

# Backward Private DSSE: Alternative Formulations of Information Leakage and Efficient Constructions

Sanjit Chatterjee, Shravan Kumar Parshuram Puria, and Akash Shah

Department of Computer Science and Automation,  
Indian Institute of Science, Bangalore, India.  
E-mail: {sanjit, shravan, shaha}@iisc.ac.in

## Abstract

Dynamic Searchable Symmetric Encryption (DSSE), apart from providing support for search operation, allows a client to perform update operations on outsourced database efficiently. Two security properties, viz., forward privacy and backward privacy are desirable from a DSSE scheme. The former captures that the newly updated entries cannot be related to previous search queries and the latter ensures that search queries should not leak matching entries after they have been deleted. These security properties are formalized in terms of the information leakage that can be incurred by the respective constructions. Existing backward private constructions either have a non-optimal communication overhead or they make use of heavy cryptographic primitives. Our main contribution consists of three efficient backward private schemes that aim to achieve practical efficiency by using light weight symmetric cryptographic components only. In the process, we also revisit the existing definitions of information leakage for backward privacy [Bost et al. CCS'17] and propose alternative formulations. Our first construction  $\Pi_{\text{BP-prime}}$  achieves a stronger notion of backward privacy whereas our next two constructions  $\Pi_{\text{BP}}$  and  $\Pi_{\text{WBP}}$  achieve optimal communication complexity at the cost of some additional leakage. The prototype implementations of our schemes depict the practicability of the proposed constructions and indicate that the cost of achieving backward privacy over forward privacy is substantially small.

## 1 Introduction

Due to a variety of crucial benefits, enterprises outsource their data to cloud resident storage. If the outsourced data is stored as plaintext on remote servers then it may be intercepted by adversaries. Hence, data is stored in encrypted form on remote servers. However, if the client has to decrypt the data in order to get results for a search query, it defeats the purpose of outsourcing data. Generic tools such as fully homomorphic encryption [19] or oblivious RAM [22, 39] can be considered to construct protocols that leak almost no information to the server. But as of now, these tools are costly for large databases and hence, are impractical.

A practical solution to this problem is Searchable Symmetric Encryption (SSE) [36, 14, 13, 11] that trades efficiency for security. Dynamic Searchable Symmetric Encryption (DSSE) [27, 10, 33] adds a vital feature to static SSE schemes, i.e., the ability for the client to efficiently perform update operations remotely on the outsourced database with the guarantee that minimal information is leaked to the server in the process. These constructions of (D)SSE that aim to achieve an acceptable balance between security and performance, explicitly describe the leakage profile and formally prove that the information leaked from the scheme is bounded by the leakage profile.

Simultaneously with the works on constructing SSE schemes with improved efficiency, security and expressiveness of queries [11, 26, 16], there is another line of work that shows the real-world consequences of these leakages [9, 43, 2]. Zhang et al. [43] through *file injection attack* showed that it is possible to reveal the contents of past search queries of DSSE schemes with a few injection of documents and Abdelraheem et al. [2] showed that the consequences of this attack is even more devastating in the case of relational databases.

Because of the file injection attack, forward privacy has garnered significant interest in the research community. The notion of forward privacy was introduced in [38], while it was first formalized in [5]. In hindsight, the first DSSE scheme that satisfied the notion of forward privacy was proposed in 2005 [12]. Along with forward privacy, Stefanov et al. [38] asserted that backward privacy should also be satisfied by a DSSE scheme. Informally, backward privacy states that search queries should not leak matching entries after they have been deleted. The notion of backward privacy was first formalized by Bost et al. [7].

## 1.1 Related Work

Kamara et al. [27] proposed the first sublinear (in the size of the database) DSSE scheme. Forward private scheme  $\Sigma\phi\phi\zeta$  [5] achieves optimal communication complexity (linear in the size of result set) but it makes use of asymmetric cryptographic primitives and does not support parallel processing. A GGM-PRF [21] based forward private DSSE scheme Diana was proposed in [7]. Diana makes use of symmetric cryptographic primitives only and supports parallel processing but doesn't have optimal computational and communication complexity. Asymptotically optimal forward private DSSE schemes that make use of symmetric cryptographic primitives only and support parallelism were proposed in [29, 15, 37]. Further, FASTIO scheme [37] ensures a reasonable locality [10], a measure of I/O efficiency.

The notion of backward privacy was first formally described in [7], through three different security definitions, ordered from most to least secure called respectively BPIP, BPUP and WBP. Bost et al. [7] proposed a generic way to achieve backward privacy from any forward private DSSE scheme. However, the communication complexity of the derived backward private scheme isn't optimal. In [7], a backward private scheme  $Diana_{del}$  based on constrained pseudo-random function (CPRF) [8, 4, 28] and a backward private Janus framework based on a puncturable encryption scheme with a particular incremental update property [25], were also proposed. The communication and computational complexity of search and update protocols of  $Diana_{del}$  are not optimal. The search protocol in Janus is single-roundtrip and has an optimal communication complexity. However, the computational complexity of search protocol is  $O(n_w \cdot d_w)$ , where  $n_w$  and  $d_w$  respectively denote the number of documents matching keyword  $w$  and delete operations performed on keyword  $w$ . As acknowledged in [7], with just a few hundred deletions per keyword, Janus will not be practical because of both computational and storage overhead reasons. Further, [7] has imposed the following restriction on  $Diana_{del}$  and Janus: reinsertion of document-keyword pairs that were previously deleted is not allowed. We refer to this constraint as *reinsertion restriction* in the rest of our paper.

While a previous version of this work was under submission in a conference, there appeared two works on backward private DSSE [20, 40]. Chamani et al. [20] proposed a practically efficient BPUP-secure scheme called Mitra. They also proposed two other backward private schemes Orion and Horus that achieve quasi-optimal (linear in  $n_w$  upto a logarithmic factor) search computation complexity but make use of Path ORAM [39] and as a result are impractical for large databases [32]. Sun et al. [40] proposed a symmetric puncturable encryption (SPE) scheme and instantiated the Janus framework with the proposed SPE scheme. We provide a comprehensive comparative analysis of the performance and security of these schemes vis à vis our proposed schemes in Section 4.5.

## 1.2 Our Contributions

We start with revisiting the notion of information leakage in the context of backward privacy. In our investigation, we identify that *Weak Backward Privacy* (WBP) notion proposed by Bost et al. [7] should only be used to argue backward privacy in reinsertion restriction setting. We propose two alternative formulations of information leakage, viz., BP-I and BP-II to capture the notion of backward privacy. BP-I is strong in the sense that, for a search query on keyword  $w$ , it allows only the timestamp of the update queries on  $w$  and the identifiers of the documents containing  $w$  to be leaked. On the other hand, BP-II relaxes the definition for backward privacy by allowing some additional leakage.

Our main contribution consists of three backward private schemes  $\Pi_{BP\text{-prime}}$ ,  $\Pi_{BP}$  and  $\Pi_{WBP}$  that are BP-I, BP-II and WBP secure respectively. We start with a simple forward private scheme  $\Pi_{FP}$ , that serves as a building block for our backward private schemes. The construction  $\Pi_{BP\text{-prime}}$  is similar to the approach of

applying the generic transformation to achieve backward privacy [7] from forward private scheme  $\Pi_{\text{FP}}$ . Due to its use of only light weight symmetric components and low leakage level,  $\Pi_{\text{BP-prime}}$  can be a suitable candidate for adoption in practice. However, one limitation of  $\Pi_{\text{BP-prime}}$  is that the communication complexity isn't optimal. We address this issue in our main construction  $\Pi_{\text{BP}}$ .  $\Pi_{\text{BP}}$  makes use of tags generated using a pseudo random permutation (PRP) which ensures that only the tags corresponding to the set of documents currently matching the keyword  $w$  are returned to client. Thus,  $\Pi_{\text{BP}}$  avoids unnecessary communication overhead, while at the same time ensuring that any information violating the notion of backward privacy is not leaked to the server. To the best of our knowledge,  $\Pi_{\text{BP}}$  is the first practical backward private scheme that has optimal update and search communication complexity and uses symmetric cryptographic primitives only. Further,  $\Pi_{\text{BP}}$  is easily parallelizable, provides reasonable locality and allows reinsertion of document-keyword pair. With a simple modification in  $\Pi_{\text{BP}}$ , we construct a single roundtrip weak backward private scheme  $\Pi_{\text{WBP}}$  that improves upon the concrete communication overhead in Search protocol by around 40%. All our constructions are forward private as well. A comparison of our schemes with some prior and concurrent works [7, 20, 40] is provided in Table 1.

Schemes	Computation		Communication			Backward Privacy
	Search	Update	Search	Update	# Rounds	
Fides [7]	$O(o'_w)$	$O(1)$	$O(o'_w)$	$O(1)$	2	BPUP
Diana <sub>del</sub> [7]	$O(a_w)$	$O(\log(a_w))$	$O(n_w + d_w \log(a_w))$	$O(1)$	2	WBP
Janus [7]	$O(n_w \cdot d_w)$	$O(1)$	$O(n_w)$	$O(1)$	1	WBP
Mitra [20]	$O(o'_w)$	$O(1)$	$O(o'_w)$	$O(1)$	2	BPUP
Orion [20]	$O(n_w \log^2(N))$	$O(\log^2(N))$	$O(n_w \log^2(N))$	$O(\log^2(N))$	$O(\log(N))$	BPIP
Horus [20]	$O(n_w \log(d_w) \log(N))$	$O(\log^2(N))$	$O(n_w \log(d_w) \log(N))$	$O(\log^2(N))$	$O(\log(d_w))$	WBP
Janus++ [40]	$O(n_w \cdot d)$	$O(d)$	$O(n_w)$	$O(1)$	1	WBP
$\Pi_{\text{BP-prime}}$ (Our Work)	$O(o'_w)$	$O(1)$	$O(o'_w)$	$O(1)$	2	BP-I
$\Pi_{\text{BP}}$ (Our Work)	$O(o'_w)$	$O(1)$	$O(n_w)$	$O(1)$	2	BP-II
$\Pi_{\text{WBP}}$ (Our Work)	$O(o'_w)$	$O(1)$	$O(n_w)$	$O(1)$	1	WBP

All the constructions are also forward private. The client storage for all the constructions is  $O(m \log(n))$  except Orion, where the corresponding complexity is  $O(1)$ . Refer Section 2.4 for the notations used.

Generally,  $o'_w < n_w \cdot d_w$ , as the former term has an additive factor whereas the latter term has a multiplicative factor.

Table 1: Comparison of backward schemes  $\Pi_{\text{BP-prime}}$ ,  $\Pi_{\text{BP}}$  and  $\Pi_{\text{WBP}}$  with some prior and concurrent works.

$\Pi_{\text{FP}}$  is the most efficient forward private scheme in literature. Our implementation results show that the performance of  $\Pi_{\text{BP}}$  is comparable to  $\Pi_{\text{FP}}$ . For example, the time taken by the search protocol of  $\Pi_{\text{BP}}$  for a search that returned 150,000 results is around 0.68 seconds as compared to 0.54 seconds in  $\Pi_{\text{FP}}$ .

We also introduce a desirable property for DSSE schemes called *inverse backward privacy*. This property ensures that a search query should leak no information about the identifier of the document that had been deleted and re-inserted later.

## 2 Notations and Definitions

The security parameter is denoted by  $\lambda$ . All procedures in our construction implicitly take  $\lambda$  as input. By efficient, we mean probabilistic polynomial-time in  $\lambda$ . All the algorithms (including adversaries and simulators) are assumed to be efficient unless otherwise specified. A function  $f: \mathbb{N} \rightarrow \mathbb{R}$  is said to be a negligible function iff for all  $c > 0$ ,  $\exists n_0 \in \mathbb{N}$  such that  $\forall n \geq n_0$ ,  $f(n) < n^{-c}$ . The function  $\text{neg}(\lambda)$  denotes a negligible function in  $\lambda$ . For a finite set  $X$ ,  $x \xleftarrow{\$} X$  means that  $x$  is uniformly sampled from  $X$  and  $|X|$  denotes the cardinality of set  $X$ .  $x \leftarrow y$  denotes that variable  $x$  is assigned the value of variable  $y$  and operator  $\parallel$  denotes concatenation. For a data structure DS,  $|\text{DS}|$  denotes the memory space occupied by the data structure in bits.  $\text{addr}(\text{D})$  denotes the address in memory at which data structure instance D is stored.  $\perp$  denotes null value. For sets  $X_1, \dots, X_n$  and  $Y$ ,  $\text{Func}(X_1 \times \dots \times X_n, Y)$  denotes the set of all functions

$\underline{\mathbf{Real}_A^{\text{prf}}(\lambda)}$ <ol style="list-style-type: none"> <li>1. <math>K \xleftarrow{\\$} \mathcal{K}</math></li> <li>2. <math>b \leftarrow \mathcal{A}^{F(K, \cdot)}</math></li> <li>3. return b</li> </ol>	$\underline{\mathbf{Ideal}_A^{\text{prf}}(\lambda)}$ <ol style="list-style-type: none"> <li>1. <math>f \xleftarrow{\\$} \text{Func}(\mathcal{I}, \mathcal{O})</math></li> <li>2. <math>b \leftarrow \mathcal{A}^{f(\cdot)}</math></li> <li>3. return b</li> </ol>
---	---

Figure 1: PRF Security Definition

$\underline{\mathbf{Real}_A^{\text{prp}}(\lambda)}$ <ol style="list-style-type: none"> <li>1. <math>K \xleftarrow{\\$} \mathcal{K}</math></li> <li>2. <math>b \leftarrow \mathcal{A}^{F(K, \cdot)}</math></li> <li>3. return b</li> </ol>	$\underline{\mathbf{Ideal}_A^{\text{prp}}(\lambda)}$ <ol style="list-style-type: none"> <li>1. <math>f \xleftarrow{\\$} \text{Perm}(X)</math></li> <li>2. <math>b \leftarrow \mathcal{A}^{f(\cdot)}</math></li> <li>3. return b</li> </ol>
---	--

Figure 2: PRP Security Definition

from  $X_1 \times \dots \times X_n$  to  $Y$ . For set  $X$ ,  $\text{Perm}(X)$  denotes the set of all permutations on  $X$ . For sets  $X_1$  and  $X_2$ ,  $\text{Perm}_{X_1}(X_2)$  denotes the set of all functions from  $X_1 \times X_2$  to  $X_2$ , where for every  $x \in X_1$ , we have a permutation on  $X_2$ .

We use pseudo random functions (PRF), pseudo random permutations (PRP) and RCPA secure symmetric key encryption schemes in our constructions.

## 2.1 Pseudorandom Function

Pseudorandom Function (PRF)  $F$  is polynomial-time computable in  $\lambda$  and is indistinguishable from a truly random function by any adversary  $\mathcal{A}$ .

**Definition 2.1.** Let  $F \in \text{Func}(\mathcal{K} \times \mathcal{I}, \mathcal{O})$  be an efficient, keyed function. For algorithms  $\mathcal{A}$ , we define the experiments  $\mathbf{Real}_A^{\text{prf}}(\lambda)$  and  $\mathbf{Ideal}_A^{\text{prf}}(\lambda)$  as shown in Figure 1.

$F$  is a pseudorandom function if for all probabilistic polynomial-time adversaries  $\mathcal{A}$ ,  $\mathbf{Adv}_{F, \mathcal{A}}^{\text{prf}}(\lambda) = |\Pr[\mathbf{Real}_A^{\text{prf}}(\lambda) = 1] - \Pr[\mathbf{Ideal}_A^{\text{prf}}(\lambda) = 1]| \leq \text{neg}(\lambda)$ .

## 2.2 Pseudorandom Permutation

Pseudorandom permutation (PRP)  $F$  is polynomial-time computable and invertible in  $\lambda$  and is indistinguishable from a truly random permutation by any adversary  $\mathcal{A}$ .

**Definition 2.2.** Let  $F \in \text{Perm}_{\mathcal{K}}(\mathcal{X})$  be an efficient, keyed function. For algorithms  $\mathcal{A}$ , we define the experiments  $\mathbf{Real}_A^{\text{prp}}(\lambda)$  and  $\mathbf{Ideal}_A^{\text{prp}}(\lambda)$  as shown in Figure 2.

$F$  is a pseudorandom permutation if for all probabilistic polynomial-time adversaries  $\mathcal{A}$ ,  $\mathbf{Adv}_{F, \mathcal{A}}^{\text{prp}}(\lambda) = |\Pr[\mathbf{Real}_A^{\text{prp}}(\lambda) = 1] - \Pr[\mathbf{Ideal}_A^{\text{prp}}(\lambda) = 1]| \leq \text{neg}(\lambda)$ .

## 2.3 Symmetric Key Encryption Scheme

A symmetric key encryption scheme  $\mathcal{E}$  consists of three algorithms: **Gen**, **Enc** and **Dec**.

- **Gen**( $\cdot$ ): It outputs a key  $k$ .
- **Enc**( $k, m$ ): It takes as input the key  $k$ , and a message  $m \in \mathcal{M}$ , where  $\mathcal{M}$  is the message space, and outputs a ciphertext  $e$ .

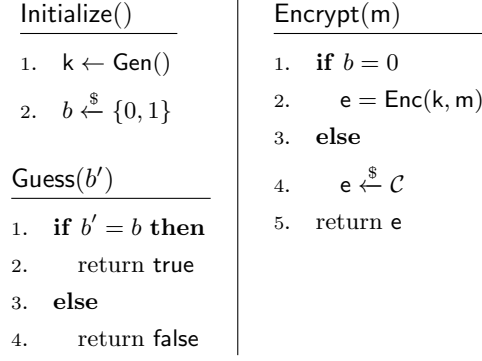


Figure 3: Security Game:  $\text{RCPA}_{\mathcal{E}}$

- $\text{Dec}(k, e)$ : It takes as input the key  $k$  and a ciphertext  $e$  and outputs a message  $m \in \mathcal{M}$  or  $\perp$ .

**Correctness:** For all  $k \leftarrow \text{Gen}()$  and for all messages  $m \in \mathcal{M}$ , it is required that  $\text{Dec}(k, \text{Enc}(k, m)) = m$ .

**RCPA Security Notion of Symmetric Key Encryption Scheme** The pseudorandom ciphertexts under chosen plaintext attack (RCPA) security notion [10] for symmetric key encryption scheme  $\mathcal{E} = (\text{Gen}, \text{Enc}, \text{Dec})$  is captured in Figure 3. Initialize() algorithm generates the key  $k$  using Gen algorithm of  $\mathcal{E}$  and picks a random challenge bit  $b$ . The adversary can then adaptively ask queries to Encrypt() oracle. The game returns true if the adversary's output  $b'$  equals the challenge bit  $b$ . In Figure 3,  $\mathcal{C}$  denotes the ciphertext space.

**Definition 2.3.** A symmetric key encryption scheme  $\mathcal{E} = (\text{Gen}, \text{Enc}, \text{Dec})$  is said to have pseudorandom ciphertexts under chosen plaintext attack (RCPA) if no adversary  $\mathcal{A}$  can win the game shown in Figure 3, except with probability at most  $\frac{1}{2} + \text{neg}(\lambda)$ . This probability is denoted by  $\Pr[\text{RCPA}_{\mathcal{E}} = 1]$ .

The advantage of  $\mathcal{A}$  in  $\text{RCPA}_{\mathcal{E}}$  game is defined as follows:  $\text{Adv}_{\mathcal{E}, \mathcal{A}}^{\text{RCPA}}(\lambda) = 2 \cdot \Pr[\text{RCPA}_{\mathcal{E}} = 1] - 1$ .

## 2.4 Dynamic Searchable Symmetric Encryption (DSSE)

Our system consists of two parties: the client  $C$  (data owner) and the server  $S$ .  $C$ , who owns the database  $DB$ , encrypts the database and outsources it to  $S$ . The encrypted copy of the database created ensures that  $S$  responds to  $C$ 's queries (search and update) in an efficient manner with the guarantee that minimal information apart from the intended output of the query operation is leaked to the server.

We primarily follow the formalization of Bost et al. [7] with certain additions. A keyword is denoted by  $w$  and a document is addressed by its document identifier  $\text{ind}$ . The database  $DB$  can be represented as:  $DB = \{(\text{ind}_i, W_i) : 1 \leq i \leq n\}$ , where  $n$  denotes the number of documents in the database,  $\text{ind}_i \in \{0, 1\}^{\ell}$  are distinct document indices and  $W_i \subseteq \{0, 1\}^*$  is a set of keywords matching document  $\text{ind}_i$ , represented by binary strings of arbitrary length. Additionally, we consider the following notations:

$$\begin{aligned}
 W &= \cup_{i=1}^n W_i, \text{ the set of keywords,} \\
 m &= |W|, \# \text{ keywords,} \\
 N &= \sum_{i=1}^n |W_i|, \# \text{ document-keyword pair,} \\
 DB(w) &= \{\text{ind}_i : w \in W_i\}, \text{ the set of documents containing } w, \\
 n_w &= |DB(w)|, \# \text{ documents containing } w \\
 a_w &, \# \text{ add operations performed on } w \\
 d_w &, \# \text{ del operations performed on } w
 \end{aligned}$$

- $o_w$ , # updates performed on  $w$
- $n'_w$ , # documents containing  $w$  in previous<sup>1</sup> search operation
- $a'_w$ , # add operations performed on  $w$  after the previous search
- $d'_w$ , # del operations performed on  $w$  after the previous search
- $o'_w = n'_w + a'_w + d'_w$ .

A DSSE scheme  $\Pi$  comprises of the following [7]:

- **Setup(DB)** is a probabilistic algorithm that takes as input the initial database DB. It outputs  $(st_C, EDB)$ , where the client's state  $st_C$  is given to C and the encrypted database EDB is given to S. Setup algorithm is executed by C.
- **Search( $q, st_C; EDB$ )**=(Search<sub>C</sub>( $q, st_C$ ), Search<sub>S</sub>(EDB)) is a protocol (possibly probabilistic) between C and S. The input of C is the search query  $q$  and client's state  $st_C$ . The input of S is the encrypted database EDB. The output to the client is the updated client's state  $st_C'$  and the set Res comprising of indices matching the search query  $q$ . The output to the server is the updated encrypted database EDB'.
- **Update( $q, st_C; EDB$ )**= (Update<sub>C</sub>( $q, st_C$ ), Update<sub>S</sub>(EDB)) is a protocol (possibly probabilistic) between C and S. The input of C is the query  $q=(op, in)$  comprising of update operation  $op \in \{add, del\}$  and the document-keyword pairs  $(w, ind)$  denoted by in, and client's state  $st_C$ . The input of S is the encrypted database EDB. The output to the client is the updated client's state  $st_C'$ . The output to the server is the updated encrypted database EDB'.

In this work, we consider the case of single keyword search, so,  $q=w$  and  $Res=DB(w)$  in Search protocol. For simplicity, we consider  $in=(ind, w)$ , i.e., a single document-keyword pair in Update protocol. Bulk-Updates can be supported by calling the Update protocol repeatedly. Similar to [11], for a given query  $q$ , we consider the output of Search protocol to be the set of document identifiers satisfying  $q$ . This allows us to decouple the storage of documents from the storage of data structures used to realize the search operation, which is the focus of this work. SSE schemes are of two types: *response-revealing* and *response-hiding* [26]. The former reveals the query response in plaintext whereas the latter does not. We use this categorization in our paper.

**Security** We consider S to be honest-but-curious. The security definition follows the real/ideal simulation paradigm [14, 11]. The definition is parameterized by a leakage profile  $\mathcal{L}$  which captures all the information that the adversary learns about the database and queries through its participation in the protocols. Hence, the view of the adversary in the real world can be simulated by  $\mathcal{L}$ .

$\mathcal{L} = \{\mathcal{L}_{Setup}, \mathcal{L}_{Search}, \mathcal{L}_{Update}\}$ , where  $\mathcal{L}_{Setup}$ ,  $\mathcal{L}_{Search}$  and  $\mathcal{L}_{Update}$  correspond to the information leaked in the Setup, Search and Update protocols respectively to the server.

**Definition 2.4.** Let  $\Pi = (Setup, Search, Update)$  be a DSSE scheme and let  $\mathcal{L} = \{\mathcal{L}_{Setup}, \mathcal{L}_{Search}, \mathcal{L}_{Update}\}$  be a stateful algorithm. For algorithms  $\mathcal{A}$  and Sim, we define the experiments  $\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda)$  and  $\mathbf{Ideal}_{\mathcal{A}, Sim}^{\Pi}(\lambda)$  as follows:

**Real $_{\mathcal{A}}^{\Pi}(\lambda)$**   $\mathcal{A}(1^\lambda)$  chooses DB. The experiment then runs  $(st_C, EDB) \leftarrow Setup(DB)$  and gives EDB to  $\mathcal{A}$ . Then  $\mathcal{A}$  makes polynomial number of adaptive queries. For each query  $q$ , if  $q$  is a search query (resp. update query), the game runs  $(Res, st_C, EDB) \leftarrow Search(q, st_C; EDB)$  (resp.  $(st_C, EDB) \leftarrow Update(q, st_C; EDB)$ ) and gives the generated transcript to  $\mathcal{A}$ . Eventually  $\mathcal{A}$  returns a bit that the game uses as its own output.

