

---

# MAKING CLOUD STORAGES SECURE AND EFFICIENT

---

Submitted to Indian Statistical Institute  
in partial fulfillment of the requirements for the Degree of  
DOCTOR OF PHILOSOPHY IN COMPUTER SCIENCE



**Avishek Majumder**  
Applied Statistics Unit (ASU)

Indian Statistical Institute  
Kolkata - 700108, India



*Dedicated to*  
***My Family***

## LIST OF PUBLICATIONS/MANUSCRIPTS

1. Debrup Chakraborty, **Avishek Majumder**, and Subhabrata Samajder “Making Searchable Symmetric Encryption Schemes Smaller and Faster.”, Published in International Journal of Information Security ([doi.org/10.1007/s10207-024-00915-y](https://doi.org/10.1007/s10207-024-00915-y)).
2. Mohamed Ahmed Abdelraheem, Sanjay Bhattacharjee, Theo Henault, and **Avishek Majumder**. “Conjunctive Dynamic Searchable Symmetric Encryption – A Generic Framework.”, The IACR Communications in Cryptology (CiC) (submitted).

## ACKNOWLEDGEMENT

I would like to take this opportunity to express my sincere gratitude to all those who have contributed to the completion of my PhD thesis.

First and foremost, I am deeply grateful to my supervisor, Dr. Debrup Chakraborty, for his unwavering support, guidance, and mentorship throughout my research journey and beyond. His valuable comments, motivating discussions, and ideas greatly helped my research work. His wisdom, values and perspective towards life have helped me become a better human being.

I am indebted to Dr. Sanjay Bhattacharjee for his guidance and support. A significant part of the work reported in this thesis has been done in collaboration with him. Furthermore, Dr. Bhattacharjee had a significant advisory role throughout my tenure as a graduate student in ISI.

A part of the thesis was done in collaboration with Dr. Subhabrata Samajder. Who was always available to listen to my ideas and enrich me with his valuable comments. I'm thankful to him for his mentorship, and I am also indebted to him for the support he provided during my visit to IIT Delhi.

I am sincerely grateful to Prof. Bimal Kumar Roy, Prof. Palash Sarkar, Prof. Mridul Nandi, Prof. Subhamoy Maitra, Prof. Sanjit Chatterjee, and Dr. Sabyasachi Karati for their invaluable help and guidance.

I am thankful to Dr. Sayantan Mukherjee, Dr. Nilanjan Datta, Dr. Avijit Dutta, Dr. Mohamed Ahmed Abdelraheem, and Dr. Dhiman Saha, who have helped me with their noble guidance, support, and encouragement as a senior and collaborator.

I would also like to thank Dr. Laltu Sardar, Mr. Bibhash Chanda Das, and Mr. Theo Henault for a wonderful journey of collaboration. I would also like to thank Sourav for helping me with my work. I would like to express my love and respect to Dr. Sayan Dey for his believe in me and continuous support and motivation throughout my journey, starting from my M.Tech.

I thank Samir for being with me in my PhD journey. I could not thank you enough to Sougata for always being there for me and encouraging me in my low times. My

journey toward a PhD would not be possible without the friendship and support of Rakesh and Chotu. Also, my sincere love and respect to Diptendu-da for always listening to me in my low time and always encouraging me.

I am also thankful to the former and current members of our research group, including Mostaf, Prabal, Nayana, Manab, Anurag, Subhra, Subhadip-da, Aniruddha, Laltu-da, Amit, Soumya, Pritam, Biswajit, Suprita, Ananda, Susanta, Aniruddha, Jyotirmoy, Animesh, Chandranan, Gourab, Anup, Suman, Suvo, Manish, Pravat with whom I have always had fruitful discussions and enjoyable times.

I would like to express my heartfelt gratitude to my entire family for their unwavering support and guidance. My dream of pursuing research would be incomplete without my parents' constant support and encouragement, Mr. Bibek Jyoti Majumdar, Mrs. Tanusree Majumdar, and sister Arijita Majumder. Last but not least, I am grateful to my wife, Koushani, for her support, which has been a constant source of inspiration for me for the past two years.



Avishek Majumder

## ABSTRACT

Over the years, *searchable symmetric encryption* (SSE) schemes have emerged as a promising tool for enabling efficient query processing over encrypted data stored in untrusted cloud servers. This thesis mainly focuses on efficiency and security enhancements of *dynamic searchable symmetric encryption* (DSSE) schemes, which support various query types and are secure against several adversarial conditions.

For any SSE scheme, its query processing, storage, and communication costs are directly related to the size of the encrypted index stored on the server. A reduction of the index size naturally leads to enhanced search efficiency and reduced storage and communication costs. We are unaware of any previous attempts to reduce the index size of SSE schemes. We introduce a novel technique to directly reduce the index size of any SSE. Our proposed method generically transforms any secure single keyword SSE into an equivalently functional and secure version with reduced storage requirements, resulting in faster search and reduced communication overhead. Our technique involves arranging the set of document identifiers  $\text{db}(w)$  related to a keyword  $w$  in the leaf nodes of a complete binary tree, eventually obtaining a succinct representation of the set  $\text{db}(w)$ . This compact representation leads to smaller index sizes. We conduct extensive theoretical analysis to prove the correctness of our scheme. Additionally, our experiments on real and synthetic data validate the effectiveness of our approach and demonstrate its practical applicability.

Among the few SSE schemes available in the literature which support complex query types like conjunctive queries, the *oblivious cross tag* (OXT) scheme from Crypto'13 is the most efficient one. OXT has the limitation that it only works for static databases. In NDSS'20, an extension of OXT called the *oblivious dynamic cross tag* (ODXT) was proposed. ODXT supports conjunctive queries with dynamic updates. However, ODXT is not *forward private*.

We propose a generic framework for designing *conjunctive dynamic SSE* (CDSSE) schemes, supporting conjunctive queries that allow dynamic updates while being both forward and backward private simultaneously. To the best of our knowledge such a

scheme does not exist till date. Our scheme assumes a restricted update model where a document with its associated keywords can be dynamically added to or deleted from the database as a whole, but the set of keywords for a document is not modified once uploaded. We define forward and backward privacy for this new setting of updates and extend the OXT scheme to make it dynamic in the new setting. We prove the security of our construction against adaptive adversaries and analyse the precise leakages to the adversarial server. Experiments show that our schemes are very efficient.

Another less studied aspect of SSE schemes is verifiability. In an SSE scheme, the server may be dishonest and may not respond to a client’s queries following the prescribed protocol. A verifiable SSE can detect such anomalous behaviour of a server. To defend against such malicious adversaries, previous approaches employ *authenticated encryption* (AE) to furnish a “*proof*” for each update. We propose a new construction where we convert any forward and backward private adaptively secure SSE scheme into a verifiable SSE. Our construction uses a new class of message authentication codes (MAC), which we call *updatable message authentication codes* (UdMAC). A UdMAC allows the verification tag for a message to be updated with each modification to the message without recomputing the entire MAC, ensuring efficiency. We establish security requirements for such a MAC and introduce two constructions, **ConCatU** and **XoRU**, which work with two different types of message updates, namely, concatenation and exclusive-or (XOR), respectively. Furthermore, we present the first generic construction for a forward and backward private *fault-tolerant verifiable* DSSE using a UdMAC construction and prove its security. Our construction converts any generic forward and backward secure SSE secure in an honest-but-curious adversarial model into an equivalently secure DSSE secure in a malicious adversarial model with faulty updates.



# Contents

<b>1</b>	<b>Introduction</b>	<b>17</b>
1.1	Searchable Symmetric Encryption . . . . .	18
1.2	Scope of This Thesis . . . . .	21
1.3	The Road-Map . . . . .	22
<b>2</b>	<b>Preliminaries</b>	<b>27</b>
2.1	General Notations . . . . .	27
2.2	Hardness Assumptions . . . . .	28
2.3	Dynamic Searchable Symmetric Encryption . . . . .	30
2.4	Conjunctive Dynamic SSE . . . . .	33
2.5	Security of DSSE . . . . .	33
2.5.1	Security Game . . . . .	35
2.5.2	Correctness of DSSE . . . . .	36
2.5.3	Soundness of DSSE . . . . .	37
2.5.4	Fault-tolerant Verifiable DSSE . . . . .	39
2.5.5	Forward and Backward Privacy (Single Keyword DSSE) . . . . .	39
2.6	Related Works . . . . .	43
2.7	Final Remarks . . . . .	48
<b>3</b>	<b>Tree Covers</b>	<b>49</b>

3.1	Binary Trees and Tree Covers . . . . .	51
3.1.1	Binary Trees . . . . .	51
3.1.2	Tree Cover . . . . .	52
3.1.3	Cover of a Configuration . . . . .	54
3.1.4	Cover of a Dynamic Configuration . . . . .	56
3.2	Cover Generation Algorithms . . . . .	58
3.2.1	Pure Cover: A Tree Cover Scheme for Pure Configurations . . . . .	58
3.2.2	Mixed Cover: Generating a Smaller Sized Cover . . . . .	61
3.2.3	Complexity of Proposed Algorithms . . . . .	71
3.2.4	Dynamic Tree Cover Scheme . . . . .	72
3.3	Expected Cover Size of a Pure Cover . . . . .	74
3.4	Experimental Results . . . . .	77
3.5	Final Remarks . . . . .	79
<b>4</b>	<b>Tree Cover Based SSE</b>	<b>81</b>
4.1	Constructing Static SSE Using Tree Cover . . . . .	82
4.1.1	Cover-based Representation of DB. . . . .	83
4.1.2	Generic Static SSE Using Keyword Cover . . . . .	84
4.2	Constructing Dynamic SSE Using Tree Cover . . . . .	87
4.3	Discussions . . . . .	92
4.4	Security of $d\Sigma$ . . . . .	94
4.5	Experimental Results . . . . .	99
4.5.1	Extra Overheads . . . . .	99
4.6	Final Remarks . . . . .	102
<b>5</b>	<b>Conjunctive Dynamic SSE – A Generic Framework</b>	<b>105</b>
5.1	Non-modifiable Database . . . . .	106
5.2	Forward and Backward Privacy in CDSSE for Non-modifiable Documents	108
5.3	Our Construction . . . . .	111
5.3.1	Our Generic CDSSE Scheme . . . . .	112
5.3.2	Design Rational . . . . .	113

5.4	Security of Our Construction	116
5.5	Leakage Analysis	122
5.6	Implementation and Experimental Results	130
5.7	Final Remarks	133
<b>6</b>	<b>Updatable Message Authentication Codes</b>	<b>135</b>
6.1	Message Authentication Codes	136
6.2	Polynomial MACs	137
6.3	Updatable MACs	139
6.3.1	Security of Updatable MACs	141
6.3.2	Some New Notations	142
6.4	ConCatU: An Updatable MAC for Concatenation	142
6.4.1	Instantiations and Efficiency	145
6.4.2	Security of ConCatU	147
6.5	XoRU: An Updatable MAC for $\oplus$ Difference	154
6.5.1	Efficiency of XoRU	156
6.5.2	Security of XoRU	158
6.6	Final Remarks	161
<b>7</b>	<b>Fault-tolerant Verifiable DSSE</b>	<b>163</b>
7.1	Generic Fault-tolerant Verifiable DSSE	164
7.1.1	A Generic FVDSSE Scheme $v\Sigma$ Using $\Psi$ and $\Sigma$	165
7.1.2	Fault-tolerance and Correctness of the Search Result	168
7.2	Comparisons and Actual Overheads	169
7.3	Security	172
7.3.1	Adaptive Security of $v\Sigma$	172
7.3.2	Forward and Backward Privacy of $v\Sigma$	175
7.3.3	Soundness of $v\Sigma$	175
7.4	Final Remark	178
<b>8</b>	<b>Conclusion</b>	<b>179</b>



## List of Figures

2-1	Correctness Game of DSSE. . . . .	36
2-2	Soundness Game of DSSE . . . . .	38
3-1	The tree $\mathcal{T}$ used in Example 1. . . . .	53
3-2	The coloring scheme for pure configurations . . . . .	59
3-3	Pure cover generation . . . . .	60
3-4	Pure cover reconstruction . . . . .	61
3-5	Coloring scheme for mixed covers . . . . .	62
3-6	An example of mixed cover: The leaf nodes represented by circles are assumed to bear the sign $+$ , the leaf nodes represented by squares bear the sign $-$ . The shaded circles represent nodes colored green and the shaded squares are nodes colored red. . . . .	63
3-7	Mixed cover generation algorithm . . . . .	68
3-8	Mixed cover reconstruction . . . . .	71
3-9	Dynamic cover generation . . . . .	73
3-10	Dynamic cover reconstruction . . . . .	74
3-11	Number of Identifiers vs Average Cover Size . . . . .	79
3-12	Expected Cover Size Comparison . . . . .	80
4-1	DB to $\widetilde{\text{DB}}$ conversion algorithm . . . . .	83
4-2	Generic static SSE using tree-cover scheme . . . . .	85

4-3	Buffer update algorithm . . . . .	86
4-4	A generic dynamic SSE using tree-cover scheme setup phase and DB to $\widetilde{\text{DB}}$ conversion algorithm . . . . .	87
4-5	A generic dynamic SSE using tree-cover scheme update phase . . . . .	89
4-6	A generic dynamic SSE using tree-cover scheme search phase . . . . .	91
4-7	A pictorial depiction of the data in Table 4.2 showing the number of tuples in the original database (represented by filled circles) and the number of tuples in the converted database (represented by filled triangles) for different database sizes. . . . .	101
4-8	Database sizes in dynamic scenario. . . . .	102
5-1	A generic dynamic single-keyword SSE (DSSE) construction $\Sigma$ . . . . .	112
5-2	A modular description of $\Gamma$ without the SKSSE part of [29] . . . . .	113
5-3	Our generic CDSSE construction $\Pi$ . . . . .	114
5-4	Search time comparison between the three instantiations of our generic CDSSE for fixed $ \text{Updates}(w_2)  = 10^5$ with varying $ \text{Updates}(w_1) $ from $10^2$ to $10^5$ . . . . .	132
5-5	Search time comparison between the three instantiations of our generic CDSSE for fixed $ \text{Updates}(w_1)  = 10$ with varying $ \text{Updates}(w_2) $ from 10 to $10^5$ . . . . .	132
6-1	The procedure <code>SeqUpdt</code> used to define correctness of an updatable MAC scheme $\Psi$ . . . . .	140
6-2	Specification of <code>ConCatU</code> using a prf $f_K : \{0, 1\}^* \rightarrow \{0, 1\}^n$ . . . . .	144
6-3	Specification of <code>XoRU.MacUpdate</code> and <code>XoRU.MacVerify</code> . . . . .	156
7-1	Setup phase of $v\Sigma$ . . . . .	166
7-2	Update protocol of $v\Sigma$ . . . . .	166
7-3	Search protocol of $v\Sigma$ . . . . .	167
7-4	Verify algorithm of $v\Sigma$ . . . . .	168
7-5	Forger game of FVDSSE. . . . .	177

## List of Tables

4.1	Characteristics of some existing SSE schemes: $N$ is the number of $(id, w)$ pairs, $ \mathcal{W} $ is the number of keywords, $i_w$ and $d_w$ are the number of insertions and deletions, and $u_w = i_w + d_w$ is the total number of updates, and $n_w = i_w - d_w$ is the number of keyword-identifier pairs currently matching the keyword $w$ . RT is the number of roundtrips for a search query. . . . .	95
4.2	The size of original database DB and the converted database $\widetilde{DB}$ in case of Enron data. . . . .	100
4.3	Cover generation and reconstruction times. . . . .	101
5.1	Comparison of our construction with previous schemes. $W$ : the set of keywords. $D$ : Set of documents. $W_d$ : Number of keywords in the document to be updated. $N$ : Number of identifier-keyword pairs in the database. $n$ : Number of keywords in a conjunction. $u_i$ : Updates performed for keyword $w_i$ . $u = \sum_{i=1}^n u_i$ : Total updates for a conjunction of $n$ keywords. $v$ : Keywords involved in the edit operation. $h$ : Average number of documents matched by any two keywords. $db(\omega)$ : Set of matching document identifiers for query $\omega = w_1 \wedge \dots \wedge w_n$ , where $n \geq 1$ . RT: Round-Trip. FP: Forward Private. BP: Backward Private. KPRP: Keyword-Pair Result Pattern Hiding. NMD: Non-Modifiable Documents. . . . .	115
5.2	Storage cost and time required for each update. . . . .	131

5.3	Search time comparison for fixed $ \text{Updates}(w_2)  = 10^5$ with varying $ \text{Updates}(w_1) $ from $10^2$ to $10^5$ . . . . .	133
7.1	Comparison of our schemes with existing VDSSE schemes. FP, BP and FT stand for forward privacy, backward privacy, and fault-tolerant, respectively. ‘Computation’ and ‘Communication’ refer to the computational and communication complexity, respectively. $W$ is the number of distinct keywords in the database, and $ D $ denotes the number of documents in the database. $N$ is the number of keyword-identifier pairs. $0 < \epsilon < 1$ is a fixed constant, and $\alpha$ denotes the number of times the queried keyword was historically added to the database. $m$ denotes the number of matching entries for the searched keyword. . .	170



## Introduction

Delegating an organization's or individual's storage requirements to a third-party server has become a common practice. Storage as a Service, commonly referred to as cloud storage, is a model in which providers offer storage solutions to clients. These solutions are advantageous as they alleviate users from costs associated with storage-related hardware, software, and maintenance. While outsourcing storage requirements offers convenience and cost-effectiveness, it also raises significant security concerns. Data protection laws like the General Data Protection Regulation (GDPR) which is in force in the European Union and the European Economic Area, require that online service providers using third-party cloud services for storing data must ensure user privacy “by design and by default”<sup>1</sup>. Similar policies are applicable to many other countries also. This makes the problem of securing the data stored in third party servers even more important, both from the perspective of the user and the service providers.

Two main dimensions of information security are data confidentiality and integrity/authenticity. These are applicable in the context of outsourced data also. For outsourced data, confidentiality means that the data stored in the server should be unreadable to the server. Integrity means that if the server deviates from the predefined protocols for answering the client's queries and/or updating the data, then the client should be able to detect such malicious behaviour of the server. Note that, once the database is outsourced, the client can no longer monitor its current status. There-

---

<sup>1</sup>As in Art. 25 of GDPR as on January 21, 2025. <https://gdpr-info.eu/art-25-gdpr/>

fore, it is essential to ensure that the server accurately stores the database, executes queries, and responds properly to user queries. Thus, the problem of authentication in the case of outsourced data is as important as the problem of confidentiality. Note that in this thesis we use both the terms data integrity and authentication interchangeably.

Historically, the challenge of maintaining data confidentiality has been addressed through encryption. Secure encryption algorithms ensure that the resulting ciphertext is indistinguishable from random strings, rendering it practically unreadable for the server. However, this solution presents a dilemma, as users also require the ability to query and update the delegated data. Encrypting the stored data with a traditional encryption algorithm would hinder the server’s capability to execute queries or updates on behalf of the client. A well accepted solution to this dilemma is provided by Searchable Symmetric Encryption (SSE) schemes, which are specialized encryption schemes that allow search and updates on encrypted data without compromising its confidentiality.

The traditional solution to the problem of data integrity/authentication in the symmetric key setting has been provided by Message Authentication Codes (MAC) and/or several variants of Authenticated Encryption (AE). These solutions, though important in the context of outsourced data also, have to be properly tailored to achieve integrity in data stored in the cloud.

In this thesis, we study several dimensions of the problem of securing outsourced data and provide some efficient and secure solutions. In this Chapter, we begin with a token introduction to SSE. Further, we discuss the scope of this thesis and finally provide a chapter wise summary of the rest of the thesis.

## 1.1 Searchable Symmetric Encryption

An SSE scheme abstracts a database as a set of documents denoted as  $\mathcal{D}$ . Each document  $d_i \in \mathcal{D}$  is associated with a unique identifier  $id_i \in \mathcal{I}$ . These documents are viewed as collections of keywords from a predefined finite set, denoted as  $\mathcal{W}$ . For each

keyword,  $w \in \mathcal{W}$ ,  $\text{db}(w)$  represents the set of identifiers of those documents containing the keyword  $w$ . For a  $w \in \mathcal{W}$ , let  $t_w = \text{db}(w) \times \{w\} = \{(\text{id}, w) : \text{id} \in \text{db}(w)\}$ . The database DB of an SSE scheme is viewed as the set of tuples  $\text{DB} = \bigcup_{w \in \mathcal{W}} t_w$ .

The functionality of searching and updating encrypted data is commonly achieved through an *inverted index* [76] that stores mappings of document identifiers with keywords. Searchable Symmetric Encryption (SSE) follows a similar approach, storing the inverted index in an encrypted form so that user queries are hidden from the server, thus preserving user privacy.

Specifically, an SSE encrypts the database DB, comprising keyword-identifier pairs, using a specialized structure often referred to as an “*encrypted multi-map*” [35] (EMM) or an inverted index. This encrypted multi-map is stored on the server and facilitates both searching and updating the encrypted database. An SSE scheme typically supports three types of operations: addition, search, and deletion, which we denote by **add**, **srch**, and **del**, respectively. To initiate a keyword search, the client provides a *search token* for the encrypted multi-map to the server. For updates (such as additions and deletions to/from the database), the client supplies an *update token* to the server, enabling modifications to the encrypted multi-map. Subsequently, these search and update tokens are used by the server to transmit encrypted search results to the client and update the database as needed, respectively. The storage of an inverted index in an encrypted form in an untrusted server incurs extra computation and communication overhead on both the client and server. In general, this overhead is directly related to the size of the inverted index.

