume-hiding. In particular, we will slightly modify the dprfMM to be differentially private while improving the query complexity significantly. In the resulting scheme, the client will have to fetch some information from the server before being able to actually issue the query and thus the scheme is not non-interactive.

Note, all volume-hiding constructions pad queries to the maximum volume resulting in $\Omega(\ell)$ query complexity. To be completely volume-hiding, this is necessary to hide the queried key with the largest volume as opposed to other queries with smaller volumes. For differentially private volume-hiding, we can relax our privacy guarantees. Instead, we will perturb the volume of each key by a small amount related to the privacy budget ϵ . As a result, the query complexity for any key is dependent only on the volume of the key as opposed to the maximum volume. We now implement these ideas to obtain our differentially private volume-hiding scheme.

High-Level Description of dpMM. To obtain dpMM, we perform a slight modification to dprfMM. Setup will store all the values of the multi-map using cuckoo hashing via a delegatable PRF. Additionally, the server will also construct an additional cuckoo hash table for a *count table*. Specifically, the count table will store $\ell(\text{key})$, denoting the number of values associated with key, for each key in the input multi-map with non-zero volume.

Queries will be two rounds as opposed to one round. First, the client downloads $\ell(\text{key})$ using cuckoo hashing from the server. Next, the client will generate tok_{key} using the dPRF. However, the client will also generate an integer that is dependent on the volume of key denoted by $\ell(\text{key})$. In particular, the client will also send the value $f(\ell(\text{key}), \lambda) := 2\ell(\text{key}) + I^*(\lambda) + \text{noise}(\text{key})$ along with tok_{key} . The value $\text{noise}(\text{key})$ is drawn according to distribution $\text{Lap}(2/\epsilon)$, which is the Laplacian distribution with parameter $2/\epsilon$. We point out that this value will be identical each time key is queried. This is accomplished by drawing from a Laplacian distribution using pseudorandom bits derived using key. We pick the value $I^*(\lambda)$ later.

Reply is modified such that the server will only expand the dPRF at values $F_K(\text{key} \parallel i \parallel 0), F_K(\text{key} \parallel i \parallel 1)$ where $1 \leq i \leq f(\ell(\text{key}), \lambda)$. Similarly, the server will only return the encrypted values at table locations for these values. As a result, the query complexity is $f(\ell(\text{key}), \lambda)$. Note that if $\text{noise}(\text{key}) = \text{Lap}(2/\epsilon) < -I^*(\lambda)$, then our construction might be lossy as some possible values associated with key stored in the cuckoo hash tables will not be returned to the client. As a result, we pick $I^*(\lambda) = \omega(\log \lambda)$ such that $\Pr[\text{Lap}(2/\epsilon) < -I^*(\lambda)]$ is negligible in λ . Therefore, all associated values will be returned except with negligible probability.

Detailed Description of dpMM. Setup will be almost identical except that the client will also store a count table in the EMM using cuckoo hashing. For an input multi-map $\{\text{key}_i, \vec{v}_i\}_{i \in [m]}$, the client will store $\text{CT} := \{\text{key}_i, \ell(\text{key}_i)\}_{i \in [m]}$ where $\ell(\text{key}_i)$ is the number of values associated with key_i . Note, dpMM will consist of four functions $\text{Query}^0, \text{Reply}^0, \text{Query}^1$ and Reply^1 which constitutes the additional round of interaction.

$(K, \text{EMM}) \leftarrow \text{dpMM.Setup}(1^\lambda, \text{MM} = \{\text{key}_i, \vec{v}_i\}_{i \in [m]}):$

- (1) Execute $(K, \text{EMM}) \leftarrow \text{dprfMM.Setup}(1^\lambda, \text{MM})$.
- (2) Create two empty arrays, CT_1, CT_2 of length $t = (1 + \alpha)n$.
- (3) Generate a $\text{Stash}_{\text{CT}} \leftarrow \emptyset$.
- (4) For each $i \in [m]$:
 - (a) Insert $(\text{key}_i, \ell(\text{key}_i))$ using the cuckoo hashing with a stash insertion algorithm where $(\text{key}_i, \ell(\text{key}_i))$ is assigned to one of $T_1[F_{K_{\text{PRF}}}(\text{CT} \parallel \text{key}_i \parallel 1)], T_2[F_{K_{\text{PRF}}}(\text{CT} \parallel \text{key}_i \parallel 1)]$ or Stash_{CT} .
- (5) If Stash_{CT} contains more than $f(n)$ items, abort.
- (6) For each location in CT_1 or CT_2 that is empty, insert $(\perp, \perp)$ into the location.
- (7) For each $i \in [t']$:
 - (a) Set $\text{CT}_1[i] \leftarrow \text{Enc}(K_{\text{Enc}}, \text{CT}_1[i])$.
 - (b) Set $\text{CT}_2[i] \leftarrow \text{Enc}(K_{\text{Enc}}, \text{CT}_2[i])$.
- (8) Set the private key as $K \leftarrow K \cup \text{Stash}_{\text{CT}}$.
- (9) Set $\text{EMM} \leftarrow \text{EMM} \cup \{\text{CT}_1, \text{CT}_2\}$.
- (10) Return (K, EMM) .

$\text{Request}_0 \leftarrow \text{dpMM.Query}^0(K, \text{key}):$

- (1) Parse K as $(K_{\text{PRF}}, K_{\text{Enc}}, \text{Stash})$.
- (2) Return $\text{tok}_{\text{CT} \parallel \text{key}} \leftarrow \text{dPRF.GenTok}(K_{\text{PRF}}, \text{CT} \parallel \text{key})$.