**Ideal $_{\mathcal{A}, Sim}^{\Pi}(\lambda)$**  The game initializes a counter  $i=0$  and an empty list  $\mathbf{q}$ .  $\mathcal{A}(1^\lambda)$  chooses DB. The experiment then runs  $EDB \leftarrow Sim(\mathcal{L}_{Setup}(DB))$  and gives EDB to  $\mathcal{A}$ . Then  $\mathcal{A}$  makes polynomial number of adaptive queries. For each query  $q$ , the game records this as  $\mathbf{q}[i]$ , increments  $i$  and if  $q$  is a search query (resp. update

<sup>1</sup>Throughout the paper by previous search on  $w$ , we mean the last search operation on  $w$  before the current search on  $w$ .

<p><b>DSSECor<sub>A</sub><sup>Π</sup>(λ)</b></p> <hr style="border: 0.5px solid black;"/> <ol style="list-style-type: none"> <li>1. <b>flag</b> = false</li> <li>2. (DB, st<sub>A</sub>) ← <math>\mathcal{A}(1^\lambda)</math></li> <li>3. (st<sub>C0</sub>, EDB<sub>0</sub>) ← Setup(DB)</li> <li>4. <b>for</b> <math>1 \leq i \leq p(\lambda)</math> <b>do</b></li> <li>5.   (q<sub>i</sub>, st<sub>A</sub>) ← <math>\mathcal{A}(\text{st}_A, \text{EDB}_0, T_{i-1})</math></li> <li>6.   <b>if</b> q<sub>i</sub> is a search query</li> <li>7.     Res<sub>i</sub>, EDB<sub>i</sub>, st<sub>C<sub>i</sub></sub> ← Search(q<sub>i</sub>, st<sub>C<sub>i-1</sub></sub>; EDB<sub>i-1</sub>)</li> </ol>	<ol style="list-style-type: none"> <li>8.     τ<sub>i</sub> ← Transcript[Search(q<sub>i</sub>, st<sub>C<sub>i-1</sub></sub>; EDB<sub>i-1</sub>)]</li> <li>9.     <b>if</b> Res<sub>i</sub> ≠ DB(w<sub>i</sub>) // q<sub>i</sub> = w<sub>i</sub></li> <li>10.      <b>flag</b> = true</li> <li>11.    <b>else</b></li> <li>12.      EDB<sub>i</sub>, st<sub>C<sub>i</sub></sub> ← Update(q<sub>i</sub>, st<sub>C<sub>i-1</sub></sub>; EDB<sub>i-1</sub>)</li> <li>13.      τ<sub>i</sub> ← Transcript[Update(q<sub>i</sub>, st<sub>C<sub>i-1</sub></sub>; EDB<sub>i-1</sub>)]</li> <li>14. <b>return flag</b></li> </ol>
--	---

Figure 4: DSSE Correctness

query), the game gives transcript generated by  $\text{Sim}(\mathcal{L}_{\text{Search}}(\text{DB}, \mathbf{q}))$  (resp.  $\text{Sim}(\mathcal{L}_{\text{Update}}(\text{DB}, \mathbf{q}))$ ). Eventually  $\mathcal{A}$  returns a bit that the game uses as its own output.

We say that  $\Pi$  is  $\mathcal{L}$ -semantically secure against adaptive attacks if for all adversaries  $\mathcal{A}$ , there exists an algorithm  $\text{Sim}$  such that  $|\Pr[\mathbf{Real}_{\mathcal{A}}^{\Pi}(\lambda) = 1] - \Pr[\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\Pi}(\lambda) = 1]| \leq \text{neg}(\lambda)$ .

**Common Leakage** We follow some of the notations of common leakages from [5, 7]. The leakage profile  $\mathcal{L}$  keeps as state the query list  $\mathbf{Q}$ , i.e., the list of all queries issued so far along with their timestamp. The entries in  $\mathbf{Q}$  are  $(u, w)$  for a search query on  $w$ , or  $(u, \text{op}, (\text{ind}, w))$  for an update query  $(\text{op}, (\text{ind}, w))$ , where  $u$  denotes the timestamp of the query. Corresponding to the search queries, the search pattern  $\text{sp}(w)$  can be defined as  $\text{sp}(w) = \{u : (u, w) \in \mathbf{Q}\}$ .

We also use the notation  $\text{Hist}(w)$  that denotes the list of all the modifications made to  $\text{DB}(w)$  over the time. It consists of  $\text{DB}_0(w)$ , the set of document indices matching  $w$  at setup, and a list  $\text{UpHist}(w)$ , comprising of the updates of documents matching  $w$ , called the *update history*. For example, consider two documents 1 and 2 matching  $w$ . Suppose, the update queries are  $(\text{add}, (1, w))$ ,  $(\text{add}, (2, w))$  and  $(\text{del}, (1, w))$  at timestamp 3, 12 and 20 respectively, then  $\text{UpHist}(w) = [(3, \text{add}, 1), (12, \text{add}, 2), (20, \text{del}, 1)]$ .

$\text{Updates}(w)$  denotes the set of timestamps of updates on  $w$ . Formally,

$$\text{Updates}(w) = \{u | (u, \text{add}, (\text{ind}, w)) \in \mathbf{Q} \text{ or } (u, \text{del}, (\text{ind}, w)) \in \mathbf{Q}\}.$$

$\text{Updates}^{\text{op}}(w)$  is exactly like  $\text{Updates}(w)$  except along with the timestamp it also stores  $\text{op}$  corresponding to the update query.

**Correctness** We say that a DSSE scheme is correct if the Search protocol returns the correct results for the keyword being searched (i.e.,  $\text{DB}(w)$ ), except with negligible probability. We follow a similar formalization to [10].

In Figure 4, the adversary  $\mathcal{A}$  makes  $p(\lambda)$  many queries, for some polynomial  $p$ .  $\text{Transcript}[\text{Protocol}]$  means the view of server in Protocol.  $T_0 = \emptyset$  and  $T_i = \{\tau_j : 1 \leq j \leq i\}$ ,  $1 \leq i \leq p(\lambda)$ . Here,  $\tau_j$  denotes the transcript of the  $j^{\text{th}}$  query.

**Definition 2.5.** Let  $\Pi = (\text{Setup}, \text{Search}, \text{Update})$  be a DSSE scheme. For algorithm  $\mathcal{A}$  we define the experiment  $\text{DSSECor}_{\mathcal{A}}^{\Pi}(\lambda)$  as shown in Figure 4.

We say that  $\Pi$  is correct if for all adversaries  $\mathcal{A}$ ,  $\Pr[\text{DSSECor}_{\mathcal{A}}^{\Pi}(\lambda) = 1] \leq \text{neg}(\lambda)$ .

### 3 Security Notions in DSSE

In this section, we perform a critical analysis of the existing notions of information leakage in backward private DSSE followed by some alternative formulations. The question of what is the *right* definition of

security is a vexed one [30]. Even for a widely used cryptographic primitive like digital signature, it has been argued that the accepted standard definition of security [24, 23] does not take into account various issues that may crop up depending upon the application scenarios [31, 35, 42]. It is but natural that for a relatively new crypto/security protocol like DSSE, one needs to look at the definitional question from various perspectives. Apart from giving better insight about the real world security assurance of the protocol, this exercise also helps to choose among alternative definitions to achieve some sort of *optimal* balance between security and efficiency for the task at hand.

Several works in the context of searchable encryption [10, 11, 13, 14, 26, 27], argue security of SSE schemes by formulating a leakage profile  $\mathcal{L}$  and proving that the leakage incurred in the proposed scheme is bounded by  $\mathcal{L}$ . However, works in the context of DSSE [5, 7] present the formulated information leakages corresponding to forward/backward privacy notion in the form of security definitions. Here, we follow the latter approach.

Bost et al. [7] made a seminal contribution in the area of DSSE by formalizing the notion of backward privacy through three different security definitions viz., BPIP, BPUP and WBP, ordered from most to least secure based on their respective information leakages. Naturally, like any other formalization of security this requires further investigation. In that vein, we first argue why *weak backward privacy* (WBP) cannot serve as a general definition for backward privacy and should be used to argue backward privacy in re-insertion restriction setting only. Second, we introduce additional definitions (Definition 3.5 and Definition 3.6) for backward privacy in terms of information leakage. Finally, we introduce a desirable property for DSSE scheme which we call *inverse backward privacy*.

For simplicity we will assume that, DB is initially empty. Thus, the Setup algorithm leaks no information. If DB is not initially empty, typically,  $\mathcal{L}_{\text{Setup}} = N$ , where  $N$  denotes the number of document-keyword pairs.

### 3.1 Forward Privacy

We recall the strongest notion of forward privacy discussed in [7]. Informally, a DSSE scheme is forward private if the Update protocol leaks no information about the updated keywords. Definition 3.1 captures that an update operation doesn't leak more than the operation  $\text{op} = \{\text{add}, \text{del}\}$  of the update query  $q$ .

**Definition 3.1.** (FP-I). An  $\mathcal{L} = \{\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}\}$ -semantically secure against adaptive attacks DSSE scheme  $\Pi$  is FP-I iff  $\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}$  can be written as:

$$\begin{aligned} \mathcal{L}_{\text{Setup}}() &= \emptyset. & \mathcal{L}_{\text{Search}}(w) &= \{\text{sp}(w), \text{Hist}(w)\}. \\ \mathcal{L}_{\text{Update}}(\text{op}, (\text{ind}, w)) &= \text{op}. \end{aligned}$$

### 3.2 Backward Privacy

The notion of backward privacy was informally introduced in [38] as “queries cannot be executed over the deleted documents”. The follow up work [7] stated that a DSSE scheme is backward private if, whenever a document-keyword pair  $(\text{ind}, w)$  is added to the database and later deleted, subsequent searches on  $w$  should not reveal  $\text{ind}$  if it cannot be inferred from the search and access pattern. Based on the above notion, we now elaborate our interpretation of backward privacy. Let  $q_{\text{prev}}$  be the previous search query on  $w$  and let  $\text{ind} \notin \text{DB}(w)'$  (here,  $\text{DB}(w)'$  denotes the set of documents matching  $w$  in the previous search on  $w$ ). If the last update operation on  $(\text{ind}, w)$  before the next search query  $q_{\text{cur}}$  is del operation, then  $\text{ind}$  should not be revealed to  $S$  as it cannot be inferred from the results of current and previous search queries. Let the set of such inds be denoted by  $I_{\text{BP}}$ . For example, let  $\text{DB}(w)' = \{1, 2\}$  for search query  $q_{\text{prev}}$  at timestamp 5. Let the update operations on  $w$  after query  $q_{\text{prev}}$  and before the next search query  $q_{\text{cur}}$  be  $(6, \text{add}, (3, w)), (12, \text{del}, (3, w))$ , then identifier 3 should not be revealed to  $S$ .

Backward privacy was formalized in [7], through three different security definitions, ordered from most to least secure called respectively BPIP, BPUP and WBP. Informally, these definitions are described below.

- *Backward Privacy with insertion pattern* (BPIP): During a search on some keyword  $w$ , BPIP schemes



leak the documents currently matching  $w$ , when they were inserted, and the total number of updates on  $w$ .

- *Backward Privacy with update pattern* (BPUP): During a search on  $w$ , BPUP schemes leak the documents currently matching  $w$ , when they were inserted, and when all the updates on  $w$  happened (but not their contents).
- *Weak backward privacy* (WBP): During a search on  $w$ , WBP schemes leak the documents currently matching  $w$ , when they were inserted, when all the updates on  $w$  happened, and which deletion update canceled which insertion update.

Let us demonstrate the differences between these notions with an example. Consider the following entries in query list  $\mathbf{Q}_1$  corresponding to  $w$ :  $(1, \text{add}, (\text{ind}_1, w))$ ,  $(4, \text{add}, (\text{ind}_2, w))$ ,  $(5, \text{del}, (\text{ind}_1, w))$ ,  $(12, \text{add}, (\text{ind}_3, w))$ . Let us consider the leakage for each definition after a search query on  $w$  at timestamp 15. The first notion reveals that  $\text{ind}_2$  and  $\text{ind}_3$  match keyword  $w$  and that this entries were added at time 4 and 12 respectively. It also reveals that there were a total of 4 updates for  $w$ . The second notion, additionally reveals that updates on  $w$  happened at time 1, 4, 5 and 12. Finally, the third definition also reveals that the index that was added for  $w$  at time 1 was deleted at time 5.

Let us recall the additional leakage functions from [7, 6] apart from the leakage functions described in Section 2.4, required to formally capture the notions of backward privacy mentioned above.

For a keyword  $w$ ,  $\text{TimeDB}(w)$  is the list of all documents matching  $w$ , excluding the deleted ones together with the timestamp of when they were inserted in the database. Formally,  $\text{TimeDB}(w)$  can be constructed from the query list  $\mathbf{Q}$  as follows:

$$\begin{aligned} \text{TimeDB}(w) = \{ & (u, \text{ind}) \mid (u, \text{add}, (w, \text{ind})) \in \mathbf{Q} \text{ and} \\ & \forall u' > u, (u', \text{del}, (w, \text{ind})) \notin \mathbf{Q} \}. \end{aligned} \quad (1)$$

The deletion history  $\text{DelHist}(w)$  of  $w$  is the list of timestamps for all deletion operations, together with the timestamp of the inserted entry it removes. Formally,  $\text{DelHist}(w)$  is constructed as:

$$\begin{aligned} \text{DelHist}(w) = \{ & (u^{\text{add}}, u^{\text{del}}) \mid \exists \text{ind s.t. } (u^{\text{add}}, \text{add}, (w, \text{ind})) \in \mathbf{Q} \text{ and} \\ & (u^{\text{del}}, \text{del}, (w, \text{ind})) \in \mathbf{Q} \}. \end{aligned} \quad (2)$$

With these tools, we can now formally define these three notions of backward privacy formally.

**Definition 3.2.** (BPIP). An  $\mathcal{L} = \{\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}\}$ -semantically secure against adaptive attacks DSSE scheme  $\Pi$  is BPIP iff  $\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}$  can be written as:

$$\begin{aligned} \mathcal{L}_{\text{Setup}}() &= \emptyset. \quad \mathcal{L}_{\text{Update}}(\text{op}, (\text{ind}, w)) = \{\text{op}\}. \\ \mathcal{L}_{\text{Search}}(w) &= \{\text{TimeDB}(w), o_w\}. \end{aligned}$$

**Definition 3.3.** (BPUP). An  $\mathcal{L} = \{\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}\}$ -semantically secure against adaptive attacks DSSE scheme  $\Pi$  is BPUP iff  $\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}$  can be written as:

$$\begin{aligned} \mathcal{L}_{\text{Setup}}() &= \emptyset. \quad \mathcal{L}_{\text{Update}}(\text{op}, (\text{ind}, w)) = \{\text{op}, w\}. \\ \mathcal{L}_{\text{Search}}(w) &= \{\text{TimeDB}(w), \text{Updates}(w)\}. \end{aligned}$$

**Definition 3.4.** (WBP). An  $\mathcal{L} = \{\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}\}$ -semantically secure against adaptive attacks DSSE scheme  $\Pi$  is WBP iff  $\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}$  can be written as:

$$\begin{aligned} \mathcal{L}_{\text{Setup}}() &= \emptyset. \quad \mathcal{L}_{\text{Update}}(\text{op}, (\text{ind}, w)) = \{\text{op}, w\}. \\ \mathcal{L}_{\text{Search}}(w) &= \{\text{TimeDB}(w), \text{DelHist}(w)\}. \end{aligned}$$

The above definitions indicate that the notion of backward privacy is more involved than that of forward privacy. Forward privacy can be formalized by ensuring that the update query doesn't leak the keyword corresponding to which the update has been made. Whereas it is more subtle in the case of backward privacy. In formulating the leakage profile for backward privacy, one must ensure that no leakage is incurred w.r.t. document identifiers in  $I_{BP}$  described above. One approach is to formulate a strong leakage profile for the security notion in hand that allows very limited information leakage to be incurred by the constructions satisfying it. This approach was followed in formulating leakage profiles in definitions BPIP and BPUP. A possible shortcoming of this approach is that there could be candidate constructions that don't satisfy these strong definitions but may still satisfy the intuitive notion of backward privacy. The other extreme could be to allow 'as much information as one can think of' to be leaked that can be allowed by the security notion in hand. This seems to be the approach followed in WBP security notion. However, in this approach one must be careful not to violate the basic security notion of the corresponding task while allowing for more information leakage. We revisit the notion of WBP from this perspective.

**Remark**  $\mathcal{L}_{Search}$  in BPIP, BPUP and WBP should be augmented with leakage function  $sp(w)$ . Moreover,  $\mathcal{L}_{Search}$  in WBP should also be augmented with leakage function  $Updates(w)$ . The rationale behind the same is described in Section 3.4.

### 3.3 Revisiting Weak Backward Privacy

In this section, we scrutinize the notion of weak backward privacy. Let us consider the following entries in query list  $\mathbf{Q}_1$  corresponding to  $w$ :  $(1, \text{add}, (\text{ind}, w))$ ,  $(3, w)$ ,  $(5, \text{del}, (\text{ind}, w))$ ,  $(6, w)$ ,  $(12, \text{add}, (\text{ind}, w))$ ,  $(14, \text{del}, (\text{ind}, w))$ ,  $(18, w)$ . Let us denote the search queries on  $w$  at timestamps 3, 6 and 18 by  $q_1$ ,  $q_2$  and  $q_3$  respectively.

Leakage of the search query  $q_1$ :

$$\text{TimeDB}(w) = \{(1, \text{ind})\}. \quad \text{DelHist}(w) = \emptyset.$$

Leakage of the search query  $q_2$ :

$$\text{TimeDB}(w) = \emptyset. \quad \text{DelHist}(w) = \{(1, 5)\}.$$

Note that,  $\text{DelHist}(w)$  leaks that the **add** operation at timestamp 1 is canceled by the **del** operation at timestamp 5. Through the content of  $\text{DB}(w)$  after queries  $q_1$  and  $q_2$  the adversary can infer that the **del** operation at timestamp 5 corresponds to document **ind**. Therefore, it adheres to the notion of backward privacy described at the beginning of this section.

Now, let us consider the search query  $q_3$ . After the search query  $q_2$ ,  $(\text{ind}, w)$  was added at timestamp 12 and later deleted at timestamp 14. Through the same intuitive notion of backward privacy, based on the state of  $\text{DB}(w)$  after queries  $q_2$  and  $q_3$ , the adversary should not infer which document does the update queries at timestamp 12 and 14 correspond to. However, the leakage of the search query  $q_3$  is:

$$\text{TimeDB}(w) = \emptyset. \quad \text{DelHist}(w) = \{(1, 5), (1, 14), (12, 5), (12, 14)\}.$$

Hence, through the leakage profile the adversary can infer that updates at timestamp 12 and 14 correspond to document **ind** as it has already inferred which document the update queries at timestamp 1 and 5 correspond to. Clearly, this goes against the intuitive notion of backward privacy.

The following restriction is imposed on the constructions  $\text{Diana}_{\text{del}}$  and  $\text{Janus}$  that are proven to be weak backward private in [7]: 'reinsertion of a document-keyword pair is not allowed after the deletion of the corresponding document-keyword pair'. The reinsertion restriction allows one to avoid scenarios such as above that violate the intuitive notion of backward privacy. Hence, WBP can be considered to argue backward privacy in reinsertion restriction setting only. However, WBP-constructions proposed in subsequent works [20, 40] do not explicitly mention that reinsertion of document-keyword pair is not allowed. Therefore, in order to avoid any ambiguity, we feel it should be clarified that WBP is applicable in such restricted scenarios only.

**Remark** The case of reinsertion of document-keyword pair, may not be a concern in certain use-cases of SSE schemes where a new document identifier can be assigned to the updated document, thereby, ensuring that the newly inserted document-keyword pairs cannot be related to older ones. But this trick may not always be applicable especially when the contents of file can change dynamically over time. Here, one needs to handle reinsertion of keyword in *existing* documents, i.e., the document identifier can't be changed. Therefore, in such scenarios one needs to support reinsertion of document-keyword pairs. As an example, consider the case of a hospital database where the patients' records are documents and the disease they are suspected to suffer from are keywords. One cannot rule out a scenario in which based on newer symptoms a patient is re-suspected to suffer from a disease, say *malignant brain tumor*, which she had been ruled out to suffer from earlier.

### 3.4 Suggested Modifications

Here, we point out subtle issues in Definitions 3.2, 3.3 and 3.4 and suggest modifications to address them. We first argue that  $\text{sp}(w)$  should be a part of  $\mathcal{L}_{\text{Search}}(w)$  in definitions of BPIP, BPUP and WBP [7, Definition 4.2]. The constructions satisfying the respective definitions leak  $\text{sp}(w)$  in Search protocol. This implies that  $\text{sp}(w)$  can be derived from the other leakage functions in the respective definitions. Now, consider the corner case where no updates corresponding to  $w$  has occurred so far and two search queries on  $w$  are executed at timestamps 5 and 8 respectively. For search query at timestamp 8,  $\text{sp}(w)=\{5\}$ . However, the state of other leakage functions at timestamp 8 are:  $\text{TimeDB}(w)=\emptyset$ ,  $\text{Updates}(w)=\emptyset$  and  $\text{DelHist}(w)=\emptyset$ . As can be observed,  $\text{sp}(w)$  cannot be derived from other leakage functions. Hence,  $\text{sp}(w)$  should be included in  $\mathcal{L}_{\text{Search}}(w)$  of BPIP, BPUP and WBP.

Next, we argue that  $\mathcal{L}_{\text{Search}}$  in WBP should also be augmented with leakage function  $\text{Updates}(w)$ . Recall that the informal notion of WBP in [7] states that a search query on  $w$  should at most leak the documents currently matching  $w$ , when they were inserted, when all the updates on  $w$  happened and the total number of updates on  $w$ . While formalizing the leakage profile of WBP [7, Definition 4.2], the leakage in Search protocol is described as:  $\mathcal{L}_{\text{Search}}(w) = (\text{TimeDB}(w), \text{DelHist}(w))$ . Note that,  $\text{Updates}(w)$  isn't explicitly a part of  $\mathcal{L}_{\text{Search}}(w)$ . This implies that, given  $\text{TimeDB}(w)$  and  $\text{DelHist}(w)$  one can construct  $\text{Updates}(w)$ . Now, consider the following entries in query list  $\mathbf{Q}_1$  corresponding to  $w$ :  $(1, \text{add}, (1, w))$ ,  $(2, \text{del}, (2, w))$ ,  $(4, \text{add}, (3, w))$ ,  $(6, \text{del}, (1, w))$ ,  $(8, w)$ . Let us denote the search query on  $w$  at timestamps 8 by  $q_1$ .

Leakage functions at search query  $q_1$ :

$$\begin{aligned} \text{TimeDB}(w) &= \{(4, 3)\}. & \text{DelHist}(w) &= \{(1, 6)\}. \\ \text{Updates}(w) &= \{1, 2, 4, 6\}. \end{aligned}$$

Note that,  $\text{TimeDB}(w)$  and  $\text{DelHist}(w)$  collectively do not capture any information about the update operation at timestamp 2. As a result,  $\text{Updates}(w)$  cannot be constructed given the leakage functions  $\text{TimeDB}(w)$  and  $\text{DelHist}(w)$ . Hence,  $\text{Updates}(w)$  should be a part of  $\mathcal{L}_{\text{Search}}(w)$  of WBP. Further, the construction Janus also leaks  $\text{Updates}(w)$  in Search protocol. Hence,  $\text{Updates}(w)$  should be included in the leakage profile of Janus construction.

Thus, the search leakage in the respective security definitions become:

$$\begin{aligned} \mathcal{L}_{\text{BPIP, Search}} &: \{\text{sp}(w), \text{TimeDB}(w), o_w\}. \\ \mathcal{L}_{\text{BPUP, Search}} &: \{\text{sp}(w), \text{TimeDB}(w), \text{Updates}(w)\}. \\ \mathcal{L}_{\text{WBP, Search}} &: \{\text{sp}(w), \text{TimeDB}(w), \text{Updates}(w), \text{DelHist}(w)\}. \end{aligned}$$

### 3.5 Alternative Formulations of Leakage

We propose two alternative formulations of information leakage in the context of backward privacy (Definition 3.5 and 3.6). As mentioned earlier, similar to [11], for a given query  $q$ , we consider the output of Search protocol to be the set of document identifiers satisfying  $q$ . This allows us to decouple the storage of documents from the storage of data structures used to realize the search operation in all our constructions and provide security definitions accordingly.