SSE schemes that support only the search operation are referred to as *static SSE*, whereas those SSEs that support both search and updates (addition and deletion) are known as *dynamic SSE* (DSSE) schemes. In the context of SSE, the input for an update protocol is a keyword-identifier pair  $(\text{id}, w)$  with an operation  $\text{op} \in \{\text{add}, \text{del}\}$ . For the search operation ( $\text{op} = \text{srch}$ ), the input can either be a single keyword or multiple keywords with a specific search rule. The output of a search is a set of document identifiers that match the query rule [60, 28]. This design allows SSE to decouple the storage of documents from the storage of the data structures used for

searching and updating the database.

Most SSE schemes to date [86, 42, 60, 28, 23, 25, 87, 49, 38, 89, 90] support equality search on single keywords in each query. However, some can also support richer queries like conjunctive or boolean queries [29, 73, 47, 58, 93, 11, 67], wild card search [85, 22, 54, 37, 96, 69], and sub-string search [36, 53, 72].

Depending upon the adversarial nature of the cloud service provider platform (CSP), an SSE adversary can be classified into two categories.

- a. **Honest-but-curious:** where the adversary follows the protocol honestly but attempts to gain knowledge about user data from the information that queries inadvertently reveal.
- b. **Malicious:** where the adversary tries to learn about the user’s data and may also cheat by either not storing the user data to save storage or not responding to a query correctly to save computation.

The primary goal of an SSE scheme is to protect users’ data from both types of adversaries mentioned above, but all existing SSEs do not provide security against both these types of adversaries.

Security of SSE schemes is an evolving topic of study and is not yet fully understood. Security guarantees are generally provided in a pre-defined security model, and as it is common in the cryptographic literature, the security of an SSE scheme is reduced to the security of one or more well studied cryptographic primitives or hardness assumptions. The usefulness of such security guarantees depends heavily on how well the security model represents the real adversarial threat. There have been several instances that a provably secure SSE scheme has been attacked as the security model failed to incorporate a real attack scenario [55, 68, 27, 97].

Over the years, two key security notions in DSSE, namely, *forward privacy* and *backward privacy* have been proposed [28], which are acceptable to the community and are widely believed to encompass a large class of real attack scenarios. Intuitively, forward private SSE schemes ensure that the server cannot link new updates with previous search operations performed on the inverted index. This protects queries

made to the inverted index with newly injected files by the adversary [97]. Backward privacy ensures that after a pair  $(id, w)$  has been deleted, the identifier  $id$  will not be leaked in future searches for  $w$ . The first security definition and a scheme achieving them were proposed in [23], and the first formal definition of backward security was proposed in [25].

The major focus, till date, has been in designing SSEs secure against honest-but-curious adversaries [86, 42, 60, 28, 23, 25, 87, 49, 38, 89, 90]. A handful of studies on SSE have focused towards the adversary being malicious. A *verifiable* SSE (VSSE) scheme prevents a malicious/dishonest server from forging the client by providing a “*proof*” for every search result. Using the proof, a user can verify the correctness of the response to a search query returned by the server [63, 30, 64, 65, 24, 99, 98, 52, 48, 71, 94].

While the works mentioned above take into account a dishonest server, very few of them can handle inappropriate behaviour from the client. The client, who is unaware of the current state of the database maintained by the server, might submit inappropriate or invalid updates. For example, a client may attempt to re-insert an already existing keyword-identifier pair or may attempt to delete a nonexistent keyword-identifier pair. An SSE scheme that can handle such malicious updates from a client and still provide correct search results is referred to as a *fault-tolerant* SSE.

## 1.2 Scope of This Thesis

This thesis aims to *make cloud storage secure and efficient*. We specifically answer the following main questions:

- a. Is it possible to develop a method for storing the encrypted inverted index in a more succinct manner, thereby reducing the computation, communication, and storage overhead of state-of-the-art SSE schemes?
- b. Is it possible to efficiently support complex queries, such as conjunctive searches, in SSE while maintaining the required security and efficiency?

- c. Is it possible to develop an SSE scheme that ensures correct search results while maintaining required security, even in the presence of a malicious server?
- d. Is it possible to maintain the desired security and correctness of SSE even if the client performs faulty updates?

In this thesis, we answer all these questions affirmatively. We propose new constructions of SSE schemes and propose modifications over existing schemes to achieve our goals. We prove the correctness and security of our proposals and, in some cases, provide extensive experimental data to validate the utility of our proposals in real-world applications. We always consider a single user model.

### 1.3 The Road-Map

In this Section, we provide a short overview of our findings and a chapter-wise summary of the rest of the thesis.

**Chapter 2** does not contain any new material. We begin by introducing the general notations used throughout the thesis and defining the hardness assumptions necessary to prove the security of our proposed constructions. Following this, we present the formal definition of SSE and the related security notions. Finally, we discuss some related works relevant to this thesis.

**Chapter 3** and **Chapter 4** presents our work related to making SSEs efficient in terms of storage, computation and communication overheads. Our main observation is that all existing SSE schemes need to store all the keyword-identifier pairs present in the database  $DB$ . Thus, the size of the encrypted index is never smaller than the size of  $DB$ , i.e.  $|DB|$ . In most cases, it is much bigger than  $|DB|$ . We propose a generic pre-processing step that can be applied to any secure SSE to obtain a new SSE, which will have a significant reduction of its index size on average but would retain the security of the original scheme. This reduction in size further reduces search time and communication requirements. The main tool we use to achieve this reduction is a novel concept that we call the *tree cover*. For each keyword  $w$  we construct a complete binary tree where the number of leaf nodes is at least  $|Z|$  and

associate each  $\text{id} \in \mathcal{I}$  with a leaf node. We label a node associated with  $\text{id}$  with  $+$  if  $\text{id} \in \text{db}(w)$  and with  $-$  if  $\text{id} \notin \text{db}(w)$ . This labeled tree uniquely represents the set  $\text{db}(w)$ . We show that this labeled tree can be represented by a few numbers of tree nodes, and we can use only these nodes to represent the set  $\text{db}(w)$  for all keywords  $w$ . This leads to a drastic reduction in the representation of  $\text{db}(w)$ . We further use this representation of  $\text{db}(w)$  to design an SSE with a reduced index size.

In Chapter 3 we introduce the idea of tree covers and discuss several algorithms to generate covers. We prove the correctness of our algorithms and present theoretical estimates on cover sizes. Finally, we present experiments to validate our theoretical estimates of the cover sizes. The characterization of tree covers and the algorithms for cover generation are further used for the construction of SSEs in Chapter 4. The material presented in Chapter 3 has little reference to SSEs. It just states and solves a combinatorial problem which can be of independent interest.

In Chapter 4 we use the tools developed in Chapter 3 to construct an SSE. We treat the cases of static and dynamic SSE separately. The main result presented in this chapter is that if we take any secure DSSE  $\Lambda$ , then we can convert it into an equivalently functional and secure DSSE, which will have much smaller index sizes on average. We prove the security of our construction and present extensive experimental data on real data sets to demonstrate the effectiveness of our scheme.

**Chapter 5** deals with SSEs, which support conjunctive search queries. The most notable existing construction of an SSE supporting conjunctive queries is the *Oblivious Cross Tag* (OXT) construction [29]. OXT supports conjunctive and general boolean queries very efficiently in large databases with reasonable leakage. However, the scheme is limited to static databases. Designing dynamic databases that support conjunctive queries while being both forward and backward private has been a long-standing open problem. The first claim of such a construction was reported in [80]. A subsequent work [100] demonstrated that the scheme in [80] fails to provide forward privacy in certain cases. Since then, several works [102, 100, 95] have proposed solutions that support conjunctive queries with forward and backward privacy, but none have achieved the efficiency close to the OXT scheme of [29].

In Chapter 5, we address the problem for a specific class of databases that we refer to as *non-modifiable databases*, which are databases where an existing document cannot be modified though new documents can be added and existing documents can be deleted. The only way to modify a document is to delete the document and then add the modified document. This model of database falls in between static databases and dynamic databases. We argue that there are numerous practical scenarios where such a model is applicable. Further, we introduce a generic technique to design an SSE scheme supporting conjunctive queries for non-modifiable databases, achieving efficiency comparable to the OXT scheme. Our approach integrates any DSSE scheme that supports equality queries and addition operations with the OXT framework of [29] in a modular, black-box fashion. We prove our scheme to be both forward and backward private if applied to non-modifiable databases.

**Chapter 6** and **Chapter 7** deal with the problem of authenticity. We consider a malicious adversarial setting, where the server not only tries to learn about user data but also differs to follow the protocol correctly. Moreover, the client, being unaware of the current state of the database, may issue incorrect updates to the database. As already stated, in such scenarios, it is crucial to have cryptographic assurance that the query responses returned by the CSP are accurate.

In Chapter 6, we develop a new class of *message authentication codes* (MAC) which we call Updatable MACs (UdMAC). Let  $\mathcal{M}$  be a message space. We fix  $\Delta : \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$  to be an associative binary operator, which we call the update function. As usual, for  $M, m \in \mathcal{M}$  we will denote  $\Delta(M, m)$  by  $M\Delta m$ . For an example, one can consider  $\Delta$  to be the concatenation operator, ie,  $M\Delta m$  represents the message  $M\|m$ . A  $\Delta$ -UdMAC is a stateful procedure, which has the following interesting property. Let  $M \in \mathcal{M}$  be a message and let its tag be  $t$ , then one can compute the tag corresponding to the message  $M\Delta m$  with only access to  $m$ ,  $t$  and some state information. Thus, to outsource a message  $M$ , a  $\Delta$ -UdMAC computes a tag  $t$  of  $M$ .  $M$  is sent to the server, and a tag  $t$ , along with some state information, is retained with the client. Further, suppose the client wishes to update  $M$  to  $M'$  where  $M' = M\Delta m$ . To do so, the client can compute the tag  $t'$  for the updated



message  $M'$  just based on the update  $m$  and the state information about  $M$  that the client retained. Thus, the client can just upload the update  $m$  to the server and retain  $t'$  with it. We define the syntax and security of UdMACs and also propose two constructions for two different update functions. We prove the security of both constructions.

In Chapter 7, we use UdMACs to construct verifiable and fault-tolerant SSEs. In particular, in Chapter 7, we present a generic construction of a single keyword DSSE scheme that ensures forward privacy, backward privacy and correctness, even in the presence of a malicious server and a client issuing faulty updates. We refer to such a DSSE scheme as a *fault-tolerant verifiable dynamic SSE* or FVDSSE. Our construction combines UdMAC with a generic forward and backward private DSSE scheme that is secure in the honest-but-curious adversarial setting to develop an equivalently secure DSSE scheme that is secure against malicious adversaries in the presence of faulty updates. To the best of our knowledge, such a scheme is the first of its kind.

**Chapter 8** summarizes our conclusions and discusses potential future directions for this area of research.



In this chapter we discuss some preliminaries which would be used throughout this thesis. We begin with introducing some general notations in Section 2.1 followed by a discussion on the basic cryptographic objects and assumptions used in this thesis in Section 2.2. In Section 2.3, we discuss in detail the syntax of a dynamic SSE which is our main object of interest. We define important concepts related to conjunctive SSEs in Section 2.4. In Section 2.5, we define the security of SSE schemes, where we define the basic  $\mathcal{L}$ -adaptive security of SSEs and further discuss the other important dimensions of security of SSEs like correctness, soundness and fault tolerance. Section 2.5 also includes a detailed discussion about leakage functions which are fundamental to the security definitions of SSEs. In Section 2.6 we summarize some important previous works on SSEs.

## 2.1 General Notations

For a finite set  $X$ ,  $|X|$  denotes the cardinality of  $X$ . Let  $X$  be any set, then  $x \stackrel{\$}{\leftarrow} X$  means  $x$  is sampled uniformly at random from the set  $X$ . Let  $X, Y$  be two finite sets then,  $X \times Y$  denotes the Cartesian products of  $X$  and  $Y$ . For a non-negative integer  $n$ ,  $[n]$  denotes the set  $\{1, 2, \dots, n\}$ , and for non-negative integers  $i, j$ ,  $i < j$ ,  $[i, j]$  represents the set  $\{i, i + 1, \dots, j\}$ . The set  $\{0, 1\}^*$  represents the set of all binary strings, including the empty string, and for a positive integer  $n$ ,  $\{0, 1\}^n$  denotes the set of all  $n$  bit strings. For  $x, y \in \{0, 1\}^*$ ,  $x||y$  denote the concatenation of the strings

$x$  and  $y$ .

We sometimes see the set  $\{0, 1\}^n$  as the finite field  $\mathbb{F}_{2^n}$  with  $2^n$  elements. For such an interpretation, for  $x, y \in \mathbb{F}_{2^n}$ , we will denote addition and multiplication of  $x, y$  in  $\mathbb{F}_{2^n}$  by  $x \oplus y$  and  $xy$  respectively. We will explicitly mention, when we treat  $\{0, 1\}^n$  as  $\mathbb{F}_{2^n}$ .

We treat adversaries as algorithms. By a ppt adversary  $\mathcal{A}$  we mean that the adversary runs in probabilistic polynomial time, i.e.,  $\mathcal{A}$  is a probabilistic algorithm which runs in polynomial time. Sometimes we describe adversaries with access to oracles. For an adversary  $\mathcal{A}$ , by  $\mathcal{A}^{\mathcal{O}} = 1$ , we mean that  $\mathcal{A}$  has oracle access to  $\mathcal{O}$ , and after interacting with  $\mathcal{O}$ ,  $\mathcal{A}$  outputs a 1.

A function  $f : \mathbb{N} \rightarrow \mathbb{R}$  is negligible if and only if for all  $c > 0$ , there exists a  $n_0 \in \mathbb{N}$  such that for all  $n \geq n_0$ ,  $f(n) < n^{-c}$ . By  $\text{negl}(\lambda)$  we denote a negligible function in  $\lambda$ . For a natural number  $n \in \mathbb{N}$ ,  $\text{poly}(n)$  denotes an arbitrary polynomial in  $n$ .

## 2.2 Hardness Assumptions

### Pseudo-Random Function (PRF).

Let  $\mathcal{K}, \mathcal{M}, \mathcal{R}$  be sets and let  $F : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{R}$  be a two input function, which we sometimes see as a function family  $F = \{F(K, \cdot)\}_{K \in \mathcal{K}}$ . Let  $\text{Func}(\mathcal{M}, \mathcal{R})$  be the set of all functions mapping  $\mathcal{M}$  to  $\mathcal{R}$ .

**Definition 2.2.1** (PRF). *Let  $F : \{0, 1\}^\lambda \times \mathcal{M} \rightarrow \mathcal{R}$  be an efficiently computable keyed function.  $F$  is said to be a pseudo-random function family if for all ppt adversary  $\mathcal{A}$  the following holds.*

$$\text{Adv}_{F, \mathcal{A}}^{\text{prf}} = |\Pr [\mathcal{A}^{F_k(\cdot)}(1^\lambda)] - \Pr [\mathcal{A}^{\rho(\cdot)}(1^\lambda)]| \leq \text{negl}(\lambda),$$

where probabilities are taken over random choices of  $k \xleftarrow{\$} \{0, 1\}^\lambda$  and  $\rho \xleftarrow{\$} \text{Func}(\mathcal{M}, \mathcal{R})$  and any randomness used by  $\mathcal{A}$ .

For denoting the PRF-advantage of an adversary  $\mathcal{A}$  attacking a function family  $F$  we sometimes use the equivalent notation  $\text{Adv}_F^{\text{prf}}(\mathcal{A})$ .

**Decisional Diffie-Hellman Assumption (ddh-Assumption) [29].**

Let  $\mathcal{G}$  be a ppt algorithm which on input  $\lambda$  outputs a description of a cyclic group  $G$  of prime order  $p$ , where  $|p| = \text{poly}(\lambda)$ , and a generator  $g \in G$ . Then the *decisional Diffie-Hellman assumption* is the following.

**Definition 2.2.2** (ddh-Assumption). *We say that the ddh problem is hard relative to  $\mathcal{G}$  if the following holds for all ppt adversary  $\mathcal{A}$ .*

$$\begin{aligned} \text{Adv}_{G,\mathcal{A}}^{\text{ddh}}(\lambda) &= \left| \Pr \left[ a, b \xleftarrow{\$} \mathbb{Z}_p^* : \mathcal{A}(g, g^a, g^b, g^{ab}) = 1 \right] \right. \\ &\quad \left. - \Pr \left[ a, b, c \xleftarrow{\$} \mathbb{Z}_p^* : \mathcal{A}(g, g^a, g^b, g^c) = 1 \right] \right| \\ &\leq \text{negl}(\lambda), \end{aligned}$$

where the probabilities are taken over the randomness of  $\mathcal{A}$  and random choices of  $a, b, c \in \mathbb{Z}_p^*$ .

**Extended Decisional Diffie-Hellman Assumption (eddh-Assumption) [80].**

Let  $\mathcal{G}$  be a ppt algorithm which on input  $\lambda$  outputs a description of a cyclic group  $G$  of prime order  $p$  (st.  $|p| = \text{poly}(\lambda)$ ), and a generator  $g \in G$ . Then the *extended decisional Diffie-Hellman assumption* is the following.

**Definition 2.2.3** (eddh-Assumption). *We say that the eddh problem is hard relative to  $\mathcal{G}$  if the following holds for all ppt adversary  $\mathcal{A}$ .*

$$\text{Adv}_{G,\mathcal{A}}^{\text{eddh}}(\lambda) = \left| \Pr [\mathcal{A}(g, M) = 1] - \Pr [\mathcal{A}(g, \widehat{M}) = 1] \right| \leq \text{negl}(\lambda),$$

where

$$M = \begin{bmatrix} g^{\alpha_1 \beta_1} & g^{\alpha_1 \beta_2} & \dots & g^{\alpha_1 \beta_n} \\ g^{\alpha_2 \beta_1} & g^{\alpha_2 \beta_2} & \dots & g^{\alpha_2 \beta_n} \\ \vdots & \vdots & \ddots & \vdots \\ g^{\alpha_m \beta_1} & g^{\alpha_m \beta_2} & \dots & g^{\alpha_m \beta_n} \end{bmatrix}, \widehat{M} = \begin{bmatrix} g^{\gamma_{1,1}} & g^{\gamma_{1,2}} & \dots & g^{\gamma_{1,n}} \\ g^{\gamma_{2,1}} & g^{\gamma_{2,2}} & \dots & g^{\gamma_{2,n}} \\ \vdots & \vdots & \ddots & \vdots \\ g^{\gamma_{m,1}} & g^{\gamma_{m,2}} & \dots & g^{\gamma_{m,n}} \end{bmatrix}$$

the probabilities are taken over the randomness of  $\mathcal{A}$  and for all  $i \in [m]$ ,  $j \in [n]$ ,  $\alpha_i, \beta_j, \gamma_{i,j} \xleftarrow{\$} \mathbb{Z}_p^*$ .

## 2.3 Dynamic Searchable Symmetric Encryption

Let  $\mathcal{D} = \{d_1, \dots, d_D\}$  denote a collection of documents that the client wants to store in the server. Let  $\mathcal{W}$  denote the set of all possible keywords. Each document  $d_i \in \mathcal{D}$ , is composed of a set of keywords  $\mathbf{w}_i \subseteq \mathcal{W}$ . Define  $\mathbf{W} = \bigcup_{i=1}^D \mathbf{w}_i$  to be the collection of all distinct keywords present in the database  $\mathcal{D}$ .

Let  $\mathcal{I} = \{\text{id}_1, \dots, \text{id}_D\} \subseteq \{0, 1\}^\lambda$ , for a fixed constant  $\lambda$  (in general considered as the security parameter), be the set of all identifiers. For each keyword,  $w \in \mathbf{W}$ ,  $\text{db}(w)$  represents the set of identifiers of those documents containing the keyword  $w$ . For a  $w \in \mathbf{W}$ , let  $t_w = \text{db}(w) \times \{w\} = \{(\text{id}, w) : \text{id} \in \text{db}(w)\}$ . The database DB of an SSE scheme is viewed as the set of tuples

$$\text{DB} = \bigcup_{w \in \mathbf{W}} t_w.$$

It is customary to abstract out a database as a collection of identifiers and their associated keywords as described above. This allows the decoupling of the storage of documents with the storage of the inverted index.

A DSSE scheme  $\Lambda = (\text{Setup}, \text{Search}, \text{Update})$  is a three tuple, comprised of three protocols between the client (C) and the server (S). In the case where the SSE supports verification of the search result, the **Search** protocol of SSE is also equipped with another algorithm called **Verify**, run by the client. A static SSE scheme does not support the **Update** protocol. A protocol  $\mathbf{P}$  between the client  $\mathbf{C}$  and the server  $\mathbf{S}$ , is denoted by  $(\dots ; \dots) \leftarrow \mathbf{P}(\dots ; \dots)$ . The semicolon separates the inputs and outputs of both parties. The first part before the semicolon is for the client and the part after the semicolon is for the server. We consider only those as output which both the parties get after the complete execution of the protocol.

During the initiation of the SSE scheme, the client runs the setup to obtain an

encrypted database EDB. The server is provided with EDB. Thereafter, the client can run three types of queries: addition, search, and deleting (`add`, `srch`, and `del` respectively) on the EDB. An SSE is primarily concerned with the storage and retrieval of the index, and not how the actual documents are stored. An update query is *always considered to be successful*. For search queries, in an honest-but-curious adversarial model, the query is always considered to be successful. In a malicious adversarial model, the client either outputs the set of identifiers matching the query rule or “reject” the result.