$\text{Response}_0 \leftarrow \text{dpMM.Reply}^1(\text{Request}_0, \text{EMM}):$

- (1) Parse Request as $\text{tok}_{\text{CT} \parallel \text{key}}$.
- (2) Get $F_{K_{\text{PRF}}}(\text{CT} \parallel \text{key} \parallel 0), F_{K_{\text{PRF}}}(\text{CT} \parallel \text{key} \parallel 1)$ using $\text{dPRF.Eval}(\text{tok}_{\text{CT} \parallel \text{key}}, 0)$ and $\text{dPRF.Eval}(\text{tok}_{\text{CT} \parallel \text{key}}, 1)$.
- (3) Return $\text{CT}_1[F_{K_{\text{PRF}}}(\text{CT} \parallel \text{key} \parallel 0)], \text{CT}_2[F_{K_{\text{PRF}}}(\text{CT} \parallel \text{key} \parallel 1)]$.

$\text{Request}_1 \leftarrow \text{dpMM.Query}^1(K, \text{key}, \text{Response}_0):$

- (1) Use Stash_{CT} and Response_0 to retrieve $\ell(\text{key})$. If $\ell(\text{key})$ is not found, set $\ell(\text{key}) \leftarrow 0$.
- (2) Parse K as $(K_{\text{PRF}}, K_{\text{Enc}}, \text{Stash})$.
- (3) Compute $\text{tok}_{\text{key}} \leftarrow \text{dPRF.GenTok}(K_{\text{PRF}}, \text{key})$.
- (4) Compute $\text{noise}(\text{key}) \leftarrow \text{Lap}_{\text{key}}(2/\epsilon)$.
- (5) Compute $X \leftarrow \ell(\text{key}) + I^*(\lambda) + \text{noise}(\text{key})$.
- (6) Return $(\text{tok}_{\text{key}}, X)$.

$\text{Response}_1 \leftarrow \text{dpMM.Reply}^1(\text{Request}_1, \text{EMM}):$

- (1) Parse Request as $(\text{tok}_{\text{key}}, X)$.
- (2) For each $i \in [X]$:
 - (a) Compute $F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 0) \leftarrow \text{dPRF.Eval}(\text{tok}_{\text{key}}, i \parallel 0)$.
 - (b) Compute $F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 1) \leftarrow \text{dPRF.Eval}(\text{tok}_{\text{key}}, i \parallel 1)$.
- (3) Return $\{T_1[F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 0)], T_2[F_{K_{\text{PRF}}}(\text{key} \parallel i \parallel 1)]\}_{i \in [X]}$.

dpMM.Result is identical to dprfMM except that the server's response will contain less than ℓ encrypted key-value pairs.

Correctness. The probability that not all values are retrieved is upper bounded by the probability that $\text{noise}(\text{key}) < -I^*(\lambda)$. Since

noise(key) is distributed according to $\text{Lap}(2/\epsilon)$, we obtain that this probability is $O(e^{l^*(\lambda)})$ which is negligible in λ , as $l^*(\lambda) = \omega(\log \lambda)$.

Security. We note that dpMM has almost identical leakage to both dprfMM and vhMM except that dpMM also leaks the value $\ell(\text{key}) + l^*(\lambda) + \text{Lap}_{\text{key}}(2/\epsilon)$ for each queried key. Additionally, dpMM leaks the number of unique keys m from the size of the cuckoo hash table for the count table. We prove that dpMM is a differentially private volume-hiding STE scheme for multi-maps in Appendix C due to lack of space.

THEOREM 16. *dpMM is ϵ -differentially private volume-hiding.*

Efficiency. Note that dpMM has almost identical efficiency to dprfMM except for two major differences. The number of encrypted key-value pairs in the server's response generated by Reply is $2\ell(\text{key}) + l^*(\lambda) + \text{Lap}_{\text{key}}(2/\epsilon)$. Note, that $l^*(\lambda) + \text{Lap}_{\text{key}}(2/\epsilon) = \omega(\log \lambda)$ except with probability negligible in λ . For keys with volume significantly smaller than the maximum volume ℓ , dpMM achieves much better query complexity. Also, the outsourced EMM consists of two cuckoo hash schemes instead of one. So, we get:

THEOREM 17. *For any $\alpha > 0$, construction dpMM requires $(2 + 2\alpha)(m + n)$ storage on the server. Each query requires $1 + \ell(\text{key}) + \omega(\log \lambda)$ communication between the client and the server except with probability negligible in λ . The client stores $O(1)$ private keys as well as at most $f(n) = \omega(1)$ key-value pairs. Both the client and server perform $O(\ell(\text{key})) + \omega(\log \lambda)$ computation. Furthermore, the construction dpMM is lossy with probability negligible in λ .*

4.1 Discussion of Differential Privacy

In this section, we present a discussion of when differentially private volume-hiding suffices for security as opposed to volume-hiding. Deciding which security definition is completely dependent on the setting. Consider the case of input multi-maps drawn from a distribution of multi-maps with extremely different volume signatures. In this setting, volume-hiding seems necessary.

On the other hand, consider multi-maps that closely follow the Zipf distribution. Here, the volumes of the same key in different possible multi-maps will not differ significantly. Therefore, differentially private volume-hiding would suffice here. Additionally, suppose we wish to hide the identity of queried keys amongst other keys with similar volumes. In this case, differential privacy may also be used as the keys with similar volume will have similar leakage.

