

Ghous Amjad\*, Seny Kamara, and Tarik Moataz

# Breach-Resistant Structured Encryption

**Abstract:** Motivated by the problem of data breaches, we formalize a notion of security for dynamic structured encryption (STE) schemes that guarantees security against a *snapshot* adversary; that is, an adversary that receives a copy of the encrypted structure at various times but does not see the transcripts related to any queries. In particular, we focus on the construction of dynamic encrypted multi-maps which are used to build efficient searchable symmetric encryption schemes, graph encryption schemes and encrypted relational databases. Interestingly, we show that a form of snapshot security we refer to as *breach resistance* implies previously-studied notions such as a (weaker version) of history independence and write-only obliviousness. Moreover, we initiate the study of *dual-secure* dynamic STE constructions: schemes that are forward-private against a persistent adversary and breach-resistant against a snapshot adversary. The notion of forward privacy guarantees that updates to the encrypted structure do not reveal their association to any query made in the past. As a concrete instantiation, we propose a new dual-secure dynamic multi-map encryption scheme that outperforms all existing constructions; including schemes that are not dual-secure. Our construction has query complexity that grows with the selectivity of the query and the number of deletes since the client executed a linear-time rebuild protocol which can be de-amortized. We implemented our scheme (with the de-amortized rebuild protocol) and evaluated its concrete efficiency empirically. Our experiments show that it is highly efficient with queries taking less than 1 microsecond per label/value pair.

**Keywords:** Searchable Encryption, Structured Encryption, Snapshot Security, Forward Privacy

DOI 10.2478/popets-2019-0014

Received 2018-05-31; revised 2018-09-15; accepted 2018-09-16.

---

**\*Corresponding Author: Ghous Amjad:** Brown University, E-mail: ghous\_amjad@brown.edu

**Seny Kamara:** Brown University, E-mail: seny@brown.edu

**Tarik Moataz:** Brown University, E-mail:

tarik\_moataz@brown.edu

## 1 Introduction

The constant occurrence of data breaches has generated a lot of interest from academia, industry and government in the subject of encrypted search and, in particular, in encrypted databases (EDB). Because EDBs can protect data at all times—even in use—they inherently provide stronger security and privacy guarantees than standard database systems. There are several ways in which a data breach can occur: (1) the server that runs the database system is compromised; (2) the database itself is somehow exfiltrated; and (3) the application is compromised and the database is retrieved using the standard database interface. The first threat can be modeled as a *persistent* adversary that controls the database server at all times. The second can be captured by a *snapshot* adversary that only gets a copy of the database at specific points in time. Note that the third cannot be addressed using cryptographic techniques because once the adversary compromises the application and gets its credentials, it is indistinguishable from the application.

Encrypted search solutions have traditionally been designed to address persistent adversaries with the understanding that security against a persistent adversary implies security against snapshot adversaries. While this is indeed the case for most *static* constructions, it is not necessarily true in the dynamic case.<sup>1</sup> In fact, solutions designed specifically against a persistent adversary could still leak information to a snapshot adversary that one would expect to be hidden. Indeed, this can happen if the query or update operation modifies the encrypted structure in a way that depends on some previous query. The resulting leakage *persists* and can therefore be observable by a snapshot adversary. We stress that there are very natural constructions that work this way including [5, 6, 16, 29]. The problem of designing snapshot-secure EDBs, therefore, is non-trivial and is even more challenging when security against both a persistent and a snapshot adversary is required.

---

**1** More precisely, it holds for static constructions with minimal setup leakage and query operations that do not modify the structure.

## 1.1 Previous Work and Challenges

**Structured encryption.** A promising approach to designing EDBs relies on *structured encryption* schemes which provide a balance of security, efficiency, expressiveness. A structured encryption scheme encrypts a data structure in such a way that it can be privately queried. Roughly speaking, an STE scheme is secure if it reveals nothing about the underlying structure and queries beyond some well-specified leakage. Special cases of STE include graph [12, 32], dictionary [10, 12] and, in particular, multi-map encryption schemes [5, 10, 13], which are going to be the focus of this work.

Multi-maps are data structures that store pairs of the form  $(\ell, \mathbf{v})$ , where  $\ell$  is a label and  $\mathbf{v} = (v_1, \dots, v_n)$  is a tuple of  $n$  values. Multi-maps support a get operation that takes as input a label and returns its associated tuple. Encrypted multi-maps (EMM) naturally yield single-keyword SSE schemes by associating to the set of encrypted documents  $(ct_1, \dots, ct_n)$ , an EMM that stores pairs of the form  $(w, \mathbf{v})$  where  $w$  is a keyword and  $\mathbf{v} = (id_1, \dots, id_m)$  is a tuple of identifiers for the files that contain  $w$ . In addition, EMMs can be used as building blocks to design graph encryption schemes [12], boolean SSE schemes [8, 23] and encrypted relational databases [25].

In many practical scenarios, encrypted data structures need to be *dynamic*; that is, able to support additions and deletions. As such, efficient dynamic SSE/EMM constructions have received a lot of attention both in terms of design [5, 6, 10, 16, 17, 21, 23, 24, 28, 33, 36–38] and cryptanalysis [7, 22, 39]. In particular, it has been shown that SSE schemes, under some assumptions, can be prone to query recovery attacks against passive [7, 22],<sup>2</sup> and dynamic [39] adversaries.

**Forward-privacy.** The most recent work on dynamic EMMs has focused on the notion of forward privacy which was first proposed by Stefanov, Papamanthou and Shi [38] and later formalized by Bost [5] with an alternative definition given by Lai and Chow in [29]. Roughly speaking, forward privacy guarantees that updates to the structure do not reveal their association to any query made in the past. While forward privacy is a useful security property in of itself, it has also been shown to mitigate the adaptive file injection attacks of Zhang, Katz and Papamanthou [39] (though not the non-adaptive attacks).

Currently, the known forward-private EMMs are the SPS construction by Stefanov et al. [38], the Sophos construction of Bost [5], the TWORAM construction by Garg, Mohassel and Papamanthou [18], the EKPE construction by Etemad, K p c , Papamanthou and Evans [16], and a recent construction by Lai and Chow [29]. Recently, Bost, Minaud and Ohrimenko [6] presented the first backward-private EMM, a notion that was informally introduced in [38]. At a high level, backward privacy guarantees that queries do not reveal their association to deleted documents. While there are no known attacks that leverage the lack of backward privacy, improving the security guarantees of EMMs is well-motivated.

**Snapshot security.** While most STE constructions are known to be adaptively-secure against a *persistent* adversary, as far as we know, no previous work has considered STE in the context of *snapshot* adversaries. Informally, a persistent adversary has access to the encrypted structure and has access to the transcripts of the interactions between the client and the server. A snapshot adversary, on the other hand, only has access to the encrypted structure (but at various times). Given that snapshot adversaries are clearly weaker, it motivates the following natural question:

*Can we design efficient and dynamic EMMs that are secure against a persistent adversary and are (almost) zero-leakage against a snapshot adversary?*

We answer this question positively and in doing so introduce a new kind of dynamic EMM which are efficient and secure against both snapshot and persistent adversaries. Specifically, these EMMs are *forward-private* against persistent adversaries and *breach-resistant* against snapshot adversaries, in the sense that they leak only the size of the structure at the time of the snapshot. Interestingly, while our constructions offer stronger security guarantees than previous work, it also achieves better asymptotic efficiency in terms of query and update complexity, token and encrypted structure size and client storage (under reasonable assumptions).

## 1.2 Our Contributions

In this work, we revisit dynamic EMMs in several ways.

**Breach resistance.** Snapshot adversaries are motivated by the threat of data breaches. In such a scenario,

<sup>2</sup> Note that the attacks [7, 22] require the knowledge of 95% and 80% of the user’s data in order to work, respectively.

the adversary is not necessarily the server itself but some other party that gets access to a copy of the encrypted structure at some point(s) in time.

We propose a formal definition of snapshot security which, intuitively, requires that a copy of the encrypted structure reveals nothing about the structure and the past operations beyond what is revealed by a well-specified leakage function we refer to as the *snapshot leakage*. We then say that an STE scheme is *breach-resistant* if its snapshot leakage reveals at most the current size of the structure. Our notion of breach resistance is similar to the notion of offline security of Lewi and Wu in the setting of property-preserving encryption (PPE) [30]. The IND-CPA-DS definition in [27] also tries to capture snapshot security in the PPE setting, but it only holds for single snapshots and semi-dynamic schemes (that do not support deletes). We also show that, breach-resistance implies some interesting and previously studied notions such as (a weaker version) of history independence [35] and write-only obliviousness [4].

We stress that our notion of breach resistance applies only to the encrypted structures we design and that, as pointed out by Grubbs, Ristenpart and Shmatikov [20], there are non-trivial implementation questions that have to be considered when they are integrated into real-world systems.

**Forward privacy.** Much like recent works, our proposal is also forward-private. Our scheme, DLS, does not make use of ORAM simulation [18, 38] or public-key operations [5]. Furthermore, its query complexity grows only with the number of delete operations performed since the client executed (linear-time) rebuild protocol; as opposed to the Sophos [5] and Diana [6] constructions whose complexity grows with the total number of deletes ever performed. Our construction is also much more efficient as it only uses simple symmetric-key operations as opposed to other schemes which use public-key operations or constrained PRFs. Compared to the EKPE construction [16], which is not snapshot-secure and forward-private only for add operations, DLS offers better asymptotics if the scheme is used in scenarios in which the structure has more frequent updates and less frequent queries.

Surprisingly, we find that forward privacy does not imply breach resistance. While this may seem counter-intuitive, it reinforces the need of rigorous security analysis of STE constructions against different types of adversaries. In fact, several known forward-private constructions are not breach-resistant [5, 6, 16, 29]. For instance, in [16], an adversary that observes multiple

snapshots can tell if a search occurred or not because entries are deleted as part of the search operation. It can also learn the exact size of the result set if snapshots were taken before and after a search operation. This leakage is already more than just the size of the data structure as required by breach resistance. For the Sophos [5] and Diana<sub>del</sub> [6] constructions, a snapshot adversary can learn if or when a delete occurs if deletes are not handled in the same structure.

**Dual security.** After formally defining snapshot security, we construct, as far as we know, the first provably-secure breach-resistant encrypted multi-map. In addition, our construction is also forward-private which means that it can be used to protect against both persistent and snapshot adversaries. We refer to such constructions as being *dual-secure*. Our scheme, DLS, achieves both notions of security, and we achieve our desirable asymptotic query complexity by an explicit rebuild protocol that must be executed by the client at certain times. We show, however, that the protocol can be de-amortized and we call this variant DLS<sup>d</sup>.<sup>3</sup>

To sum up, DLS features the best query, token size, client memory, update and storage complexity among all dual-secure schemes in literature.

**Experimental evaluation.** We implemented DLS<sup>d</sup> (the variant with de-amortized rebuilding) in Java and evaluated its performance on the Wikipedia dataset. Our experiments show that this construction is highly practical. We ran experiments on up to 83 million label/value pairs on an Amazon EC2 instance with 32 vCPUs and 60GB of memory.

DLS<sup>d</sup> is compact, producing EMMs of size 9.4GB for a 3.6GB folder composed of 554,059 files. It has a search overhead less than 1 microsecond per pair for selectivities spanning from 100 to 10,000 pairs. On the other hand, it has an update overhead of around 100 milliseconds per pair. This slowdown is mainly due to the de-amortized rebuilding protocol.

## 2 Preliminaries

**Notation.** The set of all binary strings of length  $n$  is denoted as  $\{0, 1\}^n$ , and the set of all finite binary strings as  $\{0, 1\}^*$ .  $[n]$  is the set of integers  $\{1, \dots, n\}$ , and  $2^{[n]}$

<sup>3</sup> While this can be of independent interest, DLS results in a very efficient write-only oblivious multi-map which can replace several existing constructions, e.g., the ones used in the hidden volume scheme HIVE [4] or in the oblivious file backup system ObliviSync [2].

is the corresponding power set. We write  $x \leftarrow \chi$  to represent an element  $x$  being sampled from a distribution  $\chi$ , and  $x \stackrel{\$}{\leftarrow} X$  to represent an element  $x$  being sampled uniformly at random from a set  $X$ . The output  $x$  of an algorithm  $\mathcal{A}$  is denoted by  $x \leftarrow \mathcal{A}$ . Given a sequence  $\mathbf{o}$  of  $n$  elements, we refer to its  $i$ th element as  $\mathbf{o}_i$  or  $\mathbf{o}[i]$ . If  $T$  is a set then  $\#T$  refers to its cardinality. Given strings  $x$  and  $y$ , we refer to their concatenation as either  $\langle x, y \rangle$  or  $x\|y$ .

**Basic structures.** We make use of several basic data types including dictionaries and multi-maps which we recall here. A dictionary  $\text{DX}$  of capacity  $n$  is a collection of  $n$  label/value pairs  $\{(\ell_i, v_i)\}_{i \leq n}$  and supports **Get** and **Put** operations. We write  $v_i = \text{DX}[\ell_i]$  to denote getting the value associated with label  $\ell_i$  and  $\text{DX}[\ell_i] = v_i$  to denote the operation of associating the value  $v_i$  in  $\text{DX}$  with label  $\ell_i$ . We denote by  $\mathbb{L}_{\text{DX}}$  the set of labels stored in  $\text{DX}$  and by  $\#\text{DX}$  the volume of  $\text{DX}$  which is the number of label/value pairs it holds which is  $n = \#\mathbb{L}_{\text{DX}}$ . A multi-map  $\text{MM}$  with capacity  $n$  is a collection of  $n$  label/tuple pairs  $\{(\ell_i, \mathbf{v}_i)\}_{i \leq n}$  that supports **Get** and **Put** operations. Similarly to dictionaries, we write  $\mathbf{v}_i = \text{MM}[\ell_i]$  to denote getting the tuple associated with label  $\ell_i$  and  $\text{MM}[\ell_i] = \mathbf{v}_i$  to denote operation of associating the tuple  $\mathbf{v}_i$  to label  $\ell_i$ . We denote by  $\mathbb{L}_{\text{MM}}$  the set of labels stored in  $\text{MM}$  and by  $\#\text{MM}$  the volume of  $\text{MM}$  which is  $\sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$ . Multi-maps are the abstract data type instantiated by an inverted index. In the encrypted search literature multi-maps are sometimes referred to as indexes, databases or tuple-sets (T-sets) [8, 10]. Given a sequence of operations  $\mathbf{op} = (\text{op}_1, \dots, \text{op}_n)$ , each of which can be a query, an addition or a deletion, we denote by  $\text{add}(\mathbf{op})$  the subsequence of addition operations and by  $\text{del}(\mathbf{op})$  the subsequence of delete operations. We further define the subsequence of update operations as  $\text{up}(\mathbf{op}) = \text{add}(\mathbf{op}) + \text{del}(\mathbf{op})$ .