Definition 3.5, i.e., BP-I captures that an update query should leak nothing and a search query should only reveal when the updates on  $w$  happened apart from the identifiers of the documents currently matching  $w$ . BP-I is similar to the definition of Backward privacy with update pattern (BPUP) in [7] except BP-I doesn't allow leakage of  $\text{TimeDB}(w)$ . BP-I tightly captures the leakage profile of constructions Fides, Mitra<sup>2</sup> and  $\Pi_{\text{BP-prime}}$ . This exercise of tightly capturing the leakage avoids confusion regarding the immunity of constructions against attacks which result from leakages that are indeed not incurred by them,  $\text{TimeDB}(w)$  in this case.

**Definition 3.5.** (BP-I) An  $\mathcal{L} = \{\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}\}$ -semantically secure against adaptive attacks DSSE scheme  $\Pi$  is BP-I iff  $\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}$  can be written as:

$$\begin{aligned}\mathcal{L}_{\text{Setup}}() &= \emptyset. & \mathcal{L}_{\text{Update}}(\text{op}, (\text{ind}, w)) &= \emptyset. \\ \mathcal{L}_{\text{Search}}(w) &= \{\text{sp}(w), \text{Updates}(w), \text{DB}(w)\}.\end{aligned}$$

Constructing backward private DSSE schemes with optimal communication complexity is an interesting question to explore. To achieve this goal it seems necessary for  $\mathcal{S}$  to be able to identify which insertion entry is canceled by a particular delete entry. However, the stronger definitions BPIP, BPUP and BP-I do not allow such information to be leaked to  $\mathcal{S}$ . Therefore, like WBP, one needs to craft a definition that allows some non-trivial relation among the update queries to be leaked. But unlike WBP it should be able to capture the notion of backward privacy in the general setting. We observe that there can be several alternatives to relax the formal notion of backward privacy. Next, we describe one such candidate definition, i.e., BP-II. Later, we show a natural construction (see Section 4.3), whose leakage profile is tightly captured by BP-II.

In order to describe the leakage profile of BP-II, let us introduce some new leakage functions. Let  $\text{DB}(w)$  and  $\text{DB}(w)'$  denote the set of documents matching  $w$  in the current and previous search respectively and let  $\text{LDB}(w)$  denote the list of documents matching  $w$  in the current search, in order of their insertion. An element in  $\text{LDB}(w)$  is of the form  $(i, \text{ind})$ , where  $i$  denotes the index and  $\text{ind}$  denotes the document identifier. Let the timestamp of the current and previous search be denoted as  $u_c$  and  $u_p$  respectively. We define the following leakage functions:

$$\begin{aligned}\text{PreUp}(w) &= \{(\text{ind}, u) \mid \text{ind} \in \text{DB}(w)' \text{ and } (u, \text{op}, (\text{ind}, w)) \in \mathbf{Q} \\ &\quad \text{and } u_p < u < u_c\}. \\ \text{CurUp}(w) &= \{(\text{ind}, u) \mid \text{ind} \in \text{DB}(w) \text{ and } (u, \text{op}, (\text{ind}, w)) \in \mathbf{Q} \\ &\quad \text{and } u_p < u < u_c\}. \\ \text{UpPair}(w) &= \{(u_1, u_2) \mid (u_1, \text{op}, (\text{ind}, w)) \in \mathbf{Q} \text{ and} \\ &\quad (u_2, \text{op}, (\text{ind}, w)) \in \mathbf{Q} \text{ and } u_p < u_1 < u_2 < u_c\}.\end{aligned}$$

$\text{PreUp}(w)$  captures the relation between the identifiers obtained as a result in the previous search on  $w$  and the updates that happen after the previous search on  $w$  and before the current search on  $w$ .  $\text{CurUp}(w)$  captures the relation between the identifiers obtained as a result in the current search on  $w$  and the updates that happen after the previous search on  $w$  and before the current search on  $w$ .  $\text{UpPair}(w)$  captures the relation among the updates corresponding to some  $\text{ind}$ , that happen after the previous search on  $w$  and before the current search on  $w$ .

For example, let us consider the search query on  $w$  at timestamp 16. Let the timestamp of the previous search query on  $w$  be 4. Let  $\text{DB}(w)$  and  $\text{DB}(w)'$  be  $\{1, 3\}$  and  $\{1, 2\}$  respectively. Let  $(5, \text{del}, (1, w))$ ,  $(7, \text{del}, (2, w))$ ,  $(12, \text{add}, (1, w))$  and  $(15, \text{add}, (3, w))$  be the update queries on  $w$  that occur between these two searches. The respective leakage functions will be:

$$\begin{aligned}\text{PreUp}(w) &= \{(1, 5), (2, 7), (1, 12)\}. & \text{UpPair}(w) &= \{(5, 12)\}. \\ \text{CurUp}(w) &= \{(1, 5), (1, 12), (3, 15)\}.\end{aligned}$$

Next, we describe the search leakage of BP-II. Though the leakage functions defined above may appear a bit complex, they are useful abstractions through which BP-II captures the notion of backward privacy as

<sup>2</sup>Mitra can be proven BP-I secure with a minor modification (sending document identifiers in search result  $\text{DB}(w)$  in random order).

shown below. In Search protocol, along with  $\text{sp}(w)$  and  $\text{Updates}^{\text{op}}(w)$ , leakage functions  $\text{PreUp}(w)$ ,  $\text{CurUp}(w)$ ,  $\text{UpPair}(w)$  and  $\text{LDB}(w)$  can be leaked to  $\mathcal{S}$ . As already defined in the beginning of Section 3.2, recall that  $\mathbf{I}_{\text{BP}}$  denotes the set of document identifiers relevant to the notion of backward privacy. Note that,  $\forall \text{ind} \in \mathbf{I}_{\text{BP}}$ , as  $\text{ind} \notin \text{DB}(w)'$  and as the last update operation on  $(\text{ind}, w)$  is del operation, it follows that  $\text{ind} \notin \text{DB}(w)$ . Hence, no information about the inds in  $\mathbf{I}_{\text{BP}}$  can be revealed through leakage functions  $\text{PreUp}(w)$ ,  $\text{CurUp}(w)$  and  $\text{LDB}(w)$ . As the other leakage functions viz.,  $\text{sp}(w)$  and  $\text{UpPair}(w)$ , do not leak ind corresponding to an update, no information about identifiers in  $\mathbf{I}_{\text{BP}}$  gets revealed.

**Definition 3.6.** (BP-II) An  $\mathcal{L} = \{\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}\}$ -semantically secure against adaptive attacks scheme  $\Pi$  is BP-II iff  $\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}}$  can be written as:

$$\begin{aligned} \mathcal{L}_{\text{Setup}}() &= \emptyset. & \mathcal{L}_{\text{Update}}(\text{op}, (\text{ind}, w)) &= \emptyset. \\ \mathcal{L}_{\text{Search}}(w) &= \{\text{sp}(w), \text{Updates}^{\text{op}}(w), \text{PreUp}(w), \text{CurUp}(w), \\ & \quad \text{UpPair}(w), \text{LDB}(w)\} \end{aligned}$$

Note that Definitions 3.5 (BP-I) and 3.6 (BP-II) capture the notion of forward privacy as well. Therefore, constructions that are BP-I and BP-II secure are naturally forward private.

**Inverse Backward Privacy** We identify a new desirable property for a DSSE scheme called *inverse backward privacy* which captures the complementary situation of backward privacy. Analogous to backward privacy, a DSSE scheme is inverse backward private if whenever a document-keyword pair  $(\text{ind}, w)$  is deleted and later added, subsequent search queries on  $w$  won't reveal the fact that  $(\text{ind}, w)$  was deleted unless it can be inferred by the search and access pattern of the search query. For example, let  $\text{DB}(w) = \{1, 2\}$  for search query  $\mathbf{q}$  at timestamp 5. Let the update operations on  $w$  after query  $\mathbf{q}$  and before the next search query be  $(6, \text{del}, (1, w))$ ,  $(12, \text{add}, (1, w))$ , then no information about the identifier 1 in update queries at timestamp 6 and 12 should be revealed to  $\mathcal{S}$  on the next search query on  $w$ .

Inverse Backward Privacy property could be of relevance in various use-cases. For instance, consider the employee database where the employee records correspond to document and project teams she works in correspond to keywords. An employee  $E_1$  maybe dropped and reincluded in a project team. We would like to hide the fact that employee  $E_1$  was dropped briefly, if no search on that project team had been performed during that period. Definition 3.5 (BP-I) is strong enough to capture our notion of inverse backward privacy. Therefore, all constructions that can be shown to satisfy Definition 3.5 are inverse backward private.

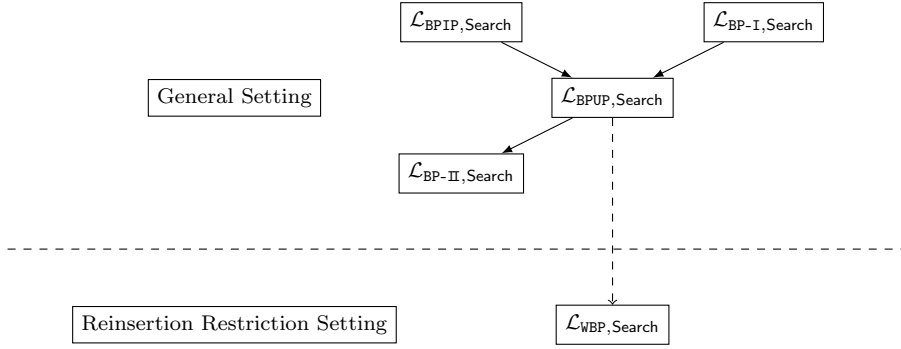
To summarize, we proposed alternative formulations of leakages in the form of two new security definitions. We argued that the two definitions capture the notion of backward privacy. An interesting open question is to analyze the real-world consequences of the leakages incurred by constructions satisfying the definitions of backward privacy, viz., BPIP, BPUP, BP-I, BP-II or even WBP (in restricted setting).

### 3.6 Comparison among Definitions

Finally we discuss the relations among the definitions of backward privacy. Precisely, we give comparison among the leakages incurred in Search protocol in the definitions of backward privacy. As we remarked earlier, the notion of WBP is only suitable to argue backward privacy in a restricted setting. Hence, it would not be meaningful to compare the definitions of backward privacy which apply to the general setting with WBP. Therefore, WBP is not considered for comparison in this section.

If  $\mathcal{L}_1$  leaks less than  $\mathcal{L}_2$ , we denote it as  $\mathcal{L}_1 \preceq \mathcal{L}_2$ . By the proposition, “ $\mathcal{L}_1$  leaks less than  $\mathcal{L}_2$ ” we mean that  $\mathcal{L}_1$  gives less information about the database and the queries to the simulator than  $\mathcal{L}_2$ , or, said otherwise, that every information given by  $\mathcal{L}_1$  can be inferred from  $\mathcal{L}_2$ . If  $\mathcal{L}_1$  leaks strictly less than  $\mathcal{L}_2$ , we denote it as  $\mathcal{L}_1 \prec \mathcal{L}_2$ . Here, we give an argument to derive the leakage functions in  $\mathcal{L}_1$  through that of  $\mathcal{L}_2$  in order to show that  $\mathcal{L}_1 \prec \mathcal{L}_2$ . And with the help of counter examples, we showcase that the leakage functions in  $\mathcal{L}_1$  cannot be derived from that of  $\mathcal{L}_2$ .

Recall the leakage functions corresponding to Search protocol in the existing definitions [7] are related in



Here,  $\mathcal{L}_1 \rightarrow \mathcal{L}_2$  denotes  $\mathcal{L}_1 \prec \mathcal{L}_2$

Figure 5: Relations among Definitions of Backward Privacy

the following way:

$$\mathcal{L}_{\text{BPIP,Search}} \prec \mathcal{L}_{\text{BPUP,Search}} \prec \mathcal{L}_{\text{WBP,Search}}. \quad (3)$$

We recall  $\mathcal{L}_{\text{Search}}$  of all the definitions below. See Section 2.4 (Common Leakages) and Section 3 for the description of the leakage functions used here.

$$\mathcal{L}_{\text{BPIP,Search}} : \{\text{sp}(w), \text{TimeDB}(w), o_w\}. \quad (4)$$

$$\mathcal{L}_{\text{BPUP,Search}} : \{\text{sp}(w), \text{TimeDB}(w), \text{Updates}(w)\}. \quad (5)$$

$$\mathcal{L}_{\text{BP-I,Search}} : \{\text{sp}(w), \text{Updates}(w), \text{DB}(w)\}. \quad (6)$$

$$\mathcal{L}_{\text{BP-II,Search}} : \{\text{sp}(w), \text{Updates}^{\text{op}}(w), \text{PreUp}(w), \text{CurUp}(w), \text{UpPair}(w), \text{LDB}(w)\}. \quad (7)$$

As can be observed from (4) and (6), for BPIP,  $\text{TimeDB}(w)$  cannot be derived from leakage functions in  $\mathcal{L}_{\text{BP-I,Search}}$  (see Example 1, described below) and for BP-I,  $\text{Updates}(w)$  cannot be derived from leakage functions in  $\mathcal{L}_{\text{BPIP,Search}}$  (see Example 2, described below). Hence,  $\mathcal{L}_{\text{BPIP,Search}}$  and  $\mathcal{L}_{\text{BP-I,Search}}$  cannot be orderly related.

From (5) and (6) it is straightforward that all the leakage functions in  $\mathcal{L}_{\text{BP-I,Search}}$  can be derived from the leakage functions in  $\mathcal{L}_{\text{BPUP,Search}}$ , as  $\text{DB}(w)$  can be easily obtained from  $\text{TimeDB}(w)$ . Further, as described Example 1 (described below),  $\text{TimeDB}(w)$  cannot be derived from leakage functions in  $\mathcal{L}_{\text{BP-I,Search}}$ . Hence,  $\mathcal{L}_{\text{BP-I,Search}} \prec \mathcal{L}_{\text{BPUP,Search}}$ .

From (5) and (7), one can conclude that all the leakage functions in  $\mathcal{L}_{\text{BPUP,Search}}$  can be derived from the leakage functions in  $\mathcal{L}_{\text{BP-II,Search}}$ , as  $\text{TimeDB}(w)$  can be derived from  $\text{CurUp}(w)$ ,  $\text{Updates}^{\text{op}}(w)$  and  $\text{DB}(w)$  (in particular  $\text{LDB}(w)$ ) of current and previous searches on  $w$ . Note that,  $\text{TimeDB}(w)$  comprises of information regarding the timestamp of the **add** updates corresponding to document identifiers in  $\text{DB}(w)$  that are not followed by their respective **del** update. This can be easily determined by keeping track of all the **add** updates corresponding to document identifiers in  $\text{DB}(w)$  that are not followed by the respective **del** update through  $\text{CurUp}(w)$  over all searches on  $w$  and  $\text{Updates}^{\text{op}}(w)$ . Further, some leakage functions in  $\mathcal{L}_{\text{BP-II,Search}}$  say,  $\text{UpPair}(w)$ , cannot be derived from  $\mathcal{L}_{\text{BPUP,Search}}$  (see Example 4, described below). Hence,  $\mathcal{L}_{\text{BPUP,Search}} \prec \mathcal{L}_{\text{BP-II,Search}}$ .

## Examples

1. Consider the entries corresponding to  $w$  in query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$  as described below. Corresponding to  $w$ ,  $(1, \text{add}, (1, w))$ ,  $(2, \text{add}, (2, w))$ ,  $(3, \text{del}, (1, w))$  and  $(6, w)$  are present in  $\mathbf{Q}_1$ . Corresponding to  $w$ ,  $(1, \text{add}, (1, w))$ ,  $(2, \text{del}, (1, w))$ ,  $(3, \text{add}, (2, w))$  and  $(6, w)$  are present in  $\mathbf{Q}_2$ .

Leakage functions corresponding to search query at timestamp 6 in  $\mathbf{Q}_1$ :

$$\begin{aligned} \text{sp}(w) &= \{6\}. & o_w &= 3. \\ \text{TimeDB}(w) &= \{(2, 2)\}. & \text{DB}(w) &= \{2\}. \\ \text{Updates}(w) &= \{1, 2, 3\}. \end{aligned}$$

Leakage functions corresponding to search query at timestamp 6 in  $\mathbf{Q}_2$ :

$$\begin{aligned} \text{sp}(w) &= \{6\}. & o_w &= 3. \\ \text{TimeDB}(w) &= \{(3, 2)\}. & \text{DB}(w) &= \{2\}. \\ \text{Updates}(w) &= \{1, 2, 3\}. \end{aligned}$$

As can be observed the leakage functions in  $\mathcal{L}_{\text{BP-I,Search}}$ , i.e.,  $\text{sp}(w)$ ,  $\text{Updates}(w)$  and  $\text{DB}(w)$  are the same for query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$ . However,  $\text{TimeDB}(w)$  is different for query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$ . Hence,  $\text{TimeDB}(w)$  cannot be uniquely determined from the leakage functions in  $\mathcal{L}_{\text{BP-I,Search}}$ .

2. Consider the entries corresponding to  $w$  in query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$  as described below. Corresponding to  $w$ ,  $(1, \text{add}, (1, w))$ ,  $(2, \text{add}, (2, w))$ ,  $(3, \text{del}, (1, w))$  and  $(6, w)$  are present in  $\mathbf{Q}_1$ . Corresponding to  $w$ ,  $(1, \text{add}, (1, w))$ ,  $(2, \text{add}, (2, w))$ ,  $(5, \text{del}, (1, w))$  and  $(6, w)$  are present in  $\mathbf{Q}_2$ .

Leakage functions corresponding to search query at timestamp 6 in  $\mathbf{Q}_1$ :

$$\begin{aligned} \text{sp}(w) &= \{6\}. & o_w &= 3. \\ \text{TimeDB}(w) &= \{(2, 2)\}. & \text{DB}(w) &= \{2\}. \\ \text{Updates}(w) &= \{1, 2, 3\}. \end{aligned}$$

Leakage functions corresponding to search query at timestamp 6 in  $\mathbf{Q}_2$ :

$$\begin{aligned} \text{sp}(w) &= \{6\}. & o_w &= 3. \\ \text{TimeDB}(w) &= \{(2, 2)\}. & \text{DB}(w) &= \{2\}. \\ \text{Updates}(w) &= \{1, 2, 4\}. \end{aligned}$$

As can be observed the leakage functions in  $\mathcal{L}_{\text{BP-IP,Search}}$ , i.e.,  $\text{sp}(w)$ ,  $\text{TimeDB}(w)$  and  $o_w$  are the same for query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$ . However,  $\text{Updates}(w)$  is different for query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$ . Hence,  $\text{Updates}(w)$  cannot be uniquely determined from the leakage functions in  $\mathcal{L}_{\text{BP-I,Search}}$ .

3. Consider the entries corresponding to  $w$  in query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$  as described below. Corresponding to  $w$ ,  $(1, \text{add}, (1, w))$ ,  $(2, \text{add}, (2, w))$ ,  $(3, \text{del}, (1, w))$  and  $(6, w)$  are present in  $\mathbf{Q}_1$ . Corresponding to  $w$ ,  $(1, \text{add}, (1, w))$ ,  $(2, \text{del}, (1, w))$ ,  $(3, \text{add}, (2, w))$  and  $(6, w)$  are present in  $\mathbf{Q}_2$ .

Leakage functions corresponding to search query at timestamp 6 in  $\mathbf{Q}_1$ :

$$\begin{aligned} \text{sp}(w) &= \{6\}. & \text{Updates}(w) &= \{1, 2, 3\}. \\ \text{DB}(w) &= \{2\}. & \text{UpPair}(w) &= \{(1, 3)\}. \end{aligned}$$

Leakage functions corresponding to search query at timestamp 6 in  $\mathbf{Q}_2$ :

$$\begin{aligned} \text{sp}(w) &= \{6\}. & \text{Updates}(w) &= \{1, 2, 3\}. \\ \text{DB}(w) &= \{2\}. & \text{UpPair}(w) &= \{(1, 2)\}. \end{aligned}$$

As can be observed the leakage functions in  $\mathcal{L}_{\text{BP-I,Search}}$ , i.e.,  $\text{sp}(w)$ ,  $\text{Updates}(w)$  and  $\text{DB}(w)$  are the same for query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$ . However,  $\text{UpPair}(w)$  is different for query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$ . Hence,  $\text{UpPair}(w)$  cannot be uniquely determined from the leakage functions in  $\mathcal{L}_{\text{BP-I,Search}}$ .

4. Consider the entries corresponding to  $w$  in query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$  as described below. Corresponding to  $w$ , (1, add, (1,  $w$ )), (2, add, (2,  $w$ )), (3, del, (1,  $w$ )), (4, del, (2,  $w$ )), (5, add, (3,  $w$ )) and (6,  $w$ ) are present in  $\mathbf{Q}_1$ . Corresponding to  $w$ , (1, add, (1,  $w$ )), (2, add, (2,  $w$ )), (3, del, (2,  $w$ )), (4, del, (1,  $w$ )), (5, add, (3,  $w$ )) and (6,  $w$ ) are present in  $\mathbf{Q}_2$ .

Leakage functions corresponding to search query at timestamp 6 in  $\mathbf{Q}_1$ :

$$\begin{aligned} \text{sp}(w) &= \{6\}. & \text{Updates}(w) &= \{1, 2, 3, 4, 5, 6\}. \\ \text{TimeDB}(w) &= \{(5, 3)\}. & \text{UpPair}(w) &= \{(1, 3), (2, 4)\}. \end{aligned}$$

Leakage functions corresponding to search query at timestamp 6 in  $\mathbf{Q}_2$ :

$$\begin{aligned} \text{sp}(w) &= \{6\}. & \text{Updates}(w) &= \{1, 2, 3, 4, 5, 6\}. \\ \text{TimeDB}(w) &= \{(5, 3)\}. & \text{UpPair}(w) &= \{(1, 4), (2, 3)\}. \end{aligned}$$

As can be observed the leakage functions in  $\mathcal{L}_{\text{BPUP,Search}}$ , i.e.,  $\text{sp}(w)$ ,  $\text{Updates}(w)$  and  $\text{TimeDB}(w)$  are the same for query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$ . However,  $\text{UpPair}(w)$  is different for query lists  $\mathbf{Q}_1$  and  $\mathbf{Q}_2$ . Hence,  $\text{UpPair}(w)$  cannot be uniquely determined from the leakage functions in  $\mathcal{L}_{\text{BPUP,Search}}$ .

## 4 Backward Private DSSE Constructions

In this section we propose three backward private schemes  $\Pi_{\text{BP-prime}}$ ,  $\Pi_{\text{BP}}$  and  $\Pi_{\text{WBP}}$  that are respectively BP-I, BP-II and WBP secure. Our starting point is a forward private DSSE scheme  $\Pi_{\text{FP}}$  which is a modified version of the scheme proposed in [15].

### 4.1 $\Pi_{\text{FP}}$ : A Warm-up Solution

The central idea in  $\Pi_{\text{FP}}$  (Figure 6) is to make updates using fresh keys. Hence, the keys disclosed in previous searches do not reveal anything about these new updates.  $\Pi_{\text{FP}}$  is described in Figure 6. The construction makes use of PRFs  $F_t, F_d: \{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$  and hash functions  $H_1: \{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$  and  $H_2: \{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^\mu$ , where  $\mu = \lambda + 1$  and  $\lambda$  is security parameter.

The **Setup** algorithm generates secret keys  $k_t$  and  $k_d$ .  $C$  initiates three maps:  $\mathbf{T}$ ,  $\mathbf{D}$  and  $\mathbf{W}$ . The maps  $\mathbf{T}$  and  $\mathbf{D}$  are stored at  $S$ 's end. Corresponding to  $w$ ,  $\mathbf{D}$  stores the pointer to  $\text{PSet}_w$ , the set of document identifiers in plaintext that were obtained as a result of the latest search operation on  $w$  and  $\mathbf{T}$  stores encrypted entries inserted after the latest search operation on  $w$ . The map  $\mathbf{W}$  is stored at  $C$ 's end. In  $\mathbf{W}$ , corresponding to  $w$ ,  $C$  stores the version  $\text{ver}_w$  (initialized to 0) and counter  $c_w$  (initialized to -1).  $\text{ver}_w$  ensures that the key  $k_w$  used in the **Update** protocol is unknown to  $S$ ,  $c_w$  stores information about the number of entries added to  $\mathbf{T}$  corresponding to  $w$  after the latest search operation.