To summarize, differentially private volume-hiding is useful when all possible input multi-maps have similar volume signatures. Differentially private volume-hiding can protect input multi-maps that do not have very different volumes. When attempting to protect significantly different input multi-maps, normal volume-hiding security seems necessary. Therefore, the choice of whether to use dprfMM or dpMM depends on the situation.

5 EXPERIMENTAL EVALUATION

In this section, we present our experimental evaluation of our two main volume-hiding STE schemes for multi-maps: dprfMM and dpMM. We start by describing the setup of our experiments as well as the choice of parameters of our constructions. Then, we compare both dprfMM and dpMM with the Densest Subgraph Transform construction described in [32].

By performing these experiments, we attempt to answer two important questions. First, are the constructions dprfMM and dpMM more concretely efficient compared to the previous most practically efficient construction (Densest Subgraph Transform [32])? Secondly, what is the total cost of dprfMM and dpMM?

We will use multipliers to describe improvements in efficiency. When we say construction A is a 2x improvement in communication over construction B, we mean that construction A uses half the communication used in construction B.

5.1 Experimental Setup

Our experiments are performed using the same machine for both the client and the server. The machines used for the client and server are Ubuntu PCs with 12 cores, 3.5 GHz Intel Xeon E5-1650 and 32 GB of RAM. All the results of our experiments have standard deviations less than 10% of their average. Network resources costs are measured at the application layer. Both the client and server binaries are built using the gRPC library [1].

Input Multi-Maps. In our experiments, we will consider multi-maps containing $n \in \{2^{16}, 2^{18}, 2^{20}, 2^{22}\}$ total values. We note that all constructions under experimentation consider general multi-map inputs. As a result, their efficiencies do not depend on the structure of the multi-map but are completely determined by the parameters n and ℓ (the maximum volume) of the input multi-map. All keys and values will be 8 byte strings.

Primitives. In all our experiments, we consider PRFs with 16 byte keys as well as 16 byte outputs. In particular, we implement our PRFs using SHA256. We also use SHA256 as the pseudorandom generator for our delegatable PRF construction. We will use AES in CBC mode as our symmetric encryption scheme. Note that encrypting a single key-value pair which is 16 bytes would result in a 32 byte ciphertext. As a result, naive encryption of the input multi-map will result in a 2x overhead in storage.

5.2 Cost of Densest Subgraph Transform [32]

In the concrete parameters described in [32], it is shown that the number of bins must be at most $O(n/\log_2 n)$ to ensure that the storage overhead remains linear in the number of total values $O(n)$. Querying in this scheme is equivalent to downloading ℓ bins. Furthermore, each of the bins are padded to the size of the maximum bin. Therefore, it is beneficial to ensure that the bin size remains small. To do this, we can attempt to ensure that the number of bins is large. For practical constants, we pick the number of bins to be $2n/(\log_2 n)$ such that the average bin size will be $\log_2 n/2$. For a concrete EMM implementation, we use a similar delegatable PRF scheme as our own. As a result, the upload consists of a single 16 byte PRF output. Downloading a single result consists of downloading an entire bucket which is padded to the maximum bucket size. For each query, ℓ results will be downloaded to ensure that the construction is volume-hiding. The experimental results of this construction can be found in the first column of Table 2.

5.3 Cost of dprfMM

We implement the dprfMM construction in C++ using OpenSSL for all the underlying cryptographic primitives (SHA256 and AES). For the cuckoo hash table sizes, we will set the constant $\alpha = 0.3$

	Densest Subgraph Transform [32]				dprfMM				dpMM			
Input Multi-Map												
Number of Values (n)	2^{16}	2^{18}	2^{20}	2^{22}	2^{16}	2^{18}	2^{20}	2^{22}	2^{16}	2^{18}	2^{20}	2^{22}
Plaintext Raw Byte Size (MB)	1.05	4.19	16.78	67.11	1.05	4.19	16.78	67.11	1.05	4.19	16.78	67.11
EMM Storage												
Server (MB)	5.53	22.74	88.25	384.40	5.45	21.81	87.24	348.97	6.81	27.26	109.05	436.21
Client Stash (KB)	N/A	N/A	N/A	N/A	0.16	0.50	1.52	4.84	0.21	0.63	1.97	6.18
Query Communication												
Upload (Bytes)	16	16	16	16	16	16	16	16	36	36	36	36
Download (Bytes Per Result)	675.2	780.8	841.6	1008.0	64	64	64	64	64	64	64	64
CPU Costs												
Query (Client ms)	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
Reply (Server ms Per Result)	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
Result (Client ms Per Result)	0.01	0.01	0.01	0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01

Table 2: Microbenchmarks for server and network costs comparing volume-hiding STE schemes. We denote n as the total number of values in the input multi-map. If ℓ is the maximum volume of any key, then the first two column constructions must download ℓ results. On the other hand, the number of results for dpMM will be significantly smaller than ℓ on average. The above results apply for any input multi-map structure, query distribution as well as any value ℓ . We denote milliseconds by ms.

such that both tables hold a total $2.6n$ encrypted key-value pairs. We set the number of maximum evictions before being placed into the stash at $5 \log_2 n$. A query consists of uploading a single 16 byte PRF output and downloading two encrypted locations per result.