With this abstraction of a database, a DSSE scheme  $\Lambda$  is defined as follows. It closely resembles the definitions of [25] and [38]. The definition of DSSE in Definition 2.3.1 captures an honest-but-curious adversary. We’ll amend the definition while considering a malicious adversary in Section 2.5.3.

**Definition 2.3.1** (DSSE). *A DSSE scheme  $\Lambda = (\text{Setup}, \text{Search}, \text{Update})$  is a three tuple of protocols between the client (C) and the server (S) defined as follows.*

- $(k, \sigma_C ; \text{EDB}) \leftarrow \text{Setup}(1^\lambda, \text{DB} ; \perp)$  : *is a protocol executed between the client and the server. It comprises the following two algorithms.*

$(k, \sigma_C, \text{EDB}) \leftarrow \text{Setup}_C(1^\lambda, \text{DB})$  : *This is a probabilistic polynomial time algorithm run by the client that takes as input the security parameter  $1^\lambda$  and the initial database DB. It outputs a key  $k$ , the client’s state  $\sigma_C$  and the encrypted database EDB. The client stores the key  $(k, \sigma_C)$  securely, and the encrypted database EDB is sent to the server.*

$\text{EDB} \leftarrow \text{Setup}_S(\text{EDB})$  : *upon receiving EDB from the client, the server saves it.*

- $(\sigma_C ; \text{EDB}) \leftarrow \text{Update}(k, \sigma_C, \text{op}, (\text{id}, w) ; \text{EDB})$  : *is a protocol executed between the client and the server. It comprises the following two algorithms.*

$(\sigma_C, \text{utk}) \leftarrow \text{Update}_C(k, \sigma_C, \text{op}, (\text{id}, w))$  : *This (possibly probabilistic) algorithm is run by the client. On input the key  $k$ , client’s state  $\sigma_C$ , an update*

query with operation  $\text{op} \in \{\text{add}, \text{del}\}$ , and a keyword-identifier pair  $(\text{id}, w)$  the algorithm outputs an update token  $\text{utk}$  and client's updated state  $\sigma_{\mathcal{C}}$ .

$\text{EDB} \leftarrow \text{Update}_{\mathcal{S}}(\text{utk}, \text{EDB})$  : This is a deterministic algorithm run by the server that takes as input the encrypted database  $\text{EDB}$  and the update token  $\text{utk}$ . It then outputs an updated encrypted database  $\text{EDB}$ .

- $(\sigma_{\mathcal{C}}, \text{db}(q) ; \text{EDB}) \leftarrow \text{Search}(k, \sigma_{\mathcal{C}}, q ; \text{EDB})$  : is a protocol executed between the client and the server. It comprises the following two algorithms.

$(\sigma_{\mathcal{C}}, \text{stk}) \leftarrow \text{Search}_{\mathcal{C}}(k, \sigma_{\mathcal{C}}, q)$  : This (possibly probabilistic) algorithm is run by the client. On input the key  $k$ , client's state  $\sigma_{\mathcal{C}}$  and search query  $q$ , the algorithm outputs a search token  $\text{stk}$  and client's updated state  $\sigma_{\mathcal{C}}$ .

$\text{res} \leftarrow \text{Search}_{\mathcal{S}}(\text{stk}, \text{EDB})$  : This is a deterministic algorithm run by the server. It takes as input the encrypted database  $\text{EDB}$  and the search token  $\text{stk}$ . It outputs  $\text{res}$  as the search results. The final search result  $\text{db}(q)$  for search query  $q$ , is computed from  $\text{res}$ .

**Remark.** In the update phase, we always consider *atomic updates*, which means a single keyword-identifier pair is updated at a time. In case of a bulk update (e.g. adding a file with many keywords), we invoke the update protocol multiple times. However, all the update tokens in such cases, can be communicated together.

**Types of SSE.** Depending upon the rule of search query and the number of keyword(s) involved in it, an SSE scheme can be classified into many types. However, we only discuss those types which are relevant to this thesis.

1. **Single-keyword SSE** supports only a single keyword per search query and returns all the matching document identifiers which match the searched keyword  $w$ . We call such a query  $q = w$  as the equality query.
2. **Conjunctive SSE** supports querying a set of keyword  $\{w_1, w_2, \dots, w_n\}$  in each search query. The scheme, as a search result, returns the set of all the document



identifiers containing all the keywords in the query. For a conjunctive query, we write  $q = \omega = w_1 \wedge w_2 \wedge \dots \wedge w_n$ , and the search result as  $\text{db}(\omega)$ .

## 2.4 Conjunctive Dynamic SSE

An SSE scheme may allow searches using multiple keywords and their conjunctions. Let  $L = \{w_1, w_2, \dots, w_n\} \subseteq \mathcal{W}$  be a subset of keywords and  $\omega(L) = w_1 \wedge w_2 \wedge \dots \wedge w_n$  be a conjunction of the elements of  $L$ . For every conjunction  $\omega(L)$ , we assume the existence of a least frequent keyword, w.l.o.g say  $w_1$ .

**Definition 2.4.1** (Least Frequent Term in a Conjunction [29]). *In a conjunctive query  $\omega(L)$ , a keyword  $w_i \in L$  is termed the least frequent term in the conjunction if, for every  $w_j \in L$  such that  $w_j \neq w_i$  and  $w_j$  appears at least in as many updates as  $w_i$  does.*

In simple terms, a least frequent keyword is a keyword in  $L$ , which has appeared in the least number of updates among all the other keywords in  $L$ .

**Definition 2.4.2** (*s-term and x-term* [29]). *In a conjunctive query  $\omega(L)$  the least frequent term (say  $w_1$ ) of the query is referred to as the “s-term” and all  $w_i \in \tilde{L} = L \setminus \{w_1\}$  are referred to as the “x-terms”.*

For convenience, we denote a conjunctive query  $\omega(L)$  as  $\omega$ .

For a conjunctive query  $\omega = w_1 \wedge w_2 \wedge \dots \wedge w_n$ , the set of identifiers matching the search query  $\omega$  is denoted by  $\text{db}(\omega)$ . Formally, for a conjunctive query  $\omega$ , we define

$$\text{db}(\omega) = \bigcap_{i=1}^n \text{db}(w_i).$$

## 2.5 Security of DSSE

The security of a DSSE scheme  $\Lambda$  is determined by a leakage profile  $\mathcal{L}_\Lambda$ . The leakage profile  $\mathcal{L}_\Lambda = (\mathcal{L}_{\Lambda, \text{Setup}}, \mathcal{L}_{\Lambda, \text{Search}}, \mathcal{L}_{\Lambda, \text{Update}})$  denote the information that the adversary learns from the setup process and each execution of the search and update protocols.

$\mathcal{L}_{\Lambda, \text{Setup}}$  denotes the leakage due to **Setup**. We note that the adversary only receives  $\text{EDB}_0$  (initial encrypted database) from the output of  $\text{Setup}_C$ , where  $\text{DB}_0$  is the initial database. So,  $\mathcal{L}_{\Lambda, \text{Setup}}$  is effectively the leakage from  $\text{EDB}_0$ .  $\mathcal{L}_{\Lambda, \text{Search}}$  denotes the leakage due to all the searches.  $\mathcal{L}_{\Lambda, \text{Update}}$  denotes the leakage due to all updates. The security of SSE schemes is generally argued by showing that an adversary cannot distinguish between the transcript of a real-world execution and an ideal-world execution (simulated using the leakage profile) of the scheme [42, 60, 28]. To formally define the security of SSE we introduce some more notations.

**Query Sequence.** A  $\ell$ -query sequence  $\text{QS}_\ell$  of a DSSE scheme  $\Lambda$  is a set with  $\ell$  element, where  $1 \leq \ell \leq \text{poly}(\lambda)$  is the total number of queries. Formally,  $\text{QS}_\ell = \{Q_1, \dots, Q_i, \dots, Q_\ell\}$ , is the list of all queries made by the client. Every element  $Q_i = (u_i, \text{op}_i, \text{in}_i)$  of a query sequence is a three tuple, where  $u_i$  is the time-stamp of the query,  $\text{op}_i \in \{\text{add}, \text{srch}, \text{del}\}$  is the type/operation of the query, and  $\text{in}_i$  is the input of the query depending upon the operation  $\text{op}_i$ . If  $\text{op}_i = \text{srch}$  is a search query, then  $\text{in}_i = q^1$ ; if  $\text{op}_i \in \{\text{add}, \text{del}\}$  is an update query, then  $\text{in}_i = (\text{id}_i, w_i)$ .

**View of the Server.** To define the view of the server in an SSE scheme, we use the formalization of [88, 38]. The “*view*” of the server in a protocol  $P$  is denoted by  $T[P]$ . For a DSSE scheme  $\Lambda$ , the first protocol invoked is always the **Setup** protocol. Then, for a query sequence  $\text{QS}_\ell$  and for every query  $Q_i \in \text{QS}_\ell$  a protocol  $P_i$  is invoked, where  $1 \leq i \leq \ell$ .  $P_i = \text{Search}$  if  $\text{op}_i = \text{srch}$  and  $P_i = \text{Update}$  if  $\text{op}_i \in \{\text{add}, \text{del}\}$ . The view of the server in every individual protocol is formally defined as follows,

$$T[\text{Setup}] = \tau_0 = \emptyset, \quad T[P_i] = \tau_i = \begin{cases} (u_i, \text{stk}_i) & \text{if } P_i = \text{Search} \\ (u_i, \text{utk}_i) & \text{if } P_i = \text{Update.} \end{cases}$$

for some  $1 \leq i \leq \ell$ . For every update query, the server gets an update token  $\text{utk}$  and for every search query, the server gets a search token  $\text{stk}$  along with their respective time stamps. So,  $\tau_i$  is either the  $(u_i, \text{utk}_i)$  or  $(u_i, \text{stk}_i)$ , depending upon whether the

---

<sup>1</sup>For example  $q = w$  in case of a single keyword search query, and  $q = \omega$  for conjunctive query.

query is an update query or a search query, respectively.

A “*transcript*” of a query sequence only contains the messages sent from the client for a search or an update query. The output of the server can efficiently be computed from the EDB and the search or update token sent by the client. This is because both the  $\text{Search}_S$  and the  $\text{Update}_S$  algorithms of Definition 2.3.1 are deterministic. With this, we finally define a transcript of a  $i$ -query sequence  $\text{QS}_i$  as follows.

$$\text{Trans}_i[\text{QS}_i] = \{\tau_j : 0 \leq j \leq i\},$$

for  $1 \leq i \leq \ell$ . For convenience, we write  $\text{Trans}_i[\text{QS}_i]$  as  $\text{Trans}_i$ .

### 2.5.1 Security Game

The security of SSE is generally defined in a real-ideal simulation paradigm [42, 60, 28]. A leakage profile  $\mathcal{L}_\Lambda$  is defined to capture all the leakages (due to the setup, search and update protocols). A simulator is given access to the leakage function to simulate the real-world execution of the scheme. Thus, the security definition is parameterised by the leakage profile.

**Definition 2.5.1** (Adaptive security of DSSE). *Let  $\Lambda = (\text{Setup}, \text{Search}, \text{Update})$  be a DSSE scheme and  $\mathcal{L}_\Lambda = (\mathcal{L}_{\Lambda, \text{Setup}}, \mathcal{L}_{\Lambda, \text{Search}}, \mathcal{L}_{\Lambda, \text{Update}})$  be a triple of functions. Then, for any ppt adversary  $\mathcal{A}$  and simulator  $\mathcal{S}$  with the leakage function  $\mathcal{L}_\Lambda$ , we define the experiment  $\text{SSEReal}_{\mathcal{A}}^\Lambda(\lambda)$ , and  $\text{SSEIdeal}_{\mathcal{A}, \mathcal{S}}^\Lambda(\lambda)$  as follows.*

$\text{SSEReal}_{\mathcal{A}}^\Lambda(\lambda)$  : *The adversary  $\mathcal{A}(1^\lambda)$  chooses a DB and the challenger runs  $(\sigma_C, \text{EDB}) \leftarrow \text{Setup}_C(1^\lambda, \text{DB})$  and returns EDB to  $\mathcal{A}$ . Then for subsequent search or update queries, the challenger runs  $(\sigma_C, \text{stk}) \leftarrow \text{Search}_C(\sigma_C, q)$  or  $(\sigma_C, \text{utk}) \leftarrow \text{Update}_C(\sigma_C, \text{op}, (\text{id}, w))$  respectively and provides the adversary with  $\text{stk}$  or  $\text{utk}$ . Finally, after the search/update queries, the adversary  $\mathcal{A}$  stops by outputting a bit, which is the output of the experiment.*

$\text{SSEIdeal}_{\mathcal{A}, \mathcal{S}}^\Lambda(\lambda)$  : *The adversary  $\mathcal{A}(1^\lambda)$  chooses a DB and the challenger runs  $\text{EDB} \leftarrow \mathcal{S}(\mathcal{L}_{\Lambda, \text{Setup}})$  and returns EDB to  $\mathcal{A}$ . Then for subsequent search or update queries, the*

challenger runs  $\text{stk} \leftarrow \mathcal{S}(\mathcal{L}_{\Lambda, \text{Search}})$  or  $\text{utk} \leftarrow \mathcal{S}(\mathcal{L}_{\Lambda, \text{Update}})$  respectively and provides the adversary with  $\text{stk}$  or  $\text{utk}$ . Finally, after the search/update queries, the adversary  $\mathcal{A}$  stops by outputting a bit, which is the output of the experiment.

A DSSE scheme  $\Lambda$  is said to be  $\mathcal{L}_{\Lambda}$ -adaptively secure if for all adversaries  $\mathcal{A}$ , there exists a simulator  $\mathcal{S}$  such that, the SSE advantage  $\left(\text{Adv}_{\mathcal{A}, \mathcal{S}}^{\text{Priv}, \Lambda}(\lambda)\right)$  of  $\mathcal{A}$  (defined below) is negligible in  $\lambda$ .

$$\text{Adv}_{\mathcal{A}, \mathcal{S}}^{\text{Priv}, \Lambda}(\lambda) = \left| \Pr [\text{SSEReal}_{\mathcal{A}}^{\Lambda}(\lambda) = 1] - \Pr [\text{SSEIdeal}_{\mathcal{A}, \mathcal{S}}^{\Lambda}(\lambda) = 1] \right| \leq \text{negl}(\lambda).$$

## 2.5.2 Correctness of DSSE

A DSSE scheme  $\Lambda$  is *correct*, informally, if the result returned by the search protocol for a search query  $q$  is the list of document identifiers matching the query, i.e.,  $(\text{db}(q))$ , except with negligible probability. To define the correctness of a DSSE scheme we use the formalization provided in [28, 38]. We assume that the adversary makes  $\ell$ -many queries where  $\ell \leq \text{poly}(\lambda)$ . The correctness game of a DSSE is defined in Figure 2-1.

DSSECorr $_{\mathcal{A}}^{\Lambda}(\lambda)$ :
01. $\text{flag} \leftarrow 0$
02. $(\text{DB}, \sigma_{\mathcal{A}}) \leftarrow \mathcal{A}(1^{\lambda})$
03. $(\sigma_{\text{C}0}, \text{EDB}_0) \leftarrow \text{Setup}(\text{DB})$
04. <b>for</b> $1 \leq i \leq \ell$ <b>do</b>
05. $(\sigma_{\mathcal{A}}, q_i) \leftarrow \mathcal{A}(\sigma_{\mathcal{A}}, \text{EDB}_0, \text{Trans}_{i-1}) \quad \triangleright q_i = (\text{op}_i, \text{in}_i)$
06. <b>if</b> $\text{op}_i = \text{srch}$
07. $(\sigma_{\text{C}i}, \text{res}_i ; \text{EDB}_i) \leftarrow \text{Search}(\sigma_{\text{C}i-1}, q_i ; \text{EDB}_{i-1})$
08. $\tau_i \leftarrow \text{T}[\text{Search}(\sigma_{\text{C}i-1}, q_i ; \text{EDB}_{i-1})]$
09. <b>if</b> $\text{res}_i \neq \text{db}(q_i)$
10. $\text{flag} \leftarrow 1$
11. <b>end if</b>
11. <b>else</b>
12. $(\sigma_{\text{C}i}, \text{EDB}_i) \leftarrow \text{Update}(\sigma_{\text{C}i-1}, q_i ; \text{utk}_i, \text{EDB}_{i-1})$
13. $\tau_i \leftarrow \text{T}[\text{Update}(\sigma_{\text{C}i-1}, q_i ; \text{utk}_i, \text{EDB}_{i-1})]$
14. <b>end if</b>
15. <b>end for</b>
16. <b>return</b> $\text{flag}$

Figure 2-1: Correctness Game of DSSE.

In the correctness game, the adversary submits an initial database  $DB_0$ , to which it receives an encrypted database  $EDB_0$ . The adversary then makes repeated search or update queries to which it receives a transcript. The DSSE scheme  $\Lambda$  is *correct* if the result returned by the search protocol is *correct* i.e., the search result is  $db(q)$ , except with negligible probability.

**Definition 2.5.2** (Correctness of DSSE). *Let,  $\Lambda = (\text{Setup}, \text{Search}, \text{Update})$  be a DSSE scheme, we say  $\Lambda$  is correct if for all ppt adversary  $\mathcal{A}$ , following holds.*

$$\text{Adv}_{\mathcal{A}}^{\text{Corr}, \Lambda}(\lambda) = \Pr[\text{DSSECorr}_{\mathcal{A}}^{\Lambda}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where the DSSECorr game is defined in Figure 2-1.

### 2.5.3 Soundness of DSSE

In Section 2.3, the Definition of SSE (Definition 2.3.1) is presented assuming an honest-but-curious adversary. Where the adversary always follows the protocol correctly and replies with the correct search result. However, the adversary can be malicious, where it differs from the protocol to save storage and/or CPU time and tries to forge the client with incorrect search results. In such cases, the client needs a guarantee from the SSE protocol that the search result is correct.

A verifiable DSSE (VDSSE) ensures that the search result returned by the server is indeed correct. In a VDSSE the Search protocol is equipped with another additional algorithm called Verify. And the output of the search protocol of the client is either  $db(q)$  or “reject”. The search protocol of a VDSSE is described as follows. All other protocols (i.e., Setup and Update) are the same as Definition 2.3.1.

$(\sigma_C, (db(q) \text{ or reject}) ; EDB) \leftarrow \text{Search}(k, \sigma_C, q ; EDB)$  : is a protocol executed between the client and the server. It comprises the following two algorithms.

$(\sigma_C, \text{stk}) \leftarrow \text{Search}_C(k, \sigma_C, q)$  : this (possibly probabilistic) algorithm is run by the client. On input the key  $k$ , client’s state  $\sigma_C$  and search query  $q$ , the algorithm outputs a search token  $\text{stk}$  and client’s updated state  $\sigma_C$ .

$\text{res} \leftarrow \text{Search}_S(\text{stk}, \text{EDB})$  : this is a deterministic algorithm run by the server. It takes as input the encrypted database  $\text{EDB}$  and the search token  $\text{stk}$ . It outputs a set  $\text{res}$  as the search results  $\text{db}(q)$  for search query  $q$  which is returned to the client.

The client then passes the search result  $\text{res}$  through another algorithm called  $\text{Verify}$ , described as follows.

$(\text{db}(q) \text{ or reject}) \leftarrow \text{Verify}(k, \sigma_C, \text{res})$  : this is a deterministic algorithm which on input the secret key  $k$ , the state  $\sigma_C$  and the result  $\text{res}$  returned by the server, outputs either the correct search result  $\text{db}(q)$  for the query  $q$  or “reject”.

$\text{DSSESound}_A^\Lambda(\lambda)$ : 01. $\text{flag} \leftarrow 0$ 02. $(\sigma_A, \text{DB}) \leftarrow \mathcal{A}(1^\lambda)$ 03. $(\sigma_{C_0}, \text{EDB}_0) \leftarrow \text{Setup}(\text{DB}_0)$ 04. <b>for</b> $1 \leq i \leq \ell$ <b>do</b> 05. $(\sigma_A, q_i) \leftarrow \mathcal{A}(\sigma_A, \text{EDB}_0, \text{Trans}_{i-1}) \quad \triangleright q_i = (\text{op}_i, \text{in}_i)$ 06. <b>if</b> $\text{op}_i = \text{srch}$ 07. $(\sigma_{C_i}, \text{res}_i ; \text{EDB}_i) \leftarrow \text{Search}(\sigma_{C_{i-1}}, q_i ; \text{EDB}_{i-1})$ 08. $\tau_i \leftarrow \text{T}[\text{Search}(\sigma_{C_{i-1}}, q_i ; \text{EDB}_{i-1})]$ 09. <b>if</b> $\text{Verify}(\text{res}_i) \neq \text{reject}$ <b>and</b> $\text{res}_i \neq \text{db}(q_i)$ 10. $\text{flag} \leftarrow 1$ 11. <b>end if</b> 11. <b>else</b> 12. $(\sigma_{C_i}, \text{EDB}_i) \leftarrow \text{Update}(\sigma_{C_{i-1}}, q_i ; \text{utk}_i, \text{EDB}_{i-1})$ 13. $\tau_i \leftarrow \text{T}[\text{Update}(\sigma_{C_{i-1}}, q_i ; \text{utk}_i, \text{EDB}_{i-1})]$ 14. <b>end if</b> 15. <b>end for</b> 16. <b>return</b> $\text{flag}$
--

Figure 2-2: Soundness Game of DSSE

The soundness game of a DSSE is defined in Figure 2-2. In the soundness game, the adversary submits an initial database  $\text{DB}_0$ , to which it receives an encrypted database  $\text{EDB}_0$ . The adversary then makes repeated search or update queries to which it receives a transcript. Informally we say a DSSE scheme  $\Lambda$  is *sound* if the result returned by the search protocol is not rejected and is *correct* i.e., the search

result is either  $\text{db}(q)$  or “reject”. To define the soundness of a DSSE scheme, we use the formalization provided in [24].