**Basic cryptographic primitives.** We make use of encryption schemes that are random-ciphertext-secure against chosen-plaintext attacks (RCPA),<sup>4</sup> and pseudo-random functions (PRF). For definitional details, we refer the readers to [10, 26].

### 3 Definitions

Structured encryption schemes [12] encrypt data structures in such a way that they can support operations

<sup>4</sup> RCPA-secure encryption can be instantiated practically using either the standard PRF-based private-key encryption scheme or, e.g., AES in counter mode.

on encrypted data. With encrypted data structures, we can distinguish between different types of operations. This includes non-interactive and interactive operations where the former require only a single message and the latter require several rounds. We can also distinguish between response-revealing and response-hiding operations, where the former reveal the answer to the query and the latter do not.

STE schemes are used as follows. During a setup phase, the client constructs an encrypted structure  $\text{EDS}$  from a plaintext structure  $\text{DS}$  under a key  $K$ . If the scheme is stateful, this setup procedure also outputs a state  $st$ . The client then sends the encrypted structure  $\text{EDS}$  to an untrusted server and keeps the state  $st$  and key  $K$  private. The client can then query  $\text{EDS}$  using the supported operations. If the operation is non-interactive, the client sends to the server a token  $\text{tk}$  constructed with its key  $K$ , state  $st$  and query  $q$ . The server then uses the token  $\text{tk}$  to query the encrypted structure  $\text{EDS}$ . If the operation is interactive, the client and server execute a two-party protocol where the former inputs  $K$ ,  $st$  and  $q$  and the latter inputs  $\text{EDS}$ .

#### Self-adjusting encrypted structures.

A data structure is self-adjusting if it re-arranges itself after being queried or updated. This is usually done to maintain correctness, consistency or to improve efficiency. We provide below the syntax of a self-adjustable STE scheme. Note that, here, the update operation is interactive which is not a requirement.

**Definition 3.1** (Self-adjusting STE). *A response-hiding dynamic structured encryption scheme  $\Sigma = (\text{Setup}, \text{Token}, \text{Query}, \text{Update}, \text{Rslv})$  with non-interactive queries and interactive updates consists of four polynomial-time algorithms and one two-party protocol between the client and server that work as follows:*

- $(K, st, \text{EDS}) \leftarrow \text{Setup}(1^k, \text{DS})$ : is a probabilistic algorithm that takes as input a security parameter  $1^k$  and a structure  $\text{DS}$ . It outputs a secret key  $K$ , a state  $st$  and an encrypted structure  $\text{EDS}$ .
- $(st', \text{tk}) \leftarrow \text{Token}(K, st, q)$ : is a (possibly) probabilistic algorithm that takes as input a secret key  $K$ , a state  $st$  and a query  $q$ . It outputs a new state  $st'$  and a query token  $\text{tk}$ .
- $(ct, \text{EDS}') \leftarrow \text{Query}(\text{EDS}, \text{tk})$ : is a (possibly) probabilistic algorithm that takes as input an encrypted structure  $\text{EDS}$  and a token  $\text{tk}$ . It outputs a message  $ct$  and an (possibly) updated encrypted structure  $\text{EDS}'$ .
- $(st', \text{EDS}') \leftarrow \text{Update}_{C,S}((K, st, u), \text{EDS})$ : is a two-party protocol between the client and the server. It takes as input from the client a key  $K$ , a state  $st$



and an update  $u$  and as input from the server an encrypted structure EDS. It outputs to the client an updated state  $st'$  and to the server a new encrypted structure EDS'.

- $r \leftarrow \text{Rslv}(K, ct)$ : is a deterministic algorithm that takes as input a secret key  $K$  and a message  $ct$ . It outputs a response  $r$ .

The syntax of response-revealing Query operation can be recovered by having it output the response  $r$  directly and omitting the Rslv algorithm.

**Rebuildable encrypted structures.** We say that a data structure is rebuildable if it supports a rebuild operation that reconstructs it. Rebuild operations are typically used to improve query efficiency or storage overhead after a sequence of operations have been performed. Whereas self-adjusting structures re-arrange themselves as part of a query or update operation, rebuildable structures support an explicit rebuild operation that is typically invoked after a sequence of operations. We provide below the syntax of a rebuildable STE scheme. Note that, here, the rebuild operation is interactive.

**Definition 3.2** (Rebuildable STE). *A response-hiding dynamic structured encryption scheme  $\Sigma = (\text{Setup}, \text{Token}, \text{Query}, \text{UToken}, \text{Update}, \text{Rebuild}, \text{Rslv})$  with non-interactive queries and interactive rebuilds consists of seven polynomial-time algorithms and one two-party protocol between the client and server. The Setup, Token and Rslv are as in Definition 3.1 while Query, UToken, Update and Rebuild work as follows:*

- $ct \leftarrow \text{Query}(\text{EDS}, tk)$ : is a (possibly) probabilistic algorithm that takes as input an encrypted structure EDS and a token  $tk$ . It outputs a message  $ct$ .
- $(st', utk) \leftarrow \text{UToken}(k, st, u)$ : is a (possibly) probabilistic algorithm that takes as input a secret key  $K$ , a state  $st$  and an update  $u$ . It outputs a new state  $st'$  and an update token  $utk$ .
- $\text{EDS}' \leftarrow \text{Update}(\text{EDS}, utk)$ : is a (possibly) probabilistic algorithm that takes as input an encrypted structure EDS and an update token  $utk$ . It outputs an updated encrypted structure EDS'.
- $(st', \text{EDS}') \leftarrow \text{Rebuild}_{\text{C,S}}((K, st), \text{EDS})$ : is a two-party protocol between the client and the server. It takes as input from the client a key  $K$  and a state  $st$  and as input from the server an encrypted structure EDS. It outputs to the client an updated state  $st'$  and to the server a new encrypted structure EDS'.

In the subsequent parts of this section, we will only focus on *rebuildable* STE. The proposed definitions can be naturally extended to *self-adjusting* STE schemes.

### 3.1 Security Against a Persistent Adversary

The standard notion of security for STE guarantees that: (1) an encrypted structure reveals no information about its underlying structure beyond the setup leakage  $\mathcal{L}_S$ ; (2) that the query algorithm reveals no information about the structure and the queries beyond the query leakage  $\mathcal{L}_Q$ ; and that (3) the update algorithm reveals no information about the structure and the update beyond the update leakage  $\mathcal{L}_U$ . Naturally, if the scheme has a rebuild protocol then we require that it reveals no information about the underlying structure beyond the rebuild leakage  $\mathcal{L}_R$ .

If this holds for non-adaptively chosen operations then the scheme is said to be non-adaptively secure. If, on the other hand, the operations can be chosen adaptively, the scheme is said to be adaptively-secure. This notion of security was first formalized by Curtmola et al. in the context of searchable encryption [13] and later generalized to structured encryption in [12].

**Definition 3.3** (Adaptive security [12, 13]). *Let  $\Sigma = (\text{Setup}, \text{Token}, \text{Query}, \text{UToken}, \text{Update}, \text{Rebuild}, \text{Rslv})$  be a rebuildable STE scheme and consider the following probabilistic experiments where  $\mathcal{A}$  is a stateful adversary,  $\mathcal{S}$  is a stateful simulator,  $\mathcal{L}_S$ ,  $\mathcal{L}_Q$ ,  $\mathcal{L}_U$  and  $\mathcal{L}_R$  are leakage profiles and  $z \in \{0, 1\}^*$ :*

**Real $_{\Sigma, \mathcal{A}}(k)$** : given  $z$  the adversary  $\mathcal{A}$  outputs a structure DS and receives EDS from the challenger, where  $(K, st, \text{EDS}) \leftarrow \text{Setup}(1^k, \text{DS})$ . The adversary then adaptively chooses a polynomial number of operations  $\text{op}_1, \dots, \text{op}_m$  such that  $\text{op}_i$  is either a query  $q_i$ , an update  $u_i$ , or a rebuild  $r_i$ . For all  $i \in [m]$ , the adversary receives  $tk_i \leftarrow \text{Token}(K, st, q_i)$  if  $\text{op}_i = q_i$  or  $utk_i \leftarrow \text{UToken}(K, st, u_i)$  if  $\text{op}_i = u_i$ . If  $\text{op}_i = r_i$ , then the client and server execute  $\text{Rebuild}_{\text{C,S}}((K, st), \text{EDS})$ . Finally,  $\mathcal{A}$  outputs a bit  $b$  that is output by the experiment.

**Ideal $_{\Sigma, \mathcal{A}, \mathcal{S}}(k)$** : given  $z$  the adversary  $\mathcal{A}$  generates a structure DS which it sends to the challenger. Given  $z$  and leakage  $\mathcal{L}_S(\text{DS})$  from the challenger, the simulator  $\mathcal{S}$  returns an encrypted structure EDS to  $\mathcal{A}$ . The adversary then adaptively chooses a polynomial number of operations  $\text{op}_1, \dots, \text{op}_m$  such that  $\text{op}_i$  is a query  $q_i$ , an update  $u_i$  or a rebuild  $r_i$ . For all

$i \in [m]$ , if the simulator receives either query leakage  $\mathcal{L}_Q(\text{DS}, q_i)$  or update leakage  $\mathcal{L}_U(\text{DS}, u_i)$ . In the former case, it returns a query token  $\text{tk}_i$  to  $\mathcal{A}$  and in the latter it returns an update token  $\text{utk}_i$  to  $\mathcal{A}$ . If  $\text{op}_i = r_i$ ,  $\mathcal{S}(\mathcal{L}_R(\text{DS}))$  and  $\mathcal{A}$  execute **Rebuild**. Finally,  $\mathcal{A}$  outputs a bit  $b$  that is output by the experiment.

We say that  $\Sigma$  is adaptively  $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U, \mathcal{L}_R)$ -secure if there exists a PPT simulator  $\mathcal{S}$  such that for all PPT adversaries  $\mathcal{A}$ , for all  $z \in \{0, 1\}^*$ , the following expression is negligible in  $k$ :

$$\left| \Pr [\mathbf{Real}_{\Sigma, \mathcal{A}}(k) = 1] - \Pr [\mathbf{Ideal}_{\Sigma, \mathcal{A}, \mathcal{S}}(k) = 1] \right|.$$

**Forward privacy.** An important property for dynamic STE schemes is *forward privacy* which was introduced in [38] to address some of the limitations of dynamic SSE constructions at the time. The informal requirement described in [38] was that forward-private SSE schemes should not reveal if the file in a file update operation (i.e., a file add or delete) has keywords that were searched for in the past. This was formalized in [5] for **AddFile** operations as

$$\mathcal{L}_{\text{AF}}(\text{MM}, f) = (\#f),$$

where  $f \subseteq \mathbb{W}$  is a file.

### 3.2 Security Against a Snapshot Adversary

In the standard notions of security for STE, the adversary is assumed to be the server itself. As such, we seek security guarantees against an adversary that sees not only the encrypted structure but all the search and update operations as well. In many real-world scenarios, however, we are concerned with a weaker adversary that, only periodically, gets access to the encrypted structure and, in particular, does not get to see any query or update operations. This adversarial model captures, for example, data breaches, malicious employees and device theft. Such an adversary is called a *snapshot adversary*.

We propose a new security definition for STE against snapshot adversaries. In our definition 3.4, the adversary has access to multiple snapshots each of which is interspersed with a batch of operations. Intuitively, we require that the encrypted structure reveals no information about the underlying structure and the sequence of operations executed prior to the multiple snapshots, beyond some snapshot leakage  $\mathcal{L}_{\text{SN}}$ .

Surprisingly, we demonstrate that for a particular class of snapshot leakage, (multiple) snapshot secure

STE schemes imply insertion independence, a variant of history independent data structures [35], and can also provide a *write-only oblivious* structure [4].

**Definition 3.4** (Snapshot security). *Let  $\Sigma = (\text{Setup}, \text{Token}, \text{Query}, \text{UToken}, \text{Update}, \text{Rebuild}, \text{Rslv})$  be a rebuildable STE scheme and consider the following probabilistic experiments where  $\mathcal{A}$  is a stateful adversary,  $\mathcal{S}$  is a stateful simulator,  $\mathcal{L}_{\text{SN}}$  is a stateful leakage function,  $z \in \{0, 1\}^*$ , and  $m \geq 1$ :*

$\mathbf{Real}_{\Sigma, \mathcal{A}}^{\text{ms}}(k, m)$ :

1. given  $z$  the adversary  $\mathcal{A}$  outputs a structure  $\text{DS}_0$ ;
2. the challenger computes  $(K, st, \text{EDS}_0) \leftarrow \text{Setup}(1^k, \text{DS}_0)$ ;
3. the adversary  $\mathcal{A}(\text{EDS}_0)$  outputs a sequence of operations  $\text{op}_1 = (\text{op}_{1,1}, \dots, \text{op}_{1,\ell})$  where  $\ell = \text{poly}(k)$ ;
4. For all  $i \in [m]$ ,
  - (a) the challenger applies all the operations in  $\text{op}_i$  to  $\text{EDS}_{i-1}$  by computing and applying the appropriate tokens. This results in  $\text{EDS}_i$ ;
  - (b) the adversary  $\mathcal{A}(\text{EDS}_i)$  outputs a sequence of operations  $\text{op}_{i+1} = (\text{op}_{i+1,1}, \dots, \text{op}_{i+1,\ell})$  where  $\ell = \text{poly}(k)$ ;
5. Finally,  $\mathcal{A}$  outputs a bit  $b$  that is returned by the experiment.