Comparing to the Densest Subgraph Transform [32], we see that dprfMM uses smaller server storage overall. Furthermore, as the multi-map increases size, dprfMM starts using significantly smaller amounts of server storage. For example, for the 2^{22} value multi-map, dprfMM uses approximately 40 MB less server storage. In terms of query overhead, dprfMM is significantly better. For each result, dprfMM downloads two encrypted locations resulting in only 64 bytes. This is a 10-16x improvement over the Densest Subgraph Transform [32] which requires 675 bytes (10 times more than dprfMM) per result for the multi-map with 2^{16} values and increases to 1 KB (16 times more than dprfMM) per result for the multi-map with 2^{22} values. Looking at the pattern, dprfMM will have even better query overhead improvements as we consider larger multi-maps. The only tradeoff that is made is that the client must store an additional stash. However, the stash is very small in practice consisting of at most 4 KB. This additional client storage is much smaller than the server storage gains of dprfMM.

In terms of CPU cost, we note that dprfMM is almost identical to the Densest Subgraph Transform [32]. The main difference is that during Reply, the Densest Subgraph Transform has to decrypt significantly larger strings as opposed to dprfMM. As a result, the CPU cost of dprfMM in Reply is smaller.

5.4 Cost of dpMM

We also implement dpMM in C++ using OpenSSL for all underlying cryptographic primitives. In particular, we modify dprfMM in two ways to obtain dpMM. First, dpMM additionally sends an integer during query as well as another PRF for retrieving an entry from the count table costing another 16 byte PRF value. For our experiments appearing in Table 2, we assume close to the worst case in our

experiments for dpMM where the number of keys is $m = n/4$ which causes a very large cuckoo hashing table for counts.

The cost of slightly increased server storage for dpMM is offset by the main benefit of dpMM, which is that the total number of results that will be downloaded is much smaller than the other two constructions. To demonstrate these gains, we consider a natural setting of a multi-map of $n = 2^{20}$ values and $m = n/8 = 2^{17}$ keys following the Zipf distribution with parameters m and 1. As a result, the volumes of the multi-map from largest to smallest will be $(\frac{n}{H_m}, \dots, \frac{n}{m \cdot H_m})$ where $H_m := \sum_{i \in [m]} i^{-1}$ is the harmonic number. In our case, $H_m \approx 12.36$ when $m = 2^{17}$.

We construct dpMM with a privacy budget of $\epsilon = 0.2$ for strong differential privacy guarantees. To pick the value $l^*(\lambda)$, we will ensure that dpMM is lossy with probability at most 2^{-64} . We know that $\Pr[\text{Lap}(2/\epsilon) \leq 2t/\epsilon] \leq e^{-t}$. By setting $t := 56.1$, we can guarantee that the probability that a single key is truncated is $\Pr[\text{Lap}(2/\epsilon) \leq 2t/\epsilon] \leq 2^{-64-17}$. Therefore, the probability that at least one of the $m = 2^{17}$ is truncated is at most 2^{-64} . By our choice, we set $l^*(\lambda) = 5610$. Note, we are being very pessimistic as we picked a small privacy budget of $\epsilon = 0.2$ and a small truncation probability of 2^{-64} . By choosing larger privacy budgets or larger truncation probabilities which might be suitable in certain settings, dpMM would incur smaller query overhead.

In Figure 1, we show the number of results downloaded by dpMM and dprfMM. It is clear that for all queried keys except the one with the largest volume, dpMM downloads significantly fewer results compared to dprfMM. Consider the average case where keys are chosen uniformly at random from the input multi-map following Zipf's distribution. In this case, the volume of the average key is 8. As a result, the average number of returned results for dpMM is 5618 while dprfMM must return more than 84,000 results (15 times more than dpMM). Since both dpMM and dprfMM communicate 64 bytes per result, dpMM is a 15x improvement in query communication over dprfMM. Additionally, we note that the

Figure 1: Comparison of number of downloaded results by dprfMM and dpMM.

Densest Subgraph Transform [32] also returns 84,000 results but each result requires communicating 841.6 bytes. As a result, dpMM is an 150-240x improvement in average query overhead over the Densest Subgraph Transform.

6 CONCLUSIONS

In this work, we present a volume-hiding scheme dprfMM that is practically more efficient than any previous volume-hiding scheme. dprfMM always uses less server storage compared to the previous construction. Furthermore, dprfMM improves the communication costs of queries by a factor of 10-16x over the previous best constructions [32] when encrypting multi-maps that occupy 1-67 MB in the plaintext and consist of 2^{16} - 2^{22} total values. From an asymptotic perspective, dprfMM is the first construction with both asymptotically optimal storage and query overhead. We also present the first formal definition of volume-hiding leakage functions.

In addition, we also introduce the notion of *differentially private volume-hiding* which strikes a better, tunable balance between privacy and efficiency. We present dpMM that is able to significantly improve the average query overhead over the previous best volume-hiding schemes [32] by a factor of 150-240x when encrypting multi-maps with 2^{16} - 2^{22} total values that occupy 1-67 MB in the plaintext.

Altogether, we significantly further the field of volume-hiding encrypted multi-maps by presenting both conceptual (formal definitions) and algorithmic (constructions that are both asymptotically and practically efficient) contributions.