**Definition 2.5.3** (Soundness of DSSE). *Let,  $\Lambda = (\text{Setup}, \text{Search}, \text{Update})$  be a DSSE scheme, we say  $\Lambda$  is sound if for all ppt adversary  $\mathcal{A}$ , following holds.*

$$\text{Adv}_{\mathcal{A}}^{\text{Sound}, \Lambda}(\lambda) = \Pr[\text{DSSESound}_{\mathcal{A}}^{\Lambda}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where the  $\text{DSSESound}$  game is defined in Figure 2-2.

## 2.5.4 Fault-tolerant Verifiable DSSE

A DSSE scheme is said to be fault-tolerant and verifiable DSSE if it archives correctness and soundness according to Definition 2.5.2 and 2.5.3 respectively even in the presence of incorrect updates, defined as follows.

**Definition 2.5.4** (Incorrect Update). *An update operation  $(\text{op}, (\text{id}, w))$  is said to be incorrect if any of the following holds*

1. if  $\text{op} = \text{del}$  and  $\text{id} \notin \text{db}(w)$
2. if  $\text{op} = \text{add}$  and  $\text{id} \in \text{db}(w)$ .

**Remark.** If a DSSE is not fault-tolerant, then the adversary in the correctness and the soundness game of Definition 2.5.2 and 2.5.3 respectively are not allowed to make any incorrect update according to Definition 2.5.4.

## 2.5.5 Forward and Backward Privacy (Single Keyword DSSE)

The two most essential security requirements (other than adaptive security) for a DSSE scheme are *forward privacy* and *backward privacy*. Intuitively, forward private SSE schemes ensure that future update queries do not leak any information about previous search operations performed on the inverted index. Similarly, backward privacy ensures that after a pair  $(\text{id}, w)$  has been deleted, the identifier  $\text{id}$  will not be

leaked in future search operations for  $w$ . Depending upon the information they leak during search and update operations, various flavours of backward privacy have been defined in the literature [25]. They are:

1. **Backward privacy with insertion pattern (BPIP):** This notion of backward privacy leaks the document identifiers currently matching the queried keyword  $w$ , the time when they were inserted into the encrypted inverted index, and the total number of updates performed on the searched keyword  $w$ . This security notion is popularly referred to in the literature as **Type I** backward privacy.
2. **Backward privacy with update pattern (BPUP):** This notion of backward privacy in addition to the BPIP leakages above, also leaks the time when all the updates on the queried keyword  $w$  happened without leaking their specific content. This security notion is popularly referred to in the literature as **Type II** backward privacy.
3. **Weak backward privacy (WBP):** This notion of backward privacy in addition to the leakages in the previous two leakages BPIP and BPUP, also leaks the deletion updates that removed corresponding insertion updates. This security notion is popularly referred to in the literature as **Type III** backward privacy.

### Some Standard Leakage Functions.

To define forward and backward privacy, we need some auxiliary definitions, which we introduce here.

$\text{sp}(w)$  : The search pattern leakage  $\text{sp}(w)$  in terms of SSE refers to the information of those time stamps when the keyword has been searched. The searched pattern for a keyword  $w$  is the set of timestamps of all search queries with the keyword  $w$ . Formally,

$$\text{sp}(w) = \{u : (u, \text{srch}, w) \in \text{QS}\}.$$

$\text{Hist}(w)$  : This leakage function captures all the modifications that were made to the set  $\text{db}(w)$ . This is the combination of two leakages, the leakage from  $\text{db}_0(w)$  (the set of document identifiers matching  $w$  at the setup) and the set  $\text{UpHist}(w)$ , defined



below.

$$\text{UpHist}(w) = \{(u, \text{op}, \text{id}) : (u, \text{op}, (\text{id}, w)) \in \text{QS} \wedge \text{op} \in \{\text{add}, \text{del}\}\}.$$

$\text{TimeDB}(w)$  : This leakage function is defined as the timestamped list of documents currently matching the queried keyword  $w$ . We consider the refined definition provided in [38] to correctly capture the timestamps of entries that have not been deleted yet.

$$\text{TimeDB}(w) = \{(u, \text{id}) : ((u, \text{add}, (\text{id}, w)) \in \text{QS}) \wedge (\forall u' > u, (u', \text{del}, (\text{id}, w)) \notin \text{QS})\}.$$

$\text{Updates}(w)$  : Similarly, we define another leakage function  $\text{Updates}(w)$  as the set of timestamps of all update queries for a given  $w$  as,

$$\text{Updates}(w) = \{u : (u, \text{op}, (\text{id}, w)) \in \text{QS}\}.$$

$\text{DelHist}(w)$  : Finally, we define  $\text{DelHist}(w)$  as the set of all pairs of timestamps of addition and deletion for a given  $w$ .

$$\text{DelHist}(w) = \{(u^{\text{add}}, u^{\text{del}}) : ((u^{\text{add}}, \text{add}, (\text{id}, w)) \in \text{QS}) \wedge ((u^{\text{del}}, \text{del}, (\text{id}, w)) \in \text{QS})\}.$$

With the above definitions in place, we are now ready to provide the definition of forward and backward privacy for a  $\mathcal{L}_\Lambda$ -adaptively secure single keyword SSE  $\Lambda$ . The definition that we present is primarily the one presented in [25] with some modifications as suggested in [38]. Before we give the definition we discuss some history of the evolution of the definitions of forward and backward security.

The notion of forward and backward privacy was first introduced in [88]. However, forward privacy was first formalised in [23]. The first formal notion of backward privacy for single keyword DSSE was introduced in [25]. The authors of [38] modified the definition of backward privacy in [25] in several ways. The primary observations of [38] regarding the definition in [25] are the following:

- The definition in [25] does not capture a specific type of leakage called *search pattern* leakage correctly.
- The definition in [25] considers that an identifier-keyword pair, once inserted and then deleted, cannot be reinserted again into the database under the *weak backward privacy* (WBP) security model. This restriction in the context of WBP security model is artificial.
- Also, the security definition of WBP in [25] missed the leakage component  $\text{Updates}(w)$ . Authors of [38] showed that the leakage  $\text{Updates}(w)$  cannot be derived from the definition of WBP provided in [25], thus, should be included explicitly in the definition of WBP

In [38] the definition in [25] was modified to take care of all the above points, i.e., the definition in [38] takes care of search pattern leakages and also lifts the artificial restriction for the WBP security model. In doing so, the authors of [38] provided an alternate definition of backward privacy called, *Backward Privacy with Link Pattern* (BPLP). We require some extra leakage function to define BPLP. However, this alternate formulation of backward privacy is not explicitly required for the works presented in this thesis. Therefore, we adopt the definitions of forward and backward privacy from [25] with the modification suggested in [38].

We will need some additional notations. For two leakage functions  $\mathcal{L}_1$  and  $\mathcal{L}_2$ ,  $\mathcal{L}_1 \preceq \mathcal{L}_2$  denotes that “ $\mathcal{L}_1$  leaks less than  $\mathcal{L}_2$ ”. In other words,  $\mathcal{L}_1$  leaks less than  $\mathcal{L}_2$  means that leakage information given by  $\mathcal{L}_1$  about the database is less than  $\mathcal{L}_2$ , or, equivalently, all the information that can be inferred from  $\mathcal{L}_1$  can be inferred from  $\mathcal{L}_2$ . Whenever,  $\mathcal{L}_1$  leaks strictly less than  $\mathcal{L}_2$ , we write  $\mathcal{L}_1 \prec \mathcal{L}_2$ .

It is a usual practice to assume the initial database of any SSE to be empty. Thus, the leakage of the setup algorithm is null. However, if the initial database is not empty, then  $\mathcal{L}_{\Lambda, \text{Setup}}(\text{DB}) = N$ , where  $N$  denotes the number of keyword document pairs in the initial database. The definition of forward and backward privacy of SSE that we present next assumes the initial database to be empty.

**Definition 2.5.5** (Forward Privacy of Single keyword DSSE). An  $\mathcal{L}_{\Lambda} = (\mathcal{L}_{\Lambda, \text{Setup}},$

$\mathcal{L}_{\Lambda, \text{Search}}, \mathcal{L}_{\Lambda, \text{Update}}$ )-adaptive secure DSSE scheme  $\Lambda$  is forward private if its leakage function  $\mathcal{L}_{\Lambda}$  can be written as follows.

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

**Definition 2.5.6** (Backward Privacy of Single keyword SSE). An  $\mathcal{L}_{\Lambda} = (\mathcal{L}_{\Lambda, \text{Setup}}, \mathcal{L}_{\Lambda, \text{Search}}, \mathcal{L}_{\Lambda, \text{Update}})$ -adaptive secure DSSE scheme  $\Lambda$  is backward private as per the notions BPIP, BPUP and WBP if its leakage functions  $\mathcal{L}_{\Lambda}$  can be written as the follows.

$$\begin{aligned}\mathcal{L}_{\Lambda, \text{Setup}}^{\text{BPIP}}(\text{DB}) &= \emptyset, & \mathcal{L}_{\Lambda, \text{Update}}^{\text{BPIP}}(\text{op}, \text{in}) &\preceq \{\text{op}\}, \\ \mathcal{L}_{\Lambda, \text{Search}}^{\text{BPIP}}(w) &\preceq \{\text{sp}(w), \text{TimeDB}(w), |\text{Updates}(w)|\}, \\ \mathcal{L}_{\Lambda, \text{Setup}}^{\text{BPUP}}(\text{DB}) &= \emptyset, & \mathcal{L}_{\Lambda, \text{Update}}^{\text{BPUP}}(\text{op}, \text{in}) &\preceq \{\text{op}, w\}, \\ \mathcal{L}_{\Lambda, \text{Search}}^{\text{BPUP}}(w) &\preceq \{\text{sp}(w), \text{TimeDB}(w), \text{Updates}(w)\}, \\ \mathcal{L}_{\Lambda, \text{Setup}}^{\text{WBP}}(\text{DB}) &= \emptyset, & \mathcal{L}_{\Lambda, \text{Update}}^{\text{WBP}}(\text{op}, \text{in}) &\preceq \{\text{op}, w\}, \\ \mathcal{L}_{\Lambda, \text{Search}}^{\text{WBP}}(w) &\preceq \{\text{sp}(w), \text{TimeDB}(w), \text{Updates}(w), \text{DelHist}(w)\}.\end{aligned}$$

In our work, specifically in Chapter 5, we are required to define forward and backward privacy for conjunctive SSEs. We modify the above definitions, which are essentially the definitions presented in [38], to make them suitable for conjunctive SSEs.

## 2.6 Related Works

**The Early Schemes.** SSE was first introduced in [86]. One of the main drawbacks of the scheme in [86] was that the search time was linear in the number of documents. This limitations were lifted in the works reported in [50] and [34]. Constructions based on an inverted index were first proposed in [50]. The use of an inverted index was significant, as its use ensured that the search time was proportional to the

number of documents matching the keyword in the collection. Authors of [50] first introduced formal security notions of SSE. They proposed the notion of indistinguishability against chosen keyword attacks (IND-CKA) and the slightly stronger IND2-CKA) notion for the security of indexes, they also provide a construction based on Bloom filters [17] and pseudo-random functions. Authors of [34] provide a construction that achieves security similar to IND2-CKA.

The first strong notion of security, which is still in use, was first proposed in [42]. However, the scheme described in [42] was only *static*. The first truly dynamic SSE with search complexity  $\mathcal{O}(\text{db}(w))$  was proposed in [60]. Since then, several dynamic single keyword SSE schemes have been proposed.

**Attacks on SSE.** User queries do leak some information to the server regarding searches and updates. Determining such leakages is important. There have been several works which describe attacks on SSE schemes which utilize the leakages during search and updates [55, 68, 27]. A unifying theme of these attacks was to use knowledge of the plain text database a priori. One of such attacks reported in [97] extended the attack by [27] to demonstrate that the server may inject only a few documents (through the client) into the database, to be able to recover a large fraction of the keywords searched by the client. Authors of [1] showed the consequences of this attack on relational databases.

**Forward and Backward Privacy.** The attacks of [97] above created a necessity to update the security definition of SSEs. To mitigate such attacks, two new notions of privacy for DSSE, namely, “*forward privacy*” and “*backward privacy*”, were first proposed in [88]. Intuitively, forward private SSE schemes ensure that the server cannot link an update operation for  $(\text{id}, w)$  with previous search operations performed on the inverted index with  $w$  or containing  $\text{id}$  in its result. This protects queries made to the inverted index with newly injected files by the adversary as was used in the attacks mounted in [97]. Backward privacy ensures that after a pair  $(\text{id}, w)$  has been deleted, the identifier  $\text{id}$  will not be leaked in future searches for  $w$ .

The first formal definition and a scheme achieving forward privacy were proposed in [23]. The scheme of [23] uses a trapdoor one-way permutation to achieve forward privacy. The first formal definition of backward privacy was defined in [25]. Depending upon the information leaked during search and update operations, various flavours of backward privacy have been defined in the literature [25], which has been followed in subsequent works [49, 90, 101, 89]. Authors of [25] defined three types of backward privacy, namely, BPIP, BPUP and WBP in increasing order of leakage (See Definition 2.5.6) to the server (decreasing order of security). These three types are popularly referred to as **Type I**, **Type II** and **Type III** backward privacy in literature. The authors of [101] gave a construction using a bit-map index with homomorphic addition to achieve **Type I** security whose leakage is even lesser than a **Type I** secure scheme. However, their storage and search complexities are high at  $\mathcal{O}(u|D|)$ , where  $u$  is the number of updates for a keyword. More efficient construction of a **Type III** backward private SSE scheme was proposed in [89] that uses *symmetric puncturable encryption*. Since then, several forward and backward private SSE schemes [49, 101, 38] have been proposed, improving security and efficiency.

In [38], the authors demonstrated that the definition of backward privacy proposed in [25] was incomplete, and it didn't cover all general scenarios. Also, the authors of [25] did miss some leakage while defining backward privacy. In [25], a relaxed version of SSE was considered in their **Type III** security definition. They considered that once a keyword-identifier pair is added and then deleted, the same pair cannot be re-inserted into the database. This restriction was lifted in [38] to provide a more general and complete definition of backward privacy for SSE with all perceivable and known leakages. Although there are no known practical attacks on SSE schemes due to a lack of backward privacy, it provides meaningful theoretical assurance.

**SSE Supporting Complex Queries.** Most SSE schemes [42, 60, 28, 23, 87, 25, 49, 89, 38, 90] allow searching for a single keyword in each query. Supporting conjunctive queries and/or general boolean queries with efficient computation and communication along with adequate security (low leakage) is a well-known challenge in the SSE

literature. There are a few SSE schemes that support conjunctive queries [29, 47, 2, 93, 58, 66, 80, 102, 100, 95]. The most notable of these is the oblivious cross-tag (OXT) system from Crypto’13 [29]. It assumes that the client knows the *least frequent keyword*  $w_1$  in a conjunction  $w_1 \wedge \dots \wedge w_n$ . This keyword  $w_1$  is called the *s-term* and all other keywords  $w_2, \dots, w_n$  are called *x-terms* [29]. OXT uses two data structures T-Set and X-Set. During a search operation in OXT, the server first conducts a single keyword search using the T-Set to find all ids containing the *s-term*  $w_1$ . For all such ids containing  $w_1$ , the X-Set is used to check if the *x-terms*  $w_i, i \geq 2$  are also in id. The OXT scheme is quite efficient with its search complexity of  $\mathcal{O}(n \cdot |\text{db}(w_1)|)$ , where  $n$  is the number of keywords (*s-term* and *x-terms*) in the conjunction and the set  $\text{db}(w_1)$  contains the ids of documents where the *s-term*  $w_1$  appears.

The main criticism of the OXT scheme was that besides being a static scheme, it leaked keyword pair result patterns (KPRP). Hiding KPRP is important in the context of the attacks presented in [97]. Moreover, OXT uses exponentiation in ddh-hard groups, which is a costly operation and removing this would lead to making OXT more efficient. Subsequently, proposed static schemes like [66, 78] were directed towards mending these shortcomings of OXT. The work in [66] used significant pre-processing on the data to restrict KPRP leakage. Their construction uses Bloom filters and symmetric hidden vector encryption in a novel way. The construction of [78] has designed a novel data structure called **ConjFiler**, which allows searching for conjunction only using symmetric key primitives. However, **ConjFiler** requires to pre-compute all possible two-conjunction to facilitate conjunctive query. This is why making **ConjFiler** dynamic is a challenging task.

The first proposed construction of a forward and backward secure CDSSE was in [80] through a new construction called Oblivious Dynamic Cross Tags (ODXT), that successfully managed the pre-processing of OXT in the dynamic setting. However, the attack in [100] showed that the construction in [80] failed to provide forward privacy in some cases. Few other forward and backward private CDSSE schemes [102, 100, 95] have followed, achieving forward privacy and different flavours of backward privacy with efficiency trade-offs. However, these schemes are far from achieving the

efficiency of the OXT scheme. The scheme of [102] performs  $\mathcal{O}(|D|)$  computation for search, where  $|D|$  is the number of documents in the database, making it prohibitively slow. The scheme of [100] uses a single keyword SSE (SKSSE) to store and fetch the elements in the conjunction which makes it very slow as well. The scheme of [95] runs in sub-linear time but is still far from the practical performance of OXT. In [95], the computation and communication complexities of the search operation are both  $\mathcal{O}(u_1 + n \cdot \text{db}(w_1))$  and involves two rounds, where  $u_1$  is the number of updates related to  $w_1$ ; the complexities of the update operation are  $\mathcal{O}(|W|/|W_d|)$  for adding a document,  $\mathcal{O}(|D|)$  for editing and  $\mathcal{O}(1)$  for deleting, where  $|W|$  is the number of keywords in the database and  $|W_d|$  denotes the number of keywords contained in the updated document. The storage requirement of [95] is  $\mathcal{O}(|W||D|)$ , as compared to  $\mathcal{O}(N)$  for our scheme, where  $N$  is the number of identifier-keyword pairs in the database. Typically,  $N$  is much smaller than  $|W||D|$  in practice. There is no additional storage requirement due to the keyword updates in [95].

**Verifiable SSE.** In SSE, it's assumed the server (seen as an adversary) is honest-but-curious. However, the servers can be "malicious". Verifiable SSE (VSSE) provides cryptographic guarantees against such servers.

The first universally composable secure VSSE scheme was introduced in [63] against non-adaptive adversaries. The first adaptive secure VSSE was proposed in [30], though both works addressed the static SSE setting. Authors of [64] extended their previous work [63] to a verifiable dynamic SSE (VDSSE), supporting updates operation, but it lacked forward privacy. The first forward private VDSSE was proposed in [24], using the incremental multi-set hashing proposed in [40]. Subsequently, Zhang et al. [98] introduced another incremental multi-set hash-based SSE using symmetric key primitives. The client-side storage required for both schemes was high. Authors of [48] addressed this issue by proposing a novel data structure called the *Accumulative Authentication Tag* (AAT), which reduces storage overhead while using symmetric key primitives.

While all these constructions supported verification of the search result, almost all

of them failed to provide correct proof in case of faulty updates from the client. There has been limited research focused on designing DSSE schemes that are fault-tolerant, verifiable, and both forward and backward private. In [88], authors proposed a DSSE scheme using an ORAM-style data structure, but it required frequent rebuild operations. This scheme was forward secure and fault-tolerant, but not verifiable. Also, the search operation in this scheme is inefficient, taking linear time (on the number of documents) in the worst case. Authors of [24] proposed three forward private, verifiable DSSE schemes GVS-Hash, GVS-Acc and GVS-Acc-RSA. The first construction is based on Merkle tree-like data structures and Multi-Set Hash, while the latter two are based on cryptographic accumulators. However, none of these constructions are fault-tolerant and they do not scale well with large databases. In addition to these constructions, authors of [24] also proposed verifiable versions of Linear SPS and Sublinear SPS introduced in [88]. Recently, an efficient FVDSSE construction was proposed in [94] that employs *authenticated encryption* (AE) for verification, avoiding the expensive use of Merkle trees or accumulators, thereby making the construction very fast in practice. Their proposed generic scheme, when instantiated with a forward private scheme [87], ensures forward privacy. However, this construction fails to provide backward privacy.

## 2.7 Final Remarks

In this Chapter, we discussed the preliminaries required to appreciate the work reported in the subsequent chapters. We listed the general notations which are heavily used throughout the thesis, some less used notations would be introduced further wherever necessary. We also discussed all necessary background for SSE schemes including their syntax, security and variants. We also provided a comprehensive discussion on the related works which trace the history of the development of SSE. Though we do not pretend that our coverage of SSE schemes in this chapter is exhaustive, we briefly discuss all works which are of relevance to the work that we present in the following chapters.



We already described in Section 2.3 that Searchable Symmetric Encryption schemes, as introduced in [86, 42], abstract a database (DB) as a set of keyword-identifier pairs. Formally,

$$\text{DB} = \bigcup_{w \in W} \{(\text{id}, w) : \text{id} \in \text{db}(w)\},$$

where  $\text{db}(w)$  is the set of identifiers corresponding to the keyword  $w$ . An SSE encrypts this set, comprising keyword-identifier pairs, using a specialized structure often referred to as an “*encrypted multi-map*” [35] (EMM) or an inverted index. This encrypted multi-map is stored on the server and facilitates both searching and updating the encrypted database.