$\mathbf{Ideal}_{\Sigma, \mathcal{A}, \mathcal{S}}^{\text{ms}}(k, m)$ :

1. given  $z$  the adversary  $\mathcal{A}$  outputs a structure  $\text{DS}_0$ ;
2. the simulator  $\mathcal{S}(z, \mathcal{L}_{\text{SN}}(\text{DS}_0, \perp))$  simulates  $\text{EDS}_0$ ;
3. the adversary  $\mathcal{A}(\text{EDS}_0)$  outputs a sequence of operations  $\text{op}_1 = (\text{op}_{1,1}, \dots, \text{op}_{1,\ell})$ ;
4. For all  $i \in [m]$ ,
  - (a) the challenger applies all the operations in  $\text{op}_i$  to  $\text{DS}_{i-1}$ , resulting in  $\text{DS}_i$ ;
  - (b) the simulator  $\mathcal{S}(\mathcal{L}_{\text{SN}}(\text{DS}_i, \text{op}_i))$  simulates  $\text{EDS}_i$ ;
  - (c) the adversary  $\mathcal{A}(\text{EDS}_i)$  outputs a sequence of operations  $\text{op}_{i+1} = (\text{op}_{i+1,1}, \dots, \text{op}_{i+1,\ell})$ ;
5. Finally,  $\mathcal{A}$  outputs a bit  $b$  that is output by the experiment.

We say that  $\Sigma$  is  $(m, \mathcal{L}_{\text{SN}})$ -snapshot secure if there exists a PPT simulator  $\mathcal{S}$  such that for all PPT adversaries  $\mathcal{A}$  and for all  $z \in \{0, 1\}^*$ , the following expression is negligible in  $k$ :

$$\left| \Pr [\mathbf{Real}_{\Sigma, \mathcal{A}}^{\text{ms}}(k, m) = 1] - \Pr [\mathbf{Ideal}_{\Sigma, \mathcal{A}, \mathcal{S}}^{\text{ms}}(k, m) = 1] \right|.$$

**Breach resistance.** Ideally, the snapshot leakage of a scheme should be as small as possible. With this in mind, we deem that an encrypted structure should be regarded as *breach-resistant* if its snapshot leakage is at most the size of the plaintext structure at the time the snapshot is taken.

**Definition 3.5** (Breach resistance). *Let  $\Sigma$  be an  $\mathcal{L}_{\text{SN}}$ -snapshot secure non-interactive dynamic STE scheme. We say that  $\Sigma$  is breach-resistant if*

$$\mathcal{L}_{\text{SN}}(\text{DS}, \text{op}_1, \dots, \text{op}_i) = \#\text{DS}_i,$$

where  $\text{DS}_i$  is the structure that results from applying  $\text{op}_1, \dots, \text{op}_i$  to  $\text{DS}$  and  $\#\text{DS}_i$  refers to its volume in the sense of the total number of “items” it stores. Note that the volume of a structure depends on its type.

For the remainder of this work, we focus on designing a multi-map encryption scheme that is secure against both persistent and snapshot adversaries. Specifically, we require that the scheme be forward-private against a persistent adversary and breach-resistant against a snapshot adversary. We refer to schemes that meet these two properties as dual-secure.

**Definition 3.6** (Dual security). *Let  $\Sigma$  be a dynamic rebuildable structured encryption scheme with leakage profiles  $\Lambda_{\text{Per}} = (\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U, \mathcal{L}_R)$  and  $\Lambda_{\text{Sna}} = \mathcal{L}_{\text{SN}}$ . We say that  $\Sigma$  is dual-secure if it is forward-private and breach-resistant.*

**Implications of breach resistance.** We demonstrate that breach resistance implies some interesting and previously studied notions of security including a weaker notion of history independence [35] and write-only obliviousness [4]. Due to lack of space, the details appear in the full version of this paper.

## 4 DLS: A Dual-Secure Multi-Map Encryption Scheme

We now describe our construction, DLS, which is a *dual-secure* rebuildable multi-map encryption scheme. Unlike SPS [38] and Sophos [5], DLS does not make use of ORAM-like techniques or public-key operations. In addition, while Sophos and Diana [6] have query complexity that is linear in the number of delete operations ever executed, DLS’s query complexity is linear only in the number of delete operations executed since the last rebuild operation. The rebuild protocol is linear in the size of the multi-map and can be de-amortized while maintaining snapshot security. In the following, we present a description of the construction followed by a detailed analysis.

### 4.1 Detailed Description

DLS makes use of a pseudo-random function  $F$  and of a private-key encryption scheme SKE. The details of the scheme are provided in Figs. 1 and 2. At a high-level, it works as follows.

**Setup.** The Setup algorithm takes as input a security parameter  $k$  and a multi-map  $\text{MM}$ . It instantiates two dictionaries: an old dictionary,  $\text{DX}_o$ , and a new dictionary,  $\text{DX}_n$ . It also instantiates the state composed of a global version  $\text{version}_g$ , a searched label set  $\mathbb{S}_e$ , and two state dictionaries, an old  $\text{DX}_o^{\text{st}}$  and a new  $\text{DX}_n^{\text{st}}$ . The global version  $\text{version}_g$  will record the number of rebuilds ever performed in the future. The set  $\mathbb{S}_e$  is a temporary one that will keep track of all searched labels within a single rebuild epoch. Dictionaries  $\text{DX}_o^{\text{st}}$  and  $\text{DX}_n^{\text{st}}$  will map a label to both its counter and version. The counter of a label  $\ell$  in  $\text{DX}_o^{\text{st}}$  and  $\text{DX}_n^{\text{st}}$  is the number of all values that have been added to  $\text{DX}_o$  and  $\text{DX}_n$  in the previous and current epoch, respectively. In other words,  $\text{DX}_o^{\text{st}}$  contains only labels that existed within the previous version, i.e., the previous rebuild epoch.

Setup outputs the old and new dictionaries as its encrypted multi-map. The old dictionary  $\text{DX}_o$  is populated as follows. In order to store  $(\ell \parallel \text{version}_g, \text{MM}[\ell])$  in  $\text{DX}_o$ , it stores the pairs  $(\ell \parallel i \parallel \text{version}_g, v_i)$ , for all  $v_i \in \text{MM}[\ell]$  and  $i \in [\#\text{MM}[\ell]]$ , where  $\text{version}_g$  is the current rebuild epoch (the rebuild epoch at setup time is initialized to 1). To store the pair in an encrypted way, a PRF evaluation is performed on the concatenation of the label  $\ell$ , its counter  $\text{count}$  and the rebuild epoch  $\text{version}_g$ . The corresponding value in  $\text{MM}[\ell]$  is first concatenated with the string  $\text{edit}^+$  and then simply encrypted. The new dictionary  $\text{DX}_n$  is only instantiated and remains empty. The output of the setup algorithm includes the encrypted structures  $(\text{DX}_o, \text{DX}_n)$ , the keys as well as the state.

**Search token.** The Token algorithm takes as input a key, a state and a label. First, it fetches from both the old and new state dictionaries the corresponding counters and rebuilding versions. Recall the old and new counters,  $\text{count}_o$  and  $\text{count}_n$ , count the number of times the label has been added, deleted from the old and new dictionaries, respectively. Based on the counters, whether old or new, it creates two sub-token vectors such that  $\text{otk} = (\text{otk}_1, \dots, \text{otk}_{\text{count}_o})$  and  $\text{ntk} = (\text{ntk}_1, \dots, \text{ntk}_{\text{count}_n})$ . The old subtokens  $\text{otk}_i$ , for  $i \in [\text{count}_o]$ , will allow the server to query the old dictionary  $\text{DX}_o$ , while the new tokens  $\text{ntk}_i$ , for  $i \in [\text{count}_n]$ , will allow the server to query the new dictionary  $\text{DX}_n$ . Finally, it updates the state by adding  $\ell$  to  $\mathbb{S}_e$  as it is

now searched. The output includes the state and the token  $\text{tk} = (\text{otk}, \text{ntk})$ .

**Query.** The Query algorithm takes as input the token and the encrypted data structure. The token is divided into two sub-token vectors,  $\text{otk}$  and  $\text{ntk}$ . Each sub-token in  $\text{otk}$  corresponds to a label in the old dictionary  $\text{DX}_o$  from which the server fetches the value and inserts in the result set,  $\text{Result}$ . The server performs the same operations for every sub-token in  $\text{ntk}$ , but on the new dictionary  $\text{DX}_n$ . The server finally outputs the result set  $\text{Result}$ . Note that the set  $\text{Result}$  does not exactly equal  $\text{MM}[\ell]$  as the client might have deleted several pairs both in the old or new dictionaries. The client, based on the meta-information included with the decrypted values, i.e.,  $\text{edit}^+$  or  $\text{edit}^-$ , can easily compute the correct  $\text{MM}[\ell]$ , where  $\text{edit}^+$  and  $\text{edit}^-$  denote added and deleted values, respectively.

**Update token.** The UToken algorithm takes as input a key, a state, and an update consisting of the type of operation  $\text{op}$ , a label  $\ell$ , and its value  $\mathbf{v}$ . It first computes  $\text{tk}_1$  which is a PRF evaluation on the concatenation of the label, its counter and the current rebuilding version. The counter represents the number of times the label  $\ell$  has been added to the new dictionary  $\text{DX}_n$  throughout the current rebuilding version, and is fetched from the new state dictionary  $\text{DX}_n^{\text{st}}$ , given the label  $\ell$ . The counter is then updated accordingly. The value  $v$  is concatenated to the operation  $\text{op}$  before being encrypted, and this represents the second part of the token,  $\text{tk}_2$ . The output of the algorithm includes  $\text{utk} = (\text{tk}_1, \text{tk}_2)$  and an updated state.

**Update.** The Update algorithm takes as input the update token and the encrypted data structure. The server will simply update the new dictionary  $\text{DX}_n$  by adding the update token,  $\text{utk} = (\text{tk}_1, \text{tk}_2)$ , to it. The output of the Update algorithm consists of the updated encrypted structure.

**Rebuilding.** The Rebuild algorithm is a two-party protocol between the client and the server. The client's input is a key and a state, while the server's input is the encrypted data structure. The goal of the rebuild operation is to merge the old dictionary into the new one. The new one will then become the old at the end of the rebuild. The rebuild differentiates between two types of labels: (1) labels that have been searched for in  $\text{DX}_o$ , and (2) labels that have not. For each of the labels in  $\mathbb{S}_e$ , the client fetches all the values corresponding to  $\ell \in \mathbb{S}_e$ , removes all values that have to be deleted and then insert the remaining ones into the new dictionary  $\text{DX}_n$ . Inserting these values into  $\text{DX}_n$  follows a similar process to the one in UToken and Update algo-

rithms. For the remaining labels, i.e., labels that have never been searched for in  $\text{DX}_o$ , the client picks a random label, fetches a value with the largest counter and inserts it into the new dictionary  $\text{DX}_n$ . The client updates the state by decreasing the counter of the selected label and removes it whenever it equals 1. Once all labels have been reinserted, both the old state dictionary  $\text{DX}_o^{\text{st}}$  and old dictionary  $\text{DX}_o$  are deleted, the set  $\mathbb{S}_e$  is reinitialized, the new state dictionary  $\text{DX}_n^{\text{st}}$  and new dictionary  $\text{DX}_n$  becomes the old state dictionary  $\text{DX}_o^{\text{st}}$  and old dictionary  $\text{DX}_o$ , and the epoch is incremented. The output of the Rebuild is an updated state for the client and an updated encrypted structure for the server.

In order to achieve snapshot security, the entire transcript of the Rebuild protocol, including its internal state, has to be kept secret. That is, a snapshot adversary must not get a snapshot of the encrypted structures while the Rebuild protocol is executing. This is mainly due to the fact that while rebuilding searched for labels, a snapshot adversary will get to know the response size of the searched for labels; which is clearly at odds with our security goals that consist of only disclosing the size of the data structure.

We show how to lift this constraint in Section 5.

**Efficiency.** The query complexity of DLS is

$$O\left(\#\text{MM}[\ell] + \text{del}(\ell, e)\right),$$

where  $\text{del}(\ell, e)$  is the number of deletes for  $\ell$  since epoch  $e$  when  $\ell$  was most recently searched for (i.e., since the last rebuild during which  $\ell$  was a searched for label). As a point of comparison, the Sophos construction of Bost [5] and the Diana construction of Bost et al. [6] have query complexity

$$O\left(\#\text{MM}[\ell] + \text{del}(\ell, 0)\right),$$

where  $\text{del}(\ell, 0)$  denotes the number of deletes for  $\ell$  since the structure was setup. The storage complexity of DLS is

$$O\left(\sum_{\ell \in \mathbb{L}_{\text{MM}}} \left(\#\text{MM}[\ell] + \text{del}(\ell, e)\right)\right),$$

while recent constructions [5, 6] have storage complexity  $O(\sum_{\ell \in \mathbb{L}_{\text{MM}}} (\#\text{MM}[\ell] + \text{del}(\ell, 0)))$ . DLS has search and update tokens of size  $O(\#\text{MM}[\ell] + \text{del}(\ell, e))$  and  $O(\#\mathbf{v})$ , respectively, and its rebuild complexity is

$$O\left(\sum_{\ell \in \mathbb{L}_{\text{MM}}} \left(\#\text{MM}[\ell] + \text{del}(\ell, e) - \text{up}(\ell, c)\right)\right),$$

where  $\#\text{up}(\ell, c)$  is the number of updates for  $\ell$  since the last rebuild operation.