REFERENCES

- [1] 2018. gRPC - An RPC library and framework. <https://github.com/grpc/grpc>.
- [2] Ghous Amjad, Seny Kamara, and Tarik Moataz. 2018. Breach-Resistant Structured Encryption. Cryptology ePrint Archive, Report 2018/195. <https://eprint.iacr.org/2018/195>.
- [3] Gilad Asharov, Ilan Komargodski, Wei-Kai Lin, Kartik Nayak, Enoch Peserico, and Elaine Shi. 2018. OptORAMA: Optimal Oblivious RAM. Cryptology ePrint Archive, Report 2018/892.
- [4] Gilad Asharov, Moni Naor, Gil Segev, and Ido Shahaf. 2016. Searchable symmetric encryption: optimal locality in linear space via two-dimensional balanced allocations. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM, 1101–1114.
- [5] Gilad Asharov, Gil Segev, and Ido Shahaf. 2018. Tight Tradeoffs in Searchable Symmetric Encryption. Cryptology ePrint Archive, Report 2018/507. <https://eprint.iacr.org/2018/507>.
- [6] Mihir Bellare, Alexandra Boldyreva, and Adam O'Neill. 2007. Deterministic and efficiently searchable encryption. In *CRYPTO '07*.
- [7] Alexandra Boldyreva, Nathan Chenette, Younho Lee, and Adam O'Neill. 2009. Order-preserving symmetric encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 224–241.
- [8] Alexandra Boldyreva, Nathan Chenette, and Adam O'Neill. 2011. Order-preserving encryption revisited: Improved security analysis and alternative solutions. In *Annual Cryptology Conference*. Springer, 578–595.
- [9] Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. 2004. Public key encryption with keyword search. In *EUROCRYPT '04*. 506–522.
- [10] Raphael Bost. 2016. Sophos: Forward Secure Searchable Encryption. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1143–1154.
- [11] Raphaël Bost, Brice Minaud, and Olga Ohrimenko. 2017. Forward and backward private searchable encryption from constrained cryptographic primitives. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1465–1482.
- [12] David Cash, Paul Grubbs, Jason Perry, and Thomas Ristenpart. 2015. Leakage-abuse attacks against searchable encryption. In *CCS '15*.
- [13] David Cash, Joseph Jaeger, Stanislaw Jarecki, Charanjit S. Jutla, Hugo Krawczyk, Marcel-Cătălin Rosu, and Michael Steiner. 2014. Dynamic searchable encryption in very-large databases: data structures and implementation.. In *NDSS*, Vol. 14. Citeseer, 23–26.
- [14] David Cash, Stanislaw Jarecki, Charanjit Jutla, Hugo Krawczyk, Marcel-Cătălin Rosu, and Michael Steiner. 2013. Highly-scalable searchable symmetric encryption with support for boolean queries. In *Annual Cryptology Conference*. Springer, 353–373.
- [15] David Cash and Stefano Tessaro. 2014. The locality of searchable symmetric encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 351–368.
- [16] Melissa Chase and Seny Kamara. 2010. Structured encryption and controlled disclosure. In *EUROCRYPT '10*. Springer, 577–594.
- [17] Reza Curtmola, Juan Garay, Seny Kamara, and Rafail Ostrovsky. 2011. Searchable symmetric encryption: improved definitions and efficient constructions. *Journal of Computer Security* (2011).
- [18] Ioannis Demertzis, Dimitrios Papadopoulos, and Charalampos Papamanthou. 2018. Searchable encryption with optimal locality: Achieving sublogarithmic read efficiency. In *Annual International Cryptology Conference*. Springer, 371–406.
- [19] Ioannis Demertzis and Charalampos Papamanthou. 2017. Fast searchable encryption with tunable locality. In *Proceedings of the 2017 ACM International Conference on Management of Data*. ACM, 1053–1067.
- [20] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). 265–284.
- [21] Eu-Jin Goh. 2003. Secure indexes. *IACR Cryptology ePrint Archive* 2003 (2003), 216.
- [22] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. 1986. How to Construct Random Functions. *J. ACM* 33, 4 (Aug. 1986), 792–807. <https://doi.org/10.1145/6490.6503>
- [23] Oded Goldreich and Rafail Ostrovsky. 1996. Software Protection and Simulation on Oblivious RAMs. *J. ACM* 43, 3 (1996).
- [24] Michael T. Goodrich and Michael Mitzenmacher. 2011. Privacy-preserving Access of Outsourced Data via Oblivious RAM Simulation. In *ICALP*. 576–587.
- [25] Paul Grubbs, Marie-Sarah Lacharité, Brice Minaud, and Kenneth G. Paterson. 2018. Pump up the Volume: Practical Database Reconstruction from Volume Leakage on Range Queries. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*. 315–331.
- [26] Paul Grubbs, Marie-Sarah Lacharité, Brice Minaud, and Kenneth G. Paterson. [n. d.]. Learning to Reconstruct: Statistical Learning Theory and Encrypted Database Attacks. Cryptology ePrint Archive, Report 2019/011.
- [27] Paul Grubbs, Richard McPherson, Muhammad Naveed, Thomas Ristenpart, and Vitaly Shmatikov. 2016. Breaking web applications built on top of encrypted data. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1353–1364.
- [28] Paul Grubbs, Kevin Sekniqi, Vincent Bindschadler, Muhammad Naveed, and Thomas Ristenpart. 2017. Leakage-abuse attacks against order-revealing encryption. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 655–672.
- [29] Mohammad Saiful Islam, Mehmet Kuzu, and Murat Kantarcioglu. 2012. Access Pattern disclosure on Searchable Encryption: Ramification, Attack and Mitigation. In *19th Annual Network and Distributed System Security Symposium, NDSS 2012, San Diego, California, USA, February 5-8, 2012*.
- [30] Seny Kamara and Tarik Moataz. 2017. Boolean searchable symmetric encryption with worst-case sub-linear complexity. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 94–124.
- [31] Seny Kamara and Tarik Moataz. 2018. SQL on structurally-encrypted databases. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 149–180.