The efficiency of an SSE scheme is measured by the time it takes to search and update. Most schemes are optimized to get better search and update times maintaining adequate security guarantees. Optimizing the size of the encrypted database has not yet received adequate attention in the literature, as it is always assumed that the server is computationally powerful and has sufficient storage. But, optimizing the size of the EMM has an immediate effect on search time and communication overhead. If the input set of tuples for an SSE is DB, then the size of the encrypted database in any existing SSE scheme, to the best of our knowledge, is at least  $|\text{DB}|$ . In fact, in most schemes (particularly in dynamic SSE schemes), the storage overhead is much more than  $|\text{DB}|$ . Similarly, the size of the result of a query  $w$  is at least  $|\text{db}(w)|$ , and in most schemes, the size of the search result which is to be communicated to the

client is much more than  $|\mathbf{db}(w)|$ . In most dynamic SSEs the size of the encrypted database and the query result increases with the number of updates performed on the database.

Recently, some effort has been made to restrict the size of the search result to at most  $\mathcal{O}(\mathbf{db}(w))$  for any given keyword  $w$  [88, 49, 43, 33]. To the best of our knowledge, no SSE scheme has considered reducing the storage size beyond  $|\mathbf{DB}|$  and the result size beyond  $|\mathbf{db}(w)|$ . In this study, we specifically address the following questions.

1. Can we build a functional and secure SSE whose storage size is smaller than  $|S|$  on average?
2. Can a search query for a keyword  $w$  on average be answered with a result size smaller than  $|\mathbf{db}(w)|$ ?

We answer both these questions in the affirmative. In order to achieve this, we introduce a generic method for transforming any single keyword SSE scheme into an equivalent secure SSE scheme where the size of the encrypted database required to store for searching is much smaller than in the original SSE scheme. This results in a more compact representation of both the encrypted database and the set  $\mathbf{db}(w)$  compared to the original SSE scheme. Thus, a reduction of the size of the encrypted database and  $\mathbf{db}(w)$  also results in a reduction of search time and communication costs.

**The Basic Technique:** Given  $\mathbf{db}(w)$  for any keyword  $w$ , we convert  $\mathbf{db}(w)$  to a new set  $c_w$ , which we call as the “cover of the keyword”  $w$ . We ensure that, on average,  $|c_w|$  is smaller than  $|\mathbf{db}(w)|$ , while retaining the ability to fully recover  $\mathbf{db}(w)$  from  $c_w$ . We provide  $\widetilde{\mathbf{DB}} = \bigcup_{w \in \mathcal{W}} c_w \times \{w\}$  as an input to any standard secure SSE. Note that  $\widetilde{\mathbf{DB}}$  is on average less than the size of  $\mathbf{DB} = \bigcup_{w \in \mathcal{W}} \mathbf{db}(w) \times \{w\}$ , which is the input to the original SSE. This generic transformation of  $\mathbf{db}(w)$  to its cover  $c_w$  produces a considerable reduction in both the storage size and result size, resulting in shorter search time and reduced communication size.

The heart of our technique lies in representing  $\mathbf{db}(w)$  for each  $w$  as a full binary

tree. For each keyword  $w \in \mathcal{W}$ , we construct a virtual full binary tree where each leaf node is associated with an identifier  $\text{id}_i$ . We label the leaf nodes with  $+$  if the identifier corresponding to it is present in  $\text{db}(w)$  and label it with  $-$  otherwise. This tree with the labels in the leaf nodes uniquely represents  $\text{db}(w)$ . We devise a scheme by which this tree can be represented uniquely by a small set of nodes of the tree, and we call this set the cover of the tree.

In this Chapter we define a tree cover and propose several algorithms to generate covers. We do an in-depth combinatorial analysis of these algorithms and prove them to be optimal in our setting. We also provide theoretical estimates of the average size of the covers generated by our algorithm(s) and experimentally validate these theoretical estimates. The cover generation algorithms proposed in this chapter are further used in Chapter 4 to design SSEs. The material in this chapter does not make reference to SSEs and has a complete combinatorial flavor, which can be of independent interest. This chapter was partly published in [31].

## 3.1 Binary Trees and Tree Covers

### 3.1.1 Binary Trees

Throughout our thesis, we'll consider a *complete binary tree*, i.e., a binary tree where all levels are present (unless specified otherwise). The *depth* of a node  $i$  of a binary tree is the number of edges in the path from  $i$  to the root. The root has a depth of 0. In a complete binary tree, there are exactly  $2^d$  nodes at depth  $d$ , and consequently, all leaf nodes are at the same depth. The *height* of a node  $i$  is the number of edges in a path from  $i$  to the deepest leaf. The height of the tree is the height of the root of the tree. A complete binary tree of height  $h$ , has exactly  $2^{h+1} - 1$  nodes, with  $2^h$  leaf nodes. The structure of a complete binary tree can be specified only using its height. Let  $\mathcal{T}$  be a tree and  $i$  a node of  $\mathcal{T}$ , then by  $\mathcal{T}(i)$  we denote the subtree rooted at  $i$ . Thus, if  $i$  is the root of  $\mathcal{T}$ , then  $\mathcal{T}(i) = \mathcal{T}$  and if  $i$  is a leaf node then  $\mathcal{T}(i)$  is a tree containing the single node  $i$ . Unless specifically mentioned, by a binary tree, we will

mean a complete binary tree.

For convenience, we will denote nodes of a complete binary tree by integers. The root will be denoted by zero. The nodes in depth  $d$  will be denoted by  $2^d$  consecutive integers starting with  $2^d - 1$ , counting left to right. Thus, for a tree  $\mathcal{T}$  of height  $h$ , the nodes of  $\mathcal{T}$  would be denoted by the set of integers  $\text{nodes}(\mathcal{T}) = \{0, 1, \dots, 2^{h+1} - 2\}$  and the leaf nodes would be denoted by the set  $\text{leaves}(\mathcal{T}) = \{2^h - 1, 2^h, \dots, 2^{h+1} - 2\}$ . For any  $i \in \text{nodes}(\mathcal{T})$ ,  $\text{leaves}(\mathcal{T}(i)) \subseteq \text{leaves}(\mathcal{T})$  will denote the leaf nodes of the subtree rooted at  $i$ .

We will sometimes use an alternative representation of the leaf nodes. Given a tree  $\mathcal{T}$  of height  $h$  with  $\text{leaves}(\mathcal{T}) = \{2^h - 1, 2^h, \dots, 2^{h+1} - 2\}$ . Let  $\phi_h : \text{leaves}(\mathcal{T}) \rightarrow [1, 2^h]$ , be a bijection defined as  $\phi_h(i) = i - 2^h + 2$ . Note that for any  $i \in \text{leaves}(\mathcal{T})$ , if  $j = \phi_h(i)$ , then the leaf node  $i$  in  $\mathcal{T}$  is the  $j^{\text{th}}$  leaf node from the left. As  $\phi$  is a bijection, we have  $\phi_h^{-1} : [1, 2^h] \rightarrow \text{leaves}(\mathcal{T})$  defined as  $\phi_h^{-1}(j) = j + 2^h - 2$ .

### 3.1.2 Tree Cover

Our main object of interest is a *complete binary tree*, i.e., a binary tree where all levels are present. We fix some basic terminology first.

The *depth* of a node  $i$  of a binary tree is the number of edges in the path from  $i$  to the root. The root has a depth of 0. In a complete binary tree, there are exactly  $2^d$  nodes at depth  $d$ , and consequently, all leaf nodes are at the same depth. The height of a node  $i$  is the number of edges in a path from  $i$  to the deepest leaf. The height of the tree is the height of the root of the tree. A complete binary tree of height  $h$ , has exactly  $2^{h+1} - 1$  nodes, with  $2^h$  leaf nodes. The structure of a complete binary tree can be specified only using its height. Let  $\mathcal{T}$  be a tree and  $i$  a node of  $\mathcal{T}$ , then by  $\mathcal{T}(i)$  we denote the subtree rooted at  $i$ . Thus, if  $i$  is the root of  $\mathcal{T}$ , then  $\mathcal{T}(i) = \mathcal{T}$  and if  $i$  is a leaf node then  $\mathcal{T}(i)$  is a tree containing the single node  $i$ . Unless specifically mentioned, by a binary tree, we will mean a complete binary tree.

For convenience, we will denote nodes of a complete binary tree by integers. The root will be denoted by zero. The nodes in depth  $d$  will be denoted by  $2^d$  consecutive integers starting with  $2^d - 1$ , counting left to right. Thus, for a tree  $\mathcal{T}$  of height  $h$ , the

nodes of  $\mathcal{T}$  would be denoted by the set of integers  $\text{nodes}(\mathcal{T}) = \{0, 1, \dots, 2^{h+1} - 2\}$  and the leaf nodes would be denoted by the set  $\text{leaves}(\mathcal{T}) = \{2^h - 1, 2^h, \dots, 2^{h+1} - 2\}$ . For any  $i \in \text{nodes}(\mathcal{T})$ ,  $\text{leaves}(\mathcal{T}(i)) \subseteq \text{leaves}(\mathcal{T})$  will denote the leaf nodes of the subtree rooted at  $i$ .

We will sometimes use an alternative representation of the leaf nodes. Given a tree  $\mathcal{T}$  of height  $h$  with  $\text{leaves}(\mathcal{T}) = \{2^h - 1, 2^h, \dots, 2^{h+1} - 2\}$ . Let  $\phi_h : \text{leaves}(\mathcal{T}) \rightarrow [1, 2^h]$ , be a bijection defined as  $\phi_h(i) = i - 2^h + 2$ . Note that for any  $i \in \text{leaves}(\mathcal{T})$ , if  $j = \phi_h(i)$ , then the leaf node  $i$  in  $\mathcal{T}$  is the  $j^{\text{th}}$  leaf node from the left. As  $\phi_h$  is a bijection, we have  $\phi_h^{-1} : [1, 2^h] \rightarrow \text{leaves}(\mathcal{T})$  defined as  $\phi_h^{-1}(j) = j + 2^h - 2$ .




Figure 3-1: The tree  $\mathcal{T}$  used in Example 1.

**Definition 3.1.1** (Configuration). Let  $\mathcal{T}$  be a complete binary tree of height  $h$  and  $\text{leaves}(\mathcal{T})$  be the set of leaf nodes of  $\mathcal{T}$ . A configuration of  $\mathcal{T}$  is a function  $\Lambda_{\mathcal{T}} : S \rightarrow \{+, -\}$ , where  $S \subseteq \text{leaves}(\mathcal{T})$ . The configuration is complete if  $S = \text{leaves}(\mathcal{T})$ , i.e., a complete configuration of  $\mathcal{T}$  is a map  $\Lambda_{\mathcal{T}} : \text{leaves}(\mathcal{T}) \rightarrow \{+, -\}$ . We will represent  $\Lambda_{\mathcal{T}}$  by a subset of  $\text{leaves}(\mathcal{T}) \times \{+, -\}$ .

Thus, a configuration puts labels from the set  $\{+, -\}$  to the leaves of a tree. If all leaves of a tree are labelled then the configuration is called a *complete configuration*. We represent a configuration  $\Lambda_{\mathcal{T}} : S \rightarrow \{+, -\}$ , where  $S \subseteq \text{leaves}(\mathcal{T})$  by the set  $L = \{(i, s) : i \in S, s = \Lambda_{\mathcal{T}}(i)\}$ .

**Example 1.** Consider the full binary tree  $\mathcal{T}$  of height 3 shown in Figure 3-1. The nodes of  $\mathcal{T}$  are elements of the set  $\text{nodes}(\mathcal{T}) = \{0, 1, \dots, 14\}$  and  $\text{leaves}(\mathcal{T}) = \{7, 8, \dots, 14\}$ .  $L = \{(7, -), (8, +), (9, +), (10, +), (11, -), (12, +), (13, +), (14, +)\}$ , is a complete configuration of  $\mathcal{T}$ , where the nodes 7 and 11 are labeled  $-$  and all other

nodes are labeled  $+$ . Whereas, the configuration  $\{(7, -), (8, +), (9, -), (10, +), (11, -)\}$  is not complete, as the leaf nodes 12, 13 and 14 are not assigned any labels.

### 3.1.3 Cover of a Configuration

Our goal is to have a succinct representation of the configuration of a binary tree. For a tree  $\mathcal{T}$  the naive representation of its configuration can be done through a subset of  $\text{leaves}(\mathcal{T}) \times \{+, -\}$ , i.e, by listing all the leaf nodes along with their labels. We are interested in a more compact representation. Informally, a cover is a succinct representation of a configuration.

We define a cover of a configuration through a pair of algorithms *cover generation* and *cover reconstruction*.

**Definition 3.1.2** (Syntax of a Tree Cover Scheme). *A tree cover scheme  $\Psi = (\text{CoverGen}, \text{ReConstruct})$  is a tuple of two algorithms defined as follows.*

- $C \leftarrow \text{CoverGen}(h, L)$ : *On input a configuration  $L$  of a tree  $\mathcal{T}$  of height  $h$ , CoverGen outputs a cover  $C \subset \text{nodes}(\mathcal{T}) \times \{+, -\}$ .*
- $L \leftarrow \text{ReConstruct}(h, C)$ : *On input the height  $h$  of a tree  $\mathcal{T}$  and a cover  $C \subset \text{nodes}(\mathcal{T}) \times \{+, -\}$ , ReConstruct outputs a configuration  $L \subset \text{leaves}(\mathcal{T}) \times \{+, -\}$ .*

**Correctness:** We say a tree cover scheme  $\Psi$  is correct if for any configuration  $L$  of a tree of height  $h$ ,

$$L = \Psi.\text{ReConstruct}(h, \Psi.\text{CoverGen}(h, L)).$$

A tree cover scheme is interesting only if the covers produced by its cover generation algorithm have a size much smaller than the configuration (which it receives as input) for most configurations. The cover generation schemes described later satisfy this property.

We do not yet provide any structural characterisation of a cover of a configuration. A cover is just what is output by a cover generation algorithm, with the guarantee

that there is a corresponding reconstruction algorithm which can reconstruct the configuration from which the cover was generated. This syntactic definition of a cover would be enough for us to develop SSE schemes with desirable properties, but at this point, it may be useful to see what a cover generation algorithm can possibly do.

Let us take the example of Fig 3-1, which represents a complete configuration  $L = \{(7, -), (8, +), (9, +), (10, +), (11, -), (12, +), (13, +), (14, +)\}$ , of a tree  $\mathcal{T}$  of height 3. Nodes drawn in circles that is the nodes  $\{8, 9, 10, 12, 13, 14\}$  are marked as  $+$  and nodes drawn in squares that is  $\{7, 11\}$  are marked as  $-$ . According to our definition so far, there can be several algorithms for the cover generation which can possibly output different covers for the same tree with the same labeled leaf nodes. We focus on two such simple instances.

1.  $L$  itself can be a cover. In this case, on input  $L$  the cover generation algorithm simply outputs  $L$  as the cover and the reconstruction algorithm on input any set  $C$ , just outputs  $C$ .
2. The set  $C = \{(7, -), (11, -)\}$ , can be a cover. Here the cover generation algorithm on input  $L$  outputs the set  $C = \{(i, -) : (i, -) \in L\}$ , i.e., outputs those nodes (along with their labels) which are labeled  $-$ . The cover reconstruction algorithm on input  $C$  labels those nodes in  $\text{leaves}(\mathcal{T})$  which are not in  $C$  with  $+$  and calls those labeled nodes as  $L_1$  and outputs  $C \cup L_1$ . It is easy to see that this reconstruction always works when  $L$  is a complete configuration.

The above two instances are trivial covers. Note that in case (2) we already have a cover whose size is much smaller than the configuration, at least for the example that we consider. In Section 3.2, we systematically study much more complex cover generation algorithms, which, on average, give covers whose sizes are smaller than the configuration. For designing dynamic SSEs a notion of covers for a dynamic configuration would be necessary which we provide in the next subsection. The design of the SSEs presented in Sections 4.1.2 and 4.2 only assumes the definitions of cover generation and reconstruction algorithms. Exact instances of cover generation

algorithms, as discussed in Section 3.2, are not required to follow the material in Sections 4.1.2 and 4.2.

### 3.1.4 Cover of a Dynamic Configuration

We will be interested in trees where the configuration is dynamic, i.e., we may start with an initial tree along with a configuration, and the tree may change by more nodes getting added to it and/or by the leaves changing labels.

For the sake of modelling dynamic trees, we will use the alternative representation of leaf nodes. Recall for a tree  $\mathcal{T}$  of height  $h$ , if  $i \in \text{leaves}(\mathcal{T}) = [2^h - 1, 2^{h+1} - 2]$ , then  $\phi_h(i) = i - 2^h + 2$ , is the position of the leaf  $i$  from the left. Let for a configuration  $L$  of tree  $\mathcal{T}$  with height  $h$ , we define

$$\phi_h(L) = \{(\phi_h(i), s) : (i, s) \in L\}.$$

Thus,  $\phi_h(L)$  is just a different representation of the configuration  $L$ , where the leaf nodes are specified by their position from the left. Similarly, for any  $S \subseteq [1, 2^h] \times \{+, -\}$ , we define  $\phi_h^{-1}(S) = \{(\phi_h^{-1}(i), s) : (i, s) \in S\}$ .

Let  $T_a, T_b$  be two complete binary trees of height  $h_a$  and  $h_b$  respectively. Let  $L_a$  and  $L_b$  be complete configurations of  $T_a$  and  $T_b$  respectively. Let  $\tilde{L}_a = \phi_{h_a}(L_a)$  and  $\tilde{L}_b = \phi_{h_b}(L_b)$ . Let  $s \in \{+, -\}$ . If  $s = +$ , then  $\bar{s} = -$  and vice versa. We define two sets  $\tilde{D}$  and  $\tilde{A}$  as

$$\begin{aligned} \tilde{D} &= \left\{ (i, s) \in \tilde{L}_b : (i, \bar{s}) \in \tilde{L}_a \right\}. \\ \tilde{A} &= \left\{ (i, s) \in \tilde{L}_b : (i, s) \notin \tilde{L}_a \wedge (i, \bar{s}) \notin \tilde{L}_a \right\}. \end{aligned}$$

Note, the set  $\tilde{D}$  contains those labeled leaf nodes in  $\tilde{L}_a$  whose labels have changed in  $\tilde{L}_b$  and  $\tilde{A}$  contains those labeled nodes in  $\tilde{L}_b$  which are not present in  $\tilde{L}_a$ . As, both  $L_a$  and  $L_b$  are complete configurations of the trees  $T_a$  and  $T_b$ , then by our definitions



of  $\tilde{D}$  and  $\tilde{A}$  they are disjoint. We define the *change in configuration* of  $L_a$  and  $L_b$  as

$$\Delta(L_a, L_b) = \phi_{h_b}^{-1} \left( \tilde{D} \cup \tilde{A} \right). \quad (3.1)$$

Consider a finite sequence of trees  $T_1, T_2, \dots, T_\ell$ , where for each  $i$  ( $1 \leq i \leq \ell - 1$ ), and their corresponding complete configurations  $L_1, \dots, L_\ell$ . Let

$$L_i^\delta = \Delta(L_i, L_{i+1}). \quad (3.2)$$

Note that,  $L_i^\delta$  represents a configuration (not necessarily, a complete one) of the tree  $T_{i+1}$ . It is easy to see that given  $T_i$  along with  $L_i^\delta$  one can reconstruct  $L_{i+1}$ , i.e., the complete configuration of  $T_{i+1}$ .

We assume a scenario where we have an initial tree  $T_1$  with a certain complete configuration and over time the tree changes where new labeled nodes are added to the tree or the labels in its leaf node change. We are given the initial tree and the changes that take place in each step, i.e., we have access to  $T_1$  along with  $L_1^\delta, L_2^\delta, \dots, L_{\ell-1}^\delta$ .

A dynamic cover generation algorithm outputs a cover on input a configuration. A dynamic cover reconstruction algorithm when given a sequence of covers generates a configuration.

**Definition 3.1.3** (Dynamic Tree Cover Scheme). *A dynamic tree cover scheme  $\Psi_d = (\text{dCoverGen}, \text{dReConstruct})$  is a tuple of two algorithms defined as follows.*

- $C \leftarrow \text{dCoverGen}(h, L)$ : *On input the height of the tree  $h$  and a configuration  $L$ ,  $\text{dCoverGen}$  outputs a cover  $C$ .*
- $L \leftarrow \text{dReConstruct}(\{h_i, C_i\}_{i \in [\ell]})$ : *On input a sequence of tree height  $h_i$  and corresponding cover  $C_i$ ,  $\text{dReConstruct}$  outputs a configuration  $L$ .*

**Correctness:** Let  $L_1, L_2, \dots, L_\ell$  be as before, let  $\emptyset$  denote the empty configuration corresponding to an empty tree and let

$$\begin{aligned} L_1^\delta &= \Delta(\emptyset, L_1), \\ L_i^\delta &= \Delta(L_i, L_{i+1}), 2 \leq i \leq \ell - 1. \end{aligned}$$

We say,  $\Psi_d$  is correct if for all  $j \leq \ell$ ,

$$\Psi_d.\text{dReConstruct}(\{h_i, \text{dCoverGen}(h_i, L_i^\delta)\}_{i \in [j]}) = L_j.$$

## 3.2 Cover Generation Algorithms

For ease of exposition, we will sometimes impose colors on the nodes of the trees. For a node  $i$ ,  $\text{color}(i)$  will denote its color. For a node  $i$ ,  $\text{leftChild}(i)$  and  $\text{rightChild}(i)$  will denote its left and right child respectively. Recall, that configuration of a tree  $\mathcal{T}$  is a set of labeled leaf nodes of  $\mathcal{T}$ , thus elements of a configuration are ordered pairs  $(i, s)$  where  $i$  is a node and  $s \in \{+, -\}$  is its label, we will sometimes denote the label of  $i$  by  $\text{sign}(i)$ .