Let  $F$  be a pseudo-random function,  $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$  be a private-key encryption scheme. Consider the dynamic encrypted multi-map  $\text{DLS} = (\text{Setup}, \text{Token}, \text{UToken}, \text{Get}, \text{Put}, \text{Rebuild})$  defined as follows:

–  $\text{Setup}(1^k, \text{MM})$ :

1. sample  $K_1, K_2 \xleftarrow{\$} \{0, 1\}^k$ ;
2. initialize an empty set  $\mathbb{S}_e$  and four empty dictionaries  $\text{DX}_o^{\text{st}}, \text{DX}_n^{\text{st}}$ , and  $\text{DX}_o$  and  $\text{DX}_n$  with capacities  $\#\mathbb{L}_{\text{MM}}$  and  $\sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$  respectively;
3. for all  $i \in \#\text{MM}$ ,
  - (a) sample a pair  $(\ell, \text{MM}[\ell])$  from  $\text{MM}$  without replacement;
  - (b) set  $\text{count} = 1$  and  $\text{version}_g = 1$ ;
  - (c) for all  $v \in \text{MM}[\ell]$ ,
    - i. compute

$\text{label} := F_{K_1}(\ell \| \text{version}_g \| \text{count})$  and  $\text{value} := \text{Enc}(K_2, v \| \text{edit}^+)$ ;

- ii. set  $\text{DX}_o[\text{label}] := \text{value}$ ;
- iii. set  $\text{DX}_o^{\text{st}}[\ell] := (\text{version}_g, \text{count})$  and increment  $\text{count}$ ;
4. increment  $\text{version}_g$ ;
5. output  $(K, st, \text{EMM})$  where  $K = (K_1, K_2)$ ,  $st = (\text{version}_g, \mathbb{S}_e, \text{DX}_o^{\text{st}}, \text{DX}_n^{\text{st}})$  and  $\text{EMM} = (\text{DX}_o, \text{DX}_n)$ .

–  $\text{Token}(K, st, \ell)$ :

1. parse  $K = (K_1, K_2)$  and  $st = (\text{version}_g, \mathbb{S}_e, \text{DX}_o^{\text{st}}, \text{DX}_n^{\text{st}})$ ;
2. if  $\text{DX}_o^{\text{st}}[\ell] \neq \perp$ , set  $(\text{version}_n, \text{count}_n) := \text{DX}_n^{\text{st}}[\ell]$ , and if  $\text{DX}_o^{\text{st}}[\ell] \neq \perp$  set  $(\text{version}_o, \text{count}_o) := \text{DX}_o^{\text{st}}[\ell]$  and add  $\ell$  to  $\mathbb{S}_e$ ;
3. set

$\text{otk} = (\text{otk}_1, \dots, \text{otk}_{\text{count}_o})$  and  $\text{ntk} = (\text{ntk}_1, \dots, \text{ntk}_{\text{count}_n})$ ,

where for all  $i \in [\text{count}_o]$  and  $j \in [\text{count}_n]$

$\text{otk}_i := F_{K_1}(\ell \| \text{version}_o \| i)$  and  $\text{ntk}_j := F_{K_1}(\ell \| \text{version}_n \| j)$ ;

4. output the update state  $st = (\text{version}_g, \mathbb{S}_e, \text{DX}_o^{\text{st}}, \text{DX}_n^{\text{st}})$  and the token  $\text{tk} = (\text{otk}, \text{ntk})$ .

–  $\text{UToken}(K, st, (\text{op}, \ell, \mathbf{v}))$ :

1. parse  $K = (K_1, K_2)$  and  $st = (\text{version}_g, \mathbb{S}_e, \text{DX}_o^{\text{st}}, \text{DX}_n^{\text{st}})$ ;
2. if  $\ell \notin \mathbb{L}_{\text{MM}}$ ,
  - (a) set  $\text{DX}_n^{\text{st}}[\ell] := (\text{version}_g, \text{count})$  where  $\text{count} = 1$ ;
  3. otherwise,
    - (a) if  $\text{DX}_n^{\text{st}}[\ell] \neq \perp$ ,
      - i. set  $(\text{version}, \text{count}) := \text{DX}_n^{\text{st}}[\ell]$  and increment  $\text{count}$ ;
      - ii. set  $\text{DX}_n^{\text{st}}[\ell] := (\text{version}_g, \text{count})$ ;
    - (b) otherwise
      - i. set  $\text{DX}_n^{\text{st}}[\ell] := (\text{version}_g, \text{count})$  where  $\text{count} = 1$ ;
  4. for all  $v \in \mathbf{v}$ ,
    - (a) compute

$\text{tk}_{v,1} := F_{K_1}(\ell \| \text{version}_g \| \text{count})$  and  $\text{tk}_{v,2} := \text{Enc}(K_2, v \| \text{op})$ ;

- (b) increment  $\text{count}$  and set  $\text{DX}_n^{\text{st}}[\ell] := (\text{version}_g, \text{count})$ ;
5. output the updated state  $st = (\text{version}_g, \mathbb{S}_e, \text{DX}_o^{\text{st}}, \text{DX}_n^{\text{st}})$  and  $\text{utk} = (\text{tk}_{v,1}, \text{tk}_{v,2})_{v \in \mathbf{v}}$ .

–  $\text{Get}(\text{tk}, \text{EMM})$ :

1. parse  $\text{EMM} = (\text{DX}_o, \text{DX}_n)$  and  $\text{tk} = (\text{otk}, \text{ntk})$  s.t.

$\text{otk} = (\text{otk}_1, \dots, \text{otk}_{\text{count}_o})$  and  $\text{ntk} = (\text{ntk}_1, \dots, \text{ntk}_{\text{count}_n})$ ;

2. instantiate an empty set  $\text{Result}$ ;
3. add  $\text{DX}_o[\text{otk}_i]$  and  $\text{DX}_n[\text{ntk}_j]$  for all  $i \in [\text{count}_o]$  and  $j \in [\text{count}_n]$  to  $\text{Result}$ ;
4. output  $\text{Result}$ .

–  $\text{Put}(\text{utk}, \text{EMM})$ :

1. parse  $\text{EMM} = (\text{DX}_o, \text{DX}_n)$  and  $\text{utk} = (\text{tk}_{v,1}, \text{tk}_{v,2})_{v \in \mathbf{v}}$ ;
2. for all  $v \in \mathbf{v}$ , set  $\text{DX}_n[\text{tk}_{v,1}] := \text{tk}_{v,2}$ ;
3. output  $\text{EMM}$ .

–  $\text{Rebuild}_{\text{C,S}}\left(\left(K, st\right), \text{EMM}\right)$ :

1.  $\text{C}$  parses  $K = (K_1, K_2)$ ,  $st = (\text{version}_g, \mathbb{S}_e, \text{DX}_o^{\text{st}}, \text{DX}_n^{\text{st}})$  and  $\text{EMM} = (\text{DX}_o, \text{DX}_n)$ ;
2. for all  $\ell \in \mathbb{S}_e$  such that  $\text{DX}_o^{\text{st}}[\ell] \neq \perp$ ,
  - (a)  $\text{C}$  sets  $(\text{version}_o, \text{count}_o) := \text{DX}_o^{\text{st}}[\ell]$  and removes  $\ell$  from  $\text{DX}_o^{\text{st}}$ ;
  - (b)  $\text{C}$  computes and sends to  $\text{S}$ ,

$\text{tk} = (\text{otk}_1, \dots, \text{otk}_{\text{count}_o})$ ,

where  $\text{otk}_i = F_{K_1}(\ell \| \text{version}_o \| i)$ ;

- (c)  $\text{S}$  computes and sends to  $\text{C}$ ,

$\text{Result} = (\text{ct}_1, \dots, \text{ct}_{\text{count}_o})$ ;

where  $\text{ct}_i = \text{DX}_o[\text{otk}_i]$  for all  $i \in [\text{count}_o]$ ;

- (d)  $\text{C}$  computes  $V = \text{Result}^+ \setminus \text{Result}^-$ , where

$\text{Result}^+ = \{v : \forall \text{ct} \in \text{Result}, v \| \text{edit}^+ = \text{Dec}(K_2, \text{ct})\}$

$\text{Result}^- = \{v : \forall \text{ct} \in \text{Result}, v \| \text{edit}^- = \text{Dec}(K_2, \text{ct})\}$

- (e) if  $\text{DX}_n^{\text{st}}[\ell] \neq \perp$ ,  $\text{C}$  sets  $(\text{version}_n, \text{count}_n) := \text{DX}_n^{\text{st}}[\ell]$ , otherwise sets  $\text{count}_n = 1$ ;
- (f) for all  $v \in V$ ,
  - i.  $\text{C}$  computes and sends,

$\text{tk}_1 := F_{K_1}(\ell \| \text{version}_g \| \text{count}_n)$   $\text{tk}_2 := \text{Enc}(K_2, v \| \text{edit}^+)$ ;

ii.  $\text{S}$  computes  $\text{DX}_n[\text{tk}_1] := \text{tk}_2$ ;

iii.  $\text{C}$  increments  $\text{count}_n$ ;

3. while  $\#\text{DX}_o^{\text{st}} > 0$ ,

(a)  $\text{C}$  picks  $\ell$  at random

(b)  $\text{C}$  sets  $(\text{version}_o, \text{count}_o) := \text{DX}_o^{\text{st}}[\ell]$  and  $\text{DX}_o^{\text{st}}[\ell] := (\text{version}_o, \text{count}_o - 1)$ ;

(c) if  $\text{count}_o - 1 < 1$ ,  $\text{C}$  then removes  $\ell$  from  $\text{DX}_o^{\text{st}}$ ;

(d)  $\text{C}$  computes and sends to  $\text{S}$   $\text{otk} = F(K_1, \ell \| \text{version}_o \| \text{count}_o)$ ;

(e)  $\text{S}$  computes and sends to  $\text{C}$   $\text{ct} = \text{DX}_o[\text{otk}]$ ;

- (f) if  $\text{DX}_n^{\text{st}}[\ell] \neq \perp$ ,
  - i.  $\text{C}$  sets  $(\text{version}_n, \text{count}_n) := \text{DX}_n^{\text{st}}[\ell]$ ;
  - ii.  $\text{C}$  computes and sends to  $\text{S}$ ,

$\text{tk}_1 := F_{K_1}(\ell \| \text{version}_n \| \text{count}_n)$   $\text{tk}_2 := \text{Enc}(K_2, \text{Dec}(K_2, \text{ct}))$ ;

iii.  $\text{S}$  computes  $\text{DX}_n[\text{tk}_1] := \text{tk}_2$ ;

iv.  $\text{C}$  sets  $\text{DX}_n^{\text{st}}[\ell] := (\text{version}_n, \text{count}_n + 1)$ ;

(g) otherwise if  $\text{DX}_n^{\text{st}}[\ell] = \perp$ , then

i.  $\text{C}$  sets  $\text{DX}_n^{\text{st}}[\ell] := (\text{version}_g, 1)$ ;

ii.  $\text{C}$  computes and sends to  $\text{S}$ ,

$\text{tk}_1 := F_{K_1}(\ell \| \text{version}_g \| 1)$   $\text{tk}_2 := \text{Enc}(K_2, \text{Dec}(K_2, \text{ct}))$

iii.  $\text{S}$  computes  $\text{DX}_n[\text{tk}_1] := \text{tk}_2$ ;

4.  $\text{C}$  sets  $\text{DX}_o^{\text{st}} := \text{DX}_o^{\text{st}}$  and initializes an empty dictionary  $\text{DX}_n^{\text{st}}$  with capacity  $2 \cdot \#\mathbb{L}_{\text{MM}}$ ;

5.  $\text{S}$  sets  $\text{DX}_o := \text{DX}_n$  and initializes an empty dictionary  $\text{DX}_n$  with capacity  $2 \cdot \#\mathbb{L}_{\text{MM}}$ ;

6.  $\text{C}$  empties  $\mathbb{S}_e$  and increments  $\text{version}_g$ ;

7.  $\text{C}$  outputs the updated states  $st = (\text{version}_g, \mathbb{S}_e, \text{DX}_o^{\text{st}}, \text{DX}_n^{\text{st}})$  and  $\text{S}$  the updated encrypted multi-map  $\text{EMM} = (\text{DX}_o, \text{DX}_n)$ .

Fig. 2. DLS (Part 2).

Fig. 1. DLS (Part 1).

DLS's amortized rebuild complexity is

$$O\left(\max_{e \in \mathbb{N}, \ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}_e[\ell]\right),$$

where  $\text{MM}_e[\ell]$  denotes the tuple associated to label  $\ell$  at epoch  $e$ . This is because for all  $e \in \mathbb{N}$ , we have

$$\text{del}(\ell, e) = O\left(\max_{e \in \mathbb{N}, \ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}_e[\ell]\right),$$

from which it follows that the total rebuild cost is

$$O\left(\#\mathbb{L}_{\text{MM}} \cdot \max_{e \in \mathbb{N}, \ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}_e[\ell]\right).$$

If rebuilds are executed every  $O(\#\mathbb{L}_{\text{MM}})$  operations, we get the amortized cost above. Note that this is a very loose asymptotic upper-bound and that, in practice, DLS's amortized cost will be much lower.

Finally, the client locally stores the state composed of both the old and new state dictionaries, and the set of searched for labels. That is, client storage is

$$O\left(\#\mathbb{L}_{\text{MM}} \cdot \log\left(\max_{\ell \in \mathbb{L}_{\text{MM}}} (\#\text{MM}[\ell] + \text{del}(\ell, e))\right)\right).$$

**Variants.** While DLS improves on the query, update and storage complexity of all previous dynamic EMM schemes, its token size is larger. This can be improved using one of the following three approaches. The first reduces the token size to be  $O(\#\text{up}(\ell, c))$ . It consists of using the counter-based approach for the *old* dictionary; that is, granting the server the ability to compute all encrypted labels for the old dictionary (as  $\pi_{\text{dyn}}$ ). Given that all new updates are going to be added to the new dictionary, the server, with a key derived from the label, can generate all the encrypted labels. That is, the client will only send a key along with the new token  $\text{ntk}$  which has size equal to the number of updates for  $\ell$  in the current epoch.