**Update:** When  $C$  wants to add/del a document-keyword pair ( $\text{ind}, w$ ), it computes key  $k_w$  using keyword  $w$  and  $\text{ver}_w$  (see line 5) and increments  $c_w$ . Based on  $k_w$  and  $c_w$ ,  $C$  computes the hash digests  $\text{label}$  and  $\text{pad}$  and sends  $(\text{label}, e = \text{pad} \oplus (b \parallel \text{ind}))$  to  $S$ .  $S$  then adds  $(\text{label}, e)$  to  $\mathbf{T}$ . For add (resp. del) operation,  $b=0$  (resp.  $b=1$ ).

Note that,  $k_w$  used in processing update queries is computed using updated  $\text{ver}_w$ . Since,  $k_w$  is output of PRF  $F_t$  at  $w \parallel \text{ver}_w$ , it is indistinguishable from random for  $S$ . As  $\text{label}$  (resp.  $\text{pad}$ ) is computed as  $H_1(k_w \parallel c_w)$  (resp.  $H_2(k_w \parallel c_w)$ ) for an update query, it is indistinguishable from random for  $S$  as  $H_1$  (resp.  $H_2$ ) is modeled as a random oracle. Hence, in the security proof, update queries can be simulated by generating random  $(\text{label}, e)$  pair.

**Search:** When  $C$  wants to perform a search query on  $w$ , it computes  $\text{label}_w$  (see line 5) and the key  $k_w$  is computed (see line 7) only if new entries are inserted in map  $\mathbf{T}$ .  $C$  sends  $\text{label}_w$ ,  $k_w$  and  $c_w$  to  $S$ .  $k_w$  gets revealed to  $S$  only if a new entry was inserted to  $\mathbf{T}$  after the previous search on  $w$ . Hence,  $C$  updates the version  $\text{ver}_w$  (see line 8).  $\text{ver}_w$  is not updated in a search query on  $w$  for which corresponding to  $w$ , no



<p><u>Setup()</u></p> <ol style="list-style-type: none"> <li>1. <math>k_t \xleftarrow{\\$} \{0, 1\}^\lambda, k_d \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>2. <math>\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow</math> empty map</li> <li>3. return ( EDB = (<math>\mathbf{D}, \mathbf{T}</math>), st<sub>c</sub> = (<math>k_t, k_d, \mathbf{W}</math>)) to (S, C)</li> </ol> <p><u>Update(op, w, ind, st<sub>C</sub>; EDB)</u></p> <p><u>Update<sub>C</sub>(op, w, ind, st<sub>C</sub>)</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, c_w) = \perp</math> <b>then</b></li> <li>3.   <math>\text{ver}_w \leftarrow 0, c_w \leftarrow -1</math></li> <li>4. <math>c_w \leftarrow c_w + 1</math></li> <li>5. <math>k_w \leftarrow F_t(k_t, w    \text{ver}_w)</math></li> <li>6. <math>\text{label} \leftarrow H_1(k_w    c_w)</math></li> <li>7. <math>\text{pad} \leftarrow H_2(k_w    c_w)</math></li> <li>8. <math>b \leftarrow 0</math> (op=add) / <math>1</math> (op=del)</li> <li>9. <math>e \leftarrow (b    \text{ind}) \oplus \text{pad}</math></li> <li>10. <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, c_w)</math></li> <li>11. Send (label, e) to S</li> </ol>	<p><u>Update<sub>S</sub>(EDB)</u></p> <ol style="list-style-type: none"> <li>12. Receive (label, e) from C</li> <li>13. <math>\mathbf{T}[\text{label}] \leftarrow e</math></li> </ol> <p><u>Search(w, st<sub>C</sub>; EDB)</u></p> <p><u>Search<sub>C</sub>(w, st<sub>C</sub>) :</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, c_w) = \perp</math> <b>then</b></li> <li>3.   return <math>\emptyset</math></li> <li>4. <b>else</b></li> <li>5.   <math>\text{label}_w \leftarrow F_d(k_d, w)</math></li> <li>6.   <b>if</b> <math>c_w \neq -1</math> <b>then</b></li> <li>7.     <math>k_w \leftarrow F_t(k_t, w    \text{ver}_w)</math></li> <li>8.     <math>\text{ver}_w \leftarrow \text{ver}_w + 1, c_w \leftarrow -1</math></li> <li>9.     <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, c_w)</math></li> <li>10.   <b>else</b></li> <li>11.     <math>k_w \leftarrow \perp</math></li> <li>12.   Send (<math>\text{label}_w, k_w, c_w</math>) to S</li> </ol>	<p><u>Search<sub>S</sub>(EDB) :</u></p> <ol style="list-style-type: none"> <li>13. Receive (<math>\text{label}_w, k_w, c_w</math>) from C</li> <li>14. <math>\text{addr}(\text{PSet}_w) \leftarrow \mathbf{D}[\text{label}_w]</math></li> <li>15. Retrieve <math>\text{PSet}_w</math> and set <math>\text{AuxSet} \leftarrow \text{PSet}_w</math></li> <li>16. <b>if</b> <math>k_w \neq \perp</math> <b>then</b></li> <li>17.   <math>c \leftarrow 0</math></li> <li>18.   <b>while</b> <math>c \leq c_w</math> <b>do</b></li> <li>19.     <math>\text{label} \leftarrow H_1(k_w    c), e \leftarrow \mathbf{T}[\text{label}]</math></li> <li>20.     <math>\text{pad} \leftarrow H_2(k_w    c), (b    \text{ind}) \leftarrow e \oplus \text{pad}</math></li> <li>21.     <b>if</b> <math>b = 0</math></li> <li>22.       <math>\text{AuxSet} \leftarrow \text{AuxSet} \cup \{\text{ind}\}</math></li> <li>23.     <b>else</b></li> <li>24.       <math>\text{AuxSet} \leftarrow \text{AuxSet} \setminus \{\text{ind}\}</math></li> <li>25.     Remove <math>\mathbf{T}[\text{label}]</math></li> <li>26.     <math>c \leftarrow c + 1</math></li> <li>27. Store <math>\text{AuxSet}</math> in disk</li> <li>28. Set <math>\mathbf{D}[\text{label}_w] \leftarrow \text{addr}(\text{AuxSet})</math></li> <li>29. Send <math>\text{AuxSet}</math> to C</li> </ol>
--	--	---

Figure 6: Scheme  $\Pi_{\text{FP}}$

updates on map  $\mathbf{T}$  were made after the previous search on  $w$ . Based on the information received from C, S computes the result set and updates  $\mathbf{D}$  with the newly computed result set.

In  $\Pi_{\text{FP}}$ ,  $\text{ver}_w$  ensures that S cannot relate later updates with previous search queries and  $c_w$  ensures that S cannot correlate update queries on  $w$  after the previous search operation on  $w$ .

**Remark** Essentially,  $\Pi_{\text{FP}}$  is same as the construction in [15] except the following: 1) The search counter corresponding to  $w$  (denoted as  $\text{ver}_w$ ) is updated differently than in [15] to avoid unnecessary increments to the search counter, 2) After a search operation on  $w$ , the revealed document identifiers are stored together ( $\text{DB}(w)$ ) in plaintext (as suggested in [7]) to provide reasonable locality without incurring any additional leakage.

**Example 1** Consider the following list of update queries: (add, (1,  $w_1$ )), (add, (1,  $w_2$ )), (add, (2,  $w_1$ )), (add, (3,  $w_1$ )), (add, (3,  $w_2$ )), (del, (1,  $w_1$ )). Figure 7.a shows the state of indexes at C and S after these updates are processed. Figure 7.b shows the state of indexes at C and S after search on  $w_1$ .

**Correctness** The scheme is correct as long as there are no repeated labels in maps  $\mathbf{T}$  and  $\mathbf{D}$ . Since,  $F_d$  is a PRF, only with negligible probability  $\text{label}_w$  in  $\mathbf{D}$  is same for two distinct keywords. The input to  $H_1$  is repeated only with negligible probability as  $F_t$  is a PRF. If we consider  $H_1$  to be a collision resistant hash function, only with negligible probability  $\text{label}$  in  $\mathbf{T}$  is repeated.

In Appendix A, we provide complete proofs of correctness and FP-I-security (see Definition 3.1) of  $\Pi_{\text{FP}}$ .

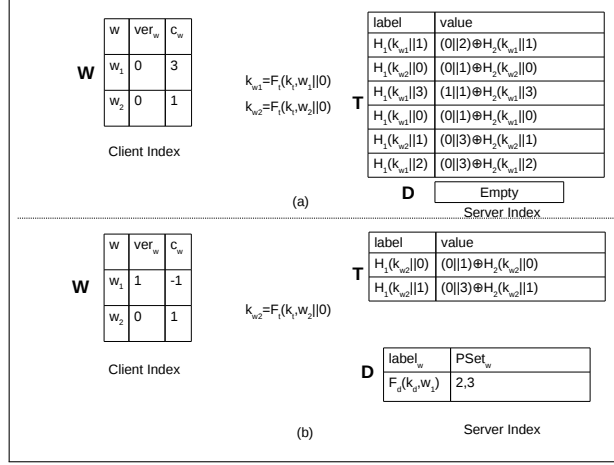


Figure 7: Example 1: Indexes at C and S before and after search on  $w_1$  in construction  $\Pi_{FP}$ .  $\mathbf{W}$  = Client Index,  $\mathbf{D}$  = Index that stores search results of previous search query at S and  $\mathbf{T}$  = Index that stores updates after the last search on  $w$  at S

## 4.2 $\Pi_{BP}$ -prime : Realizing Strong Privacy

As  $\Pi_{FP}$  reveals the deleted entries, it isn't backward private. We propose backward private DSSE scheme  $\Pi_{BP}$ -prime that makes use of light weight symmetric primitives only. The scheme in Figure 8 is similar to applying the generic two-roundtrip backward-private scheme transformation [7] on  $\Pi_{FP}$  but with the following essential modification: the document identifiers revealed by search query are stored together in plaintext thus providing reasonable locality by avoiding re-encryption of revealed identifiers (as mentioned in Section 4.1). A high degree of locality ensures that the document identifiers are stored in contiguous memory locations which results in significant I/O efficiency. The construction makes use of PRFs  $F_t, F_d: \{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ , symmetric encryption scheme  $\mathcal{E}$  and hash function  $H_1: \{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$ .

Unlike  $\Pi_{FP}$ , encryption of  $b||ind$  is stored in  $\mathbf{T}$  at S in Update protocol (see line 8). Padding in  $\Pi_{FP}$  using  $H_2$  was necessary to prevent S from learning anything about the update. However, in  $\Pi_{BP}$ -prime, encryption of the updated entry is stored which is random in the view of S, thus revealing no information about the updated entry.

The Search protocol in  $\Pi_{BP}$ -prime comprises of two rounds. Round 1 of Search is similar to the Search protocol of  $\Pi_{FP}$ . At the end of round 1 of Search, S sends  $PSet_w$ , the result set of previous search on  $w$  and  $AList$ , the list of ciphertexts corresponding to all the updates on  $w$  after the previous search on  $w$ , in order of their insertion. While preparing  $AList$ , S only learns the timestamps at which updates corresponding to  $w$  have happened because the entries in  $AList$  are encrypted. Based on the information received from S, C computes  $AuxSet$  (see line 28-35 of Search protocol).  $AuxSet$  consists of all the document identifiers currently matching keyword  $w$  in random order. In the second round of communication,  $AuxSet$  is forwarded to S, who stores  $AuxSet$  in  $\mathbf{D}[label_w]$ .

**Correctness** As mentioned in correctness of  $\Pi_{FP}$ , the scheme is correct as long as there are no repeated labels in maps  $\mathbf{T}$  and  $\mathbf{D}$ . Since, these labels are generated in the same manner as they were generated in  $\Pi_{FP}$ , correctness of  $\Pi_{BP}$ -prime immediately follows from that of  $\Pi_{FP}$ .

**Asymptotic Complexity** The communication complexity in round 1 of the Search protocol is  $O(o'_w)$  and that of round 2 is  $O(n_w)$ . The computational complexity of the Search protocol is  $O(o'_w)$ . The communication and computation cost of the Update protocol is  $O(1)$ . Space complexity at the server's end is  $O(N + D')$ , where  $D' = \sum_{w \in \mathcal{W}} d'_w$  and at the client's end is  $O(m \log(n))$ .

Setup()

1.  $k_t \xleftarrow{\$} \{0, 1\}^\lambda$ ,  $k_d \xleftarrow{\$} \{0, 1\}^\lambda$
2.  $k_e \leftarrow \mathcal{E}.Gen()$
3.  $\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow$  empty map
4. return ( EDB = (  $\mathbf{D}, \mathbf{T}$  ),  
     $st_c = (k_e, k_t, k_d, \mathbf{W})$ ) to (S, C)

Update(op, w, ind, st<sub>c</sub>; EDB)Update<sub>C</sub>(op, w, ind, st<sub>c</sub>)

1.  $(ver_w, c_w) \leftarrow \mathbf{W}[w]$
2. **if**  $(ver_w, c_w) = \perp$  **then**
3.      $ver_w \leftarrow 0$ ,  $c_w \leftarrow -1$
4.      $c_w \leftarrow c_w + 1$
5.      $k_w \leftarrow F_t(k_t, w || ver_w)$
6.     label  $\leftarrow H_1(k_w || c_w)$
7.      $b \leftarrow 0$  (op=add) /  $1$  (op=del)
8.      $e \leftarrow \mathcal{E}.Enc(k_e, b || ind)$
9.      $\mathbf{W}[w] \leftarrow (ver_w, c_w)$
10.     Send (label, e) to S

Update<sub>S</sub>(EDB)

11. Receive (label, e) from C
12.  $\mathbf{T}[label] \leftarrow e$

Search(w, st<sub>c</sub>; EDB)**Round 1**Search<sub>C</sub>(w, st<sub>c</sub>) :

1.  $(ver_w, c_w) \leftarrow \mathbf{W}[w]$
2. **if**  $(ver_w, c_w) = \perp$  **then**
3.     return  $\emptyset$
4. **else**
5.     label<sub>w</sub>  $\leftarrow F_d(k_d, w)$
6.     **if**  $c_w \neq -1$  **then**
7.          $k_w \leftarrow F_t(k_t, w || ver_w)$
8.          $ver_w \leftarrow ver_w + 1$
9.          $c_w \leftarrow -1$
10.          $\mathbf{W}[w] \leftarrow (ver_w, c_w)$
11.     **else**
12.          $k_w \leftarrow \perp$  //No need of round 2
13.     Send (label<sub>w</sub>, k<sub>w</sub>, c<sub>w</sub>) to S

Search<sub>S</sub>(EDB) :

14. Receive (label<sub>w</sub>, k<sub>w</sub>, c<sub>w</sub>) from C
15.  $addr(PSet_w) \leftarrow \mathbf{D}[label_w]$
16. Retrieve PSet<sub>w</sub>
17. AList  $\leftarrow \emptyset$
18. **if**  $k_w \neq \perp$  **then**

19.      $c \leftarrow 0$
20.     **while**  $c \leq c_w$  **do**
21.         label  $\leftarrow H_1(k_w || c)$
22.          $e \leftarrow \mathbf{T}[label]$
23.         Append e to AList
24.         Remove  $\mathbf{T}[label]$
25.          $c \leftarrow c + 1$
26.     Send (PSet<sub>w</sub>, AList) to C
27.     **Round 2**
28.     Search<sub>C</sub>(w, st<sub>c</sub>) :
29.     Receive (PSet<sub>w</sub>, AList) from S
30.      $(ver_w, c_w) \leftarrow \mathbf{W}[w]$
31.     AuxSet  $\leftarrow PSet_w$
32.     **for**  $c \leftarrow 0$  to |AList| **do**
33.          $(b || ind) \leftarrow \mathcal{E}.Dec(k_e, AList[c])$
34.         **if**  $b = 0$
35.             AuxSet  $\leftarrow AuxSet \cup \{ind\}$
36.         **else**
37.             AuxSet  $\leftarrow AuxSet \setminus \{ind\}$
38.     return AuxSet
39.     Search<sub>S</sub>(EDB) :
40.     Receive AuxSet from C
41.      $\mathbf{D}[label_w] \leftarrow AuxSet$

Figure 8: Scheme  $\Pi_{BP}$ -prime

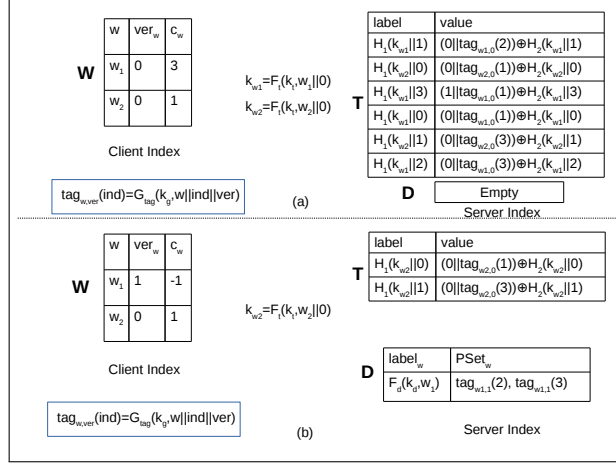


Figure 9: Example 2: Indexes at C and S before and after search on  $w_1$  in  $\Pi_{BP}$ .  $\mathbf{W}$  = Client Index,  $\mathbf{D}$  = Index that stores tags of search results of previous search query at S and  $\mathbf{T}$  = Index that stores updates after the last search on  $w$  at S

**Security** We prove  $\Pi_{BP}$ -prime is BP-I secure in the random oracle model. The proof relies on the pseudo randomness of  $F_d$ ,  $F_t$  and RCPA security of encryption scheme  $\mathcal{E}$ . The complete proof appears in Appendix B.

In a concurrent work [20], a backward private construction *Mitra* was proposed independently. *Mitra* was further optimized to obtain *Mitra\** in [20].  $\Pi_{BP}$ -prime is similar to *Mitra\** except the following minor difference: the revealed document identifiers in  $\Pi_{BP}$ -prime are stored in plaintext whereas in *Mitra\** they are re-encrypted. This property of  $\Pi_{BP}$ -prime results in better locality. The leakage profile of constructions  $\Pi_{BP}$ -prime and *Mitra\** are same.

### 4.3 $\Pi_{BP}$ : Realizing Optimal Communication Complexity

As can be observed from the asymptotic analysis of  $\Pi_{BP}$ -prime, the communication complexity of Search protocol is  $O(o'_w)$  which is not optimal i.e.,  $O(n_w)$ . Further, as C has to process each ciphertext it receives from S, the computation complexity at C's end also becomes  $O(o'_w)$  due to the above communication overhead. In order to obtain optimal communication complexity, S should send entries corresponding to only the set of documents currently matching  $w$  ( $DB(w)$ ) to C. One approach to satisfy the above requirement is to associate a *tag* corresponding to each update entry. Using these tags, S, while performing a search on  $w$  will be able to correlate the update queries on  $w$  corresponding to the same *ind*. As in *Diana<sub>del</sub>* and *Janus*, if the tags are generated deterministically using just *ind* and  $w$ , it leaks  $DelHist(w)$  and hence, will not satisfy BP-II. In  $\Pi_{BP}$ , we leverage the version  $ver_w$  to generate the tags in a clever yet simple manner to ensure that the leakage is bounded by Definition 3.6.

Scheme  $\Pi_{BP}$  is described in Figure 10. For a keyword  $w$ ,  $\mathbf{D}$  stores the pointer to  $PSet_w$ , the set of tags corresponding to document identifiers that were obtained as a result of the latest search operation on  $w$ . The construction makes use of PRFs  $F_t$ ,  $F_d$ :  $\{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ , PRP  $G_{tag}$ :  $\{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$  and hash functions  $H_1$ :  $\{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$  and  $H_2$ :  $\{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^\mu$ , where  $\mu = \lambda + 1$ .

For an update query ( $op, (ind, w)$ ), a *tag* corresponding to  $(ind, w)$  is generated using the current version  $ver_w$ . Then,  $b||tag$  is stored in  $\mathbf{T}$  at S in Update protocol (see line 10).

Round 1 of Search protocol is similar to the Search protocol of  $\Pi_{FP}$ . At the end of round 1 of Search, S sends  $TS$ , the set of tags corresponding to the document identifiers currently matching  $w$  in order of their insertion. For every *tag* in  $TS$ , C computes  $G_{tag}^{-1}$  to get the document identifier *ind* which it adds to  $AuxSet$  and re-computes *tag* using the updated version which it stores in  $PSet$ .  $AuxSet$  consists of all the document identifiers currently matching keyword  $w$  and  $PSet$  consists of the updated tags, in order of their insertion.

<p><u>Setup()</u></p> <ol style="list-style-type: none"> <li>1. <math>k_t \xleftarrow{\\$} \{0, 1\}^\lambda, k_d \xleftarrow{\\$} \{0, 1\}^\lambda,</math>  <math>k_g \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>2. <math>\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow</math> empty map</li> <li>3. return ( EDB = ( <math>\mathbf{D}, \mathbf{T}</math> ),  <math>\text{st}_c = (k_t, k_d, k_g, \mathbf{W})</math>) to (S, C)</li> </ol> <p><u>Update(op, w, ind, st<sub>C</sub>; EDB)</u></p> <p><u>Update<sub>C</sub>(op, w, ind, st<sub>C</sub>)</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, c_w) = \perp</math> <b>then</b></li> <li>3.   <math>\text{ver}_w \leftarrow 0, c_w \leftarrow -1</math></li> <li>4.   <math>c_w \leftarrow c_w + 1</math></li> <li>5.   <math>k_w \leftarrow F_t(k_t, w    \text{ver}_w)</math></li> <li>6.   <math>\text{label} \leftarrow H_1(k_w    c_w)</math></li> <li>7.   <math>\text{pad} \leftarrow H_2(k_w    c_w)</math></li> <li>8.   <math>\text{tag} \leftarrow G_{\text{tag}}(k_g, w    \text{ind}    \text{ver}_w)</math></li> <li>9.   <math>b \leftarrow 0</math> (op=add) / <math>1</math> (op=del)</li> <li>10.   <math>e \leftarrow (b    \text{tag}) \oplus \text{pad}</math></li> <li>11.   <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, c_w)</math></li> <li>12.   Send (label, e) to S</li> </ol> <p><u>Update<sub>S</sub>(EDB)</u></p> <ol style="list-style-type: none"> <li>13. Receive (label, e) from C</li> <li>14. <math>\mathbf{T}[\text{label}] \leftarrow e</math></li> </ol>	<p><u>Search(w, st<sub>C</sub>; EDB)</u></p> <p><b>Round 1</b></p> <p><u>Search<sub>C</sub>(w, st<sub>C</sub>) :</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, c_w) = \perp</math> <b>then</b></li> <li>3.   return <math>\emptyset</math></li> <li>4. <b>else</b></li> <li>5.   <math>\text{label}_w \leftarrow F_d(k_d, w)</math></li> <li>6.   <b>if</b> <math>c_w \neq -1</math> <b>then</b></li> <li>7.     <math>k_w \leftarrow F_t(k_t, w    \text{ver}_w)</math></li> <li>8.     <math>\text{ver}_w \leftarrow \text{ver}_w + 1</math></li> <li>9.     <math>c_w \leftarrow -1</math></li> <li>10.    <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, c_w)</math></li> <li>11.   <b>else</b></li> <li>12.     <math>k_w \leftarrow \perp</math> //No need of round 2</li> <li>13.    Send (label<sub>w</sub>, k<sub>w</sub>, c<sub>w</sub>) to S</li> </ol> <p><u>Search<sub>S</sub>(EDB) :</u></p> <ol style="list-style-type: none"> <li>14. Receive (label<sub>w</sub>, k<sub>w</sub>, c<sub>w</sub>) from C</li> <li>15. <math>\text{addr}(\text{PSet}_w) \leftarrow \mathbf{D}[\text{label}_w]</math></li> <li>16. Retrieve PSet<sub>w</sub></li> <li>17. <math>\text{TS} \leftarrow \text{PSet}_w</math></li> <li>18. <b>if</b> <math>k_w \neq \perp</math> <b>then</b></li> <li>19.   <math>c \leftarrow 0</math></li> <li>20.   <b>while</b> <math>c \leq c_w</math> <b>do</b></li> </ol>	<ol style="list-style-type: none"> <li>21.   <math>\text{label} \leftarrow H_1(k_w    c)</math></li> <li>22.   <math>e \leftarrow \mathbf{T}[\text{label}]</math></li> <li>23.   <math>\text{pad} \leftarrow H_2(k_w    c)</math></li> <li>24.   <math>(b    \text{tag}) \leftarrow e \oplus \text{pad}</math></li> <li>25.   <b>if</b> <math>b = 0</math></li> <li>26.     <math>\text{TS} \leftarrow \text{TS} \cup \{\text{tag}\}</math></li> <li>27.   <b>else</b></li> <li>28.     <math>\text{TS} \leftarrow \text{TS} \setminus \{\text{tag}\}</math></li> <li>29.    Remove <math>\mathbf{T}[\text{label}]</math></li> <li>30.    <math>c \leftarrow c + 1</math></li> <li>31.    Send TS to C</li> </ol> <p><b>Round 2</b></p> <p><u>Search<sub>C</sub>(w, st<sub>C</sub>) :</u></p> <ol style="list-style-type: none"> <li>32. Receive TS from S</li> <li>33. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>34. <math>\text{PSet} \leftarrow \emptyset, \text{AuxSet} \leftarrow \emptyset</math></li> <li>35. <b>for</b> all tag <math>\in</math> TS <b>do</b></li> <li>36.   <math>(w    \text{ind}    \text{ver}') \leftarrow G_{\text{tag}}^{-1}(k_g, \text{tag})</math></li> <li>37.   <math>\text{AuxSet} \leftarrow \text{AuxSet} \cup \{\text{ind}\}</math></li> <li>38.   <math>\text{PSet} \leftarrow \text{PSet} \cup \{G_{\text{tag}}(k_g, w    \text{ind}    \text{ver}_w)\}</math></li> <li>39. return PSet, AuxSet</li> </ol> <p><u>Search<sub>S</sub>(EDB) :</u></p> <ol style="list-style-type: none"> <li>40. Receive PSet, AuxSet from C</li> <li>41. <math>\mathbf{D}[\text{label}_w] \leftarrow \text{PSet}</math></li> </ol>
---	--	---

Figure 10: Scheme  $\Pi_{\text{BP}}$

C sends PSet and AuxSet to S, who stores PSet at  $\mathbf{D}[\text{label}_w]$  and can use AuxSet to fetch the matching documents. Note that, recomputed tags enables S to consistently handle future searches and updates as the tags corresponding to subsequent updates on  $w$  are made using the same value of  $\text{ver}_w$ .