We call a configuration  $L$  a *pure configuration* if, for every  $(i, s) \in L$ ,  $s$  is the same, i.e., in a pure configuration, every node is either labeled  $+$  or  $-$ . A configuration which is not pure is called a *mixed configuration*.

### 3.2.1 Pure Cover: A Tree Cover Scheme for Pure Configurations

For the algorithms that follow, we assume that each tree node  $i$  is endowed with two fields  $\text{color}(i)$  and  $\text{sign}(i)$ .

We start with a simple scheme which generates covers only for pure configurations. Let  $L_p$  be a pure configuration of a tree  $\mathcal{T}$ , and we want to construct a cover of  $L_p$ . As a first step, we color the nodes of the tree according to the scheme described in the algorithm in Figure 3-2. We start from the leaf nodes and color each node which

is in the configuration with the color “GREEN”. We proceed with the non-leaf nodes starting from the level just above the leaf nodes and color a node GREEN if both of its children are colored GREEN.

<b>PureColoring</b> ( $h, L_p$ ): 01. Initialize a tree $\mathcal{T}$ with height $h$ , where the nodes of the tree has no color. 02. <b>for</b> $i \leftarrow 2^h - 1$ to $2^{h+1} - 2$ $\triangleright$ leaf nodes 03. <b>if</b> $(i, s) \in L_p$ 04. $\text{color}(i) \leftarrow \text{GREEN}, \text{sign}(i) \leftarrow s$ 05. <b>for</b> $i \leftarrow 2^h - 1$ to 0 of $\mathcal{T}$ $\triangleright$ non-leaf nodes 06. <b>if</b> $\text{color}(\text{leftChild}(i)) = \text{color}(\text{rightChild}(i)) = \text{GREEN}$ 07. $\text{color}(i) \leftarrow \text{GREEN}, \text{sign}(i) \leftarrow \text{sign}(\text{leftChild}(i))$ 08. <b>return</b> $\mathcal{T}$
---

Figure 3-2: The coloring scheme for pure configurations

If we apply the algorithm  $\text{PureColoring}(h, L_p)$  on tree  $\mathcal{T}$  with a pure configuration  $L_p$ , its nodes either get colored GREEN or they are without color.

**Definition 3.2.1** (Top Node). *A colored node in a tree  $\mathcal{T}$  is called a top node if the path from that node to the root does not contain any other colored node.*

We call the set of top nodes in a tree  $\mathcal{T}$  with a pure configuration  $L_p$  as the cover of the configuration.

With the above characterization of a cover of a pure configuration, we formulate an algorithm to construct one in Figure 3-3. Initially, the algorithm assigns an empty set  $X_i$  to each node  $i$  of the tree. For any leaf node  $j$ , if the leaf node is colored GREEN, then  $X_j$  is set to the singleton set  $\{(j, \text{sign}(j))\}$ , otherwise  $X_j$  remains the empty set. For any non-leaf node  $i$ , if the node is colored GREEN, then  $X_i$  is set to  $\{(i, \text{sign}(i))\}$ , else  $X_i$  is set as  $X_{\text{leftChild}(i)} \cup X_{\text{rightChild}(i)}$ . Finally, the algorithm returns  $X_0$ , i.e., the set associated with the root node.

**Proposition 3.2.1.**  *$\text{PureCoverGen}(h, L_p)$  returns the top nodes of a tree of height  $h$  with configuration  $L_p$ .*

*Proof.* According to our coloring scheme and the definition of a top node, the following is true for any top node  $i$ .

<pre> PureCoverGen(<math>h, L_p</math>): 01. <math>\mathcal{T} \leftarrow \text{PureColoring}(h, L_p)</math> 02. <b>for</b> <math>i \leftarrow 0</math> to <math>2^{h+1} - 2</math> 03.   <math>X_i \leftarrow \emptyset</math> 04. <b>for</b> <math>i = 2^h - 1</math> to <math>2^{h+1} - 2</math> <math>\triangleright</math> leaf nodes 05.   <b>if</b> <math>\text{color}(i) = \text{GREEN}</math> 06.     <math>X_i \leftarrow \{(i, \text{sign}(i))\}</math> 07. <b>for</b> <math>i \leftarrow 2^h - 2</math> to <math>0</math> <math>\triangleright</math> every non-leaf node <math>i</math> 08.   <b>if</b> <math>\text{color}(i) = \text{GREEN}</math> 09.     <math>X_i \leftarrow \{(i, \text{sign}(i))\}</math> 10.   <b>else</b> 11.     <math>X_i \leftarrow X_{\text{leftChild}(i)} \cup X_{\text{rightChild}(i)}</math> 12. <b>return</b> <math>X_0</math> </pre>
--

Figure 3-3: Pure cover generation

1. Both children of  $i$  are GREEN. Our coloring scheme guarantees this.
2. The sibling of  $i$  is not GREEN, as otherwise, our coloring scheme will make the immediate ancestor of  $i$  also GREEN, which violates the condition that  $i$  is a top node.
3. No ancestor of  $i$  is GREEN, as  $i$  is a top node.

Based on the above observations, it follows that if  $i$  is a top node, then  $X_i = \{(i, s)\}$  (see lines 6 and 9 of the Algorithm in Figure 3-3). Further, for any ancestor  $j$  of  $i$ ,  $X_j$  contains  $\{(i, s)\}$  (see line 11 of Figure 3-3). As the root (i.e., node 0) is also an ancestor of  $i$  hence  $X_0$  contains  $(i, s)$ . Thus, the set returned by  $\text{PureCoverGen}(h, L_p)$  contains all top nodes. Conversely, following the same arguments, it is easy to see that if  $(i, s) \in X_0$ , then  $i$  is a top node.  $\square$

We list some additional properties of covers generated by the cover generation algorithm  $\text{PureCoverGen}$  which are immediate from the algorithm.

**Proposition 3.2.2.** *Let  $L_p$  be a pure configuration of a tree  $\mathcal{T}$  of height  $h$ , and let  $C_p = \text{PureCoverGen}(h, L_p)$  then the following are true.*

1. If  $(i, s) \in C_p$  and  $i$  is a leaf node then  $(i, s) \in L_p$ .

2. If  $(i, s) \in C_p$  and  $i$  is not a leaf node then every leaf node of the subtree rooted at  $i$  occurs in  $L_p$  with sign  $s$ .
3. If  $i$  and  $i'$  be siblings in the tree  $\mathcal{T}$  then both  $(i, s)$ ,  $(i', s)$  cannot be in  $C_p$ .

<p><b>PureReConstruct</b>(<math>h, C_p</math>):</p> <ol style="list-style-type: none"> <li>01. Initialize a tree <math>\mathcal{T}</math> with height <math>h</math>, where the nodes of the tree has no color.</li> <li>02. Initialize <math>L_p \leftarrow \emptyset</math></li> <li>03. <b>for</b> every node <math>i = 0</math> to <math>2^h - 2</math>   ▷ non-leaf nodes</li> <li>04.   <b>if</b> <math>(i, s) \in C_p</math></li> <li>05.     <b>for</b> all <math>j \in \text{leaves}(\mathcal{T}(i))</math></li> <li>06.       <math>L_p \leftarrow L_p \cup \{(j, s)\}</math></li> <li>07. <b>return</b> <math>L_p</math></li> </ol>
--

Figure 3-4: Pure cover reconstruction

Now, given a pure cover  $C_p$ , corresponding to a pure configuration  $L_p$ , a reconstruction algorithm works as follows (see Figure 3-4). For every non-leaf node  $i$  in the cover, the reconstruction algorithm assigns the same sign to all the leaf nodes of the sub-tree rooted at node  $i$  and adds those leaf nodes along with their sign to the configuration  $L_p$ . And for every leaf node in the cover, it adds the node with its sign to the configuration.

The following proposition, which is easy to verify, asserts that the cover generation scheme is correct.

**Proposition 3.2.3.** *Let  $L_p$  be a pure configuration and  $C_p = \text{PureCoverGen}(h, L_p)$  and  $L'_p = \text{PureReConstruct}(h, C_p)$ , then  $L_p = L'_p$ .*

### 3.2.2 Mixed Cover: Generating a Smaller Sized Cover

In the previous section we discussed a scheme to generate covers assuming the configuration to be pure. In this section, we remove this restriction. Let  $L$  be a complete configuration for a tree  $\mathcal{T}$  of height  $h$ . Hence,  $L$  contains all leaf nodes of  $\mathcal{T}$  along with their sign, i.e.,

$$L = \{(i, \text{sign}(i)) : i \in \text{leaves}(\mathcal{T})\},$$

where  $\text{sign}(i) \in \{+, -\}$ . Such a complete configuration can be naturally decomposed into pure configurations  $L_g$  and  $L_r$  where,

$$L_g = \{(i, +) : (i, +) \in L\}, \quad L_r = \{(i, -) : (i, -) \in L\}.$$

As  $L_g$  and  $L_r$  are pure configurations, we can use our cover generation algorithm for computing the covers of the pure configurations  $L_g$  and  $L_r$  as  $C_g = \text{PureCoverGen}(h, L_g)$ ,  $C_r = \text{PureCoverGen}(h, L_r)$ . Note, that any one of  $C_g$  or  $C_r$  can be used as a cover of the complete configuration  $L$ . The reconstruction would be simple. We would reconstruct the pure configuration  $L_g$  (respectively  $L_r$ ) from  $C_g$  (respectively  $C_r$ ) and assign the opposite sign to all other leaf nodes of the tree. Thus, we can choose the one with smaller number of elements among  $C_g$  and  $C_r$  as the cover of  $L$ .

The above procedure would generate a correct cover for  $L$ , but we seek to find a more succinct cover. The above procedure yields a cover which contains nodes of the same sign, and hence we call such a cover as a *pure cover*. The procedure that we are about to describe will contain nodes with both signs and hence we name this algorithm a *mixed cover* algorithm.

MixedColoring( $h, L_c$ ):
01. Initialize a tree $\mathcal{T}$ with height $h$ , where the nodes of the tree has no color.
02. <b>for</b> $i \leftarrow 2^h - 1$ to $2^{h+1} - 2$
03. <b>if</b> $(i, +) \in L_c$
04. $\text{color}(i) \leftarrow \text{GREEN}$
05. <b>else</b>
06. $\text{color}(i) \leftarrow \text{RED}$
07. <b>for</b> $i \leftarrow 2^h - 2$ to 0
08. <b>if</b> $\text{color}(\text{leftChild}(i)) = \text{color}(\text{rightChild}(i)) = \text{GREEN}$
09. $\text{color}(i) \leftarrow \text{GREEN}$
10. <b>if</b> $\text{color}(\text{leftChild}(i)) = \text{color}(\text{rightChild}(i)) = \text{RED}$
11. $\text{color}(i) \leftarrow \text{RED}$
12. <b>return</b> $\mathcal{T}$

Figure 3-5: Coloring scheme for mixed covers

As before, the heart of the algorithm is a coloring scheme  $\text{MixedColoring}(h, L_c)$  which is described in Figure 3-5. The algorithm takes in a complete configuration

( $L_c$ ) for a tree  $\mathcal{T}$  of height  $h$ , and colors the nodes of  $\mathcal{T}$  as GREEN or RED. The algorithm examines the nodes level wise starting from the lowest level, i.e., the leaf nodes. All leaf nodes with sign  $+$  are colored GREEN and leaf nodes with sign  $-$  are colored RED. A non-leaf node gets the color GREEN if both its children are colored GREEN and gets the color RED if both its children are RED. Other nodes remain uncolored.

We discuss the main idea of our cover generation scheme next. It may be helpful to consider an example shown in Figure 3-6 all along. The number of leaf nodes in the tree  $n = 16$ . Consider that the nodes labeled  $+$  are  $L_g = \{15, 16, 18, 19, 20, 23\}$ , and the nodes labeled  $-$  are  $L_r = \{17, 21, 22, 24, 25, 26, 27, 28, 29, 30\}$ . According to our coloring algorithm of Figure 3-5, the nodes  $\{15, 16, 18, 19, 20, 23\}$  will receive the color GREEN (denoted by shaded circles in the Figure 3-6) and the nodes  $\{17, 21, 22, 24, 25, 26, 27, 28, 29, 30\}$  will be colored RED (shown by shaded squares in the figure).

Let  $\mathcal{T}$  be a tree of height  $h$  with a complete configuration  $L_c$  and colored using  $\text{MixedColoring}(h, L_c)$ . Recall that a node  $i$  is a top node of the colored tree  $\mathcal{T}$  if  $i$  is colored and there is no colored node in the path from  $i$  to the root. In our example in Figure 3-6 the nodes 7, 17, 18, 9, 10, 23, 24, 12, and 6 are top nodes. Note that, all leaf nodes of the subtree rooted at a top node  $i$  have the same color as that of  $i$ .

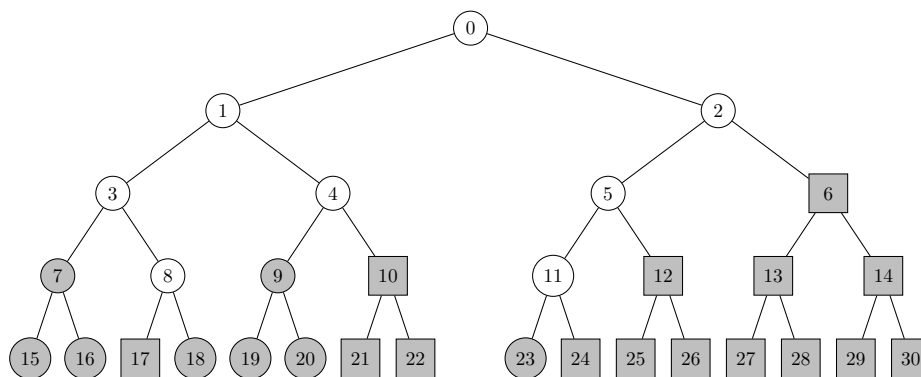


Figure 3-6: An example of mixed cover: The leaf nodes represented by circles are assumed to bear the sign  $+$ , the leaf nodes represented by squares bear the sign  $-$ . The shaded circles represent nodes colored green and the shaded squares are nodes colored red.

The main intuition of our mixed cover generation algorithm is the following. Let  $i$  be a top node for a colored tree  $\mathcal{T}$ . Hence, the color of  $i$  is enough to determine the color of all the leaf nodes of the subtree rooted at  $i$ . In fact, in many cases, a top node  $i$  along with its color can be used to uniquely color the leaf nodes of a tree much bigger than the subtree rooted at  $i$ . For example, in Figure 3-6, node 10 can be used to reconstruct the tree rooted at node 4: we will assign the sign of 10 to all of its leaf nodes and the opposite sign to all the other leaf nodes of the subtree rooted at node 4. We want to make this intuition more concrete. We introduce some definitions for this purpose.

**Definition 3.2.2** (Representable Top Nodes). *Let  $\mathcal{T}$  be a tree of height  $h$  with complete configuration  $L_c$  and colored by  $\text{MixedColoring}(h, L_c)$ . A set of top nodes  $S$  of  $\mathcal{T}$  is called representable, if the following holds*

1. *All nodes in  $S$  are of same color  $c \in \{\text{GREEN}, \text{RED}\}$ .*
2. *There exists a subtree  $\mathcal{T}'$  (of  $\mathcal{T}$ ) containing all nodes in  $S$ , such that all leaf nodes of  $\mathcal{T}'$  other than the leaf nodes of the trees rooted at nodes in  $S$  have a color different from  $c$ .*

*If  $S$  is representable then any  $\mathcal{T}'$  satisfying the property (2) above is said to represent  $S$ .*

**Definition 3.2.3** (Representative Tree). *Let  $\mathcal{T}$  be a tree of height  $h$  with complete configuration  $L_c$  and colored by  $\text{MixedColoring}(h, L_c)$ .  $S$  be a representable set of top nodes of  $\mathcal{T}$ . The largest subtree of  $\mathcal{T}$  which represents  $S$  is called the representative tree of  $S$ . By  $\text{rep}_{\mathcal{T}}(S)$  we will denote the root of the representative tree of  $S$ . If  $S$  is a singleton set containing only  $i$ , we will denote the root of their representative tree of  $S$  by  $\text{rep}_{\mathcal{T}}(i)$ .*

From the definition, it is immediate that every representable set of top nodes will have a representative tree, moreover every singleton set containing a top node is representable and thus has a representative tree.



In our example tree  $\mathcal{T}$  in Figure 3-6 we have the set  $\{17, 10\}$  is representable and  $\text{rep}_{\mathcal{T}}(\{17, 10\}) = 1$ , whereas  $\{17, 24\}$  is not representable. Moreover, we have  $\text{rep}_{\mathcal{T}}(7) = 7$ ,  $\text{rep}_{\mathcal{T}}(17) = 3$ ,  $\text{rep}_{\mathcal{T}}(9) = \text{rep}_{\mathcal{T}}(10) = 4$ ,  $\text{rep}_{\mathcal{T}}(23) = 2$ ,  $\text{rep}_{\mathcal{T}}(12) = 12$  and  $\text{rep}_{\mathcal{T}}(6) = 6$ .

**Definition 3.2.4** (Independent Nodes). *Let  $S_1, S_2$  be two sets of representable top nodes of a colored tree  $\mathcal{T}$  and let  $\text{rep}_{\mathcal{T}}(S_1) = i_1$  and  $\text{rep}_{\mathcal{T}}(S_2) = i_2$ . We say  $S_1, S_2$  are independent if  $\mathcal{T}(i_1)$  and  $\mathcal{T}(i_2)$  are disjoint.*

In our example, the sets  $\{7\}, \{17\}$  are not independent whereas  $\{17\}, \{10\}$  and  $\{23\}, \{17\}$  are independent.

**Definition 3.2.5** (Span). *Let  $S$  be a set of top nodes of a tree  $\mathcal{T}$  colored by MixedColoring  $(h, L)$ . We say that  $S$  spans  $\mathcal{T}$ , if  $S$  can be decomposed into  $S = S_1 \cup S_2 \cup \dots \cup S_k$ , for some  $k \leq |S|$  such that the following holds*

1. For all  $i, j \in [k]$ ,  $S_i \cap S_j = \emptyset$ .
2. For all  $i \in [k]$ ,  $S_i$  is representable.
3. For all  $i, j \in [k]$ ,  $i \neq j$ ,  $S_i$  and  $S_j$  are independent.
4. The union of the leaf nodes of the representative trees of the sets  $S_1, \dots, S_k$ , gives the leaves of the tree  $\mathcal{T}$ , i.e.,

$$\bigcup_{i \in [k]} \text{leaves}(\mathcal{T}(\text{rep}_{\mathcal{T}}(S_i))) = \text{leaves}(\mathcal{T}).$$

Observe that if  $X$  spans a colored tree  $\mathcal{T}$ , then the nodes in  $X$  along with their colors contain enough information to reconstruct the complete configuration of  $\mathcal{T}$ . Thus, for a given colored tree our goal is to find a set  $X$  of top nodes of minimum size which spans the tree. For the reconstruction, it is necessary to just decompose the span into representable sets and determine the representative tree for each such representable set.

Consider any node  $i$  of a colored tree  $\mathcal{T}$ . If  $i$  is colored and is of color  $c \in \{\text{RED}, \text{GREEN}\}$ , then  $i$  along with its color/sign spans the tree  $\mathcal{T}(i)$ . If  $i$  is un-colored, then the following statements hold and are easy to verify.

**Proposition 3.2.4.** *Let  $i$  be any un-colored node in a colored tree  $\mathcal{T}$  and  $\rho$  and  $\lambda$  be the right child and left child of  $i$ , respectively. Let  $X_\rho$  and  $X_\lambda$  span  $\mathcal{T}(\rho)$  and  $\mathcal{T}(\lambda)$  respectively. Then either both  $X_\rho$  and  $X_\lambda$  individually spans  $\mathcal{T}(i)$  or  $X_\lambda \cup X_\rho$  spans  $\mathcal{T}(i)$ .*

**Proposition 3.2.5.** *Let  $i$  be any un-colored node in a colored tree  $\mathcal{T}$  and  $\rho$  and  $\lambda$  be the right child and left child of  $i$ , respectively. Let  $X$  spans  $\mathcal{T}(i)$  and let  $X = X_\lambda \cup X_\rho$ , where  $X_\lambda \subseteq \text{nodes}(\mathcal{T}(\lambda))$  and  $X_\rho \subseteq \text{nodes}(\mathcal{T}(\rho))$ . For  $j \in \{\lambda, \rho\}$ , if  $X_j \neq \emptyset$ , then  $X_j$  spans  $\mathcal{T}(j)$ . Moreover, if  $X_\lambda = \emptyset$  then  $X_\rho$  contains nodes of the same color and all leaves of  $\mathcal{T}(\lambda)$  are of color different from that of the nodes of  $X_\rho$ .*

The above observation is central to our algorithm `MixedCoverGen`( $h, L_c$ ) for generating covers described in Figure 3-7. The algorithm finds a cover for a tree  $\mathcal{T}$  of height  $h$  with a complete configuration  $L_c$ . The algorithm assigns to each node  $i$ , three sets namely  $X_{i,g}$ ,  $X_{i,r}$ ,  $X_{i,m}$ .  $X_{i,g}$  and  $X_{i,r}$  contains the GREEN and RED top nodes of the subtree rooted at  $i$ .  $X_{i,m}$  is empty if  $i$  is a leaf node, otherwise we set

$$Y \leftarrow \text{MixedUnion}(i).$$

The procedure `MixedUnion`( $i$ ), when  $i$  is not a leaf node, is described in the right column of Figure 3-7. `MixedUnion`( $i$ ) computes the following collection of sets

$$\mathcal{C} = \left\{ X_{\text{leftChild}(i),x} \cup X_{\text{rightChild}(i),y} : x, y \in \{r, g, m\}; \right. \\ \left. x \neq y; X_{\text{leftChild}(i),x}, X_{\text{rightChild}(i),y} \neq \emptyset \right\},$$

and returns the set  $Y$  of minimum cardinality from  $\mathcal{C}$ . If  $|Y|$  is smaller than both  $|X_{i,g}|$  and  $|X_{i,r}|$ , then  $X_{i,m}$  is set to  $Y$  otherwise it remains empty. Finally, the algorithm outputs  $C = \text{minCard}(X_{0,r}, X_{0,g}, X_{0,m})$ . The function `minCard`( $X_1, \dots, X_n$ ), for finite

sets  $X_1, X_2, \dots, X_n$  gives a nonempty set of minimum cardinality among the sets  $X_1, X_2, \dots, X_n$ .