The second approach leverages constrained pseudo-random functions, introduced in Diana [6]. The client can simply send two constrained keys for the required ranges for both the old and new counters. This approach reduces the token size to be

$$O\left(\log(\#\text{MM}[\ell] + \text{del}(\ell, e))\right),$$

when using GGM [19] as the constrained PRF.

The third approach is the combination of the first two approaches. The client sends a token composed of: (1) a key for a standard PRF with which the server computes all encrypted labels in the old dictionary; and (2) a constrained key for the appropriate range in the

new dictionary. The size of the search token in this case is

$$O\left(\log(\#\text{up}(\ell, c))\right),$$

when using GGM as the constrained PRF.

This last variant of DLS outperforms all previous constructions in query complexity, update complexity, token size, query round complexity and client and server storage—all while being secure in the standard model. DLS can also be easily modified to have constant client memory by storing the state on the server in a zero-leakage encrypted multi-map, e.g., using ORAM at the cost of an additive poly-logarithmic overhead per query.

## 4.2 Security

In the following, we detail the leakage of DLS against standard and snapshot adversaries, respectively.

**Against a persistent adversary.** The setup leakage of DLS consists of the size of the multi-map  $\text{MM}$ . The query leakage of DLS for a label  $\ell$  consists of the search, response length and operation patterns. The search pattern captures if and when the label has been searched for in the past. As DLS is response-hiding, it does not reveal the access pattern but only the response length which is the cardinality of the result. The operation pattern reveals if an operation for a label  $\ell$  was an update (i.e., an add or delete) or not. The update leakage of DLS is the size of the tuple to be updated. The rebuild leakage is the size of the updated multi-map and, for each label that was searched for in the current epoch, the number of deletes in  $\text{DX}_o$ . We now give a precise description of DLS's leakage profile and show that it is adaptively-secure with respect to it. Its setup leakage is

$$\mathcal{L}_S(\text{MM}) = \sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell].$$

The query leakage is

$$\mathcal{L}_Q(\text{MM}, \ell) = \left(\text{QP}, \text{RL}, \text{OP}\right).$$

Here, QP is the query pattern which reveals if and when a query is repeated. More formally, it is defined as  $\text{QP} = B$ , where  $B$  is a binary square matrix of size  $t$  and  $t$  is the total number of operations that have been made.  $B$  is such that  $B_{i,j} = 1$  if the  $i$ th and  $j$ th queries are the same, and 0 otherwise. The response length pattern is

$$\text{RL}(\text{MM}, \ell) = \#\text{MM}[\ell].$$

The operation pattern is

$$\text{OP}(\text{MM}, \ell) = \mathbf{m},$$

where  $\mathbf{m}$  is a binary vector of size  $t$ , where  $t$  is the total number of operations. For all  $i \in [t]$ ,  $m_i = 1$  if the  $i$ th operation is an update and  $m_i = 0$  otherwise. Its update leakage is

$$\mathcal{L}_U(\text{MM}, (\text{op}, \ell, \mathbf{v})) = \#\mathbf{v},$$

for all  $\text{op} \in \{\text{edit}^+, \text{edit}^-\}$ . Its rebuild leakage is

$$\mathcal{L}_R(\text{MM}) = (\#\text{del}_\ell^o)_{\ell \in \mathbb{S}_e},$$

where  $\#\text{del}_\ell^o$  is number of pairs with label  $\ell$  removed from the old dictionary  $\text{DX}_o$  and  $\mathbb{S}_e \subseteq \mathbb{L}$  is the set of labels that were searched for in the current epoch.

**Theorem 4.1.** *If SKE is an RCPA-secure encryption scheme and  $F$  is a pseudo-random function, then DLS is  $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U, \mathcal{L}_R)$  secure.*

The proof of Theorem 4.1 appears in Appendix A.

**Leakage against a snapshot adversary.** The snapshot leakage  $\mathcal{L}_{\text{SN}}$  of DLS is

$$\mathcal{L}_{\text{SN}}(\text{MM}, \text{op}_1, \dots, \text{op}_i) = \sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}_i[\ell]$$

where  $\text{MM}_i$  is the current version of the multi-map. Note that, given  $\mathcal{L}_{\text{SN}}$ , DLS is breach-resistant based on Definition 3.5.

**Theorem 4.2.** *If SKE is an RCPA-secure encryption scheme and  $F$  is a pseudo-random function, then DLS is  $(m, \mathcal{L}_{\text{SN}})$ -snapshot secure, for  $m \in \text{poly}(k)$ .*

The proof of Theorem 4.2 is in Appendix C.

**Corollary 4.3.** *DLS is dual-secure.*

## 5 DLS<sup>d</sup>: DLS with De-amortized Rebuilding

In Section 4, we introduced an instantiation of the rebuilding protocol with a worst-case complexity linear in the size of the old dictionary, i.e., in

$$O\left(\sum_{\ell \in \mathbb{L}_{\text{MM}}} \left(\#\text{MM}[\ell] + \text{del}(\ell, e) - \text{up}(\ell, c)\right)\right),$$

where  $\#\text{up}(\ell, c)$  is the number of updates for  $\ell$  since the last rebuild operation. Moreover, DLS only achieves dual security under the assumption that all intermediate

steps of the rebuild protocol are hidden from the snapshot adversary. In this section, we show how to overcome these two challenges, resulting in a variant of our construction called DLS<sup>d</sup>. DLS<sup>d</sup> uses a de-amortized rebuild process, that maintains dual security with minimal leakage while providing optimal worst-case efficiency under *no assumptions* on when snapshots occur.

**Intuition.** DLS<sup>d</sup> is a variant of DLS with de-amortized rebuilding, i.e., it has the same setup, search token and get algorithms, but differs on how the rebuilding process is performed. We now provide a gradual description of this process starting from our previous rebuild protocol.

We can de-amortize the Rebuild protocol of DLS by simply operating on one label at a time in the case of a searched label, or on a pair in the other case. This solution provides a de-amortized rebuilding and clearly does not introduce any additional leakage against a persistent adversary. Unfortunately, this solution still induces additional leakage against a snapshot adversary because it can differentiate between the two types of labels based on the number of pairs that have been inserted into the new dictionary. Remember that our goal is to have snapshot leakage composed only of the size of the current structure. So far, this is not the case as we have an additional leakage.

We solve the issue above by allowing some client storage. That is, instead of inserting an arbitrary number of pairs in the case of searched for labels, only  $\lambda$  pairs are inserted and the remaining pairs are kept in the client side. For the unsearched for labels, the same number of pairs,  $\lambda$ , is similarly inserted in the structure. This will make differentiating between the two types of labels impossible for a snapshot adversary.

**De-amortization.** An important design decision we have to make is to figure out when to rebuild. There are three possible approaches, including:

- **at update time:** this approach runs the de-amortized rebuild steps as a sub-routine of the update protocol. That is, whenever the client updates the encrypted structure, the client simultaneously performs a partial rebuilding that operates on  $\lambda$  pairs. Note that this choice is natural as a snapshot adversary already knows that a data structure has been modified due to the update operation. Therefore performing a partial-rebuilding that is triggered by an update will not leak more than the number of rebuilt pairs, which is  $\lambda$ .
- **at query time:** this approach consists of running the de-amortized rebuild step as a sub-routine of the query protocol. In this case, it is easy to see that

a snapshot adversary would learn that a search occurred by just looking to the encrypted structure, which violates breach-resistance and, therefore, dual security.

- **continuously**: this approach runs the de-amortized rebuild steps continuously in the background. This is quite similar to the previous rebuild protocol except that the client can always query or update the structure independently of the rebuilding process. In other words, the rebuild does not affect the correctness of the query operation. This approach provides the client with the flexibility to rebuild the data structure whenever it is required (which is not the case of the previous approaches).

Our preferred approach is to execute rebuild steps at update time. In the following, we provide the algorithmic details. Note that the syntax of the STE scheme will change slightly because we are now dealing with a self-adjusting EMM. We refer the reader to Definition 3.1 for more details.

**Details.**  $\text{DLS}^d$  is a self-adjusting dynamic encrypted multi-map composed of four non-interactive algorithms and one interactive protocol such that  $\text{DLS}^d = (\text{Setup}, \text{Token}, \text{Get}, \text{Update}, \text{Rslv})$ . Below, we only detail the Update protocol as Setup, Token, Get and Rslv remain the same as in DLS. The pseudo-code details of  $\text{DLS}^d$  appear in the full version of the paper.

Update is a two-party protocol between the client and the server. The client’s input consists of a key  $K$ , a state  $st$  and an update  $u$ . The server’s input consists of the encrypted multi-map EMM. The state is the same as the state of DLS except that it is augmented with the de-amortization rate  $\lambda$  and two lists  $L_{sr}$  and  $L_{un}$ .  $L_{sr}$  is a stash that stores the pairs corresponding to the searched for labels, while  $L_{un}$  is a stash that stores the pairs for the unsearched for labels.

The client starts by creating the update token and sending it to the server which then inserts the updates in the new dictionary using DLS UToken and Put algorithms, respectively. The client then samples a bit  $b$ . If  $b = 0$ , the client rebuilds a searched label in  $L_{sr}$ , otherwise it rebuilds an unsearched label in  $L_{un}$ . If  $\#L_{sr} < \lambda$  or  $\#L_{un} < \lambda$ , it means the client does not hold enough pairs in the stashes to rebuild and needs to prepare additional pairs as follows.

If it is rebuilding searched labels and  $\#L_{sr} < \lambda$ , then the client picks a label  $\ell \in \mathbb{S}_e$ , where  $\mathbb{S}_e$  contains all the searched for labels. The client fetches the corresponding pairs from the old dictionary  $\text{DX}_o$ , decrypts them, removes all the deleted pairs, and appends the remain-

ing pairs to  $L_{sr}$ . The client then removes the label  $\ell$  from  $\mathbb{S}_e$ . When it is time to send pairs to the server, it generates  $\lambda$  freshly encrypted pairs from  $L_{sr}$  to send. Similarly to the UToken algorithm, a fresh pair is composed of two sub-tokens. The first,  $\text{tk}_1$ , is the evaluation of the pseudo-random function on the label  $\ell$ , the global version  $\text{version}_g$ , and a counter  $\text{count}_n$ . The second,  $\text{tk}_2$ , is generated by encrypting the value  $v$  corresponding to the label  $\ell$  concatenated to the string  $\text{edit}^+$ .

Otherwise, if it is rebuilding unsearched labels and  $\#L_{un} < \lambda$ , the client retrieves an unsearched label uniformly at random from  $\mathbb{L}_{MM} \setminus \mathbb{S}_e$  and only retrieves one pair using the label’s counter. The client then decrements the counter and updates the state accordingly. If the counter is less than one, it removes the label from the state,  $\text{DX}_o^{\text{st}}$ . The client refreshes the pair similarly to the operations described earlier, but then appends it to  $L_{un}$ . The client repeats this process as long as  $\#L_{un} < \lambda$ . Finally if both the stashes are empty and each pair in  $\text{DX}_o$  is refreshed, the server deletes the old dictionary and the new dictionary becomes the old dictionary. The client similarly deletes the state of the old dictionary and replaces it with the state of the new dictionary. A new empty dictionary and state are initialized for future updates. Finally the client increments the global version.

**Efficiency.**  $\text{DLS}^d$  has the same query complexity, storage complexity and token size as DLS. Below, we only detail the client memory and update complexity.  $\text{DLS}^d$  introduces two new data structures that are stored at the client. The first,  $L_{sr}$ , stores the pairs for the searched labels and has size

$$O\left(\max\left\{\lambda, \max_{\ell \in \mathbb{L}_{MM}} (\#\text{MM}[\ell] + \text{del}(\ell, e))\right\}\right).$$

The second,  $L_{un}$ , stores the pairs for the unsearched labels and has size  $O(\lambda)$ . In addition, similarly to DLS, the client also stores the state which is composed of both the old and new state dictionaries, and the set of searched for labels. This results in  $O(\#\mathbb{L}_{MM} \cdot \log(\max_{\ell \in \mathbb{L}_{MM}} (\#\text{MM}[\ell] + \text{del}(\ell, e))))$  storage at the client. The overall client storage is then

$$O\left(\#\mathbb{L}_{MM} \cdot \log\left(\max_{\ell \in \mathbb{L}_{MM}} (\#\text{MM}[\ell] + \text{del}(\ell, e))\right) + \max\left\{\lambda, \max_{\ell \in \mathbb{L}_{MM}} (\#\text{MM}[\ell] + \text{del}(\ell, e))\right\}\right).$$

The update complexity is  $O(\lambda + \#\mathbf{v})$  with two rounds of interactions.<sup>5</sup> The first round fetches the pairs and

<sup>5</sup> One might think that the update complexity should be equal to  $O(\#\mathbf{v} + \max\{\lambda, \max_{\ell \in \mathbb{L}_{MM}} (\#\text{MM}[\ell] + \text{del}(\ell, e))\})$ . However,



sends the update token to the server and the second inserts the freshly encrypted pairs from the states along with the updates.

## 5.1 Security

In the following, we describe the leakage of  $\text{DLS}^d$  against persistent and snapshot adversaries, respectively.

**Leakage against a persistent adversary.** The setup and query leakage of  $\text{DLS}^d$  is the same as DLS. The Update leakage consists of the update leakage of DLS and the de-amortization rate which is public. Note that, contrary to DLS, there is no explicit rebuild leakage since the rebuild process is de-amortized and executed as a sub-routine of the update operation.

The update with rebuild leakage is then

$$\mathcal{L}_{U_r}(\text{MM}, u) = \left( \mathcal{L}_U(\text{MM}, u), \mathcal{L}_{R_d}(\text{MM}) \right),$$

where  $u = (\text{op}, \ell, \mathbf{v})$ ,  $\mathcal{L}_U(\text{MM}, u) = \#\mathbf{v}$ ,  $\mathcal{L}_{R_d}(\text{MM}) = \left( \lambda, \left( \#\text{del}_\ell^o \right)_{\ell \in \mathbb{S}_e} \right)$ ,  $\lambda$  is the rebuild rate of the Update protocol, and  $\mathbb{S}_e \subseteq \mathbb{L}$  is the set of labels that were searched for in the current epoch.