- [32] Seny Kamara and Tarik Moataz. 2019. Computationally Volume-Hiding Structured Encryption. In *EUROCRYPT 2019*, Yuval Ishai and Vincent Rijmen (Eds.), 183–213.
- [33] Seny Kamara, Tarik Moataz, and Olga Ohrimenko. 2018. Structured Encryption and Leakage Suppression. In *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part I*. 339–370.
- [34] Seny Kamara, Charalampos Papamanthou, and Tom Roeder. 2012. Dynamic searchable symmetric encryption. In *Proceedings of the 2012 ACM conference on Computer and communications security*. ACM, 965–976.
- [35] Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. 2016. Generic Attacks on Secure Outsourced Databases. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*. 1329–1340.
- [36] Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. 2017. Accessing Data while Preserving Privacy. *CoRR* abs/1706.01552 (2017). <http://arxiv.org/abs/1706.01552>
- [37] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. 2013. Delegatable Pseudorandom Functions and Applications. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security (CCS '13)*. 669–684.
- [38] Adam Kirsch, Michael Mitzenmacher, and Udi Wieder. 2009. More Robust Hashing: Cuckoo Hashing with a Stash. *SIAM J. Comput.* 39, 4 (2009), 1543–1561.
- [39] Eyal Kushilevitz, Steve Lu, and Rafail Ostrovsky. 2012. On the (in)Security of Hash-based Oblivious RAM and a New Balancing Scheme. In *SODA*. 143–156.
- [40] Marie-Sarah Lacharité, Brice Minaud, and Kenneth G. Paterson. 2018. Improved Reconstruction Attacks on Encrypted Data Using Range Query Leakage. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*. 297–314.
- [41] Kasper Green Larsen and Jesper Buus Nielsen. 2018. Yes, There is an Oblivious RAM Lower Bound!. In *CRYPTO '18*.
- [42] Ian Miers and Payman Mohassel. 2016. IO-DSSE: Scaling Dynamic Searchable Encryption to Millions of Indexes By Improving Locality. *IACR Cryptology ePrint Archive* 2016 (2016), 830.
- [43] Muhammad Naveed, Seny Kamara, and Charles V. Wright. 2015. Inference Attacks on Property-Preserving Encrypted Databases. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*. 644–655.
- [44] Muhammad Naveed, Manoj Prabhakaran, and Carl A Gunter. 2014. Dynamic searchable encryption via blind storage. In *2014 IEEE Symposium on Security and Privacy*. IEEE, 639–654.
- [45] Rasmus Pagh and Flemming Friche Rodler. 2004. Cuckoo Hashing. *J. Algorithms* 51, 2 (2004), 122–144.
- [46] Sarvar Patel, Giuseppe Persiano, Mariana Raykova, and Kevin Yeo. 2018. PanORAMa: Oblivious RAM with logarithmic overhead. In *FOCS '18*.
- [47] Sarvar Patel, Giuseppe Persiano, and Kevin Yeo. 2017. Symmetric Searchable Encryption with Sharing and Unsharing. *IACR Cryptology ePrint Archive* 2017 (2017), 973.
- [48] Giuseppe Persiano and Kevin Yeo. 2019. Lower Bounds for Differentially Private RAMs. In *EUROCRYPT 2019*, Yuval Ishai and Vincent Rijmen (Eds.), Vol. 11476. 404–434.
- [49] Raluca Ada Popa, Catherine Redfield, Nikolai Zeldovich, and Hari Balakrishnan. 2011. CryptDB: protecting confidentiality with encrypted query processing. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*. ACM, 85–100.
- [50] David Prouliot and Charles V Wright. 2016. The shadow nemesis: Inference attacks on efficiently deployable, efficiently searchable encryption. In *CCS '16*.
- [51] D. Song, D. Wagner, and A. Perrig. 2000. Practical techniques for searches on encrypted data. In *Proceeding 2000 IEEE Symposium on Security and Privacy. S&P 2000*. 44–55.
- [52] Emil Stefanov, Charalampos Papamanthou, and Elaine Shi. 2014. Practical Dynamic Searchable Encryption with Small Leakage.. In *NDSS*, Vol. 71. 72–75.
- [53] Emil Stefanov, Marten van Dijk, Elaine Shi, Christopher Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. 2013. Path ORAM: An Extremely Simple Oblivious RAM Protocol. In *CCS '13*. 299–310.
- [54] Yupeng Zhang, Jonathan Katz, and Charalampos Papamanthou. 2016. All Your Queries Are Belong to Us: The Power of File-Injection Attacks on Searchable Encryption.. In *USENIX Security Symposium*. 707–720.

A SECURITY PROOF OF vhMM

PROOF OF LEMMA 1. We construct a stateful simulator $\mathcal{S}$ for both Setup and Query.

$(\text{state}, \text{EMM}) \leftarrow \mathcal{S}.\text{SimSetup}(1^\lambda, \text{dsize}(\text{MM}))$:

- (1) Set $n \leftarrow \text{dsize}(\text{MM})$.

- (2) Construct two empty arrays T_1, T_2 of length $t = (1 + \alpha)n$.
- (3) Fill each array location with a uniformly random value from $\{0, 1\}^\lambda$.
- (4) Return $(\perp, (T_1, T_2))$.