Let  $\mathbf{U}$  be the set of update queries corresponding to  $w$  after the previous search and before the current search on  $w$ . The tags are generated using the same  $\text{ver}_w$  exclusively  $\forall q_u \in \mathbf{U}$  and  $\forall \text{ind} \in \text{DB}(w)'$ . Since, the tags are computed using PRP  $G_{\text{tag}}$  taking  $w, \text{ind}$  and  $\text{ver}_w$  as input, S can only learn the relation between the tags corresponding to the same  $\text{ind}$  among the update queries in  $\mathbf{U}$  and  $\text{DB}(w)'$ . As the inds in PSet are stored in order of insertion, S can link the update queries in  $\mathbf{U}$  with these inds. The above leakage is precisely captured in BP-II via leakage functions UpPair, PreUp and CurUp. Moreover, S cannot relate update queries in  $\mathbf{U}$  with the update queries that are made before the previous search on  $w$  and after the current search on  $w$  as the tags are generated using different  $\text{ver}_w$ .

**Example 2** Consider the following list of update queries: (add, (1,  $w_1$ )), (add, (1,  $w_2$ )), (add, (2,  $w_1$ )), (add, (3,  $w_1$ )), (add, (3,  $w_2$ )), (del, (1,  $w_1$ )). Figure 9.a shows the state of indexes at C and S after these updates are processed. Figure 9.b shows the state of indexes at C and S after search on  $w_1$ .

In conclusion,  $\Pi_{\text{BP}}$  achieves optimal communication complexity and uses symmetric primitives only.  $\mathbf{D}$  provides reasonable locality as it stores the *tags* corresponding to previous search results together.  $\Pi_{\text{BP}}$  is easily parallelizable and doesn't impose reinsertion restriction.

Further, in order to obtain a response-hiding scheme, C sends only PSet to S in line 39 in second round

<p><u>Setup()</u></p> <ol style="list-style-type: none"> <li>1. <math>\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow</math> empty map</li> <li>2. <math>\text{bad} \leftarrow \text{false}</math></li> <li>3. return ( <math>\text{EDB} = (\mathbf{D}, \mathbf{T})</math>, <math>\text{st}_c = \mathbf{W}</math>) to (S, C)</li> </ol> <p><u>Update(op, w, ind, st<sub>c</sub>; EDB)</u></p> <p><u>Update<sub>c</sub>(op, w, ind, st<sub>c</sub>)</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, \text{c}_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, \text{c}_w) = \perp</math> <b>then</b></li> <li>3.   <math>\text{ver}_w \leftarrow 0, \text{c}_w \leftarrow -1</math></li> <li>4.   <math>\text{c}_w \leftarrow \text{c}_w + 1</math></li> <li>5.   <math>\text{k}_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> <li>6.   <math>\text{val}_r \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>7.   <b>if</b> <math>\mathbf{H}_1[\text{k}_w, \text{c}_w] \neq \perp</math> <b>then</b></li> <li>8.     <math>\text{bad} \leftarrow \text{true}, \boxed{\text{val}_r \leftarrow \mathbf{H}_1(\text{k}_w \  \text{c}_w)}</math></li> <li>9.   <math>\text{Hash}_1[w, \text{ver}_w, \text{c}_w] \leftarrow \text{val}_r</math></li> <li>10.   <math>\text{pad} \leftarrow \text{H}_2(\text{k}_w \  \text{c}_w)</math></li> <li>11.   <math>\text{tag} \leftarrow \text{Key}_{\text{tag}}[w, \text{ind}, \text{ver}_w]</math></li> <li>12.   <math>\text{b} \leftarrow 0</math> (op=add) / <math>1</math> (op=del)</li> <li>13.   <math>\text{e} \leftarrow (\text{b} \  \text{tag}) \oplus \text{pad}</math></li> </ol>	<ol style="list-style-type: none"> <li>14. <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, \text{c}_w)</math></li> <li>15. Send <math>(\text{val}_r, \text{e})</math> to S</li> </ol> <p><u>Search(w, st<sub>c</sub>; EDB)</u></p> <p><b>Round 1</b></p> <p><u>Search<sub>c</sub>(w, st<sub>c</sub>) :</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, \text{c}_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, \text{c}_w) = \perp</math> <b>then</b></li> <li>3.   return <math>\emptyset</math></li> <li>4. <b>else</b></li> <li>5.   <math>\text{label}_w \leftarrow \text{Key}_d[w]</math></li> <li>6.   <b>if</b> <math>\text{c}_w \neq -1</math> <b>then</b></li> <li>7.     <math>\text{k}_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> <li>8.     <b>for</b> <math>i = 0</math> to <math>\text{c}_w</math> <b>do</b></li> <li>9.       <math>\mathbf{H}_1[\text{k}_w, i] \leftarrow \text{Hash}_1[w, \text{ver}_w, i]</math></li> <li>10.      <math>\text{ver}_w \leftarrow \text{ver}_w + 1</math></li> <li>11.      <math>\text{c}_w \leftarrow -1</math></li> <li>12.      <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, \text{c}_w)</math></li> <li>13.   <b>else</b></li> <li>14.      <math>\text{k}_w \leftarrow \perp</math> //No need of round 2</li> <li>15.   Send <math>(\text{label}_w, \text{k}_w, \text{c}_w)</math> to S</li> </ol>	<p><b>Round 2</b></p> <p><u>Search<sub>c</sub>(w, st<sub>c</sub>) :</u></p> <ol style="list-style-type: none"> <li>16. Receive TS from S</li> <li>17. <math>(\text{ver}_w, \text{c}_w) \leftarrow \mathbf{W}[w]</math></li> <li>18. <math>\text{PSet} \leftarrow \emptyset, \text{AuxSet} \leftarrow \emptyset</math></li> <li>19. <b>for</b> all <math>\text{tag} \in \text{TS}</math> <b>do</b></li> <li>20.   <math>(w, \text{ind}, \text{ver}') \leftarrow \text{Key}_{\text{tag}}^{-1}[\text{tag}]</math></li> <li>21.   <math>\text{AuxSet} \leftarrow \text{AuxSet} \cup \{\text{ind}\}</math></li> <li>22.   <math>\text{tag}' \leftarrow \text{Key}_{\text{tag}}[w, \text{ind}, \text{ver}_w]</math></li> <li>23.   <math>\text{PSet} \leftarrow \text{PSet} \cup \{\text{tag}'\}</math></li> <li>24. return <math>\text{PSet}, \text{AuxSet}</math></li> </ol> <p><u>H<sub>1</sub>(k  c)</u></p> <ol style="list-style-type: none"> <li>1. <math>\text{val} \leftarrow \mathbf{H}_1[\text{k}, \text{c}]</math></li> <li>2. <b>if</b> <math>\text{val} = \perp</math> <b>then</b></li> <li>3.   <math>\text{val} \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>4.   <b>if</b> <math>\exists w, \text{ver}</math> s.t. <math>\text{ver} = \mathbf{W}[w].\text{ver}_w \wedge</math> <math>\text{k} = \text{Key}_t[w, \text{ver}] \wedge \text{c} \leq \mathbf{W}[w].\text{c}_w</math> <b>then</b></li> <li>5.     <math>\text{bad} \leftarrow \text{true}, \boxed{\text{val} \leftarrow \text{Hash}_1[w, \text{ver}, \text{c}]}</math></li> <li>6.   <math>\mathbf{H}_1[\text{k}, \text{c}] \leftarrow \text{val}</math></li> <li>7. return <math>\text{val}</math></li> </ol>
--	--	---

Figure 11: Games  $G_3$  and  $G'_3$  (Theorem 4.1).  $G'_3$  includes the box code and  $G_3$  does not.

of search protocol. We can eliminate the second round of communication using standard piggybacking technique [15, 18] and upload the updated tag set PSet with the next search query, thus, achieving a *single roundtrip* response-hiding backward private DSSE protocol.

**Correctness** Like  $\Pi_{\text{BP-prime}}$ , correctness of  $\Pi_{\text{BP}}$  immediately follows from that of  $\Pi_{\text{FP}}$ .

**Asymptotic Complexity** The communication complexity of the Search protocol is  $O(n_w)$ . The computational complexity of the Search protocol is  $O(o'_w)$ . The communication and computational cost of the Update protocol is  $O(1)$ . Space complexity at the server's end is  $O(N + D')$ , where  $D' = \sum_{w'} d'_w$  and at the client's end is  $O(m \log(n))$ .

**Security** We prove  $\Pi_{\text{BP}}$  is BP- $\Pi$  in the random oracle model. The proof relies on pseudo randomness of  $F_d, F_t$  and  $G_{\text{tag}}$ .

**Theorem 4.1.** *If  $F_d, F_t$  are secure PRFs,  $G_{\text{tag}}$  is a secure PRP and  $H_1, H_2$  are hash functions modeled as random oracles outputting  $2\lambda$  and  $\mu$  bits respectively then  $\Pi_{\text{BP}}$  is BP- $\Pi$  secure (Definition 3.6).*

*Proof.* We structure our proof using a sequence of games  $G_0$  to  $G_5$ .  $G_0$  will compute a distribution identical to  $\mathbf{Real}_{\mathcal{A}}^{\Pi_{\text{BP}}}(\lambda)$  and  $G_5$  will compute a distribution that can be simulated perfectly given the leakage profile  $\mathcal{L}$ , i.e., its distribution is identical to  $\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\Pi_{\text{BP}}}(\lambda)$ .

Game  $G_0$ :  $G_0$  is exactly identical to  $\mathbf{Real}_{\mathcal{A}}^{\Pi_{\text{BP}}}(\lambda)$ .

$$\Pr[\mathbf{Real}_{\mathcal{A}}^{\Pi_{\text{BP}}}(\lambda) = 1] = \Pr[G_0 = 1]. \quad (8)$$

Game  $G_1$ : In  $G_1$ , every call to PRFs  $F_t$  and  $F_d$  are answered using tables  $\text{Key}_t$  and  $\text{Key}_d$  respectively. The entries in table  $\text{Key}_t$  are referred by  $(w, \text{ver})$  and entries in table  $\text{Key}_d$  are referred by  $w$ . Conventionally, when an entry is being accessed for the first time, it is chosen at random and then used thereafter, which is followed in the rest of the paper unless mentioned explicitly. If there exists an adversary  $\mathcal{A}$  that is able to distinguish between games  $G_0$  and  $G_1$ , we can construct an adversary  $\mathcal{B}_1$  that can distinguish  $F_t$  from a truly random function and/or an adversary  $\mathcal{B}_2$  that can distinguish  $F_d$  from a truly random function. Formally, there exist adversaries  $\mathcal{B}_1$  and  $\mathcal{B}_2$ , such that

$$\begin{aligned} |\Pr[G_0 = 1] - \Pr[G_1 = 1]| &\leq \mathbf{Adv}_{F_t, \mathcal{B}_1}^{\text{prf}}(\lambda) + \mathbf{Adv}_{F_d, \mathcal{B}_2}^{\text{prf}}(\lambda) \\ &\leq \text{neg}(\lambda). \end{aligned} \tag{9}$$

Game  $G_2$ : In  $G_2$ , every call to PRP  $G_{\text{tag}}$  is answered using table  $\text{Key}_{\text{tag}}$ . The entries in table  $\text{Key}_{\text{tag}}$  are referred by  $(w, \text{ind}, \text{ver})$ . If the randomly generated  $\text{tag}$  has been selected earlier in  $\text{Key}_{\text{tag}}$ ,  $G_2$  is aborted. Since, the number of queries to PRP  $G_{\text{tag}}$ , say  $q$ , is a polynomial in security parameter, by the birthday bound we can conclude that the probability that the two tags are equal is at most  $\frac{q^2}{2^\lambda}$ , i.e.,  $\text{neg}(\lambda)$ . Therefore,  $G_2$  aborts with negligible probability.

Now, if there exists an adversary  $\mathcal{A}$  that is able to distinguish between games  $G_1$  and  $G_2$ , we can construct an adversary  $\mathcal{B}$  that can distinguish  $G_{\text{tag}}$  from a truly random permutation. Formally, there exists an adversary  $\mathcal{B}$ , such that

$$\begin{aligned} |\Pr[G_1 = 1] - \Pr[G_2 = 1]| &\leq \mathbf{Adv}_{G_{\text{tag}}, \mathcal{B}}^{\text{prp}}(\lambda) + \frac{q^2}{2^\lambda} \\ &\leq \text{neg}(\lambda). \end{aligned} \tag{10}$$

Game  $G_3$ : In  $G_3$ , instead of calling  $H_1$  to generate label in the **Update** protocol, we pick random strings. Then, during the **Search** protocol, the random oracle  $H_1$  is programmed accordingly, to ensure consistency.

Tables  $\text{Hash}_1$  and  $\mathbf{H}_1$  are used to simulate the random oracle  $H_1$ , the entries in the table  $\text{Hash}_1$  are referred by  $(w, \text{ver}, c)$  and in the table  $\mathbf{H}_1$  by  $(k, c)$ .

Figure 11 formally describes  $G_3$ , and an intermediate game  $G'_3$ . In  $G'_3$ ,  $H_1$  is never programmed to two different values for the same input, thus, ensuring consistency. Instead of storing the randomly picked value in table  $\text{Hash}_1$  at position  $(w, \text{ver}_w, c_w)$ , one first checks whether  $H_1$  is already programmed at value  $(k_w, c_w)$  which can happen if there was a query to the random oracle  $H_1$  with input  $(k_w \| c_w)$ . If the check is true, the value  $H_1(k_w \| c_w)$  is stored in  $\text{Hash}_1[w, \text{ver}_w, c_w]$  else the randomly picked value is stored in  $\text{Hash}_1[w, \text{ver}_w, c_w]$ . The random oracle when needed in the **Search** protocol in line 9 or by an adversary's query to random oracle  $H_1$  in line 5 is lazily programmed in  $G_3$ , so that the outputs are consistent throughout.

The only difference between game  $G_2$  and  $G'_3$  is how we model the random oracle  $H_1$ . The outputs of  $H_1$  is perfectly indistinguishable in both these games, therefore,

$$\Pr[G'_3 = 1] = \Pr[G_2 = 1]. \tag{11}$$

Let us denote the event 'the flag **bad** is set to true' in  $G'_3$  by  $\mathbb{E}_1$ . The games  $G'_3$  and  $G_3$  are also perfectly identical unless the event  $\mathbb{E}_1$  occurs, and we can apply *identical-until-bad* technique [3] to bound the distinguishing advantage between  $G'_3$  and  $G_3$ .

$$|\Pr[G'_3 = 1] - \Pr[G_3 = 1]| \leq \Pr[\mathbb{E}_1]. \tag{12}$$

The event  $\mathbb{E}_1$  occurs in line 8 of **Update** protocol and in line 5 of  $H_1$  algorithm. The former captures the fact that adversary has already queried random oracle  $H_1$  at input  $(k_w \| c_w)$  and the latter captures the fact that the adversary queries random oracle  $H_1$  on a valid input  $(k \| c)$  and the value  $k$  is not currently revealed to adversary. Since, the value  $k_w$  is picked uniformly at random, the probability with which event  $\mathbb{E}_1$  occurs is negligible. Using (11) and (12), we can conclude:

$$|\Pr[G_2 = 1] - \Pr[G_3 = 1]| \leq \text{neg}(\lambda). \tag{13}$$

<p><u>Setup()</u></p> <ol style="list-style-type: none"> <li>1. <math>\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow</math> empty map</li> <li>2. <math>u = 0</math></li> <li>3. <math>\mathbf{Update}, \mathbf{STs} \leftarrow</math> empty table</li> <li>4. <math>\mathbf{STs}[w][0] \leftarrow -1, \forall w \in \mathbf{W}</math></li> <li>5. return ( EDB = ( <math>\mathbf{D}, \mathbf{T}</math> ), <math>\mathbf{st}_c = \mathbf{W}</math> ) to ( <math>\mathbf{S}, \mathbf{C}</math> )</li> </ol> <p><u>Update(op, w, ind, st<sub>c</sub>; EDB)</u></p> <p><u>Update<sub>c</sub>(op, w, ind, st<sub>c</sub>)</u></p> <ol style="list-style-type: none"> <li>1. Append (u, op, ind) to <math>\mathbf{Update}[w]</math></li> <li>2. <math>\mathbf{Hash}_1[u] \xleftarrow{\\$} \{0, 1\}^{2\lambda}</math></li> <li>3. <math>\mathbf{Hash}_2[u] \xleftarrow{\\$} \{0, 1\}^\mu</math></li> <li>4. Send ( <math>\mathbf{Hash}_1[u], \mathbf{Hash}_2[u]</math> ) to <math>\mathbf{S}</math></li> <li>5. <math>u \leftarrow u + 1</math></li> </ol> <p><u>Search(w, st<sub>c</sub>; EDB)</u></p> <p><b>Round 1</b></p> <p><u>Search<sub>c</sub>(w, st<sub>c</sub>) :</u></p> <ol style="list-style-type: none"> <li>1. Append u to <math>\mathbf{STs}[w]</math></li> <li>2. ( empty, <math>\mathbf{ver}_w, c_w, \{\mathbf{tag}_c, \mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}</math> ) <math>\leftarrow</math> <u>GetData<sub>r1</sub></u>( <math>\mathbf{Update}, \mathbf{STs}, \mathbf{Hash}_1, \mathbf{Hash}_2, \mathbf{Tag}</math> )</li> <li>3. <b>if</b> empty = true <b>then</b></li> </ol>	<ol style="list-style-type: none"> <li>4. return <math>\emptyset</math></li> <li>5. <b>else</b></li> <li>6. <math>\mathbf{label}_w \leftarrow \mathbf{Key}_d[w]</math></li> <li>7. <b>if</b> <math>c_w \neq -1</math> <b>then</b></li> <li>8. <math>k_w \leftarrow \mathbf{Key}_t[w, \mathbf{ver}_w]</math></li> <li>9. <b>for</b> <math>i = 0</math> to <math>c_w</math> <b>do</b></li> <li>10. <math>\mathbf{H}_1[k_w, i] \leftarrow \mathbf{H}_{1,w,i}</math></li> <li>11. <math>\mathbf{H}_2[k_w, i] \leftarrow \mathbf{H}_{2,w,i} \oplus \mathbf{tag}_i</math></li> <li>12. <b>else</b></li> <li>13. <math>k_w \leftarrow \perp</math></li> <li>14. Send ( <math>\mathbf{label}_w, k_w, c_w</math> ) to <math>\mathbf{S}</math></li> </ol> <p><b>Round 2</b></p> <p><u>Search<sub>c</sub>(w, st<sub>c</sub>) :</u></p> <ol style="list-style-type: none"> <li>15. Receive TS from <math>\mathbf{S}</math></li> <li>16. <math>\mathbf{PSet}_w, \mathbf{AuxSet}_w \leftarrow</math> <u>GetData<sub>r2</sub></u>( <math>\mathbf{Tag}, \mathbf{TS}</math> )</li> <li>17. return <math>\mathbf{PSet}_w, \mathbf{AuxSet}_w</math></li> <li>18. <math>u \leftarrow u + 1</math></li> </ol> <p><u>GetData<sub>r2</sub>(Tag, TS)</u></p> <ol style="list-style-type: none"> <li>1. <math>\mathbf{PSet} \leftarrow \emptyset</math></li> <li>2. <math>\mathbf{Tag}_{\mathbf{temp}} \leftarrow</math> empty map</li> <li>3. <b>for</b> <math>\forall \mathbf{tag} \in \mathbf{TS}</math> <b>do</b></li> <li>4. <math>(w, \mathbf{ind}) \leftarrow \mathbf{Tag}[w]^{-1}[\mathbf{tag}]</math></li> <li>5. <math>\mathbf{PSet} \leftarrow \mathbf{PSet} \cup \{\mathbf{Tag}_{\mathbf{temp}}[\mathbf{ind}]\}</math></li> <li>6. <math>\mathbf{AuxSet} \leftarrow \mathbf{AuxSet} \cup \{\mathbf{ind}\}</math></li> </ol>	<ol style="list-style-type: none"> <li>7. <math>\mathbf{Tag}[w] \leftarrow \mathbf{Tag}_{\mathbf{temp}}</math></li> <li>8. return <math>\mathbf{PSet}, \mathbf{AuxSet}</math></li> </ol> <p><u>GetData<sub>r1</sub>(Update, STs, Hash<sub>1</sub>, Hash<sub>2</sub>, Tag)</u></p> <ol style="list-style-type: none"> <li>1. <b>if</b> <math>\mathbf{Update}[w] = \perp</math></li> <li>2. empty <math>\leftarrow</math> true</li> <li>3. <b>else</b></li> <li>4. empty <math>\leftarrow</math> false</li> <li>5. <math>\mathbf{ver}_w \leftarrow -1</math></li> <li>6. <b>for</b> <math>i = 1</math> to <math> \mathbf{STs}[w]  - 1</math> <b>do</b></li> <li>7. <b>if</b> <math>\exists u, \mathbf{op}, \mathbf{ind}</math> s.t. <math>(u, \mathbf{op}, \mathbf{ind}) \in \mathbf{Update}[w] \wedge u &gt; \mathbf{STs}[w][i - 1] \wedge u &lt; \mathbf{STs}[w][i]</math> <b>then</b></li> <li>8. <math>\mathbf{ver}_w \leftarrow \mathbf{ver}_w + 1</math></li> <li>9. <math>c_w \leftarrow -1, ul \leftarrow  \mathbf{STs}[w]  - 1</math></li> <li>10. <b>for</b> <math>u = \mathbf{STs}[w][ul - 1]</math> to <math>\mathbf{STs}[w][ul]</math> <b>do</b></li> <li>11. <b>if</b> <math>\exists \mathbf{op}, \mathbf{ind}</math> s.t. <math>(u, \mathbf{op}, \mathbf{ind}) \in \mathbf{Update}[w]</math> <b>then</b></li> <li>12. <math>c_w \leftarrow c_w + 1</math></li> <li>13. <math>\mathbf{b} \leftarrow 0</math> (op=add)/ 1 (op=del)</li> <li>14. <math>\mathbf{tag}_{c_w} \leftarrow \mathbf{b} \parallel \mathbf{Tag}[w, \mathbf{ind}]</math></li> <li>15. <math>\mathbf{H}_{1,w,c_w} \leftarrow \mathbf{Hash}_1[u]</math></li> <li>16. <math>\mathbf{H}_{2,w,c_w} \leftarrow \mathbf{Hash}_2[u]</math></li> <li>17. return ( empty, <math>\mathbf{ver}_w, c_w, \{\mathbf{tag}_c, \mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}</math> )</li> </ol>
---	--	---

Figure 12: Game  $\mathbf{G}_5$  (Theorem 4.1)

Game  $\mathbf{G}_4$ : In  $\mathbf{G}_4$ , what we did for  $\mathbf{H}_1$  in game  $\mathbf{G}_3$ , we do for  $\mathbf{H}_2$ . Using the same arguments, we can conclude:

$$|\Pr[\mathbf{G}_3 = 1] - \Pr[\mathbf{G}_4 = 1]| \leq \text{neg}(\lambda). \quad (14)$$

Game  $\mathbf{G}_5$ : In  $\mathbf{G}_5$  (see Figure 12), we abstract out the information that needs to be simulated by the simulator in order to output transcripts identical to  $\mathbf{G}_4$ . Using GetData<sub>r1</sub> and GetData<sub>r2</sub> algorithms in  $\mathbf{G}_5$ , one keeps track of the randomly generated  $\mathbf{tag}$ ,  $\mathbf{label}$  and  $\mathbf{pad}$  differently than in  $\mathbf{G}_4$ . In Search protocol, the random oracles are programmed identically to that in  $\mathbf{G}_4$ . Queries to random oracles  $\mathbf{H}_1$  and  $\mathbf{H}_2$  can be simulated by outputting random values.