**Theorem 5.1.** *If SKE is an RCPA-secure encryption scheme and  $F$  is a pseudo-random function, then  $\text{DLS}^d$  is  $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_{U_r})$  secure.*

The proof of Theorem 5.1 appears in Appendix B.

**Leakage against a snapshot adversary.** The snapshot leakage  $\mathcal{L}_{\text{SN}}$  of  $\text{DLS}^d$  is

$$\mathcal{L}_{\text{SN}}(\text{MM}, \text{op}) = \left( \lambda, \sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}(\ell) \right).$$

Note that  $\lambda$  is a public parameter so  $\text{DLS}^d$  is breach-resistant.

**Theorem 5.2.** *If SKE is an RCPA-secure encryption scheme and  $F$  is a pseudo-random function, then  $\text{DLS}^d$  is  $(m, \mathcal{L}_{\text{SN}})$ -snapshot secure, for  $m \in \text{poly}(k)$ .*

The proof of Theorem 5.2 appears in Appendix D.

**Corollary 5.3.**  $\text{DLS}^d$  is dual-secure.

---

the quantity  $\max_{\ell \in \mathbb{L}_{\text{MM}}} (\#\text{MM}[\ell])$  can be de-amortized over  $\lambda^{-1}$ .  $\max_{\ell \in \mathbb{L}_{\text{MM}}} (\#\text{MM}[\ell])$  updates, as the client will not fetch any pair as long as the state  $\mathbb{L}_{\text{sr}}$  contains more than  $\lambda$  pairs.

## 6 Empirical Evaluation

We now evaluate how our construction performs in practice. We implemented  $\text{DLS}^d$ , the de-amortized variant of DLS, in Java using the Clusion encrypted search library [34]. It consists of 1114 lines of codes excluding 180 lines for testing purposes calculated using CLOC [14]. We set  $\lambda$  to 3 for all experiments except for one where we vary  $\lambda$  to study its effect on update time.

**Parsing and indexing.** We used the parsing and indexing functionality of the Clusion library to process data (which is itself based on the Lucene parser [31]). Through Clusion, our implementation handles pdf files, Microsoft Office files (doc, docx, pptx, xlsx), basic text files. For media files, it only indexes the file names.

**Cryptographic primitives.** We use the cryptographic primitives provided by Clusion (themselves based on Bouncy Castle [11]). For symmetric encryption, we use AES in CTR mode with a 256 bit key. We use of HMAC-SHA256/512 for PRFs.

**Experimental setup.** We ran our experiments on an Amazon EC2 instance running Ubuntu Linux (c3.8xlarge) with an Intel Xeon E5-2680 v2 (Ivy Bridge) Processor with 32 vCPU and 60 GB of RAM. For all our experiments, we used the Wikipedia data dumps. The total uncompressed size of our dataset was 26.5GB. There are a total number of 2,681,795 files in this dataset. We partitioned these files into different folders. The first folder had 17,600 files with a total size of about 250 MB and the last folder had 554,059 files with a total size of 3.6GB. In our empirical evaluation, we want to quantify the following characteristics of  $\text{DLS}^d$ :

1. The time to set up the EMM as a function of the number of pairs;
2. The size of each EMM and of the client state as a function of the number of pairs;
3. The time taken to respond to a query for labels with different selectivity as a function of the number of pairs. The selectivity of a label is the number of values associated to it;
4. The time taken for an update operation as a function of the number of pairs in the EMM. We also measure how different rebuild parameters  $\lambda$  affect the time taken;
5. The effect of de-amortized rebuilding on the time taken by a query operation in  $\text{DLS}^d$  specifically when there are deletes involved.

**Setup time and storage overhead.** Fig. 3 describes the time taken to set up an EMM as a function of

the number of label/value pairs stored. We created EMMs with number of pairs ranging from 2,758,254 to 83,239,341. The setup phase takes under 13 minutes with a multi-threaded implementation.<sup>6</sup> Fig. 4 shows the size on disk of both the client state and the EMM for different number of pairs. We observe that even for about 83 million pairs, the client state is only 210MB for an 11GB EMM. Recall that in scenarios in which client memory is expensive (e.g., in the case of email clients [3]), the state can always be outsourced using ORAM at the cost of a poly-log overhead. Also note that DLS’s state is (asymptotically) the same as all recent forward-private schemes [5, 6, 16, 29].

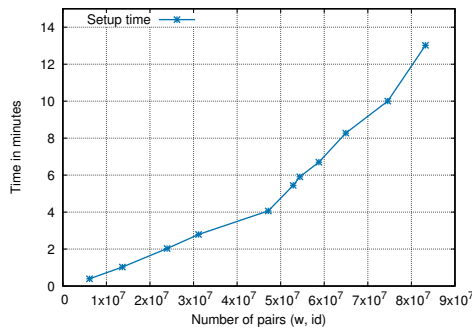


Fig. 3. Setup time.




Fig. 4. Encrypted multi-map and state sizes.

**Search and update operations.** Fig. 5 describes the time taken to search for labels of different selectivities ( $MM(w)$ ). We also measure time taken by an update operation. In our experiments, the client and server are running on the same machine. We measure the effect of increasing the EMM size on the query time, which

<sup>6</sup> Setup time would be roughly  $\times 32$  slower with a single-threaded execution.

seems negligible. We do not send any update tokens between the query operations during the experiment. For each EMM we first search for labels of different selectivities (100, 1000 and 10,000). The search time for all selectivities is less than 1 microsecond per pair. We ran every search data point corresponding to every number of pairs in the x-axis 500 times. We re-ran the whole experiment 10 times, then we took the median. From Fig. 6 we can see that the update operation is more costly as it takes around 100 milliseconds when  $\lambda$  is set to 3. This can be attributed to the fact that the rebuild is performed with the update algorithm. Whenever we do an update, we also execute  $\lambda$  de-amortized rebuild steps. One can see that as we increase  $\lambda$ , the update time increases proportionally. Note also that both the search and update operations are independent of the number of pairs in the EMM.




Fig. 5. Search time for 100, 1000 and 10000 search selectivities.


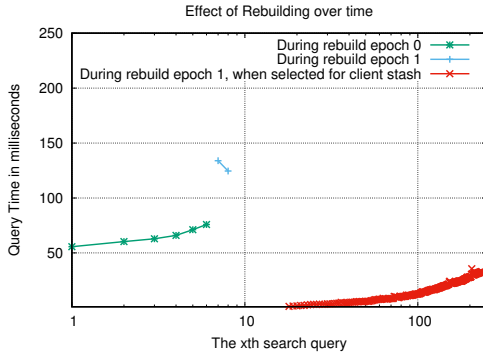


Fig. 6. Update time for different rebuilding parameters.

**Effect of rebuilding.** Figure 7 describes the query time as a function of the number of delete opera-



**Fig. 7.** Query times for the same keyword at different times. During epoch 0, the keyword starts with a selectivity of 100,000 and 8000 delete operations are performed between each query operation. No further deletes happen in the next epoch.

tions. The keyword being queried initially has selectivity 100,000 which was achieved by updating the EMM with 100,000 update tokens. The consequence of this is that these values are initially in  $DX_n$ . The first point of epoch 0 (the green line) represents the first query. After that, we send 8000 delete tokens, execute a search, send 8000 more deletes, execute a search and so on. This continues until we send a total of 40,000 delete tokens (the end of the green line). At this stage, the selectivity of the keyword is 60,000 but the total number of pairs in  $DX_n$  associated with the keyword is 140,000. During epoch 0 we can clearly see that query time increases as the number of deletes increases.

At some point, the rebuild operation ends and epoch 1 begins and  $DX_n$  becomes  $DX_o$ . At the beginning of epoch 1 two queries occur before the keyword has been selected to be rebuilt. The gap in query time between epoch 0 and the first two queries of epoch 1 is due to the fact that the pairs are stored in different sized dictionaries in the two epochs. In epoch 0 they are stored in a relatively empty  $DX_n$  whereas in epoch 1 they are stored in  $DX_o$  along with 50 million other pairs (remember that we need to keep updating the EMM to make the rebuild process progress). Once the keyword is selected to be rebuilt (red line), it is moved to the client’s local stash; specifically, to  $L_{sr}$  since it was searched. At this stage, query time is negligible but increases as the pairs are inserted into  $DX_n$ . When all pairs have been inserted,  $DX_n$  only holds 60,000 pairs associated to this keyword (i.e., all the delete pairs have been removed).

One can clearly see the difference in query time before and after the keyword has been rebuilt. Before it has been rebuilt (i.e., epoch 0 and beginning of epoch 1 which are the green and blue lines, respectively), the

query times range from 55.6 ms to 75.8 ms. After being rebuilt (i.e., the end of epoch 1 which is the red line), query times range from 0.1 ms to 33 ms.

Also, note that all query times stay well below 1 microsecond per pair.

**Comparison with previous constructions.** Given the asymptotic overhead of the SPS construction, we do not compare it to  $DLS^d$  and focus mainly on the Sophos and Diana schemes of [5] and [6]. The empirical evaluations of these constructions are based on C/C++ implementations whereas our implementation of  $DLS^d$  is in Java. In addition, the Sophos implementation is multi-threaded. Our implementation of  $DLS^d$ , on the other hand, is only multi-threaded for setup. And while our experiments are in memory and the ones of [5] and [6] are on disk, they are conducted using SSDs which have efficiency comparable to memory. Even with these advantages, the performance of  $DLS^d$  is better or, at the very least, comparable to those of Sophos and Diana. [5] reports search times ranging from 24 microseconds to 7 microseconds per pair depending on the selectivity of the keyword. On similar datasets, the Diana implementation in [6] is 10 times faster. Similarly to our  $DLS^d$  implementation, the hardware accelerated C/C++ implementation of Diana takes less than 1 microsecond per pair.

An important distinction between our constructions and Sophos and Diana is in how their query time is affected by delete operations. While the query time of Sophos and Diana will increase with the total number of deletes ever performed, the query time of  $DLS^d$  will only increase with the number of deletes since the last rebuild (or the current epoch). As shown above, this makes a significant difference in the query time of  $DLS^d$ .

**Locality.** Several works [1, 9, 10, 15] have considered the design of locality-friendly SSE schemes. While  $DLS^d$  achieves near-optimal search complexity, it has a non-optimal locality. This results from the way the plaintext multi-map tuples are stored in the underlying dictionary (each tuple is split and stored in multiple non-contiguous memory locations). While locality had no impact in our experiments (since they were conducted in memory) one might observe a slowdown if the encrypted structure is stored on disk.

It is worth mentioning, however, that  $DLS^d$  could be made local using techniques from [10]. In fact, the design and approach of  $DLS^d$  could be used to extend the locality-friendly 2Lev construction of [10] instead of  $\pi_{bas}$ .

## References

- [1] G. Asharov, M. Naor, G. Segev, and I. Shahaf. Searchable symmetric encryption: Optimal locality in linear space via two-dimensional balanced allocations. In *ACM Symposium on Theory of Computing (STOC '16)*, STOC '16, pages 1101–1114, New York, NY, USA, 2016. ACM.
- [2] Adam J. Aviv, Seung Geol Choi, Travis Mayberry, and Daniel S. Roche. Oblivisync: Practical oblivious file backup and synchronization. In *Network and Distributed System Security Symposium (NDSS '16)*, 2016.
- [3] Wei Bai, Ciara Lynton, Michelle L. Mazurek, and Charalampos Papamanthou. Understanding user tradeoffs for search in encrypted communication. *EuroSP*, 2018.
- [4] E.-O. Blass, T. Mayberry, G. Noubir, and K. Onarlioglu. Toward robust hidden volumes using write-only oblivious RAM. In *ACM Conference on Computer and Communications Security (CCS '14)*, pages 203–214, 2014.
- [5] R. Bost. Sophos - forward secure searchable encryption. In *ACM Conference on Computer and Communications Security (CCS '16)*, 20016.
- [6] R. Bost, B. Minaud, and O. Ohrimenko. Forward and backward private searchable encryption from constrained cryptographic primitives. In *ACM Conference on Computer and Communications Security (CCS '17)*, 2017.
- [7] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. Leakage-abuse attacks against searchable encryption. In *ACM Conference on Communications and Computer Security (CCS '15)*, pages 668–679. ACM, 2015.
- [8] D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *Advances in Cryptology - CRYPTO '13*. Springer, 2013.
- [9] D. Cash and S. Tessaro. The locality of searchable symmetric encryption. In *Advances in Cryptology - EUROCRYPT 2014*, 2014.
- [10] David Cash, Joseph Jaeger, Stanislaw Jarecki, Charanjit Jutla, Hugo Krawczyk, Marcel Rosu, and Michael Steiner. Dynamic searchable encryption in very-large databases: Data structures and implementation. In *Network and Distributed System Security Symposium (NDSS '14)*, 2014.
- [11] Bouncy Castle. Crypto API. In <http://www.bouncycastle.org>.
- [12] M. Chase and S. Kamara. Structured encryption and controlled disclosure. In *Advances in Cryptology - ASIACRYPT '10*, volume 6477 of *Lecture Notes in Computer Science*, pages 577–594. Springer, 2010.
- [13] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky. Searchable symmetric encryption: Improved definitions and efficient constructions. In *ACM Conference on Computer and Communications Security (CCS '06)*, pages 79–88. ACM, 2006.
- [14] Al Danial. Cloc. In <http://www.cloc.sourceforge.net>.
- [15] I. Demertzis and C. Papamanthou. Fast searchable encryption with tunable locality. In *ACM International Conference on Management of Data (SIGMOD '17)*, SIGMOD '17, pages 1053–1067, New York, NY, USA, 2017. ACM.
- [16] Mohammad Etemad, Alptekin Küpçü, Charalampos Papamanthou, and David Evans. Efficient dynamic searchable encryption with forward privacy. *PoPETs '18, Issue 1*, 2018.
- [17] B. A. Fisch, B. Vo, F. Krell, A. Kumarasubramanian, V. Kolesnikov, T. Malkin, and S. M. Bellare. Malicious-client security in blind seer: a scalable private dbms. In *IEEE Symposium on Security and Privacy*, pages 395–410. IEEE, 2015.
- [18] S. Garg, P. Mohassel, and C. Papamanthou. TWORAM: efficient oblivious RAM in two rounds with applications to searchable encryption. In *Advances in Cryptology - CRYPTO 2016*, pages 563–592, 2016.
- [19] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. In *IEEE Symposium on the Foundations of Computer Science (FOCS '84)*, pages 464–479. IEEE Computer Society, 1984.
- [20] P. Grubbs, T. Ristenpart, and V. Shmatikov. Why your encrypted database is not secure. In *Workshop on Hot Topics in Operating Systems (HotOS '17)*, pages 162–168, New York, NY, USA, 2017. ACM.
- [21] F. Hahn and F. Kerschbaum. Searchable encryption with secure and efficient updates. In *ACM Conference on Computer and Communications Security (CCS '14)*, CCS '14, pages 310–320, New York, NY, USA, 2014. ACM.
- [22] M. Saiful Islam, M. Kuzu, and M. Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In *Network and Distributed System Security Symposium (NDSS '12)*, 2012.
- [23] S. Kamara and T. Moataz. Boolean searchable symmetric encryption with worst-case sub-linear complexity. In *Advances in Cryptology - EUROCRYPT '17*, 2017.
- [24] S. Kamara and C. Papamanthou. Parallel and dynamic searchable symmetric encryption. In *Financial Cryptography and Data Security (FC '13)*, 2013.
- [25] Seny Kamara and Tarik Moataz. SQL on structurally-encrypted databases. *IACR Cryptology ePrint Archive*, 2016:453, 2016.
- [26] J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC, 2008.
- [27] Florian Kerschbaum and Anselme Tuono. An efficiently searchable encrypted data structure for range queries. *CoRR*, abs/1709.09314, 2017.
- [28] K. Kurosawa and Y. Ohtaki. How to update documents verifiably in searchable symmetric encryption. In *International Conference on Cryptology and Network Security (CANS '13)*, pages 309–328, 2013.
- [29] Russell W. F. Lai and Sherman S. M. Chow. Forward-secure searchable encryption on labeled bipartite graphs. In *Applied Cryptography and Network Security - 15th International Conference, ACNS 2017, Kanazawa, Japan, July 10-12, 2017, Proceedings*, pages 478–497, 2017.
- [30] K. Lewi and D. Wu. Order-revealing encryption: New constructions, applications, and lower bounds. In *ACM Conference on Computer and Communications Security (CCS '16)*, 2016.
- [31] Lucene. Parser. In <http://lucene.apache.org>.
- [32] X. Meng, S. Kamara, K. Nissim, and G. Kollios. GreCs: Graph encryption for approximate shortest distance queries. In *ACM Conference on Computer and Communications Security (CCS '15)*, 2015.
- [33] I. Miers and P. Mohassel. Io-dsse: Scaling dynamic searchable encryption to millions of indexes by improving local-