$(\text{state}, \text{Response}) \leftarrow \mathcal{S}.\text{SimQuery}(1^\lambda, \text{req}(\text{key}_1, \dots, \text{key}_i), \text{mrlen}(\text{MM}))$:

- (1) Using $\text{req}(\text{key}_1, \dots, \text{key}_i)$, determine if $\text{key}_i = \text{key}_j$ for some $j < i$.
- (2) If $\text{key}_i = \text{key}_j$ for some $j < i$:
 - (a) Return $(\text{state}, \text{state}[\text{key}_j])$.
- (3) Suppose key_i is unique from all previous queried keys:
 - (a) Set $\text{Response} \leftarrow (r_1, \dots, r_{2\ell})$ where each r_i is drawn uniformly at random from $\{0, 1\}^\lambda$.
 - (b) Set $\text{state}[\text{key}_i] \leftarrow \text{Response}$.
 - (c) Return $(\text{state}, \text{Response})$.

We now show that for all PPT adversaries $\mathcal{A}$, the probability $\text{Real}_{\text{vhMM}, \mathcal{A}}(1^\lambda)$ outputs 1 is negligibly different from the probability that $\text{Ideal}_{\text{vhMM}, \mathcal{A}, \mathcal{S}}(1^\lambda)$ outputs 1. To do this, we use the following sequence of games:

- **Game₀** is identical to $\text{Real}_{\text{vhMM}, \mathcal{A}}(1^\lambda)$.
- **Game₁** remove Step 5 from $\text{vhMM}.\text{Setup}$.
- **Game₂** replaces the IND-CPA encryption scheme Enc with a random function.
- **Game₃** replaces the outputs of both random functions with uniformly random chosen values.
- **Game₄** replaces the PRF F with a random function.

Note **Game₀** and **Game₁** are only negligibly-distinguishable by a computational adversary as cuckoo hashing with a stash fails with negligible probability when $f(n) = \omega(1)$. It can be shown that **Game₄** is identical to $\text{Ideal}_{\text{vhMM}, \mathcal{A}, \mathcal{S}}(1^\lambda)$. $\square$

PROOF OF LEMMA 2. To prove that $\mathcal{L} = (\text{dsize}, (\text{req}, \text{mrlen}))$ is volume-hiding, we consider any two multi-map signatures with the same number of values n and maximum volume ℓ . Note, $\mathcal{L}_{\text{Setup}}$ for vhMM consists of $\text{dsize}(\text{MM}) = n$. Similarly, $\mathcal{L}_{\text{Query}}$ consists of only $\text{mrlen}(\text{MM}) = \ell$ and query equality for the queried keys. The leakage of dsize and mrlen are identical for both signatures. Query equality leakage is independent of the input multi-maps. As a result, the input to the adversary in both games with different signatures is identical completing the proof. $\square$

B SECURITY PROOF OF dprfMM

PROOF OF THEOREM 13. The stateful simulator for construction dprfMM is identical to vhMM for SimSetup . We now present SimQuery for dprfMM .

$\mathcal{S}.\text{SimQuery}(1^\lambda, \text{req}(\text{key}_1, \dots, \text{key}_i), \text{mrlen}(\text{MM}))$:

- (1) Using $\text{req}(\text{key}_1, \dots, \text{key}_i)$, determine if $\text{key}_i = \text{key}_j$ for some $j < i$.
- (2) If $\text{key}_i = \text{key}_j$ for some $j < i$:
 - (a) Return $(\text{state}, \text{state}[\text{key}_j])$.
- (3) Suppose key_i is unique from all previous queried keys:
 - (a) Set $\text{Response} \leftarrow r$ where each r is drawn uniformly at random from $\{0, 1\}^\lambda$.
 - (b) Set $\text{state}[\text{key}_i] \leftarrow \text{Response}$.
 - (c) Return $(\text{state}, \text{Response})$.

Note the only difference is that Response consists of a single random value as opposed to 2ℓ random values.

The sequence of games is identical except that the dPRF (instead of the PRF) is replaced by a random function. The two games being computationally indistinguishable follows from the proof of security in [37]. $\square$

PROOF OF THEOREM 14. The leakage of dprfMM and vhMM are identical. As a result, the proof follows identically to Theorem 2. $\square$

C SECURITY PROOF OF dpMM

Before proving Theorem 16, we note that extra $l^*(\lambda) + \text{noise}(\text{key})$ downloads are performed by the client, where $\text{noise}(\text{key}) \leftarrow \text{Lap}(2/\epsilon)$. We stress though that the same value of $\text{noise}(\text{key})$ must be used for each query for key and this can be achieved in two possible ways:

- (1) We consider the sanitizer San_ϵ that, for each key of multi-set $\mathcal{D}$, samples $\text{noise}(\text{key}) \leftarrow \text{Lap}(2/\epsilon)$ and adds $l^*(\lambda) + \text{noise}(\text{key})$ dummy values to key, before executing Setup over it.

This approach has the drawback of a multiplicative overhead of $\omega(\log N)$ for server storage, for a set of size $N = \text{poly}(\lambda)$.

- (2) We do not sanitize the multi-set $\mathcal{D}$ before executing Setup but each time we query for key we sample $\text{noise}(\text{key})$. To ensure that for every query of key we obtain the same value of $\text{noise}(\text{key})$, we fix the random bits to be used by the sampling algorithm as the output of a PRF whose seed is part of the client secret key.

We stress that, as far as differential privacy and the volume-hiding property are concerned, the two approaches are equivalent. So we use the second approach in our construction and the first one in our proof.

The proof of Theorem 16 consists of two parts. First, we prove that our construction is $\mathcal{L}$ -secure for an appropriate leakage function $\mathcal{L}$ and then we show that $\mathcal{L}$ is (ϵ, δ) -differentially private volume-hiding when coupled with San_ϵ .