As we output fresh random strings in Update protocol, the transcripts of Update protocol is identical to that of Update protocol in  $\mathbf{G}_4$ .

Next, we describe the Search protocol in  $\mathbf{G}_5$ . Based on tables **Update**, **STs** (search timestamps) and **Tag** map, the value of the following components:  $\mathbf{empty}_w, \mathbf{ver}_w, c_w$  and  $\{\mathbf{tag}_c, \mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}$  are determined using GetData<sub>r1</sub> algorithm. Flag  $\mathbf{empty}$  is set to 1 if  $\mathbf{Update}[w]$  is empty.  $\mathbf{ver}_w$  is the count of searches for which there was an update on map  $\mathbf{T}$  corresponding to keyword  $w$  after the previous search. The loop in line 6 of GetData<sub>r1</sub> algorithm determines the number of times version number is updated, i.e., value of  $\mathbf{ver}_w$ . Here,  $c_w$  denotes the count of updates on map  $\mathbf{T}$  corresponding to keyword  $w$  after the previous search.  $\mathbf{tag}$  is picked from the map  $\mathbf{Tag}$ , therefore, if the indices are same, the tags are equal, thus ensuring the consistency. The values  $\{\mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}$  are used to simulate the random oracles consistently with the response given at the time of update queries. The loop in line 10 of GetData<sub>r1</sub> algorithm



computes the values of  $c_w$  and  $\{\text{tag}_c, \mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}$ .  $\text{GetData}_{r_2}$  algorithm outputs  $\text{LDB}(w)$  and new tags corresponding to the document identifiers currently present in  $\text{DB}(w)$ . These tags are ordered based on the order of insertion of documents they correspond to. As the values of components are computed correctly and consistently w.r.t. previous queries, the transcripts of  $\text{Search}$  protocol is identical to that of  $\text{Search}$  protocol in  $\mathbf{G}_4$ . Therefore, we conclude that:

$$\Pr[\mathbf{G}_5 = 1] = \Pr[\mathbf{G}_4 = 1]. \quad (15)$$

**Simulator Sim:** Finally, we construct a simulator that given the leakage profile  $\mathcal{L}$  simulates game  $\mathbf{G}_5$  correctly.  $\text{Sim}$  can simulate  $\text{Update}$  protocol correctly as in  $\mathbf{G}_5$ . Instead of using keyword  $w$ ,  $\text{Sim}$  uses the counter  $\bar{w} = \min \text{sp}(w)$  uniquely mapped from  $w$  using  $\mathcal{L}_{\text{Search}}$  in simulating the  $\text{Search}$  protocol (line 6 and line 8). In line 2 of  $\text{Search}$  protocol,  $\text{Sim}$  uses  $\text{sp}(w)$ ,  $\text{Updates}^{\text{op}}(w)$ ,  $\text{PreUp}(w)$  and  $\text{UpPair}(w)$  instead of  $\text{STs}$ ,  $\text{Update}$  and  $\text{Tag}$  as input to the  $\text{GetData}_{r_1}$  algorithm. In  $\text{GetData}_{r_1}$ , we use the timestamps of search and update queries and make use of indices to generate  $\text{tag}$ . However, the indices are used just to identify when same tags need to be generated. This can be ensured using,  $\text{PreUp}(w)$  and  $\text{UpPair}(w)$ .

Also, the output of  $\text{GetData}_{r_2}$  is  $\text{LDB}(w)$  and an ordered list of freshly generated random tags which can be simulated easily using  $\text{CurUp}(w)$  and  $\text{LDB}(w)$ . In  $\text{GetData}_{r_2}$ , the indices are used to just associate an order to the generated tags, which can be ensured by the components  $\text{CurUp}(w)$  and  $\text{LDB}(w)$  of the leakage profile. Thus,  $\text{Sim}$  is able to produce transcripts of output of  $\text{Search}$  and  $\text{Update}$  protocols identical to  $\mathbf{G}_5$ . Hence, we conclude that:

$$\Pr[\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\Pi_{\text{BP}}}(\lambda) = 1] = \Pr[\mathbf{G}_5]. \quad (16)$$

By connecting all the games, we conclude

$$\left| \Pr[\mathbf{Real}_{\mathcal{A}}^{\Pi_{\text{BP}}}(\lambda) = 1] - \Pr[\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\Pi_{\text{BP}}}(\lambda) = 1] \right| \leq \text{neg}(\lambda). \quad (17)$$

□

#### 4.4 $\Pi_{\text{WBP}}$ : A Weak Backward Private Variant

As mentioned in Section 3.3, there are various use-cases such as storing a collection of text files in which reinsertion restriction is not a concern. Here, we propose a simple one roundtrip response-hiding backward private scheme  $\Pi_{\text{WBP}}$  in the *reinsertion restriction* setting. Our construction  $\Pi_{\text{WBP}}$  is essentially a simple modification to  $\Pi_{\text{BP}}$ . Further, it also inherits all the salient features of  $\Pi_{\text{BP}}$ .  $\Pi_{\text{WBP}}$  improves upon the concrete efficiency of  $\text{Search}$  protocol in  $\Pi_{\text{BP}}$ , by eliminating the computation and transmission of newly generated tags in the second round of communication. This shows that one can construct an efficient WBP scheme with optimal communication complexity using simple primitives only.

In  $\Pi_{\text{BP}}$ , we use  $\text{ver}_w$  in computation of tags to securely handle reinsertions. The tags corresponding to fresh updates after a search are computed using the updated  $\text{ver}_w$ . Hence, these tags cannot be related to the deleted entries before this search. This ensures backward privacy in cases where reinsertion is allowed. However, when we consider the reinsertion restriction setting, an add query is not allowed after a delete query corresponding to the same document-keyword pair. Hence,  $\text{ver}_w$  is not required in computation of tags in this setting. Therefore, for an update query,  $\text{tag}$  in line 8 of  $\text{Update}$  protocol in  $\Pi_{\text{BP}}$  (see Figure 10) is computed as  $\text{tag} \leftarrow G_{\text{tag}}(k_g, w \parallel \text{ind})$  in  $\Pi_{\text{WBP}}$ .

$\text{Search}$  protocol in  $\Pi_{\text{WBP}}$  is identical to Round 1 of  $\text{Search}$  protocol in  $\Pi_{\text{BP}}$ . In line 31 of  $\text{Search}$  protocol in  $\Pi_{\text{BP}}$ ,  $\text{S}$  along with sending  $\text{TS}$  to  $\text{C}$ , stores  $\text{TS}$  at  $\mathbf{D}[\text{label}_w]$ . From  $\text{TS}$ ,  $\text{C}$  retrieves the search results in similar fashion as in  $\Pi_{\text{BP}}$ . Recomputation of tags in  $\text{Search}$  protocol in  $\Pi_{\text{BP}}$  is needed because  $\text{ver}_w$  gets updated. While in  $\Pi_{\text{WBP}}$ , since the tags are independent of  $\text{ver}_w$ , recomputation of tags is not required. Hence, the second round of communication is not needed.

The changes we make in  $\Pi_{\text{BP}}$  in constructing  $\Pi_{\text{WBP}}$  induces additional leakage which can be shown to be bounded by the leakage profile of WBP. The proof of Theorem 4.1 can be easily adapted to prove security of  $\Pi_{\text{WBP}}$ .

The crucial observation from our constructions is that efficient backward private schemes can be realized without involving complex cryptographic primitives. The simplicity of design in our backward private constructions is an appealing feature, particularly from the implementation perspective.

## 4.5 Comparative Analysis

In Table 1, we provided a comparison of our schemes with some prior and concurrent works [7, 20, 40]. On that line, we conclude this section with a comparative study of the currently available BP-secure DSSE schemes. The goal is to figure out the scenarios in which each of these constructions would fit best. Let us first consider the scenario where the requirement is to achieve minimal information leakage. The candidate constructions are Orion [20] and Moneta [7] as they satisfy strong notion of backward privacy (BPIP). As the search time of Orion is quasi-optimal in  $n_w$  (linear in  $n_w$  upto a logarithmic factor), it may appear to be more suitable than Moneta in such scenarios. However, Orion may not be practical for very large databases as Path-ORAM [39] is used as a building block in its construction, which limits the applicability of Orion in such scenarios [32]. Constructions Mitra [20], Fides [7] and  $\Pi_{\text{BP-prime}}$  (Section 4.2) satisfy the next level of backward privacy. Mitra and  $\Pi_{\text{BP-prime}}$  are essentially the same, make use of symmetric primitives only and as a result are very efficient in practice. But, these constructions still suffer from significant communication overhead.

In order to overcome the above limitations, one can look to trade security a bit for performance, while at the same time ensure that the notion of forward and backward privacy is preserved. Constructions  $\text{Diana}_{\text{del}}$  [7], Janus [7], Janus++ [40], Horus [20] and  $\Pi_{\text{WBP}}$  (Section 4.2) satisfy the notion of weak backward privacy (WBP). The communication and computation complexity of search and update protocols of  $\text{Diana}_{\text{del}}$  isn't optimal (See Table 1). To improve upon the communication overhead of Search protocol, a single round-trip Janus framework [7] was proposed. It was instantiated using asymmetric and symmetric puncturable encryption scheme in Janus [7] and Janus++ [40] respectively. However, the computational complexity of search protocol in Janus framework is  $O(n_w \cdot d_w)$ , which is unreasonably high ( $n_w \cdot d_w \gg o'_w$ ). Horus, a modified version of Orion, was proposed in order to improve the number of round trips in the Search protocol ( $O(\log(N))$  to  $O(\log(d_w))$ ). But Horus suffers from the same scalability issue as Orion.  $\Pi_{\text{WBP}}$  is a single-round trip DSSE scheme that achieves optimal communication complexity, makes use of symmetric primitives only and is very efficient in practice. However, the computation complexity of the Search protocol is not quasi-optimal in  $n_w$ . Moreover, the notion of weak backward privacy can only be used in scenarios where reinsertion of keyword-document is not allowed.

BP-II (Definition 3.6) along with satisfying the intuitive notion of backward privacy, allows the reinsertion of keyword-document pairs. The corresponding construction,  $\Pi_{\text{BP}}$  (Section 4.3) is the first DSSE scheme that satisfies the notion of backward privacy in the general setting that achieves optimal communication complexity, makes use of symmetric primitives only and is very efficient in practice. The only limitation is that the search time in  $\Pi_{\text{BP}}$  is not quasi-optimal in  $n_w$  which seems to be the cost that one has to pay to achieve optimal communication complexity in Update and Search protocol.

## 5 Implementation Results

In this section, we discuss the implementation results of schemes  $\Pi_{\text{FP}}$  and  $\Pi_{\text{BP}}$ .  $\Pi_{\text{FP}}$  being currently the most efficient forward private scheme in literature, serves as a benchmark to evaluate the performance of other DSSE schemes. The implementation results for  $\Pi_{\text{BP}}$  gives a fair indication about the performance of  $\Pi_{\text{BP-prime}}$  and  $\Pi_{\text{WBP}}$  as their asymptotic complexity are same and they also make use of light-weight symmetric primitives. Further, the concrete computation cost of  $\Pi_{\text{WBP}}$  reduces further as the tags need not be recomputed. We have implemented the schemes in C++. For pseudo random functions  $F_d, F_t$  and pseudo random permutation  $G_{\text{tag}}$ , we use AES, and for hash functions  $H_1$  and  $H_2$ , we use SHA-256. We use the AES and SHA-256 function available in OpenSSL library [41] in our code. Maps  $\mathbf{T}$  and  $\mathbf{W}$  are stored using RocksDB [17].

All our experiments were performed on a desktop computer with an Intel Core i5 4460 3.20GHz CPU and

Implementation	Time (s)	Time per pair * ( $\mu$ s)	Strg. at S (GB)	Strg. at C (MB)
$\Pi_{FP}$ (in main memory)	164.66	4.49	2.8	2.4
$\Pi_{FP}$ (in ext. disk)	172.66	4.71	2.8	2.4
$\Pi_{BP}$ (in main memory)	182.90	4.98	2.8	2.4
$\Pi_{BP}$ (in ext. disk)	193.74	5.28	2.8	2.4

\* - document-keyword pair, Strg. = Storage

Table 2: EDB Creation




Figure 13: Average per entry search time (Main memory)

8GB RAM running Ubuntu 16.04 LTS. Our code is designed to run as a single program as we are interested in determining the performance of **Search** and **Update** protocols of our constructions.

We used Enron email dataset [1] to create EDB on which we perform our search and update operations. We wrote a python code to extract keywords from the mails in Enron email dataset using NLTK library [34]. The number of documents, number of keywords and number of document-keyword pairs in our dataset are 517,401, 212,020 and 36,688,028 respectively.

**EDB Creation** EDB was created to store all the document-keyword pairs extracted from the Enron email dataset. The computational works and I/O latency required for EDB creation are parallelized using thread pool. Table 2 shows the time taken to create an EDB, the time taken to process each document-keyword pair and the size of EDB and **W** just after EDB creation. For in memory results, only map **T** was mounted on to main memory.

For each entry in Table 2, we ran our experiment 10 times and computed the average. The time taken to create EDB in main memory and external disk is similar for both the schemes. The time taken to create EDB for  $\Pi_{BP}$  is reasonably close to that of  $\Pi_{FP}$ . The per document-keyword processing time is very less for both the schemes as only symmetric primitives are used in our constructions.




Figure 14: Average per entry search time (External disk)




Figure 15: Average per entry search time (# Threads)

**EDB Search** To evaluate the search performance, we searched all the keywords extracted from the Enron email dataset just after EDB creation and measured the overall time taken for the Search protocol. The performance of both the schemes when  $\mathbf{T}$  is mounted on to main memory and when  $\mathbf{T}$  is stored on external disk was evaluated. Figure 13 describes the search time per matched entry ( $\text{stpme}$ ) based on the number of documents returned in the search results for both the schemes when  $\mathbf{T}$  is mounted on to main memory. It illustrates the performance in single threaded and multi-threaded environment. Figure 14 describes the same when  $\mathbf{T}$  is stored on the external hard disk. In Figures 13 and 14,

$$RS(i) = \begin{cases} \text{DB}(w) = 1 & \text{if } i = 0 \\ 2^{i-1} < \text{DB}(w) \leq 2^i & \text{if } 1 \leq i \leq 18. \end{cases}$$

Figure 15 illustrates that the  $\text{stpme}$  for both the schemes is affected by the number of threads used to perform the search operation. For searches matching less number of documents, the cost is high because of one time computations such as storage access in computation of token at  $\mathbf{C}$ 's end, creation of threads, access to map  $\mathbf{D}$ , etc., which is amortized for searches matching large number of documents. The performance of both the schemes improve significantly in multi-threaded environment, thus, highlighting the potential for parallelism in practice. The schemes performed the best when the number of threads were around the number of cores in the processor, i.e., 4. The performance of  $\Pi_{BP}$  is comparable to the performance of  $\Pi_{FP}$ . The  $\text{stpme}$  in  $\Pi_{BP}$  was 4.50 (resp. 4.89)  $\mu\text{s}$  for queries whose result set size was in the interval  $(2^{17}, 2^{18}]$  and the average  $\text{stpme}$  was 5.59 (resp. 5.95)  $\mu\text{s}$  for all queries when  $\mathbf{T}$  was mounted on to main memory (resp. stored on external disk).

**EDB Dynamic Environment** To show the performance of search queries in dynamic environment, i.e., where search queries are interspersed with update queries, we identified the set of keywords matching more than 80k documents from the extracted document-keyword pairs, say  $\mathbf{S}_{80k}$ . An initial EDB was constructed using extracted document-keyword pairs apart from those corresponding to keywords in  $\mathbf{S}_{80k}$ . The document-keyword pairs corresponding to keywords in  $\mathbf{S}_{80k}$  are then utilized to perform update queries dynamically. Figure 16 describes the  $\text{stpme}$  with regard to the probability ( $p$ ) of search queries on keywords in  $\mathbf{S}_{80k}$ . This implies that the update queries occur with probability  $1-p$ , of which, with probability 0.1 it is a del query. The performance evaluation includes the time required to process del operation in search queries. For both the schemes, Figure 16 illustrates that the average  $\text{stpme}$  decreases in scenarios where search queries are frequent, as both the schemes exploit the locality introduced by  $\mathbf{D}$ . The average  $\text{stpme}$  in  $\Pi_{BP}$  turns out to be 1.53  $\mu\text{s}$ , when the probability of search query is 0.0005.

The prototype implementations of  $\Pi_{FP}$  and  $\Pi_{BP}$  indicate that the cost of achieving backward privacy over and above forward privacy is substantially small. This makes  $\Pi_{BP}$  a suitable candidate for practical deployment.

**Communication Cost** As all our constructions make use of symmetric primitives only, the communication cost becomes the main performance bottleneck. We now compare the communication cost of  $\Pi_{BP}$ -prime,  $\Pi_{BP}$  and  $\Pi_{WBP}$  as a function of the nature of update queries. Figure 17 depicts the communication cost of the trio based on the number of documents returned in the search results and the probability with which an




Figure 16: Average per entry search time (Dynamic Environment)

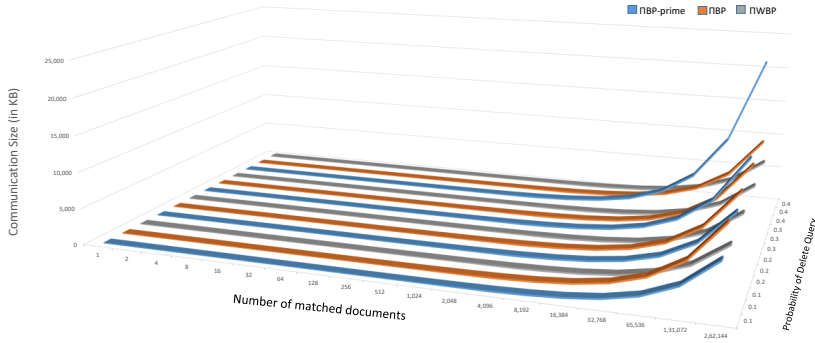


Figure 17: Communication Cost: Number of matching documents vs. Probability of delete queries

update query is a del query. As can be observed, for  $\Pi_{BP}$  and  $\Pi_{WBP}$  the communication cost depends only upon the result set size and is independent of the nature of update queries. Therefore, the communication complexity is the same for all types of update query distribution for  $\Pi_{BP}$  and  $\Pi_{WBP}$ . However, that is not the case for  $\Pi_{BP}$ -prime, as can be observed from Figure 17, the communication complexity increases with increase in the percentage of delete queries. But note that if the percentage of delete query is less, then the communication cost of  $\Pi_{BP}$ -prime is less than  $\Pi_{BP}$  in practice. Concretely,  $\Pi_{WBP}$  has the least communication overhead among all our backward private constructions in practice. As already mentioned in Section 4.2,  $\Pi_{BP}$ -prime avoids re-encryption of entries and thus improves upon the communication cost of *Mitra*<sup>\*</sup>. Hence, the communication cost of our constructions fair reasonably well against the most efficient construction until this work.

## 6 Conclusion

The main contribution of this paper is to propose three efficient backward private DSSE schemes, viz.,  $\Pi_{BP}$ -prime,  $\Pi_{BP}$  and  $\Pi_{WBP}$ . We start with revisiting the existing definitions of backward privacy and propose alternative formulations of leakage for backward privacy, viz., BP-I and BP-II. The proposed constructions achieve practical efficiency by using light weight symmetric cryptographic components only. In particular, our construction  $\Pi_{BP}$  is the first backward private scheme in the general setting that achieves optimal communication complexity using symmetric cryptographic primitives only. The main takeaway from this work is that efficient backward private schemes can be realized without involving complex cryptographic primitives. The simplicity of their design make our backward private constructions even more appealing, particularly from the implementation perspective. On the definition front, an interesting question arising out of this study is to analyze the real-world consequences of the leakages incurred by constructions satisfying the definitions of backward privacy, viz., BP-I, BP-II or even WBP (in restricted setting). On

the construction front, an interesting problem to pursue is to design an efficient, single roundtrip, response revealing backward private scheme in the general setting ideally with optimal communication complexity.

## References

- [1] “Enron Email Dataset,” <https://www.cs.cmu.edu/enron/>, Accessed: 2018-05-14.
- [2] M. A. Abdelraheem, T. Andersson, and C. Gehrman, “Inference and Record-Injection Attacks on Searchable Encrypted Relational Databases,” *IACR Cryptology ePrint Archive*, vol. 2017, p. 24, 2017.
- [3] M. Bellare and P. Rogaway, “The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs,” in *EUROCRYPT*, ser. LNCS, vol. 4004. Springer, 2006, pp. 409–426.
- [4] D. Boneh and B. Waters, “Constrained Pseudorandom Functions and Their Applications,” in *ASIACRYPT*, ser. LNCS, vol. 8270. Springer, 2013, pp. 280–300.
- [5] R. Bost, “Σοφος: Forward Secure Searchable Encryption,” in *ACM CCS*. ACM Press, 2016, pp. 1143–1154.
- [6] —, “Searchable Encryption: New Constructions of Encrypted Databases,” Ph.D. dissertation, 2018.
- [7] R. Bost, B. Minaud, and O. Ohrimenko, “Forward and Backward Private Searchable Encryption from Constrained Cryptographic Primitives,” in *ACM CCS*. ACM Press, 2017, pp. 1465–1482.
- [8] E. Boyle, S. Goldwasser, and I. Ivan, “Functional Signatures and Pseudorandom Functions,” in *PKC*, ser. LNCS, vol. 8383. Springer, 2014, pp. 501–519.
- [9] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart, “Leakage-Abuse Attacks Against Searchable Encryption,” in *ACM CCS*. ACM Press, 2015, pp. 668–679.
- [10] D. Cash, J. Jaeger, S. Jarecki, C. S. Jutla, H. Krawczyk, M. Rosu, and M. Steiner, “Dynamic Searchable Encryption in Very-Large Databases: Data Structures and Implementation,” in *NDSS*. The Internet Society, 2014.
- [11] D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Roşu, and M. Steiner, “Highly-Scalable Searchable Symmetric Encryption with Support for Boolean Queries,” in *CRYPTO*, ser. LNCS. Springer, 2013, vol. 8042, pp. 353–373.
- [12] Y. Chang and M. Mitzenmacher, “Privacy preserving keyword searches on remote encrypted data,” in *ACNS*, ser. Lecture Notes in Computer Science, vol. 3531, 2005, pp. 442–455.
- [13] M. Chase and S. Kamara, “Structured Encryption and Controlled Disclosure,” in *ASIACRYPT*, ser. LNCS, vol. 6477. Springer, 2010, pp. 577–594.
- [14] R. Curtmola, J. A. Garay, S. Kamara, and R. Ostrovsky, “Searchable Symmetric Encryption: Improved Definitions and Efficient Constructions,” in *ACM CCS*. ACM Press, 2006, pp. 79–88.
- [15] M. Etemad, A. Küpçü, C. Papamanthou, and D. Evans, “Efficient Dynamic Searchable Encryption with Forward Privacy,” *PoPETs*, vol. 2018, no. 1, pp. 5–20, 2018.
- [16] S. Faber, S. Jarecki, H. Krawczyk, Q. Nguyen, M. Rosu, and M. Steiner, “Rich Queries on Encrypted Data: Beyond Exact Matches,” in *ESORICS*, ser. LNCS, vol. 9327. Springer, 2015, pp. 123–145.
- [17] Facebook, “RocksDB: A persistent key-value store for fast storage environment,” <https://rocksdb.org/>, Accessed: 2018-05-14.
- [18] S. Garg, P. Mohassel, and C. Papamanthou, “TWRORAM: Efficient Oblivious RAM in Two Rounds with Applications to Searchable Encryption,” in *CRYPTO*, ser. LNCS, vol. 9816. Springer, 2016, pp. 563–592.