- ity. Cryptology ePrint Archive, Report 2016/830, 2016. <http://eprint.iacr.org/2016/830>.
- [34] T. Moataz. Clusion. <https://github.com/encryptedsystems/Clusion>.
- [35] M. Naor and V. Teague. Anti-persistence: history independent data structures. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 492–501, New York, NY, USA, 2001. ACM.
- [36] M. Naveed, M. Prabhakaran, and C. Gunter. Dynamic searchable encryption via blind storage. In *IEEE Symposium on Security and Privacy (S&P '14)*, 2014.
- [37] V. Pappas, F. Krell, B. Vo, V. Kolesnikov, T. Malkin, S.-G. Choi, W. George, A. Keromytis, and S. Bellovin. Blind seer: A scalable private dbms. In *Security and Privacy (SP), 2014 IEEE Symposium on*, pages 359–374. IEEE, 2014.
- [38] E. Stefanov, C. Papamanthou, and E. Shi. Practical dynamic searchable encryption with small leakage. In *Network and Distributed System Security Symposium (NDSS '14)*, 2014.
- [39] Y. Zhang, J. Katz, and C. Papamanthou. All your queries are belong to us: The power of file-injection attacks on searchable encryption. In *USENIX Security Symposium*, 2016.

## A Proof of Theorem 4.1

**Theorem 4.1.** *If SKE is an RCPA-secure encryption scheme and  $F$  is a pseudo-random function, then DLS is  $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U, \mathcal{L}_R)$  secure.*

*Proof.* Consider the simulator that works as follows.

1. **Setup simulation.** Given the setup leakage  $\mathcal{L}_S(\text{MM}) = (\sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell])$ , the simulator initializes two dictionaries  $\text{DX}_0$  and  $\text{DX}_1$  of size  $\sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$ . It then initializes a *state multi-map*  $\text{MM}_0$  and three vectors  $\text{vec}_i, \text{vec}_o, \text{vec}_n$ . For  $1 \leq i \leq \sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$ :
  - (a) compute  $\text{str}_i, \text{ctr}_i \xleftarrow{\$} \{0, 1\}^k$  where  $\text{str}_i, \text{ctr}_i$  are of appropriate lengths;
  - (b) set  $\text{DX}_0[\text{str}_i] = \text{ctr}_i$ ;
  - (c) add  $\text{str}_i$  to  $\text{vec}_i$ .

It then outputs  $\text{EMM} = (\text{DX}_0, \text{DX}_1)$ .

2. **UToken simulation.** Given the update leakage

$$\mathcal{L}_U(\text{MM}, \ell, \mathbf{v}, \text{op}) = \#\mathbf{v},$$

for all  $\text{op} \in \{\text{edit}^+, \text{edit}^-\}$ . For  $1 \leq i \leq \#\mathbf{v}$ ,

- (a) compute  $\text{str}_i, \text{value}_i \xleftarrow{\$} \{0, 1\}^k$  where  $\text{str}_i, \text{value}_i$  are of appropriate lengths;
- (b) add  $\text{str}_i$  to  $\text{vec}_n$ .

It then outputs  $\text{UToken} = (\text{str}_i, \text{value}_i)_{i \in \bar{\mathbf{v}}}$ .

3. **Token simulation.** Given the query leakage

$$\mathcal{L}_Q(\text{MM}, \ell) = \left( \text{QP}, \text{RL}, \text{OP} \right),$$

where QP is the search pattern,  $\text{RL} = (\#\text{DX}_o[\ell], \#\text{DX}_n[\ell])$  is the access pattern and OP is the operation pattern. It adds  $\ell$  to a set of searched labels  $\mathbb{S}$  and executes the following steps.

- (a) if  $\text{MM}_0[\ell] = \perp$  (which occurs if  $\ell$  is being searched for the first time),
  - i. set  $\text{count} = \#\text{DX}_o[\ell]$  from RL and initialize an empty set  $S_l$ ;
  - ii. using OP, extract and remove all labels that were updates of  $\ell$  from  $\text{vec}_o$  and add them to  $S_l$ . This step is performed in a case if where a rebuild has already occurred and some updates are in the old dictionary.
  - iii. from  $i = \#S_l$  to  $\text{count}$ , pick and remove a label from  $\text{vec}_i$  and add it to  $S_l$ .
  - iv. assign  $S_l$  to  $\text{MM}_0[\ell]$
- (b) else if  $\text{MM}_0[\ell] < \#\text{DX}_o[\ell]$  (which happens when there is a rebuild and some new updates are now in the old dictionary),
  - using OP, extract and remove all labels that were updates of  $\ell$  from  $\text{vec}_o$  and prepend them to  $\text{MM}_0[\ell]$ .
- (c) set  $\text{count} = \#\text{DX}_n[\ell]$  from RL and initialize an empty set  $S_l$ ;
- (d) Using OP, add all labels that were updates of  $\ell$  from  $\text{vec}_n$  (all updates in the new dictionary) and add them to  $S_l$ .

It then outputs  $\text{otk} = \text{MM}_0[\ell]$  and  $\text{ntk} = S_l$ .

4. **Rebuild simulation.** Given the rebuild leakage

$$\mathcal{L}_R(\text{MM}) = (\#\text{del}_\ell^o)_{\ell \in \mathbb{S}},$$

where  $\#\text{del}_\ell^o$  is the number of pairs with label  $\ell$  that have been removed from the old dictionary since the last rebuild. It executes the following steps,

- (a) to remove the appropriate number of tuples it does the following. For  $\ell \in \mathbb{S}$ ,
  - i. if  $\text{MM}_0[\ell] \neq \perp$ ,
    - A. initialize  $v_{lab}$ ;
    - B. for  $\text{count} = 0$  to  $\#\text{MM}_0[\ell] - \#\text{del}_\ell^o$ ;
      - compute  $\text{str}, \text{ctr} \xleftarrow{\$} \{0, 1\}^k$  where  $\text{str}, \text{ctr}$  are of appropriate lengths;
      - set  $v_{lab}[\text{count}] = \text{str}$  and  $\text{DX}_1[\text{str}] = \text{ctr}$ ;
      - send the pair  $(\text{str}, \text{ctr})$  to adversary.
    - C. set  $\text{MM}_0[\ell] := v_{lab}$ .
  - (b) To freshly encrypt the remaining pairs in the old dictionary, it sets  $\text{counter}$  to  $\#\text{vec}_i + \#\text{vec}_o + \sum_{\ell \in \mathbb{K}(\text{MM}_0) \setminus \mathbb{S}} \#\text{MM}_0[\ell]$  where  $\text{vec}_i$  has all the still

unsearched for labels from setup,  $vec_o$  has the unsearched for updates which happened before the last rebuild and  $\mathbb{K}(\text{MM}_0) \setminus \mathbb{S}$  is a set of all  $\ell$  that were searched for before the last rebuild but not since (where  $\mathbb{K}(\text{MM}_0)$  is a set that contains all the labels of the state multi-map  $\text{MM}_0$ ). It then does the following:

- (c) initialize  $new_l$  of size  $counter$ ;
- (d) for  $i = 1$  to  $counter$ ,
  - i.  $str, ctr \xleftarrow{\$} \{0, 1\}^k$  where  $str, ctr$  are of appropriate lengths;
  - ii. append  $(str, ctr)$  to  $new_l$ ;
  - iii. send  $(str, ctr)$  to adversary.
 Now that the simulator has all fresh encryptions, it just needs to replace them in its internal state.
- (e) for all  $\ell \in (\mathbb{K}(\text{MM}_0) \setminus \mathbb{S})$ ,
  - i. initialize  $v_{lab}$  and for  $count = 0$  to  $\#\text{MM}_0[\ell]$ ;
    - A. pick and remove a pair  $(str, ctr)$  from  $new_l$  at random;
    - B. append  $str$  to  $v_{lab}$ ;
    - C. set  $\text{DX}_1[str] = ctr$ .
  - ii. set  $\text{MM}_0[\ell] := v_{lab}$ .
- (f) initialize  $v_{lab}$  and for  $\ell \in vec_l$ ,
  - i. pick and remove a pair  $(str, ctr)$  from  $new_l$  at random;
  - ii. append  $str$  to  $v_{lab}$ ;
  - iii. set  $\text{DX}_1[str] = ctr$ .
- (g)  $vec_l := v_{lab}$ ;
- (h) initialize  $v_{lab}$  and for  $\ell \in vec_o$ ,
  - i. pick and remove a pair  $(str, ctr)$  from  $new_l$  at random;
  - ii. append  $str$  to  $v_{lab}$ ;
  - iii. set  $\text{DX}_1[str] = ctr$ ;
- (i) set  $vec_o := v_{lab}$ ;
- (j) prepend  $vec_n$  to  $vec_o$  and delete everything from  $vec_n$ ;
- (k) set  $\text{DX}_0 := \text{DX}_1$  and delete everything from  $\text{DX}_1$ ;
- (l) delete all labels from  $\mathbb{S}$ .

Now, we have to show that for PPT adversaries  $\mathcal{A}$ , the output of the real experiment and ideal experiment are indistinguishable. This can be shown by standard sequence of games argument that shows that EMM,  $utk$  and  $tk$  are indistinguishable from the real ones due to the RCPA security of SKE and the pseudo-randomness of  $F$ .

**Game<sub>0</sub>:** It is the same as a real experiment.

**Game<sub>1</sub>:** It is the same as **Game<sub>0</sub>** except that we replace the function  $F$  by a call to a random function  $G$ . This is indistinguishable because of the pseudo-randomness of  $F$ .

**Game<sub>2</sub>:** It is the same as **Game<sub>1</sub>** except that we do not generate any keys and replace encryption steps to simply producing a random string. During rebuild step for searched for labels, we simply remove encrypted values at random from the result set till the result set is of the leaked size. RCPA security of SKE guarantees indistinguishability between a ciphertext and a randomly generated string.

**Game<sub>3</sub>:** It is the same as **Game<sub>2</sub>** except that we replace the random function  $G$  and use random strings for labels. In setup, keep track of all labels (random strings) generated. When generating search tokens, we pick these labels at random for a fresh query as the number of labels originally in the old dictionary for the query is leaked. We further pick appropriate labels from both dictionaries as the total number of updates in each dictionary and when they were received is leaked. We assign these labels to this particular query for future repetitions. During rebuild step for unsearched for labels, we pick them one by one at random and then replace them with a new random string. This is the same as **Game<sub>2</sub>** as the output of  $G$  and a random string are indistinguishable.

**Game<sub>3</sub>** is the same as an ideal experiment.

This concludes our proof. ■

## B Proof of Theorem 5.1

**Theorem B.1.** *If SKE is an RCPA-secure encryption scheme and  $F$  is a pseudo-random function, then  $\text{DLS}^d$  is  $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U)$  secure.*

Consider the simulator that works as follows.