**Theorem 3.2.1.** *Let  $\mathcal{T}$  be a tree of height  $h$ , and  $L_c$  be a complete configuration of  $\mathcal{T}$ . Let  $\text{MixedCoverGen}(h, L_c)$  compute the sets  $X_{i,g}, X_{i,r}, X_{i,m}$  for each node  $i$  of  $\mathcal{T}$  as described in Figure 3-7. Let  $C_i = \{X_{i,g}, X_{i,r}, X_{i,m}\}$ . Then the following are true.*

1.  $C_i$  contains at least one non-empty set.
2. Any non-empty set  $X \in C_i$  spans  $\mathcal{T}(i)$ .
3.  $\text{minCard}(C_i)$  is the smallest set of nodes which spans  $\mathcal{T}(i)$ .

*Proof.* The proof of (1) is immediate, as at least one of  $X_{i,r}$  or  $X_{i,g}$  must be non-empty as  $L_c$  is a complete configuration of  $\mathcal{T}$  and hence all the leaves of  $\mathcal{T}$  are colored.

(2) Directly follows from the description of the algorithm and Proposition 3.2.4.

We prove (3) by induction on the height of a node. Let  $i$  be of height 0, i.e., a leaf node and without loss of generality let  $i$  be colored GREEN, thus  $X_{i,g} = \{(i, +)\}$  and  $X_{i,r} = X_{i,m} = \emptyset$ , and  $X_{i,g}$  is the smallest set which spans  $\mathcal{T}(i)$  as  $|X_{i,g}| = 1$ . This serves as the base case. As induction hypothesis consider that (3) is true for all nodes at height less than  $\ell$ . Consider, a node  $i$  at height  $\ell$ . As  $\ell > 0$ ,  $i$  have two children say  $\lambda$  and  $\rho$ . By our induction hypothesis,  $X_\lambda = \text{minCard}(X_{\lambda,r}, X_{\lambda,g}, X_{\lambda,m})$  and  $X_\rho = \text{minCard}(X_{\rho,r}, X_{\rho,g}, X_{\rho,m})$  are the smallest sets which spans  $\mathcal{T}(\lambda)$  and  $\mathcal{T}(\rho)$  respectively. Now we have a few cases to consider:

**Case 1:**  $i$  is colored. Without loss of generality let  $\text{color}(i) = \text{GREEN}$ , then  $X_{i,g} = \{(i, +)\}$  and  $X_{i,r}$  and  $X_{i,m}$  are empty. Moreover, as  $X_{i,g} = \{(i, +)\}$  contains a single node which is the root of  $\mathcal{T}(i)$ , thus  $X_{i,g}$  is the smallest set which spans  $\mathcal{T}(i)$ .

**Case 2:**  $i$  is not colored. Let  $X = \text{minCard}(C_i)$ , and for the sake of contradiction let  $Y$  span  $\mathcal{T}(i)$  and let  $|Y| < |X|$ . Let  $Y = Y_\lambda \cup Y_\rho$  where  $Y_j$  contains nodes in  $\mathcal{T}(j)$  for  $j \in \{\lambda, \rho\}$ . We consider two subcases:

- (a) Both  $Y_\lambda$  and  $Y_\rho$  are non-empty. By Proposition 3.2.5,  $Y_\lambda$  and  $Y_\rho$  spans  $\mathcal{T}(\lambda)$  and  $\mathcal{T}(\rho)$  respectively. Now, we claim that either  $|Y_\lambda| < |X_\lambda|$  or  $|Y_\rho| < |X_\rho|$ , and this contradicts our induction hypothesis. Finally, to see why our claim is correct, notice that by our algorithm  $X_\lambda \cup X_\rho \in C_i$ , and as  $|Y| < |X|$  and  $X = \text{minCard}(C_i)$  it cannot be that  $|Y_\lambda| \geq |X_\lambda|$  and  $|Y_\rho| \geq |X_\rho|$ .
- (b) One of  $Y_\lambda$  and  $Y_\rho$  is empty. Without loss of generality, let  $Y_\rho = \emptyset$ , i.e.,  $Y = Y_\lambda$ . Then by Proposition 3.2.5,  $Y_\lambda$  must contain nodes of the same color, say GREEN, and all leaves of  $\mathcal{T}(\rho)$  must be RED. Hence, node  $\rho$  must be RED and all nodes in  $\text{leaves}(\mathcal{T}(\lambda)) \setminus Y_\lambda$  must also be RED. In this configuration, the set  $X_{i,g}$  computed by `MixedCoverGen`, will contain the same nodes as in  $Y_\lambda$ . Hence, we have  $Y = Y_\lambda \in C_i$ , which contradicts  $|Y| < |X|$ , as  $X = \text{minCard}(C_i)$ .

□

MixedCoverGen( $h, L_c$ ):	MixedUnion( $i$ ):
01. $\mathcal{T} \leftarrow \text{MixedColoring}(h, L_c)$	01. $\text{cnt} \leftarrow 0$
02. <b>for</b> every node $i$ of tree $\mathcal{T}$	02. <b>for</b> $x \in \{r, g, m\}$
03. $X_{i,g}, X_{i,r}, X_{i,m} \leftarrow \emptyset$	03. <b>for</b> $y \in \{r, g, m\}$
04. <b>for</b> $i \leftarrow 2^h - 1$ to 0	04. $\lambda \leftarrow \text{leftChild}(i)$
05. <b>if</b> $\text{color}(i) = \text{GREEN}$	05. $\rho \leftarrow \text{rightChild}(i)$
06. $X_{i,g} \leftarrow X_{i,g} \cup \{(i, +)\}$	06. <b>if</b> $(x \neq y) \wedge X_{\lambda,x} \neq \emptyset \wedge X_{\rho,y} \neq \emptyset$
07. <b>else-if</b> $\text{color}(i) = \text{RED}$	07. $Y_{\text{cnt}} \leftarrow X_{\lambda,x} \cup X_{\rho,y}$
08. $X_{i,r} \leftarrow X_{i,r} \cup \{(i, -)\}$	08. $\text{cnt} \leftarrow \text{cnt} + 1$
09. <b>else</b>	09. $X \leftarrow \text{minCard}(Y_0, Y_1, \dots, Y_{\text{cnt}})$
10. $X_{i,g} \leftarrow X_{\text{leftChild}(i),g} \cup X_{\text{rightChild}(i),g}$	10. <b>return</b> $X$
11. $X_{i,r} \leftarrow X_{\text{leftChild}(i),r} \cup X_{\text{rightChild}(i),r}$	
12. $Y \leftarrow \text{MixedUnion}(i)$	
13. <b>if</b> $ Y  < \min\{ X_{i,g} ,  X_{i,r} \}$	
14. $X_{i,m} \leftarrow Y$	
15. <b>return</b> $\text{minCard}(X_{0,g}, X_{0,r}, X_{0,m})$	

Figure 3-7: Mixed cover generation algorithm

Theorem 3.2.1 asserts that `MixedCoverGen`( $h, L_c$ ) outputs the smallest possible set  $X$  which spans a tree of height  $h$  with a complete configuration  $L_c$ . Now our goal is

to regenerate the configuration of a tree of height  $h$ , given  $h$  and a set  $C$  which spans the tree and was produced by  $\text{MixedCoverGen}(h, L_c)$ . The reconstruction algorithm is presented in Figure 3-8.

First, note that it may be the case that the algorithm  $\text{MixedCoverGen}(h, L_c)$  generates a pure cover, i.e., it generates  $C$  where all nodes in  $C$  are labeled with the same sign. In such a case either  $C = X_{0,r}$  or  $C = X_{0,g}$ , i.e.,  $C$  contains all top nodes of a single color. In this case, the reconstruction is simple, as the representative tree for  $C$  is the complete tree, i.e.,  $\text{rep}_{\mathcal{T}}(C) = 0$ . For reconstruction, the following procedure would suffice: For every  $(i, \text{sign}(i)) \in C$ , the leaves of the sub-tree  $\mathcal{T}(i)$  are labeled with  $\text{sign}(i)$  and the rest of the leaf nodes are labeled with the opposite sign. These steps are done in lines 5 to 8 of the algorithm described in Figure 3-8.

If  $C$  is a mixed cover, i.e.,  $C$  contains nodes of both signs then  $C = X_{0,m}$ . In this case, the reconstruction procedure is a bit more involved. For convenience, we introduce some additional notations. For any node  $i$  in a tree  $\mathcal{T}$ , let  $\text{path}(i)$  denote the sequence of nodes in the unique path from  $i$  to the root of  $\mathcal{T}$ . For any node  $j \neq i$  of  $\text{path}(i)$ ,  $\text{prev}_i(j)$  denotes the node preceding  $j$  in the sequence  $\text{path}(i)$ .

Reconstructing the configuration, given a mixed cover  $C$  boils down to decomposing  $C$  into disjoint subsets such that each subset is representable. Once such representable subsets are obtained the leaves of the corresponding representative trees can be assigned colors (signs) and this assignment of signs would yield the desired complete configuration.

Let  $C$  be the output of  $\text{MixedCoverGen}(h, L_c)$  and let  $C$  be a mixed cover. Let  $(j, s) \in C$ , where  $s \in \{+, -\}$  and let  $\bar{s}$  be the sign opposite to  $s$ . Let,  $\ell_j$  be the first node in  $\text{path}(j)$  which intersects with  $\text{path}(i)$  for some node  $i$  such that  $(i, \bar{s}) \in C$ . As  $C$  is a mixed cover, hence  $C$  contains at least two elements with different signs and thus for every  $(j, s) \in C$ ,  $\ell_j$  is well-defined.

Let  $\lambda$  and  $\rho$  be the left and right children of  $\ell_j$ . Then, if  $i$  belongs to  $\mathcal{T}(\lambda)$  then  $j$  belongs to  $\mathcal{T}(\rho)$  and vice versa. Without loss of generality, let  $i$  belong to  $\mathcal{T}(\lambda)$  also

let  $J$  be those nodes in  $\mathcal{T}(\rho)$  which belongs to  $C$ , i.e.,

$$J = \{(k, \text{sign}(k)) : (k, \text{sign}(k)) \in C, k \in \mathcal{T}(\rho)\}. \quad (3.3)$$

It is easy to see that all nodes in  $J$  bears the same sign as that of  $j$ , as otherwise  $\ell_j$  cannot be the first node in  $\text{path}(j)$  which intersects with the path of another node in  $C$  with a sign different from the sign of  $j$ . We observe the following

**Proposition 3.2.6.** *The set  $J$  as defined in Eq. (3.3) is representable and  $\rho = \text{rep}_{\mathcal{T}}(J)$ .*

*Proof.* As all nodes in  $J$  are top nodes and bear the same sign, then in the corresponding colored tree  $\mathcal{T}$  all nodes in  $J$  have the same color, say GREEN. With reference to the algorithm  $\text{MixedCoverGen}((h, L_c))$ , it is immediate that  $J \subseteq X_{\rho,g}$ . We claim that  $J = X_{\rho,g}$ , i.e.,  $J$  contains all GREEN top nodes in  $\mathcal{T}(\rho)$ . As  $C$  is a cover of  $\mathcal{T}$  and  $J \subseteq C$ , by Theorem 3.2.1,  $C$  spans  $\mathcal{T}$ , and the nodes in  $J$  are the only GREEN top nodes in  $\mathcal{T}(\rho)$  which are in  $C$ . If there are GREEN top nodes in  $\mathcal{T}(\rho)$  which are not in  $J$  then  $C$  cannot span  $\mathcal{T}$ . Thus,  $J$  is representable and  $\mathcal{T}(\rho)$  represents  $J$ .

We are left to show that  $\mathcal{T}(\rho)$  is the largest tree that represents  $J$ , i.e.,  $\text{rep}_{\mathcal{T}}(J) = \rho$ . The smallest tree larger than  $\mathcal{T}(\rho)$  which contains  $\mathcal{T}(\rho)$  is  $\mathcal{T}(\ell_j)$ . We will argue that  $\ell_j \neq \text{rep}_{\mathcal{T}}(J)$ . Note that,  $i \in C$  and  $i$  belongs to the left subtree  $\mathcal{T}(\lambda)$  of  $\mathcal{T}(\ell_j)$  and  $i$  has color different from  $j$ , i.e.,  $i$  is colored RED. If  $\ell_j = \text{rep}_{\mathcal{T}}(J)$ , then a set containing  $J$  and another set containing  $i$  cannot be independent, which implies that both nodes in  $J$  and the node  $i$  cannot be in  $C$  which by Theorem 3.2.1 spans  $\mathcal{T}$ .  $\square$

The above proposition is central to the reconstruction algorithm presented in Figure 3-8 when  $C$  is a mixed cover. Lines 10-14 of the algorithm presented in Figure 3-8 finds the node  $\ell_j$  (as described above) for each  $j$  in  $C$  and thus decomposes  $C$  into representative sets  $J$  as asserted in Proposition 3.2.6. This process is iterated until all leaves of the tree are labeled.

<p><b>MixedReConstruct</b>(<math>h, C</math>):</p> 01. Initialize a tree $\mathcal{T}$ with height $h$ , where the nodes of the tree has no color. 02. $L \leftarrow \emptyset$ 03. <b>for</b> all $(i, s) \in C$ 04. <b>for</b> all $k \in \text{leaves}(\mathcal{T}(i))$ 05. $L \leftarrow L \cup \{(k, s)\}$ 06. <b>if</b> $C$ is a pure cover with sign $s \in \{+, -\}$ 07. <b>for</b> all $k' \in \text{leaves}(\mathcal{T})$ 08. <b>if</b> $(k', s) \notin L$ 09. $L \leftarrow L \cup \{(k', \bar{s})\}$ 10. <b>else</b> 11. <b>for</b> each $(i, s) \in C$ 12.             find the first ancestor node $j$ of $i$ in $\text{path}(i)$ such that $j \in \text{path}(l)$ for some $(l, \bar{s}) \in C$ 13. <b>for</b> all the leaf nodes $k$ of the sub-tree rooted at the node $\text{prev}_i(j)$ 14. <b>if</b> $(k, s) \notin L$ 15. $L \leftarrow L \cup \{(k, \bar{s})\}$ 16. <b>return</b> $L$
--

Figure 3-8: Mixed cover reconstruction

### 3.2.3 Complexity of Proposed Algorithms

In the `MixedCoverGen` algorithm (Figure 3-7), the loop in line 4 iterates  $N = 2^{h+1} - 1$  times, i.e., for all nodes in the tree. For each iteration, the algorithm computes a minimum-sized cover of the node using the `MixedUnion` algorithm. This algorithm operates on a single node and determines the minimum-sized mixed cover by taking the union of all possible covers from the left and right children of the target node. Thus, the overall complexity of the mixed cover generation algorithm is determined by the complexity of the `MixedUnion` algorithm.

The `MixedUnion` algorithm has a constant-time computational complexity and requires linear space. To achieve this, the algorithm maintains the RED, GREEN, and MIXED covers of each node using linked lists. Each linked list is structured such that its head contains the length of the list and an additional pointer to the tail. This setup allows us to efficiently compute the cover with the minimum cardinality using the length information stored at the head, and a union operation can be performed by manipulating the head and tail pointers of the linked lists being merged. Thus, the `MixedUnion` operations have constant time complexity. In addition, the size of the linked lists is linear with respect to the number of tree nodes. The algorithm processes

the tree level-wise, and at each level, the total size of all linked lists is linear in the number of tree nodes. Covers for all nodes below the current level can be discarded. Therefore, the time and space complexity of the proposed algorithm `MixedCoverGen` is  $\mathcal{O}(N)$ .

The `MixedReConstruct` algorithm of Figure 3-8 runs for at most the number of elements in the input cover  $C$ . The size of the input cover is at most half the number of leaf nodes, i.e.,  $n/2 = 2^{h-1}$ . For each node in the cover, the algorithm examines the path from that node to the root. It searches for the first intersection of this path with another path from a different node in the cover with opposite color. This first intersection determines the representative tree for the target node (the details are described in the algorithm `MixedReConstruct` of Figure 3-8).

This process can be optimized by pre-computing all paths from the elements in the cover, which requires at most  $\mathcal{O}(h)$  time per element. As a result, the overall time complexity of the reconstruction algorithm is  $\mathcal{O}(C \cdot h)$ . Lastly, the reconstruction algorithm has a space complexity  $\mathcal{O}(N)$ , since the entire process can be performed directly on the tree structure.

### 3.2.4 Dynamic Tree Cover Scheme

Recall the setting of a dynamic tree cover scheme as discussed in Section 3.1.4. We have a sequence of trees  $T_1, T_2, \dots, T_\ell$ , with their corresponding complete configurations  $L_1, L_2, \dots, L_\ell$ , and for  $1 \leq i \leq \ell - 1$ ,  $L_i^\delta$  is defined as in Equations (3.1) and (3.2). Note each  $L_i^\delta$  represents a configuration, not necessarily a complete one, of a tree of height  $h_i$ , where  $h_i$  is the height of the shortest complete binary tree which contains all nodes in  $L_i^\delta$ . The dynamic cover generation algorithm takes  $L_i^\delta$  and generates a cover  $C_i$ . The property which is required is that, knowing the complete configuration of  $T_1$  and the sequence of covers  $C_2, C_3, \dots, C_\ell$  corresponding to the configurations  $L_2^\delta, L_3^\delta, \dots, L_{\ell-1}^\delta$ , the cover reconstruction algorithm can generate the complete configuration of  $T_n$ .

The dynamic cover generation algorithm just takes a configuration and produces its cover. If the input configuration  $L$  is a complete configuration then the mixed



cover generation algorithm is used to generate the cover. Otherwise,  $L$  is decomposed into two sets  $L_g$  and  $L_r$ , where  $L_g$  contains the nodes labeled  $+$  and  $L_r$  contains the nodes labeled  $-$ , and the pure cover generation algorithm is used to generate covers  $C_g$  and  $C_r$  for the configurations  $L_g$  and  $L_r$  respectively. Finally,  $C = C_g \cup C_r$  is produced as the output. The details are shown in Figure 3-9.

<p><b>dCoverGen</b>(<math>h, L</math>):</p> <ol style="list-style-type: none"> <li>01. Initialize an empty tree <math>\mathcal{T}</math> of height <math>h</math></li> <li>02. <b>if</b> <math> L  = 2^h</math></li> <li>03.   <math>C \leftarrow \text{MixedCoverGen}(h, L)</math></li> <li>04. <b>else</b></li> <li>05.   <math>C, L_g, L_r \leftarrow \emptyset</math></li> <li>06.   <b>for</b> all <math>(i, +) \in L</math></li> <li>07.     <math>L_g \leftarrow L_g \cup \{(i, +)\}</math></li> <li>08.   <math>C_g \leftarrow \text{PureCoverGen}(h, L_g)</math></li> <li>09.   <b>for</b> all <math>(i, -) \in L</math></li> <li>10.     <math>L_r \leftarrow L_r \cup \{(i, -)\}</math></li> <li>11.   <math>C_r \leftarrow \text{PureCoverGen}(h, L_r)</math></li> <li>12.   <math>C \leftarrow C_g \cup C_r</math></li> <li>13. <b>return</b> <math>C</math></li> </ol>
--

Figure 3-9: Dynamic cover generation

To explain the reconstruction procedure we introduce a new operation on configurations. Let  $L_\alpha$  and  $L_\beta$  be two configurations for trees with heights  $h_\alpha$  and  $h_\beta$ . Let  $h_{\max} = \max\{h_\alpha, h_\beta\}$ ,  $\tilde{L}_\alpha = \phi_{h_\alpha}(L_\alpha)$  and  $\tilde{L}_\beta = \phi_{h_\beta}(L_\beta)$ . We define a new configuration for a tree of height  $h_{\max}$  as

$$L_\alpha \triangleright_{(h_\alpha, h_\beta)} L_\beta = \phi_{h_{\max}}^{-1}(X \cup Y \cup Z), \quad (3.4)$$

where  $X, Y, Z$  are defined as follows:

$$\begin{aligned} X &= \{(i, s) \in \tilde{L}_\alpha : (i, s) \notin \tilde{L}_\beta \text{ and } (i, \bar{s}) \notin \tilde{L}_\beta\} \\ Y &= \{(i, s) \in \tilde{L}_\beta : (i, s) \notin \tilde{L}_\alpha, (i, \bar{s}) \notin \tilde{L}_\alpha\} \\ Z &= \{(i, s) \in \tilde{L}_\beta : (i, \bar{s}) \in \tilde{L}_\alpha\}. \end{aligned}$$

$L_\alpha \triangleright_{(h_\alpha, h_\beta)} L_\beta$  is essentially the configuration obtained by overwriting  $L_\alpha$  by  $L_\beta$ .