- [19] C. Gentry, “A Fully Homomorphic Encryption Scheme,” Ph.D. dissertation, Stanford, CA, USA, 2009, aAI3382729.
- [20] J. Ghareh Chamani, D. Papadopoulos, C. Papamanthou, and R. Jalili, “New constructions for forward and backward private symmetric searchable encryption,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18. ACM, pp. 1038–1055.
- [21] O. Goldreich, S. Goldwasser, and S. Micali, “How to Construct Random Functions (Extended Abstract),” in *FOCS*. IEEE Computer Society Press, 1984, pp. 464–479.
- [22] O. Goldreich and R. Ostrovsky, “Software Protection and Simulation on Oblivious RAMs,” *J. ACM*, no. 3, pp. 431–473, 1996.
- [23] S. Goldwasser, S. Micali, and R. L. Rivest, “A ”paradoxical” solution to the signature problem (extended abstract),” in *FOCS*. IEEE Computer Society, 1984, pp. 441–448.
- [24] —, “A digital signature scheme secure against adaptive chosen-message attacks,” *SIAM J. Comput.*, vol. 17, no. 2, pp. 281–308, 1988.
- [25] M. D. Green and I. Miers, “Forward Secure Asynchronous Messaging from Puncturable Encryption,” in *IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 2015, pp. 305–320.
- [26] S. Kamara and T. Moataz, “Boolean Searchable Symmetric Encryption with Worst-Case Sub-linear Complexity,” in *EUROCRYPT*, ser. LNCS, vol. 10212. Springer, 2017, pp. 94–124.
- [27] S. Kamara, C. Papamanthou, and T. Roeder, “Dynamic Searchable Symmetric Encryption,” in *ACM CCS*. ACM Press, 2012, pp. 965–976.
- [28] A. Kiayias, S. Papadopoulos, N. Triandopoulos, and T. Zacharias, “Delegatable Pseudorandom Functions and Applications,” in *ACM CCS*. ACM Press, 2013, pp. 669–684.
- [29] K. S. Kim, M. Kim, D. Lee, J. H. Park, and W. Kim, “Forward Secure Dynamic Searchable Symmetric Encryption with Efficient Updates,” in *ACM CCS*. ACM Press, 2017, pp. 1449–1463.
- [30] N. Kobitz and A. Menezes, “Another look at security definitions,” *Adv. in Math. of Comm.*, vol. 7, no. 1, pp. 1–38, 2013.
- [31] A. Menezes and N. P. Smart, “Security of signature schemes in a multi-user setting,” *Des. Codes Cryptography*, vol. 33, no. 3, pp. 261–274, 2004.
- [32] M. Naveed, “The fallacy of composition of oblivious RAM and searchable encryption,” *IACR Cryptology ePrint Archive*, vol. 2015, p. 668, 2015.
- [33] M. Naveed, M. Prabhakaran, and C. A. Gunter, “Dynamic Searchable Encryption via Blind Storage,” in *IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 2014, pp. 639–654.
- [34] NLTK Project, “Natural Language Toolkit,” <https://www.nltk.org/>, Accessed: 2018-05-14.
- [35] T. Pornin and J. P. Stern, “Digital signatures do not guarantee exclusive ownership,” in *ACNS*, ser. LNCS, vol. 3531, 2005, pp. 138–150.
- [36] D. X. Song, D. Wagner, and A. Perrig, “Practical Techniques for Searches on Encrypted Data,” in *IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 2000, pp. 44–55.
- [37] X. Song, C. Dong, D. Yuan, Q. Xu, and M. Zhao, “Forward Private Searchable Symmetric Encryption with Optimized I/O Efficiency,” *IEEE Transactions on Dependable and Secure Computing*, 2018.
- [38] E. Stefanov, C. Papamanthou, and E. Shi, “Practical Dynamic Searchable Encryption with Small Leakage,” in *NDSS*. The Internet Society, 2014.
- [39] E. Stefanov, M. van Dijk, E. Shi, C. W. Fletcher, L. Ren, X. Yu, and S. Devadas, “Path ORAM: An Extremely Simple Oblivious RAM Protocol,” in *ACM CCS*. ACM Press, 2013, pp. 299–310.

- [40] S.-F. Sun, X. Yuan, J. K. Liu, R. Steinfeld, A. Sakzad, V. Vo, and S. Nepal, “Practical backward-secure searchable encryption from symmetric puncturable encryption,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18. ACM, 2018, pp. 763–780.
- [41] The OpenSSL Project, “OpenSSL Cryptography and SSL/TLS Toolkit,” <https://www.openssl.org/>, Accessed: 2018-05-14.
- [42] S. Vaudenay, “Digital signature schemes with domain parameters: Yet another parameter issue in ECDSA,” in *ACISP*, ser. LNCS, vol. 3108. Springer, 2004, pp. 188–199.
- [43] 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*. USENIX Association, 2016, pp. 707–720.

## A Correctness and Security of $\Pi_{\text{FP}}$

### A.1 Proof of Correctness

**Theorem A.1.** *If  $F_d, F_t$  are secure PRFs and  $H_1$  is a collision-resistant hash function then  $\Pi_{\text{FP}}$  is correct.*

*Proof.* We use the games  $G_0$  and  $G_1$ . The games are identical to  $\Pi_{\text{FP}}$  (see Figure 6) except for the following changes:

1. In the Setup algorithm, we initialize the sets  $\text{LabSet}_1, \text{LabSet}_2$  to null set and set boolean variable `bad` to false.
2. We remove line 5 and 6 of Update protocol and after line 4 in the Update protocol, we add the following code:

$G_0,$	$G_1$
<ol style="list-style-type: none"> <li>5. <math>k_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> <li>6. <math>\text{label} \leftarrow H_1(k_w \  c_w)</math></li> <li>7. <b>if</b> <math>\text{label} \in \text{LabSet}_1</math> <b>then</b></li> <li style="padding-left: 20px;">8. <math>\text{bad} \leftarrow \text{true}, \text{label} \xleftarrow{\\$} \{0, 1\}^\lambda \setminus \text{LabSet}_1</math></li> <li style="padding-left: 20px;">9. <math>\text{LabSet}_1 \leftarrow \text{LabSet}_1 \cup \{\text{label}\}</math></li> </ol>	

3. We replace line 5 in the Search protocol with the following code:

$G_0,$	$G_1$
<ol style="list-style-type: none"> <li>5. <math>\text{label}_w \leftarrow \text{Key}_d[w]</math></li> <li>6. <b>if</b> <math>\exists w' \text{ s.t. } (w', \text{label}_w) \in \text{LabSet}_2 \wedge w' \neq w</math> <b>then</b></li> <li style="padding-left: 20px;">7. <math>\text{bad} \leftarrow \text{true}, \text{label}_w \xleftarrow{\\$} \{0, 1\}^\lambda \setminus \text{LabSet}_2</math></li> <li style="padding-left: 20px;">8. <math>\text{LabSet}_2 \leftarrow \text{LabSet}_2 \cup \{(w, \text{label}_w)\}</math></li> </ol>	

4. We replace line 7 in the Search protocol with the following code:

$G_0,$	$G_1$
<ol style="list-style-type: none"> <li style="padding-left: 20px;">7. <math>k_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> </ol>	

In the modification snippets,  $G_1$  includes the box code and  $G_0$  does not. The first game  $G_0$  will output 1, only if `bad` is set, as repeated labels in maps  $\mathbf{T}$  and  $\mathbf{D}$  are the only source of incorrectness.  $G_0$  produces an



identical distribution to real game when `bad` is not set. If the value assigned to `label` is repeated,  $G_0$  replaces it with new value which hasn't been assigned to any `label` up till now.

Let us denote the event ‘the flag `bad` is set to `true`’ in  $G_0$  by  $\mathbb{E}_0$ . This gives,

$$\Pr[\mathbf{DSSECor}_{\mathcal{A}}^{\Pi_{\text{FP}}}(\lambda) = 1] \leq \Pr[\mathbb{E}_0]. \quad (18)$$

In  $G_1$ , every call to PRFs  $F_t$  and  $F_d$  are answered using tables  $\text{Key}_t$  and  $\text{Key}_d$  respectively. The entries in table  $\text{Key}_t$  are referred by  $(w, \text{ver})$  and entries in table  $\text{Key}_d$  are referred by  $w$ .

Let us denote the event ‘the flag `bad` is set to `true`’ in  $G_1$  by  $\mathbb{E}_1$ .

There exists an adversary  $\mathcal{B}_1$  that can distinguish  $F_t$  from a truly random function and/or an adversary  $\mathcal{B}_2$  that can distinguish  $F_d$  from a truly random function such that:

$$\begin{aligned} |\Pr[\mathbb{E}_1] - \Pr[\mathbb{E}_0]| &\leq \mathbf{Adv}_{F_t, \mathcal{B}_1}^{\text{prf}}(\lambda) + \mathbf{Adv}_{F_d, \mathcal{B}_2}^{\text{prf}}(\lambda) \\ &\leq \text{neg}(\lambda). \end{aligned} \quad (19)$$

The event  $\mathbb{E}_1$  occurs only when the newly picked label value was already picked earlier in  $G_1$ .  $\mathbb{E}_1$  occurs in line 7 of modification 3, if the same label is generated for more than one keyword. Since, the labels are picked uniformly at random in line 5 and the number of keywords,  $m$ , is a polynomial in security parameter, by the birthday bound we can conclude that the probability that the two labels for map  $\mathbf{D}$  are equal is at most  $\frac{m^2}{2^\lambda}$ , i.e.,  $\text{neg}(\lambda)$ .

Further, the key  $k_w$  is picked uniformly at random (see line 5 of modification 2) and the number of updates on  $\mathbf{T}$ , say  $q$ , is a polynomial in security parameter. By the birthday bound we can conclude that the probability that the two keys are equal is at most  $\frac{q^2}{2^\lambda}$ , i.e.,  $\text{neg}(\lambda)$ . Therefore, only with negligible probability the input to  $H_1$  is repeated. So, if  $\mathbb{E}_1$  occurs in line 8 of modification 2, one can find collision in the hash function  $H_1$ . Since,  $H_1$  is collision resistant this happens with negligible probability. Therefore,

$$\Pr[\mathbb{E}_1] \leq \text{neg}(\lambda). \quad (20)$$

From (18), (19) and (20), we get

$$\Pr[\mathbf{DSSECor}_{\mathcal{A}}^{\Pi_{\text{FP}}}(\lambda) = 1] \leq \text{neg}(\lambda). \quad (21)$$

□

## A.2 Security Proof

**Theorem A.2.** *If  $F_d, F_t$  are secure PRFs and  $H_1, H_2$  are hash functions modeled as random oracles outputting  $2\lambda$  and  $\mu$  bits respectively then  $\Pi_{\text{FP}}$  is FP-I secure (Definition 3.1).*

*Proof.* We structure our proof using a sequence of games  $G_0$  to  $G_4$ .  $G_0$  will compute a distribution identical to  $\mathbf{Real}_{\mathcal{A}}^{\Pi_{\text{FP}}}(\lambda)$  and  $G_4$  will compute a distribution that can be simulated perfectly given  $\mathcal{L}$ , i.e., its distribution is identical to  $\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\Pi_{\text{FP}}}(\lambda)$  and the intermediate games are hybrids.

Game  $G_0$ :  $G_0$  is exactly identical to  $\mathbf{Real}_{\mathcal{A}}^{\Pi_{\text{FP}}}(\lambda)$ .

$$\Pr[\mathbf{Real}_{\mathcal{A}}^{\Pi_{\text{FP}}}(\lambda) = 1] = \Pr[G_0 = 1]. \quad (22)$$

Game  $G_1$ : In  $G_1$ , every call to PRFs  $F_t$  and  $F_d$  are answered using tables  $\text{Key}_t$  and  $\text{Key}_d$  respectively. The entries in table  $\text{Key}_t$  are referred by  $(w, \text{ver})$  and entries in table  $\text{Key}_d$  are referred by  $w$ . If there exists an adversary  $\mathcal{A}$  that is able to distinguish between games  $G_0$  and  $G_1$ , we can construct an adversary  $\mathcal{B}_1$  that

<p><u>Setup()</u></p> <ol style="list-style-type: none"> <li>1. <math>\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow</math> empty map</li> <li>2. <math>\text{bad} \leftarrow \text{false}</math></li> <li>3. return ( EDB = ( <math>\mathbf{D}, \mathbf{T}</math> ),  <math>\text{st}_c = (k_t, k_d, \mathbf{W})</math>) to (S, C)</li> </ol> <p><u>Update(op, w, ind, st<sub>C</sub>; EDB)</u></p> <p><u>Update<sub>C</sub>(op, w, ind, st<sub>C</sub>)</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, c_w) = \perp</math> <b>then</b></li> <li>3.     <math>\text{ver}_w \leftarrow 0, c_w \leftarrow -1</math></li> <li>4.     <math>c_w \leftarrow c_w + 1, k_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> <li>5.     <math>\text{val}_r \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>6.     <b>if</b> <math>\mathbf{H}_1[k_w, c_w] \neq \perp</math> <b>then</b></li> <li>7.         <math>\text{bad} \leftarrow \text{true}, \boxed{\text{val}_r \leftarrow \mathbf{H}_1(k_w \  c_w)}</math></li> <li>8.     <math>\text{Hash}_1[w, \text{ver}_w, c_w] \leftarrow \text{val}_r</math></li> </ol>	<ol style="list-style-type: none"> <li>9.   <math>\text{b} \leftarrow 0(\text{op} = \text{add})/1(\text{op} = \text{del})</math></li> <li>10.   <math>\text{pad} \leftarrow \mathbf{H}_2(k_w \  c_w)</math></li> <li>11.   <math>\text{e} \leftarrow \text{b} \  \text{ind} \oplus \text{pad}</math></li> <li>12.   <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, c_w)</math></li> <li>13.   Send <math>(\text{val}_r, \text{e})</math> to S</li> </ol> <p><u>Search(w, st<sub>C</sub>; EDB)</u></p> <p><u>Search<sub>C</sub>(w, st<sub>C</sub>) :</u></p> <ol style="list-style-type: none"> <li>1.   <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>2.   <b>if</b> <math>(\text{ver}_w, c_w) = \perp</math> <b>then</b></li> <li>3.     return <math>\emptyset</math></li> <li>4.   <b>else</b></li> <li>5.     <math>\text{label}_w \leftarrow \text{Key}_d[w]</math></li> <li>6.     <b>if</b> <math>c_w \neq -1</math> <b>then</b></li> <li>7.         <math>k_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> <li>8.         <b>for</b> <math>i = 0</math> to <math>c_w</math> <b>do</b></li> <li>9.             <math>\mathbf{H}_1[k_w, i] \leftarrow \text{Hash}_1[w, \text{ver}_w, i]</math></li> </ol>	<ol style="list-style-type: none"> <li>10.   <math>\text{ver}_w \leftarrow \text{ver}_w + 1, c_w \leftarrow -1</math></li> <li>11.   <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, c_w)</math></li> <li>12.   <b>else</b></li> <li>13.     <math>k_w \leftarrow \perp</math></li> <li>14.   Send <math>(\text{label}_w, k_w, c_w)</math> to S</li> </ol> <p><u>H<sub>1</sub>(k  c)</u></p> <ol style="list-style-type: none"> <li>1.   <math>\text{val} \leftarrow \mathbf{H}_1[k, c]</math></li> <li>2.   <b>if</b> <math>\text{val} = \perp</math> <b>then</b></li> <li>3.     <math>\text{val} \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>4.     <b>if</b> <math>\exists w, \text{ver}</math> s.t. <math>\text{ver} = \mathbf{W}[w].\text{ver}_w \wedge</math>  <math>k = \text{Key}_t[w, \text{ver}] \wedge c \leq \mathbf{W}[w].c_w</math> <b>then</b></li> <li>5.         <math>\text{bad} \leftarrow \text{true}, \boxed{\text{val} \leftarrow \text{Hash}_1[w, \text{ver}, c]}</math></li> <li>6.     <math>\mathbf{H}_1[k, c] \leftarrow \text{val}</math></li> <li>7.   return <math>\text{val}</math></li> </ol>
---	---	--

Figure 18: Games  $G_2$  and  $G'_2$  (Theorem A.2).  $G'_2$  includes the box code and  $G_2$  does not.

can distinguish  $F_t$  from a truly random function and/or an adversary  $\mathcal{B}_2$  that can distinguish  $F_d$  from a truly random function. Formally, there exist adversaries  $\mathcal{B}_1$  and  $\mathcal{B}_2$ , such that

$$\begin{aligned}
|\Pr[G_0 = 1] - \Pr[G_1 = 1]| &\leq \text{Adv}_{F_t, \mathcal{B}_1}^{\text{prf}}(\lambda) + \text{Adv}_{F_d, \mathcal{B}_2}^{\text{prf}}(\lambda) \\
&\leq \text{neg}(\lambda).
\end{aligned} \tag{23}$$

Game  $G_2$ : In  $G_2$ , instead of calling  $\mathbf{H}_1$  to generate label in the Update protocol, we pick random strings. Then, during the Search protocol, the random oracle  $\mathbf{H}_1$  is programmed accordingly to ensure consistency.

Table  $\text{Hash}_1$  and  $\mathbf{H}_1$  are used to simulate the random oracle  $\mathbf{H}_1$ , the entries in the table  $\text{Hash}_1$  are referred by  $(w, \text{ver}, c)$  and in the table  $\mathbf{H}_1$  by  $(k, c)$ .

Figure 18 formally describes  $G_2$ , and an intermediate game  $G'_2$ . In  $G'_2$ ,  $\mathbf{H}_1$  is never programmed to two different values for the same input, thus, ensuring consistency. Instead of storing the randomly picked value in table  $\text{Hash}_1$  at position  $(w, \text{ver}_w, c_w)$ , it first checks whether  $\mathbf{H}_1$  is already programmed at value  $(k_w, c_w)$  which can happen if there was a query to the random oracle  $\mathbf{H}_1$  with input  $(k_w \| c_w)$ . If the check is true, the value  $\mathbf{H}_1(k_w \| c_w)$  is stored in  $\text{Hash}_1[w, \text{ver}_w, c_w]$  else the randomly picked value is stored in  $\text{Hash}_1[w, \text{ver}_w, c_w]$ . The random oracle when needed in the Search protocol in line 9 or by an adversary's query to the random oracle  $\mathbf{H}_1$  in line 5 is lazily programmed in  $G_2$ , so that the outputs are consistent throughout.

The only difference between game  $G_1$  and  $G'_2$  is how we model the random oracle  $\mathbf{H}_1$ . The outputs of  $\mathbf{H}_1$  is perfectly indistinguishable in both these games, therefore,

$$\Pr[G'_2 = 1] = \Pr[G_1 = 1]. \tag{24}$$

Let us denote the event ‘the flag  $\text{bad}$  is set to true’ in  $G'_2$  by  $\mathbb{E}_1$ . The games  $G'_2$  and  $G_2$  are also perfectly identical unless the event  $\mathbb{E}_1$  occurs, and we can apply *identical-until-bad* technique to bound the distinguishing advantage between  $G'_2$  and  $G_2$ .

$$|\Pr[G'_2 = 1] - \Pr[G_2 = 1]| \leq \Pr[\mathbb{E}_1]. \tag{25}$$

<p><u>Setup()</u></p> <ol style="list-style-type: none"> <li>1. <math>\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow</math> empty map</li> <li>2. <math>u \leftarrow 0</math></li> <li>3. <math>\mathbf{Update}, \mathbf{STs} \leftarrow</math> empty table</li> <li>4. <math>\mathbf{STs}[w][0] \leftarrow -1, \quad \forall w \in W</math></li> <li>5. return ( EDB = (<math>\mathbf{D}, \mathbf{T}</math>),     <math>\text{st}_c = (\mathbf{W})</math>) to (S, C)</li> </ol> <p><u>Update(op, w, ind, st<sub>c</sub>; EDB)</u></p> <p><u>Update<sub>c</sub>(op, w, ind, st<sub>c</sub>)</u></p> <ol style="list-style-type: none"> <li>1. Append (u, op, ind) to <math>\mathbf{Update}[w]</math></li> <li>2. <math>\text{Hash}_1[u] \xleftarrow{\\$} \{0, 1\}^{2\lambda}</math></li> <li>3. <math>\text{Hash}_2[u] \xleftarrow{\\$} \{0, 1\}^\mu</math></li> <li>4. Send (<math>\text{Hash}_1[u], \text{Hash}_2[u]</math>) to S</li> <li>5. <math>u \leftarrow u + 1</math></li> </ol> <p><u>Search(w, st<sub>c</sub>; EDB)</u></p> <p><u>Search<sub>c</sub>(w, st<sub>c</sub>) :</u></p> <ol style="list-style-type: none"> <li>1. Append u to <math>\mathbf{STs}[w]</math></li> </ol>	<ol style="list-style-type: none"> <li>2. (<math>\text{empty}, \text{ver}_w, c_w</math>     <math>\{\text{ind}_{w,c}, \mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}</math>) <math>\leftarrow</math>     GetData(<math>\mathbf{Update}, \mathbf{STs}, \text{Hash}_1, \text{Hash}_2</math>)</li> <li>3. if <math>\text{empty} = \text{true}</math> then</li> <li>4. return <math>\emptyset</math></li> <li>5. else</li> <li>6. <math>\text{label}_w \leftarrow \text{Key}_d[w]</math></li> <li>7. if <math>c_w \neq -1</math> then</li> <li>8. <math>k_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> <li>9. for <math>i = 0</math> to <math>c_w</math> do</li> <li>10. <math>\mathbf{H}_1(k_w, i) \leftarrow \mathbf{H}_{1,w,i}</math></li> <li>11. <math>\mathbf{H}_2(k_w, i) \leftarrow \mathbf{H}_{2,w,i} \oplus \text{ind}_{w,i}</math></li> <li>12. else</li> <li>13. <math>k_w \leftarrow \perp</math></li> <li>14. Send (<math>\text{label}_w, k_w, c_w</math>) to S</li> <li>15. <math>u \leftarrow u + 1</math></li> </ol> <p><u>GetData(<math>\mathbf{Update}, \mathbf{STs}, \text{Hash}_1, \text{Hash}_2</math>)</u></p> <ol style="list-style-type: none"> <li>1. if <math>\mathbf{Update}[w] = \perp</math></li> <li>2. <math>\text{empty} \leftarrow \text{true}</math></li> <li>3. else</li> </ol>	<ol style="list-style-type: none"> <li>4. <math>\text{empty} \leftarrow \text{false}</math></li> <li>5. <math>\text{ver}_w \leftarrow -1</math></li> <li>6. for <math>i = 1</math> to <math> \mathbf{STs}[w]  - 1</math> do</li> <li>7. if <math>\exists u, \text{op}, \text{ind}</math> s.t. <math>(u, \text{op}, \text{ind}) \in</math>     <math>\mathbf{Update}[w] \wedge u &gt; \mathbf{STs}[w][i - 1] \wedge</math>     <math>u &lt; \mathbf{STs}[w][i]</math> then</li> <li>8. <math>\text{ver}_w \leftarrow \text{ver}_w + 1</math></li> <li>9. <math>c_w \leftarrow -1</math></li> <li>10. <math>ul \leftarrow  \mathbf{STs}[w]  - 1</math></li> <li>11. <math>\text{Buf}_w \leftarrow \emptyset</math></li> <li>12. for <math>u = \mathbf{STs}[w][ul - 1]</math> to <math>\mathbf{STs}[w][ul]</math> do</li> <li>13. if <math>\exists \text{op}, \text{ind}</math> s.t. <math>(u, \text{op}, \text{ind})</math>     <math>\in \mathbf{Update}[w]</math> then</li> <li>14. <math>c_w \leftarrow c_w + 1</math></li> <li>15. <math>b \leftarrow 0(\text{op} = \text{add})/1(\text{op} = \text{del})</math></li> <li>16. <math>\text{ind}_{w,c_w} \leftarrow b \parallel \text{ind}</math></li> <li>17. <math>\mathbf{H}_{1,w,c_w} \leftarrow \text{Hash}_1[u]</math></li> <li>18. <math>\mathbf{H}_{2,w,c_w} \leftarrow \text{Hash}_2[u]</math></li> <li>19. return (<math>\text{empty}, \text{ver}_w, c_w,</math>     <math>\{\text{ind}_{w,c}, \mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}</math>)</li> </ol>
--	---	--

Figure 19: Game  $G_4$  (Theorem A.2)

The event  $\mathbb{E}_1$  occurs in line 7 of  $\mathbf{Update}$  protocol and in line 5 of  $\mathbf{H}_1$  algorithm. The former captures the fact that adversary has already queried random oracle  $\mathbf{H}_1$  at input  $(k_w \parallel c_w)$  and the latter captures the fact that the adversary queries random oracle  $\mathbf{H}_1$  on a valid input  $(k \parallel c)$  and the value  $k$  is not currently revealed to adversary. Since, the value  $k_w$  is picked uniformly at random, the probability with which event  $\mathbb{E}_1$  occurs is negligible. Using (24) and (25), we can conclude:

$$|\Pr[G_1 = 1] - \Pr[G_2 = 1]| \leq \text{neg}(\lambda). \quad (26)$$

Game  $G_3$ : In  $G_3$ , what we did for  $\mathbf{H}_1$  in game  $G_2$ , we do for  $\mathbf{H}_2$ . Using the same arguments, we can conclude:

$$|\Pr[G_2 = 1] - \Pr[G_3 = 1]| \leq \text{neg}(\lambda). \quad (27)$$

Game  $G_4$ : In  $G_4$  (see Figure 19), we abstract out the information that needs to be simulated by the simulator in order to output transcripts identical to  $G_3$ . Using  $\mathbf{GetData}$  algorithm in  $G_4$ , one keeps track of the randomly generated  $\text{label}$  and  $\text{pad}$  differently than in  $G_3$ . In  $\mathbf{Search}$  protocol, the random oracles are programmed identically to that in  $G_3$ . Queries to random oracles  $\mathbf{H}_1$  and  $\mathbf{H}_2$  can be simulated by outputting random values.