1. **Setup simulation.** The simulator takes as input the setup leakage  $\mathcal{L}_S(\text{MM}) = \sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$  and initializes dictionaries  $\text{DX}_0$  and  $\text{DX}_1$  of size  $\sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$ . It also initializes a state multi-map  $\text{MM}_0$  and three vectors  $vec_l, vec_o, vec_n$ . Then For  $1 \leq i \leq \sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$ :
  - (a) compute  $str_i, ctr_i \xleftarrow{\$} \{0, 1\}^k$  where  $str_i, ctr_i$  are of appropriate lengths;
  - (b) set  $\text{DX}_0[str_i] = ctr_i$ ;
  - (c) add  $str_i$  to  $vec_l$ .

It then outputs  $\text{EMM} = (\text{DX}_0, \text{DX}_1)$

2. **Update Simulation.** The simulator first simulates an update token  $utk$  and then outputs the rebuilt label,

value pairs for the adversary. To simulate the  $\text{utk}$ , it takes as input the update leakage which is equal to  $\mathcal{L}_U(\text{MM}, \ell, \mathbf{v}, \text{op}) = \#\mathbf{v}$  for all  $\text{op} \in \{\text{edit}^+, \text{edit}^-\}$ . To output the rebuilt pairs, it takes as input the public parameter  $\lambda$  and also when a rebuild gets completed. Then, for  $1 \leq v \leq \#\mathbf{v}$ , It does the following:

- (a)  $\text{str}_v, \text{value}_v \xleftarrow{\$} \{0, 1\}^k$  where  $\text{str}_v, \text{value}_v$  are of appropriate lengths;
- (b) add  $\text{str}_v$  to  $\text{vec}_n$ .

It outputs  $\text{utk} = (\text{str}_v, \text{value}_v)_{v \in \bar{v}}$ , and executes the following steps if the rebuild was completed:

- (a) delete  $\text{vec}_l$  and  $\text{MM}_0$ ;
- (b) set  $\text{vec}_o := \text{vec}_n$  and reset  $\text{vec}_n$  where  $\text{vec}_n$  had all the updates and re-inserts during current rebuild epoch.

If the rebuild was not complete, then for  $v \in [\lambda]$ ,

- (a)  $\text{str}_v, \text{value}_v \xleftarrow{\$} \{0, 1\}^k$  where  $\text{str}_v, \text{value}_v$  are of appropriate lengths
- (b) add  $\text{str}_v$  to  $\text{vec}_n$  and output  $(\text{str}_v, \text{value}_v)$ .

**3. Token Simulation.** The simulator takes as input the query leakage which is equal to  $\mathcal{L}_Q(\text{MM}, \ell) = (\text{QP}, \text{RL}, \text{OP})$  where QP is the search pattern, RL =  $\#\text{MM}[\ell]$  is the response length pattern and OP is the operation pattern which captures the update tokens and rebuild insertions of  $\ell$ . It first extracts from RL and OP the old and new response lengths,  $\#\text{DX}_o[\ell]$  and  $\#\text{DX}_n[\ell]$  respectively, such that  $\#\text{MM}[\ell] = \#\text{DX}_o[\ell] + \#\text{DX}_n[\ell]$ . It initializes empty vectors  $S_1$  and  $S_2$  and if the first rebuild hasn't been completed yet it executes the following steps:

- (a) if  $\text{MM}_0[\ell] = \perp$ :
  - i. set  $\text{count} = \#\text{DX}_o[\ell]$ ;
  - ii. from  $i = 1$  to  $\text{count}$ , pick and remove a label from  $\text{vec}_l$  and add it to  $S_1$ ;
  - iii. assign  $S_1$  to  $\text{MM}_0[\ell]$ .
- (b) Using OP, extract all labels that were updates of  $\ell$  from  $\text{vec}_n$ , and add them to  $S_2$ .

But if the first rebuild was completed, it simply does the following:

- (a) using OP, extract all labels that were updates or re-inserts of  $\ell$  from  $\text{vec}_o$  and add to  $S_1$ ;
- (b) using OP, extract all labels that were updates or re-inserts of  $\ell$  from  $\text{vec}_n$  and add to  $S_2$ .

It then outputs  $\text{otk} = S_1$  and  $\text{ntk} = S_2$ .

Now, we have to show that for all PPT adversaries  $\mathcal{A}$ , the output of the real experiment and ideal experiment are indistinguishable. This can be shown by standard sequence of games argument that shows that EMM,  $\text{utk}$  and  $\text{tk}$  are indistinguishable from the real ones due to

the RCPA security of SKE and the pseudo-randomness of  $F$ .

**Game<sub>0</sub>:** It is the same as a real experiment.

**Game<sub>1</sub>:** It is the same as **Game<sub>0</sub>** except that we replace the function  $F$  by a call to random function  $G$ . This is indistinguishable because of the pseudo-randomness of  $F$ .

**Game<sub>2</sub>:** It is the same as **Game<sub>1</sub>** except that we do not generate any keys and replace encryption steps to simply producing a random string. During rebuild step for a searched for label, we simply remove encrypted values at random from the result set till the result set is of the leaked size. RCPA security of SKE guarantees indistinguishability between a ciphertext and a randomly generated string.

**Game<sub>3</sub>:** We now get rid of  $G$  and use random strings for labels. In setup, keep track of all labels (random strings) generated. When generating search tokens, we pick these labels at random for a fresh query as the number of labels originally in the old dictionary for the query is leaked if the first rebuild hasn't been completed. We assign these labels to this particular query for future repetitions till the first rebuild is completed. If it has, these labels have already been re-inserted so using OP, we can retrieve them. We further pick appropriate labels from both dictionaries as the total number of updates and re-inserts in each dictionary and when they were received is also leaked. During rebuild step for unsearched labels, we pick  $\lambda$  labels one by one at random and then replace them with a new random string. This is the same as **Game<sub>2</sub>** because the output of  $G$  and a random string are indistinguishable.

**Game<sub>3</sub>:** It is the same as an ideal experiment. The rebuild step is now essentially equivalent to sending  $\lambda$  label, value pairs (which are both random strings now) whenever Update protocol is executed.

## C Proof of Theorem 4.2

**Theorem 4.2.** *If SKE is an RCPA-secure encryption scheme and  $F$  is a pseudo-random function, then DLS is  $(\mathcal{L}_{\text{SN}}, m, \ell)$ -multi-snapshot secure, for  $m, \ell \in \text{poly}(k)$ .*

*Proof.* Consider the simulator that works as follows.

1. The snapshot leakage is composed of the total number of label, value pairs in the old dictionary. So the simulator initializes  $\text{DX}_0$  and  $\text{DX}_1$  of the size  $\sum_{\ell \in \mathbb{L}_{\text{MM}}} \#\text{MM}[\ell]$ .

For  $1 \leq i \leq \sum_{\ell \in \mathbb{L}_{\text{MM}}} \# \text{MM}[\ell]$ , it performs the following steps:

- (a) compute  $str_i, ctr_i \xleftarrow{\$} \{0, 1\}^k$  where  $str_i, ctr_i$  are of appropriate lengths;
- (b) set  $\text{DX}_0[str_i] = ctr_i$ .

Then output  $\text{EMM} = (\text{DX}_0, \text{DX}_1)$  and set a variable  $size_1 = 0$  which captures the size of  $\text{DX}_1$ .

2. The snapshot leakage is composed of the size of the dictionaries ( $\# \text{DX}_o, \# \text{DX}_n$ ). The simulator then sets  $count$  to be  $\# \text{DX}_n - size_1$  where  $size_1$  captures the size of  $\text{DX}_n$  before the update. Then for  $1 \leq v \leq count$ , it does the following:

- (a) compute  $str, value \xleftarrow{\$} \{0, 1\}^k$  where  $str, value$  are of appropriate lengths;
- (b) set  $\text{DX}_1[str] = ctr$ .

It outputs  $\text{EMM} = (\text{DX}_0, \text{DX}_1)$  and sets  $size_1$  to  $\# \text{DX}_n$ .

3. The rebuild leakage for a snapshot adversary is, similarly, composed of the size of the dictionaries ( $\# \text{DX}_o, \# \text{DX}_n$ ) after the rebuild. The simulator sets  $count$  to be  $\# \text{DX}_o - size_1$  as the  $\text{DX}_n$  is now  $\text{DX}_o$ . Then for  $i$  from 1 to  $count$ , it does the following:

- (a) compute  $str, ctr \xleftarrow{\$} \{0, 1\}^k$  where  $str, ctr$  are of appropriate lengths;
- (b) set  $\text{DX}_1[str] = ctr$ .

It then sets  $size_1 := 0$ ,  $\text{DX}_0 := \text{DX}_1$ , deletes everything from  $\text{DX}_1$  and outputs  $\text{EMM} = (\text{DX}_0, \text{DX}_1)$ .

Now, we have to show that for PPT adversaries  $\mathcal{A}$ , the output of the real experiment and ideal experiment are indistinguishable. This can be shown by standard sequence of games argument that shows that snapshots of EMM after different protocols are indistinguishable from the real ones due to the RCPA security of SKE and the pseudo-randomness of  $F$ .

**Game<sub>0</sub>:** It is the same as a real experiment.

**Game<sub>1</sub>:** It is the same as **Game<sub>0</sub>** except that we replace the function  $F$  by a call to random function  $G$ . This is indistinguishable because of the pseudo-randomness of  $F$ .

**Game<sub>2</sub>:** It is the same as **Game<sub>1</sub>** except that we do not generate any keys and replace encryption steps to simply producing a random string. During rebuild step for searched for labels, we now use the snapshot leakage capturing the decrease in size ( $\#del$ ) of the old dictionary, and just create label, value pairs that are exactly  $\#del$  less than the original size in number. It does not matter if we remove from each individual result sets of the searched for labels accurately as long as the overall size decrease is satisfied.

RCPA security of SKE guarantees indistinguishability between a ciphertext and a randomly generated string.

**Game<sub>3</sub>:** It is the same as **Game<sub>2</sub>** except that we now remove the random function  $G$  and replace it with random strings for labels and during rebuild. We simply generate random label, value pairs that are equal in number to the new size of the old dictionary and add them to the new dictionary. This is the same as **Game<sub>2</sub>** because output of  $G$  and a random string are indistinguishable.

**Game<sub>3</sub>** is the same as an ideal experiment. ■

## D Proof of Theorem 5.2

**Theorem D.1.** *If SKE is an RCPA-secure encryption scheme and  $F$  is a pseudo-random function, then  $\text{DLS}^d$  is  $(m, \mathcal{L}_{\text{SN}})$ -snapshot secure, for  $m, \ell \in \text{poly}(k)$ .*

*Proof.* Consider the simulator that works as follows.

1. The snapshot leakage consists of the total number of label, value pairs in the old dictionary. The simulator takes this as input and initializes a dictionary  $\text{DX}_0$  of size  $\sum_{\ell \in \mathbb{L}_{\text{MM}}} \# \text{MM}[\ell]$  and an empty dictionary  $\text{DX}_1$ . Then for  $1 \leq i \leq \sum_{\ell \in \mathbb{L}_{\text{MM}}} \# \text{MM}[\ell]$ , it does the following:

- (a)  $str_i, ctr_i \xleftarrow{\$} \{0, 1\}^k$  where  $str_i, ctr_i$  are of appropriate lengths;
- (b) set  $\text{DX}_0[str_i] = ctr_i$ .

It then outputs  $\text{EMM} = (\text{DX}_0, \text{DX}_1)$  and sets  $size_1 = 0$  which captures the size of  $\text{DX}_1$ .

2. We can derive from the snapshot leakage the size of the old and new dictionaries,  $\# \text{DX}_o$  and  $\# \text{DX}_n$  respectively. The simulator takes this derived leakage as input and simulates **Update** as follows:

- (a) if  $\# \text{DX}_n$  is 0 (this is when a rebuild just got completed), it sets  $count$  to be  $\# \text{DX}_o - size_1$  and for  $1 \leq v \leq count$  it does the following,
  - i.  $str, value \xleftarrow{\$} \{0, 1\}^k$  where  $str, value$  is of appropriate lengths;
  - ii. set  $\text{DX}_1[str] = ctr$ .

If  $size_1 > 0$ , it then sets  $size_1 := 0$ ,  $\text{DX}_0 := \text{DX}_1$ , deletes everything from  $\text{DX}_1$ . It then outputs  $\text{EMM} = (\text{DX}_0, \text{DX}_1)$ .

- (b) Otherwise, it sets  $count$  to be  $\# \text{DX}_n - size_1$  and for  $1 \leq v \leq count$  it does the following,



- i.  $str, value \xleftarrow{\$} \{0,1\}^k$  where  $str, value$  is of appropriate lengths;
- ii. set  $DX_1[ctr] = ctr$ .  
It then outputs  $EMM = (DX_0, DX_1)$  and sets  $size_1$  to  $\#DX_n$ .

Now, we have to show that for all PPT adversaries  $\mathcal{A}$ , the output of the real experiment and ideal experiment are indistinguishable. This can be shown by standard sequence of games argument that shows that snapshots of EMM after different protocols are indistinguishable from the real ones due to the RCPA security of SKE and the pseudo-randomness of  $F$ .

**Game<sub>0</sub>**: It is the same as a real experiment.

**Game<sub>1</sub>**: It is the same as **Game<sub>0</sub>** except that we replace the function  $F$  by a call to random function  $G$ . This is indistinguishable because of the pseudo-randomness of  $F$ .

**Game<sub>2</sub>**: It is the same as **Game<sub>1</sub>** except that we do not generate any keys and replace encryption steps to simply producing a random string. During rebuild step for searched for labels, as  $\lambda$  is leaked we simply output  $\lambda$  pairs without caring about deletes. When the rebuild is complete, we would automatically send the the right number of pairs. RCPA security of SKE guarantees indistinguishability between a ciphertext and a randomly generated string.

**Game<sub>3</sub>**: We now get rid of  $G$  and use random strings for labels and during rebuild, we simply generate random label, value pairs that are equal in number to the new size of the old dictionary and add them to the new dictionary. This is the same as **Game<sub>2</sub>** because output of  $G$  and a random string are indistinguishable.

**Game<sub>3</sub>** is the same as an ideal experiment.

This concludes our proof. ■

## E Acknowledgements

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.