We consider the leakage function $\mathcal{L} = (\text{dsize}, (\text{qep}, \text{dprlen}))$, where $\text{dprlen}(\mathcal{D}, \text{key}_1, \dots, \text{key}_i)$ is simply rlen , the response length leakage function, applied to the sanitized multi-map $\text{San}(\mathcal{D})$. The proof of $\mathcal{L}$ -security follows similarly to the proof of security of vhMM for any San_ϵ .

We next prove that, for every ϵ , $(\text{San}_\epsilon, \mathcal{L})$ is $(\epsilon, 0)$ -differential privacy volume-hiding. We remind the reader that sanitizer San_ϵ , described at bullet (1) above, adds extra $l^*(\lambda) + \text{noise}(\text{key})$ values for key, where $\text{noise}(\text{key}) \leftarrow \text{Lap}(2/\epsilon)$.

We use the following well known property of the Laplacian distribution. Specifically, let us denote by h the density function of the Laplacian distribution with parameter γ . Then, for every y, y' we have that

$$h(y) \leq h(y') \cdot e^{\frac{|y-y'|}{\gamma}}.$$

For an adversary $\mathcal{A}$, we denote by $\mathcal{T}_{\mathcal{A}}(S)$ the random variable of the leakage transcript obtained by $\mathcal{A}$ in the game $\text{dpVH}^{\mathcal{A}, \mathcal{L}, \text{San}_\epsilon}$ on input signature S . We will show that for any transcript $T = (t_1, \dots, t_q)$ and for any two neighboring signatures, S_0 and S_1 ,

$$\text{Prob}[\mathcal{T}_{\mathcal{A}}(S_0) = T] \leq e^\epsilon \cdot \text{Prob}[\mathcal{T}_{\mathcal{A}}(S_1) = T].$$

This will suffice to prove differential privacy.

First of all, we observe that if query $Q_i = Q_j$ and $t_i \neq t_j$, then

$$\text{Prob}[\mathcal{T}_{\mathcal{A}}(S_0) = T] = \text{Prob}[\mathcal{T}_{\mathcal{A}}(S_1) = T] = 0.$$

We can thus restrict ourselves to transcripts that are *consistent* with Q ; that is, transcripts for which $Q_i = Q_j$ implies $t_i = t_j$. For $b = 0, 1$, we write

$$\text{Prob}[\mathcal{T}_{\mathcal{A}}(S_b) = T] = \prod_{i=1}^q \text{Prob}[\mathcal{T}_{\mathcal{A}}(S_b)_i = t_i \mid t_1, \dots, t_{i-1}]$$

where $\mathcal{T}_{\mathcal{A}}(S_b)_i$ denotes the i -th component of the random variable of the transcript. We also denote by $\text{noise}_0(\text{key})$ and $\text{noise}_1(\text{key})$ the random variable relative to the distributions $\mathcal{T}_{\mathcal{A}}(S_0)$ and $\mathcal{T}_{\mathcal{A}}(S_1)$, respectively. Consider the following cases.

- (1) If S_0 and S_1 do not differ for key Q_i then

$$\text{Prob}[\mathcal{T}_{\mathcal{A}}(S_0)_i = t_i \mid t_1, \dots, t_{i-1}] = \text{Prob}[\mathcal{T}_{\mathcal{A}}(S_1)_i = t_i \mid t_1, \dots, t_{i-1}].$$

- (2) Suppose that S_0 and S_1 differ for key $Q_i = \text{key}$ (that is, $\ell_0(\text{key}) \neq \ell_1(\text{key})$) and suppose that the i -th query is the first query for key; that is, for all $j < i$, $Q_j \neq \text{key}$. Then,

$$\begin{aligned} \text{Prob}[\mathcal{T}_{\mathcal{A}}(S_0)_i = t_i \mid t_1, \dots, t_{i-1}] &= \text{Prob}[\mathcal{T}_{\mathcal{A}}(S_0)_i = t_i] \\ &= \text{Prob}[\text{noise}_0(\text{key}) = t_i - \ell_0(\text{key}) - \ell^*(\lambda)] \\ &\leq e^{\frac{\epsilon}{2} |\ell_0(\text{key}) - \ell_1(\text{key})|} \cdot \text{Prob}[\text{noise}_1(\text{key}) = t_i - \ell_1(\text{key}) - \ell^*(\lambda)] \\ &\leq e^{\frac{\epsilon}{2}} \cdot \text{Prob}[\text{noise}_1(\text{key}) = t_i - \ell_1(\text{key}) - \ell^*(\lambda)], \\ &\quad (\text{since } |\ell_0(\text{key}) - \ell_1(\text{key})| = 1) \\ &= e^{\frac{\epsilon}{2}} \cdot \text{Prob}[\mathcal{T}_{\mathcal{A}}(S_1)_i = t_i \mid t_1, \dots, t_{i-1}] \end{aligned}$$

- (3) If instead S_0 and S_1 differ for key $Q_i = \text{key}$ but for some $j < i$, $Q_j = \text{key}$, then, by the consistency of the transcript, we have

$$\text{Prob}[\mathcal{T}_{\mathcal{A}}(S_b)_i = t_i \mid t_1, \dots, t_{i-1}] = 1,$$

for $b = 0, 1$.

Since S_0 and S_1 differ for at most two keys we have that

$$\text{Prob}_{\mathcal{A}}[\mathcal{T}(S_0) = T] \leq e^\epsilon \cdot \text{Prob}_{\mathcal{A}}[\mathcal{T}(S_1) = T].$$