The transcripts outputted by  $\mathbf{Update}$  protocol is identical to that of  $\mathbf{Update}$  protocol in  $G_3$ , as we output fresh random strings in the  $\mathbf{Update}$  protocol.

Next, we describe the  $\mathbf{Search}$  protocol in  $G_4$ . Based on tables  $\mathbf{Update}$  and  $\mathbf{STs}$ , the value of the following components:  $\text{empty}, \text{ver}_w, c_w$  and  $\{\text{ind}_{w,c}, \mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}$  is determined using  $\mathbf{GetData}$  algorithm. Flag  $\text{empty}$  is set to 1 if  $\mathbf{Update}$  is empty.  $\text{ver}_w$  is the count of searches for which there was an update on map  $\mathbf{T}$  corresponding to keyword  $w$  after the previous search. The loop in line 6 of  $\mathbf{GetData}$  determines the number of times version number is updated, i.e., value of  $\text{ver}_w$ . Here,  $c_w$  denotes the count of

updates on map  $\mathbf{T}$  corresponding to keyword  $w$  after the previous search.  $\{\text{ind}_{w,c}\}_{0 \leq c \leq c_w}$  are the document identifier values along with operation bit that have been added to  $\mathbf{T}$  after the previous search operation and the values  $\{\mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}$  are used to simulate the random oracles consistently with the response given at the time of update queries. The loop in line 12 of `GetData` computes the values of  $c_w$  and  $\{\text{ind}_{w,c}, \mathbf{H}_{1,w,c}, \mathbf{H}_{2,w,c}\}_{0 \leq c \leq c_w}$ . As the value of components are computed correctly and consistently w.r.t. previous queries, the transcripts of `Search` protocol is identical to that of `Search` protocol in  $G_3$ . Therefore, we conclude that:

$$\Pr[G_4 = 1] = \Pr[G_3 = 1]. \quad (28)$$

**Simulator Sim:** Finally, we construct a simulator that given the leakage profile  $\mathcal{L}$  simulates game  $G_4$  correctly. It is easy to see that, `Sim` can simulate `Update` protocol correctly. Instead of using keyword  $w$ , `Sim` uses the counter  $\bar{w} = \min \text{sp}(w)$  uniquely mapped from  $w$  using  $\mathcal{L}_{\text{Search}}$  in simulating the `Search` protocol (line 6 and line 8). In line 2 of `Search` protocol, `Sim` uses  $\text{sp}(w)$  and  $\text{Hist}(w)$  instead of STs and `Update` as input to the `GetData` algorithm. Thus, `Sim` is able to produce transcripts of output of `Search` and `Update` protocols identical to  $G_4$ . Hence, we conclude that:

$$\Pr[\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\text{IFP}}(\lambda) = 1] = \Pr[G_4]. \quad (29)$$

By connecting all the games, we conclude

$$\left| \Pr[\mathbf{Real}_{\mathcal{A}}^{\text{IFP}}(\lambda) = 1] - \Pr[\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\text{IFP}}(\lambda) = 1] \right| \leq \text{neg}(\lambda).$$

□

## B Security Proof of $\Pi_{\text{BP-prime}}$

**Theorem B.1.** *If  $F_d, F_t$  are secure PRFs,  $\mathcal{E}$  is RCPA secure and  $H_1$  is a hash function modeled as random oracle outputting  $2\lambda$  bits then  $\Pi_{\text{BP-prime}}$  is BP-I secure (Definition 3.5).*

*Proof.* We structure our proof using a sequence of games  $G_0$  to  $G_4$ .  $G_0$  will compute a distribution identical to  $\mathbf{Real}_{\mathcal{A}}^{\text{BP-prime}}(\lambda)$  and  $G_4$  will compute a distribution that can be simulated perfectly given  $\mathcal{L}$ , i.e., its distribution is identical to  $\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\text{BP-prime}}(\lambda)$  and the intermediate games are hybrids.

**Game  $G_0$ :**  $G_0$  is exactly identical to  $\mathbf{Real}_{\mathcal{A}}^{\text{BP-prime}}(\lambda)$ .

$$\Pr[\mathbf{Real}_{\mathcal{A}}^{\text{BP-prime}}(\lambda) = 1] = \Pr[G_0 = 1]. \quad (30)$$

**Game  $G_1$ :** In  $G_1$ , every call to PRF  $F_t$  and  $F_d$  are answered using tables  $\text{Key}_t$  and  $\text{Key}_d$  respectively. The entries in table  $\text{Key}_t$  are referred by  $(w, \text{ver})$  and entries in table  $\text{Key}_d$  are referred by  $w$ . If there exists an adversary  $\mathcal{A}$  that is able to distinguish between games  $G_0$  and  $G_1$ , we can construct an adversary  $\mathcal{B}_1$  that can distinguish  $F_t$  from a truly random function and/or an adversary  $\mathcal{B}_2$  that can distinguish  $F_d$  from a truly random function. Formally, there exists adversaries  $\mathcal{B}_1$  and  $\mathcal{B}_2$ , such that

$$\begin{aligned} |\Pr[G_0 = 1] - \Pr[G_1 = 1]| &\leq \mathbf{Adv}_{F_t, \mathcal{B}_1}^{\text{prf}}(\lambda) + \mathbf{Adv}_{F_d, \mathcal{B}_2}^{\text{prf}}(\lambda) \\ &\leq \text{neg}(\lambda). \end{aligned} \quad (31)$$

**Game  $G_2$ :** In  $G_2$ , instead of calling  $H_1$  to generate label in the `Update` protocol, we pick random strings. Then, during the `Search` protocol, the random oracle  $H_1$  is programmed accordingly, to ensure consistency.

Table  $\text{Hash}_1$  and  $\mathbf{H}_1$  are used to simulate the random oracle  $H_1$ , the entries in the table  $\text{Hash}_1$  are referred by  $(w, \text{ver}, c)$  and in the table  $\mathbf{H}_1$  by  $(k, c)$ .

Figure 20 formally describes  $G_2$ , and an intermediate game  $G'_2$ . In  $G'_2$ ,  $H_1$  is never programmed to two different values for the same input, thus, ensuring consistency. Instead of storing the randomly picked value in table  $\text{Hash}_1$  at position  $(w, \text{ver}_w, c_w)$ , it first checks whether  $H_1$  is already programmed at value  $(k_w, c_w)$

<p><u>Setup()</u></p> <ol style="list-style-type: none"> <li>1. <math>k_e \leftarrow \mathcal{E}.\text{Gen}()</math></li> <li>2. <math>\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow</math> empty map</li> <li>3. <math>\text{bad} \leftarrow \text{false}</math></li> <li>4. return (<math>\text{EDB} = (\mathbf{D}, \mathbf{T})</math>, <math>\text{st}_c = (k_e, k_t, k_d, \mathbf{W})</math>) to <math>(S, C)</math></li> </ol> <p><u>Update(op, w, ind, st<sub>C</sub>; EDB)</u></p> <p><u>Update<sub>C</sub>(op, w, ind, st<sub>C</sub>)</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, c_w) = \perp</math> <b>then</b></li> <li>3.   <math>\text{ver}_w \leftarrow 0, c_w \leftarrow -1</math></li> <li>4.   <math>c_w \leftarrow c_w + 1</math></li> <li>5.   <math>k_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> <li>6.   <math>\text{val}_r \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>7.   <b>if</b> <math>\mathbf{H}_1[k_w, c_w] \neq \perp</math> <b>then</b></li> <li>8.     <math>\text{bad} \leftarrow \text{true}, \boxed{\text{val}_r \leftarrow \mathbf{H}_1(k_w \  c_w)}</math></li> <li>9.   <math>\text{Hash}_1[w, \text{ver}_w, c_w] \leftarrow \text{val}_r</math></li> <li>10. <math>b \leftarrow 0</math> (op=add) / <math>1</math> (op=del)</li> <li>11. <math>e \leftarrow \mathcal{E}.\text{Enc}(k_e, b \  \text{ind})</math></li> </ol>	<ol style="list-style-type: none"> <li>12. <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, c_w)</math></li> <li>13. Send <math>(\text{val}_r, e)</math> to <math>S</math></li> </ol> <p><u>Search(w, st<sub>C</sub>; EDB)</u></p> <p><b>Round 1</b></p> <p><u>Search<sub>C</sub>(w, st<sub>C</sub>) :</u></p> <ol style="list-style-type: none"> <li>1. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>2. <b>if</b> <math>(\text{ver}_w, c_w) = \perp</math> <b>then</b></li> <li>3.   return <math>\emptyset</math></li> <li>4. <b>else</b></li> <li>5.   <math>\text{label}_w \leftarrow \text{Key}_d[w]</math></li> <li>6.   <b>if</b> <math>c_w \neq -1</math> <b>then</b></li> <li>7.     <math>k_w \leftarrow \text{Key}_t[w, \text{ver}_w]</math></li> <li>8.     <b>for</b> <math>i = 0</math> to <math>c_w</math> <b>do</b></li> <li>9.       <math>\mathbf{H}_1[k_w, i] \leftarrow \text{Hash}_1[w, \text{ver}_w, i]</math></li> <li>10.      <math>\text{ver}_w \leftarrow \text{ver}_w + 1</math></li> <li>11.      <math>c_w \leftarrow -1</math></li> <li>12.      <math>\mathbf{W}[w] \leftarrow (\text{ver}_w, c_w)</math></li> <li>13.   <b>else</b></li> <li>14.      <math>k_w \leftarrow \perp</math> //No need of round 2</li> <li>15.   Send <math>(\text{label}_w, k_w, c_w)</math> to <math>S</math></li> </ol>	<p><b>Round 2</b></p> <p><u>Search<sub>C</sub>(w, st<sub>C</sub>) :</u></p> <ol style="list-style-type: none"> <li>16. Receive <math>(\text{PSet}_w, \text{AList})</math> from <math>S</math></li> <li>17. <math>(\text{ver}_w, c_w) \leftarrow \mathbf{W}[w]</math></li> <li>18. <math>\text{AuxSet} \leftarrow \text{PSet}_w</math></li> <li>19. <b>for</b> <math>c \leftarrow 0</math> to <math> \text{AList} </math> <b>do</b></li> <li>20.   <math>(b \  \text{ind}) \leftarrow \mathcal{E}.\text{Dec}(k_e, \text{AList}[c])</math></li> <li>21.   <b>if</b> <math>b = 0</math></li> <li>22.     <math>\text{AuxSet} \leftarrow \text{AuxSet} \cup \{\text{ind}\}</math></li> <li>23.   <b>else</b></li> <li>24.     <math>\text{AuxSet} \leftarrow \text{AuxSet} \setminus \{\text{ind}\}</math></li> <li>25. return <math>\text{AuxSet}</math></li> </ol> <p><u>H<sub>1</sub>(k  c)</u></p> <ol style="list-style-type: none"> <li>1. <math>\text{val} \leftarrow \mathbf{H}_1[k, c]</math></li> <li>2. <b>if</b> <math>\text{val} = \perp</math> <b>then</b></li> <li>3.   <math>\text{val} \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>4.   <b>if</b> <math>\exists w, \text{ver}</math> s.t. <math>\text{ver} = \mathbf{W}[w].\text{ver}_w \wedge</math> <math>k = \text{Key}_t[w, \text{ver}] \wedge c \leq \mathbf{W}[w].c_w</math> <b>then</b></li> <li>5.     <math>\text{bad} \leftarrow \text{true}, \boxed{\text{val} \leftarrow \text{Hash}_1[w, \text{ver}, c]}</math></li> <li>6.   <math>\mathbf{H}_1[k, c] \leftarrow \text{val}</math></li> <li>7. return <math>\text{val}</math></li> </ol>
---	---	---

Figure 20: Games  $G_2$  and  $G'_2$  (Theorem B.1).  $G'_2$  includes the box code and  $G_2$  does not.

which can happen if there was a query to the random oracle  $\mathbf{H}_1$  with input  $(k_w \| c_w)$ . If the check is true, the value  $\mathbf{H}_1(k_w \| c_w)$  is stored in  $\text{Hash}_1[w, \text{ver}_w, c_w]$  else the randomly picked value is stored in  $\text{Hash}_1[w, \text{ver}_w, c_w]$ . The random oracle when needed in the Search protocol in line 9 or by an adversary's query to random oracle  $\mathbf{H}_1$  in line 5 is lazily programmed in  $G_2$ , so that the outputs are consistent throughout.

The only difference between game  $G_1$  and  $G'_2$  is how we model the random oracle  $\mathbf{H}_1$ . The outputs of  $\mathbf{H}_1$  is perfectly indistinguishable in both these games, therefore,

$$\Pr[G'_2 = 1] = \Pr[G_1 = 1]. \quad (32)$$

Let us denote the event 'the flag bad is set to true' in  $G'_2$  by  $\mathbb{E}_1$ . The games  $G'_2$  and  $G_2$  are also perfectly identical unless the event  $\mathbb{E}_1$  occurs, and we can apply *identical-until-bad* technique to bound the distinguishing advantage between  $G'_2$  and  $G_2$ .

$$|\Pr[G'_2 = 1] - \Pr[G_2 = 1]| \leq \Pr[\mathbb{E}_1]. \quad (33)$$

The event  $\mathbb{E}_1$  occurs in line 8 of Update protocol and in line 5 of  $\mathbf{H}_1$  algorithm. The former captures the fact that the adversary has already queried random oracle  $\mathbf{H}_1$  at input  $(k_w \| c_w)$  and the latter captures the fact that the adversary queries random oracle  $\mathbf{H}_1$  on a valid input  $(k \| c)$  and the value  $k$  is not currently revealed to adversary. Since, the value  $k_w$  is picked uniformly at random, the probability with which event  $\mathbb{E}_1$  occurs is negligible. Using (32) and (33), we can conclude:

$$|\Pr[G_1 = 1] - \Pr[G_2 = 1]| \leq \text{neg}(\lambda). \quad (34)$$

Game  $G_3$ : In  $G_3$ , instead of computing ciphertexts corresponding to actual op bit 'b' and document identifier ind, we replace them with randomly chosen ciphertexts from the ciphertext space. We also store the plaintext-ciphertext pairs in order to simulate the view of round 2 of the Search protocol correctly. If there exists an

<u>Setup()</u> 1. $\mathbf{W}, \mathbf{T}, \mathbf{D} \leftarrow$ empty map 2. $u \leftarrow 0$ 3. <b>Update</b> , STs $\leftarrow$ empty table 4. $\text{STs}[w][0] \leftarrow -1, \forall w \in W$ 5. return ( EDB = ( $\mathbf{D}, \mathbf{T}$ ), $\text{st}_c = (k_e, \mathbf{W})$ ) to ( $\mathbf{S}, \mathbf{C}$ )  <u>Update(op, w, ind, st<sub>c</sub>; EDB)</u> <b>Update<sub>c</sub>(op, w, ind, st<sub>c</sub>)</b> 1. Append (u, op, ind) to <b>Update</b> [w] 2. $\text{Hash}_1[u] \xleftarrow{\$} \{0, 1\}^{2\lambda}$ 3. $e \xleftarrow{\$} \mathcal{C}$ 4. Send ( $\text{Hash}_1[u], e$ ) to $\mathbf{S}$ 5. $u \leftarrow u + 1$  <u>GetData<sub>r1</sub>(Update, STs, Hash<sub>1</sub>)</u> 1. <b>if</b> <b>Update</b> [w] = $\perp$ 2.   empty $\leftarrow$ true 3. <b>else</b> 4.   empty $\leftarrow$ false 5. $\text{ver}_w \leftarrow -1$ 6. <b>for</b> $i = 1$ to $ \text{STs}[w]  - 1$ <b>do</b> 7. <b>if</b> $\exists u, \text{op}, \text{ind}$ s.t. $(u, \text{op}, \text{ind}) \in$ <b>Update</b> [w] $\wedge u > \text{STs}[w][i - 1]$	$\wedge u < \text{STs}[w][i]$ <b>then</b> 8. $\text{ver}_w \leftarrow \text{ver}_w + 1$ 9. $c_w \leftarrow -1$ 10. $ul \leftarrow  \text{STs}[w]  - 1$ 11. <b>for</b> $u = \text{STs}[w][ul - 1]$ to $\text{STs}[w][ul]$ <b>do</b> 12. <b>if</b> $\exists \text{op}, \text{ind}$ s.t. $(u, \text{op}, \text{ind})$ $\in$ <b>Update</b> [w] <b>then</b> 13. $c_w \leftarrow c_w + 1$ 14. $\mathbf{H}_{1,w,c_w} \leftarrow \text{Hash}_1[u]$ 15.   return (empty, $\text{ver}_w, c_w,$ $\{\mathbf{H}_{1,w,c}\}_{0 \leq c \leq c_w}$ )  <u>Search(w, st<sub>c</sub>; EDB)</u> <b>Round 1</b> <b>Search<sub>c</sub>(w, st<sub>c</sub>) :</b> 1. Append $u$ to $\text{STs}[w]$ 2. (empty, $\text{ver}_w, c_w, \{\mathbf{H}_{1,w,c}\}_{0 \leq c \leq c_w}$ ) $\leftarrow$ <b>GetData<sub>r1</sub>(Update, STs, Hash<sub>1</sub>)</b> 3. <b>if</b> empty = true <b>then</b> 4.   return $\emptyset$ 5. <b>else</b> 6. $\text{label}_w \leftarrow \text{Key}_d[w]$ 7. <b>if</b> $c_w \neq -1$ <b>then</b> 8. $k_w \leftarrow \text{Key}_t[w, \text{ver}_w]$	9. <b>for</b> $i = 0$ to $c_w$ <b>do</b> 10. $\mathbf{H}_1[k_w, i] \leftarrow \mathbf{H}_{1,w,i}$ 11. <b>else</b> 12. $k_w \leftarrow \perp$ 13.   Send ( $\text{label}_w, k_w, c_w$ ) to $\mathbf{S}$ <b>Round 2</b> <b>Search<sub>c</sub>(w, st<sub>c</sub>) :</b> 14. Receive ( $\text{PSet}_w, \text{AList}$ ) from $\mathbf{S}$ 15. $\text{AuxSet} \leftarrow$ <b>GetData<sub>r2</sub>(PSet<sub>w</sub>, Update, STs)</b> 16. return $\text{AuxSet}$ 17. $u \leftarrow u + 1$  <u>GetData<sub>r2</sub>(PSet<sub>w</sub>, Update, STs)</u> 1. $\text{AuxSet} \leftarrow \text{PSet}_w$ 2. $ul \leftarrow  \text{STs}[w]  - 1$ 3. <b>for</b> $u = \text{STs}[w][ul - 1]$ to $\text{STs}[w][ul]$ <b>do</b> 4. <b>if</b> $\exists \text{op}, \text{ind}$ s.t. $(u, \text{op}, \text{ind})$ $\in$ <b>Update</b> [w] <b>then</b> 5. <b>if</b> op = add <b>then</b> 6. $\text{AuxSet} \leftarrow \text{AuxSet} \cup \{\text{ind}\}$ 7. <b>else</b> 8. $\text{AuxSet} \leftarrow \text{AuxSet} \setminus \{\text{ind}\}$ 9.   return $\text{AuxSet}$
---	---	---

Figure 21: Game  $\mathbf{G}_4$  (Theorem B.1)

adversary  $\mathcal{A}$  that is able to distinguish between games  $\mathbf{G}_2$  and  $\mathbf{G}_3$ , we can construct an adversary  $\mathcal{B}$  that breaks the RCPA security of  $\mathcal{E}$ .

**Game  $\mathbf{G}_4$ :** In  $\mathbf{G}_4$  (see Figure 21), we abstract out the information that needs to be simulated by the simulator in order to output transcripts identical to  $\mathbf{G}_3$ . Using  $\text{GetData}_{r1}$  and  $\text{GetData}_{r2}$  algorithms in  $\mathbf{G}_4$ , one keeps track of label differently than in  $\mathbf{G}_4$ . In **Search** protocol, the random oracle is programmed identically to that in  $\mathbf{G}_4$ . Queries to random oracle  $\mathbf{H}_1$  can be simulated by outputting random values.

As we output fresh random strings in **Update** protocol, the transcripts of **Update** protocol is identical to that of **Update** protocol in  $\mathbf{G}_3$ .

Next, let us describe the **Search** protocol in  $\mathbf{G}_4$ . Based on tables **Update** and STs, the value of following components: empty,  $\text{ver}_w$ ,  $c_w$  and  $\{\mathbf{H}_{1,w,c}\}_{0 \leq c \leq c_w}$  are determined using  $\text{GetData}_{r1}$  algorithm. Flag empty is set to 1 if **Update** is empty.  $\text{ver}_w$  is the count of searches for which there was an update on map  $\mathbf{T}$  corresponding to keyword  $w$  after the previous search. The loop in line 6 of  $\text{GetData}_{r1}$  determines the number of times version number is updated, i.e., value of  $\text{ver}_w$ . Here,  $c_w$  denotes the count of updates on map  $\mathbf{T}$  corresponding to keyword  $w$  after the previous search and the values  $\{\mathbf{H}_{1,w,c}\}_{0 \leq c \leq c_w}$  are used to simulate the random oracle consistently with the response given at the time of update queries. The loop in line 11 of  $\text{GetData}_{r1}$  computes the values of  $c_w$ , and  $\{\mathbf{H}_{1,w,c}\}_{0 \leq c \leq c_w}$ . Given  $\text{PSet}_w$ , STs and **Update**, we determine the current set of documents matching the keyword  $w$ , i.e.,  $\text{AuxSet} = \text{DB}(w)$  in  $\text{GetData}_{r2}$  which is then used to simulate the output of round 2 of **Search** protocol. As the value of components are determined

correctly and consistently, the transcripts of **Search** protocol is identical to that of **Search** protocol in  $G_3$ . Therefore, we conclude that:

$$\Pr[G_4 = 1] = \Pr[G_3 = 1]. \quad (35)$$

**Simulator Sim:** Finally, we construct a simulator that given the leakage profile  $\mathcal{L}$  simulates game  $G_4$  correctly. **Sim** can simulate **Update** protocol correctly as in  $G_4$ . Instead of using keyword  $w$ , **Sim** uses the counter  $\bar{w} = \min \text{sp}(w)$  uniquely mapped from  $w$  using  $\mathcal{L}_{\text{Search}}$  in simulating the **Search** protocol (line 6 and line 8). In line 2 of **Search** protocol, **Sim** uses  $\text{sp}(w)$  and  $\text{Updates}(w)$  instead of **STs** and **Update** as input to the  $\text{GetData}_{r_1}$  algorithm. Note that in  $\text{GetData}_{r_1}$ , we only use the timestamps of the search and update queries in order to perform the simulation. Also, the output of  $\text{GetData}_{r_2}$  (line 15) is the current set of documents matching the keyword  $w$ , i.e.,  $\text{DB}(w)$  which is a part of leakage profile,  $\mathcal{L}_{\text{Search}}$ . Thus, **Sim** is able to produce transcripts of output of **Search** and **Update** protocols identical to  $G_4$ . Hence, we conclude that:

$$\Pr[\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\Pi_{\text{BP}}\text{-prime}}(\lambda) = 1] = \Pr[G_4]. \quad (36)$$

By connecting all the games, we conclude

$$\left| \Pr[\mathbf{Real}_{\mathcal{A}}^{\Pi_{\text{BP}}\text{-prime}}(\lambda) = 1] - \Pr[\mathbf{Ideal}_{\mathcal{A}, \text{Sim}}^{\Pi_{\text{BP}}\text{-prime}}(\lambda) = 1] \right| \leq \text{neg}(\lambda).$$

□