Notice that the set  $X$  contains those labeled nodes in  $L_\alpha$  which are not present in  $L_\beta$ . Similarly,  $Y$  contains those nodes in  $L_\beta$  which are not present in  $L_\alpha$  and  $Z$  contains those nodes in  $L_\beta$  which are present in  $L_\alpha$  but with the opposite label.

For reconstruction, the sequence of covers, along with the corresponding tree heights, is used. Suppose  $C_1$  be the cover of  $L_1$  and for  $2 \leq i \leq \ell - 1$ ,  $C_i$  be the cover of  $L_i^\delta$ . The cover reconstruction algorithm first reconstructs the configuration  $L_i$  for each cover  $C_i$  and then outputs the configuration

$$L = (((L_2 \triangleright L_2) \triangleright L_3) \triangleright \cdots) \triangleright L_{\ell-1}.$$

The details are in Figure 3-10.

dCoverReConstruct( $\{(h_i, C_i)\}_{i \in [0, \ell-1]}$ ):
01. $L_0 \leftarrow \text{MixedReConstruct}(h_0, C_0)$
02. <b>for</b> $i = 1$ to $\ell - 1$
03.   Initialize $X_i, C_p, C_r \leftarrow \emptyset$
04. <b>for</b> all $(j, +) \in C_i$
05. $C_p \leftarrow C_p \cup \{(j, +)\}$
06. $L_p \leftarrow \text{PureReConstruct}(h_i, C_p)$
07. <b>for</b> all $(j, -) \in C_i$
08. $C_r \leftarrow C_r \cup \{(j, -)\}$
09. $L_r \leftarrow \text{PureReConstruct}(h_i, C_r)$
10. $X_i \leftarrow L_p \cup L_r$
11. $h \leftarrow h_0, L \leftarrow L_0$
12. <b>for</b> $i = 1$ to $\ell - 1$ ,
13. $L \leftarrow L \triangleright_{(h, h_i)} X_i$
14. $h \leftarrow \max\{h, h_i\}$
15. <b>return</b> $L$

Figure 3-10: Dynamic cover reconstruction

### 3.3 Expected Cover Size of a Pure Cover

In this section, we provide an analysis of the expected size of a cover returned by the Algorithm in Figure 3-3. Let  $\mathcal{T}$  be a full binary tree of height  $h$  and thus  $\mathcal{T}$  has a total of  $n = 2^h$  leaf nodes. Let  $L_p$  be a pure configuration of  $\mathcal{T}$  such that  $|L_p| = r$ , and without loss of generality, we assume that all nodes in  $L_p$  bear the sign  $+$ . There

are  $\binom{n}{r}$  such configurations possible, and we are interested in the average (expected value of) cover size over all these  $\binom{n}{r}$  configurations.

Consider a sequence of  $2n - 1$  binary random variables  $P_0, P_1, \dots, P_{2n-2}$  corresponding to each node  $i$  of  $\mathcal{T}$ . Define

$$P_i = \begin{cases} 1; & \text{if } \text{color}(i) = \text{GREEN} \\ 0; & \text{otherwise.} \end{cases}$$

That is, according to our coloring scheme,  $P_i = 1$  denotes if the node  $i$  is colored GREEN or not. Consider any node  $i$  at the  $\ell^{\text{th}}$  level of the tree  $\mathcal{T}$ , i.e.,  $i \in \{2^\ell - 1, \dots, 2^{\ell+1} - 2\}$ . Then, the subtree rooted at node  $i$  contains  $2^{h-\ell}$  many leaf nodes. The event “ $\{P_i = 1\}$ ” can then be viewed as choosing  $2^{h-\ell}$  many BLACK balls from a bag containing a total of  $n$  balls, where  $r$  many BLACK balls and remaining  $n - r$  are WHITE balls. Therefore,

$$\Pr[P_i = 1] = \begin{cases} \frac{\binom{r}{2^{h-\ell}}}{\binom{n}{2^{h-\ell}}} = \eta_{2^{h-\ell}}(n, r); & \text{if } 2^{h-\ell} \leq r \\ 0; & \text{otherwise,} \end{cases}$$

where  $\eta_\rho(\nu, \xi) = \frac{\binom{\xi}{\rho}}{\binom{\nu}{\rho}}$  denotes the probability of choosing  $\rho$  many BLACK balls from a bag containing  $\xi$  many are BLACK balls and the  $\nu - \xi$  many WHITE balls. In order to avoid writing the boundary conditions every time, we extend the definition of  $\eta_\rho(\nu, \xi)$  in the following way.

$$\eta_\rho(\nu, \xi) \triangleq \begin{cases} \frac{\binom{\xi}{\rho}}{\binom{\nu}{\rho}}; & \text{if } \nu, \xi, \rho \geq 0 \text{ and } \rho \leq \xi \\ 0; & \text{otherwise.} \end{cases}$$

Define another sequence of  $2n - 1$  binary random variables  $X_0, X_1, \dots, X_{2n-2}$  in the

following manner.

$$X_i = \begin{cases} 1; & \text{if } \begin{cases} i = \text{odd and } P_i = 1 \text{ and } P_{i+1} = 0 \\ i = \text{even and } P_i = 1 \text{ and } P_{i-1} = 0 \end{cases} \\ 0 & \text{otherwise.} \end{cases}$$

That is, the random variable  $X_i$  corresponding to the node  $i$  contributes 1 to the size of the cover if  $i$  is a GREEN node, but its sibling node is not a GREEN node. Assume that  $i$  is odd and  $2^\ell < i < 2^{\ell+1}$ . Then,

$$\begin{aligned} \Pr[X_i = 1] &= \Pr[\{P_i = 1\} \cap \{P_{i+1} = 0\}] \\ &= \Pr[P_{i+1} = 0 | P_i = 1] \cdot \Pr[P_i = 1] \\ &= (1 - \Pr[P_{i+1} = 1 | P_i = 1]) \cdot \Pr[P_i = 1] \\ &= \left(1 - \frac{\binom{r-2^{h-\ell}}{2^{h-\ell}}}{\binom{n-2^{h-\ell}}{2^{h-\ell}}}\right) \cdot \eta_{2^{h-\ell}}(n, r) \\ &= (1 - \eta_{2^{h-\ell}}(n - 2^{h-\ell}, r - 2^{h-\ell})) \cdot \eta_{2^{h-\ell}}(n, r). \end{aligned}$$

Let the random variable  $X$  denote the cover size. Then,  $X$  can be expressed as

$$X = X_0 + X_1 + \cdots + X_{2^{h+1}-2}.$$

Assume that  $2^\alpha \leq r < 2^{\alpha+1}$ . This implies that  $P_i = 0$  for all  $0 \leq i \leq 2^{h-\alpha} - 2$  with the convention that if  $2^{h-\alpha} - 2 < 0$  then such an  $i$  does not exist. Then the cover size  $X$  is given by,

$$\begin{aligned} X &= X_{2^{h-\alpha}-1} + X_{2^{h-\alpha}} + \cdots + X_{2^{h+1}-2} \\ &= \sum_{i=h-\alpha}^h (X_{2^i-1} + X_{2^i} + \cdots + X_{2^{i+1}-2}). \end{aligned}$$

Therefore, by linearity of expectation, the expected cover size is given by

$$\begin{aligned}
E[X] &= \sum_{i=h-\alpha}^h (E[X_{2^{i-1}}] + E[X_{2^i}] + \cdots + E[X_{2^{i+1}-2}]) \\
&= \sum_{i=h-\alpha}^h \{2^i \cdot (1 - \eta_{2^{h-i}}(n - 2^{h-i}, r - 2^{h-i})) \cdot \\
&\quad \eta_{2^{h-i}}(n, r)\}, \tag{3.5}
\end{aligned}$$

where  $r$  denote the size of the pure configuration such that  $2^\alpha \leq r < 2^{\alpha+1}$ .

### 3.4 Experimental Results

In this section, we test the performance of our proposed algorithms experimentally. Consider a tree with  $n$  leaf nodes, where  $n = 2^k$  for some  $k$ . Let  $r$  leaves of the tree would be labeled  $+$  and  $n - r$  leaves would be labeled  $-$ . We are interested in finding the cover size of such a tree. Note that if we fix a tree of size  $n$  and  $r$  of the leaf are  $+$ , then the tree has  $\binom{n}{r}$  different configurations, and each configuration will give rise to a cover of a different size. In our first experiment, we keep  $n$  fixed to 32, and for each  $0 \leq r \leq 32$ , we generate  $\binom{n}{r}$  configurations and compute the cover of all these configurations. For each configuration, we generate three types of covers:

1. Pure cover with the nodes labeled  $+$ , by using the algorithm `PureCoverGen( $\cdot, \cdot$ )` described in Figure 3-3.
2. Pure cover with the nodes labeled  $-$ , also using algorithm `PureCoverGen( $\cdot, \cdot$ )` of Figure 3-3.
3. Mixed cover using Algorithm `MixedCoverGen( $\cdot, \cdot$ )` described in Figure 3-7.

For each  $r$ , we compute the average cover size over all the configurations. These results are summarized in Figure 3-11. The  $X$ -axis of the graph shown in Figure 3-11 represents the value  $r$ , i.e., the number of identifiers which are present corresponding to the keyword  $w$ . The  $Y$ -axis represents the average cover size of all  $\binom{n}{r}$  configurations corresponding to  $r$ . The size of the three types of covers is represented by three

different types of lines. In addition, we have also plotted the line  $Y = X$  for the sake of reference.

The following can be immediately observed from the results demonstrated in Figure 3-11:

1. For all three types of covers, the average size of the covers increases with  $r$ , and then it decreases.
2. For all values of  $r$ , the size of the mixed cover is smaller than the size of the pure covers. The difference is significant when the number of identifiers present is around half of the total number of identifiers.
3. The curves representing the average size of pure cover with nodes labeled  $\dagger$  and the mixed cover always lies below the line  $Y = X$

We have theoretically calculated the expected cover size for pure configuration (see Section 3.3), given by Equation (3.5). To validate our theoretical result (Equation (3.5)), we ran the pure cover generation algorithm for all possible configurations of the set  $\binom{n}{r}$  in a tree with 32 leaf nodes. The experimental results for the expected cover size match very closely with the values predicted by the expression (Equation 3.5), confirming the tightness of our expression.

Next, we repeat the same experiment with different database sizes for  $n = 512, 1024, 2048, \text{ and } 4096$ . In this setting, it would be computationally prohibitive to compute the cover size of all configurations corresponding to a value of  $r$ , for all  $0 \leq r \leq n$ . Hence, for each  $r$ , we generate 5000 uniform random configurations (with replacement), then compute their mixed cover and the average cover size over these 5000 configurations. The results of this experiment are depicted in Figure 3-12, which shows the variation of the average cover size with  $r$  for different values of  $n$ . In each figure, the line  $Y = X$  is also drawn for reference. For pure covers, the curves were drawn using Equation (3.5). The results in Figure 3-12 show the same pattern as in Figure 3-11.


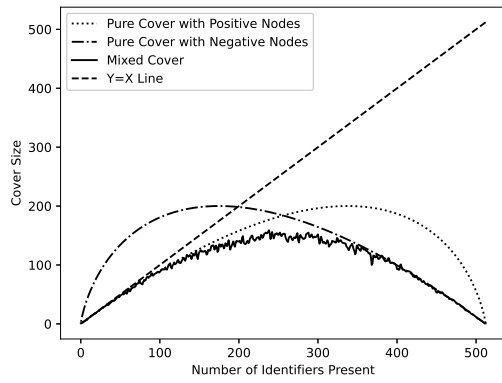


Figure 3-11: Number of Identifiers vs Average Cover Size


### 3.5 Final Remarks

We introduced the concept of a cover of a tree, which is a novel method for representing trees with labeled leaves more concisely. This representation has interesting combinatorial properties, which we explored in detail. We described several algorithms for generating covers and the corresponding algorithms for reconstructing the configuration of a tree. We also proved that our algorithm(s) produce optimal sized covers. We also provided an estimation of the expected size of a cover and experimentally validated our estimates.

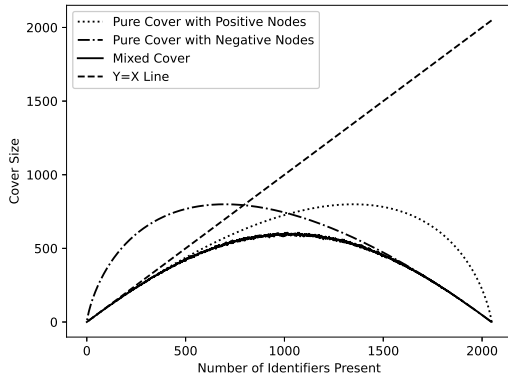
In the following chapter, we demonstrate how the techniques developed in this chapter can be applied to transform a generic SSE scheme into an equivalently secure and efficient version.




(a) Expected Cover Size (512)



(b) Expected Cover Size (1024)



(c) Expected Cover Size (2048)



(d) Expected Cover Size (4096)

Figure 3-12: Expected Cover Size Comparison



## Tree Cover Based SSE

In this chapter, we show how to use our formulation of tree cover and the algorithms for cover generation discussed in Chapter 3 to design a cover based SSE scheme. Our method is generic, i.e., given any secure SSE scheme we can apply our method to convert it into a cover based scheme which will make the resulting scheme more efficient in terms of storage, search time and communication costs. This conversion of any base SSE scheme to cover based SSE scheme is security preserving, i.e., the new scheme retains the security of the base scheme.

To demonstrate our idea, we start by describing our technique for a static SSE in Section 4.1. Further in Section 4.2 we extend our scheme for dynamic SSEs. In Section 4.3, we describe in detail our methodology and point out how our scheme can result in savings if paired with an existing SSE. In Section 4.4, we also provide a detailed security argument showing that the new scheme resulting from pairing our scheme with a base SSE scheme enjoys the same security as the base scheme.

In Section 4.5 we validate our efficiency claims with extensive experimentation on both real and synthetic data. We report the results of our schemes when applied to the Enron Email database [45]. Our proposed scheme achieves a significant reduction (between 35% to 60%) in the size of the encrypted database over a base SSE scheme. We also simulate the dynamic setting in SSE using a synthetic database, and we demonstrate a significant advantage of our scheme in the dynamic setting as well. Furthermore, we provide results on extra overhead incurred by adopting our technique and show that the overhead of our scheme is fairly reasonable in a practical context.

The basic codes and data used for our experiment are publicly available at [77]. This chapter was partly published in [31].

## 4.1 Constructing Static SSE Using Tree Cover

A static database is where the number of documents in the database is a fixed natural number  $D$ . Let,  $\mathcal{D} = \{d_1, \dots, d_D\}$  be the set of documents and  $\mathcal{I} = \{id_1, \dots, id_D\}$  be the set of identifiers associated with those documents. We assume a natural ordering of the identifiers in the database, where  $id_i$  represents the  $i$ -th identifier of the database. Let  $w \in \mathcal{W}$  be an arbitrary keyword and  $m_w$  be the largest integer such that  $id_{m_w} \in \text{db}(w)$ , and let  $h_w$  be the smallest integer such that  $m_w \leq 2^{h_w}$ .

Let  $\mathcal{T}_w$  be a complete binary tree of height  $h_w$ . We will represent  $\text{db}(w)$  by a configuration of the tree  $\mathcal{T}_w$ . We associate each identifier  $id_i$ ,  $i \leq m_w$  with a leaf node of the tree  $\mathcal{T}_w$ , through the injective map  $\varphi^{-1} : \mathcal{I} \rightarrow \text{nodes}(\mathcal{T}_w)$ , where

$$\varphi^{-1}(id_i) = \phi^{-1}(i) = i + 2^{h_w} - 2.$$

Note, the  $\phi^{-1}$  function was introduced in Section 3.1.1. With the above specification,  $\varphi^{-1}(id_i)$  represents the  $i^{\text{th}}$  leaf node from the left of the tree  $\mathcal{T}_w$ .

We label the leaf nodes of  $\mathcal{T}_w$  as follows. For every  $i \leq m_w$ , the leaf node  $\varphi^{-1}(id_i)$  is labeled  $+$  if  $id_i \in \text{db}(w)$  and is labeled  $-$  if  $id_i \notin \text{db}(w)$ . Moreover, if  $m_w < 2^{h_w}$ , then for all  $i$ ,  $m_w < i \leq 2^{h_w}$ , the leaf nodes  $\varphi^{-1}(id_i)$  are labeled  $-$ . This labelling of the leaves of  $\mathcal{T}_w$  yields a complete configuration of  $\mathcal{T}_w$ , and we call this configuration  $L_w$ . Note that this configuration  $L_w$  uniquely represents the set  $\text{db}(w)$ .

We now fix an efficient tree cover scheme  $\Psi = (\text{CoverGen}, \text{ReConstruct})$ , and let  $C_w \leftarrow \Psi.\text{CoverGen}(h_w, L_w)$ , where  $C_w$  be the cover of the configuration  $L_w$  of the tree  $\mathcal{T}_w$  of height  $h_w$ . Thus, from  $C_w$ , the configuration  $L_w$  and further the set  $\text{db}(w)$  can be uniquely reconstructed. This interpretation of  $\text{db}(w)$  as a configuration of  $\mathcal{T}_w$ , which can be represented by a cover, will help us construct an efficient SSE scheme. It is important to note that the height of a tree completely specifies it. It is not

required by any of our schemes described later to store the tree explicitly.

### 4.1.1 Cover-based Representation of DB.

As a first step of constructing a cover-based SSE, we need to convert the given database  $\text{DB}$  to a different representation  $\widetilde{\text{DB}}$ . We call this the *converted database*. This  $\widetilde{\text{DB}}$  then acts as the input to the existing SSE scheme. For each  $w \in \mathbf{W}$ , we generate a configuration  $L_w$  from the set  $\text{db}(w)$  as described before. Let  $C_w \leftarrow \Psi.\text{CoverGen}(L_w)$ . We define

$$\widetilde{\text{db}}(w) = \{(c, s, h_w) : (c, s) \in C_w\}. \quad (4.1)$$

The elements of  $\widetilde{\text{db}}(w)$  are the  $\lambda$ -bit encoding of the tuple  $(c, s, h_w)$ . With this we define

$$\widetilde{\text{DB}} = \bigcup_{w \in \mathbf{W}} \widetilde{\text{db}}(w) \times \{w\}.$$

The conversion scheme from  $\text{DB}$  to  $\widetilde{\text{DB}}$  is summarized in `dbConversion` Algorithm in Figure 4-1.

dbConversion (DB)	Conversion (db(w))
01. set $\widetilde{\text{DB}} \leftarrow \emptyset$	01. set $\widetilde{\text{db}}(w), L_w \leftarrow \emptyset$
02. <b>for</b> each keyword $w \in \mathbf{W}$	02. let $m_w$ be the largest integer such that $\text{id}_{m_w} \in \text{db}(w)$
03. $\widetilde{\text{db}}(w) \leftarrow \text{Conversion}(\text{db}(w))$	03. let $h_w$ be the smallest integer such that $m_w \leq 2^{h_w}$
04. <b>for</b> all $\widetilde{\text{id}} \in \widetilde{\text{db}}(w)$	04. initialize an empty full binary tree $\mathcal{T}_w$ of height $h_w$
05. $\widetilde{\text{DB}} \leftarrow \widetilde{\text{DB}} \cup \{(\widetilde{\text{id}}, w)\}$	05. define $\mathcal{P} = \{\varphi^{-1}(\text{id}_i) : \text{id}_i \in \text{db}(w)\}$
06. <b>return</b> $\widetilde{\text{DB}}$	06. define $\mathcal{R} = \text{leaves}(\mathcal{T}_w) \setminus \mathcal{P}$
	07. <b>for</b> all $p \in \mathcal{P}$
	08. $L_w \leftarrow L_w \cup \{(p, +)\}$
	09. <b>for</b> all $r \in \mathcal{R}$
	10. $L_w \leftarrow L_w \cup \{(r, -)\}$
	11. $C_w \leftarrow \Psi.\text{CoverGen}(h_w, L_w)$
	12. <b>for</b> all $(c, s) \in C_w$
	13. $\widetilde{\text{db}}(w) \leftarrow \widetilde{\text{db}}(w) \cup \{(c, s, h_w)\}$
	14. <b>return</b> $\widetilde{\text{db}}(w)$

Figure 4-1: DB to  $\widetilde{\text{DB}}$  conversion algorithm

### 4.1.2 Generic Static SSE Using Keyword Cover

Let us consider a database

$$\text{DB} = \bigcup_{w \in W} \{(\text{id}, w) : \text{id} \in \text{db}(w)\},$$

and any secure static SSE scheme  $\Sigma = (\Sigma.\text{Setup}, \Sigma.\text{Search})$  as defined in Definition 2.3.1 (without the update protocol). Our goal is to covert  $\Sigma$  into a new SSE  $s\Sigma = (s\Sigma.\text{Setup}, s\Sigma.\text{Search})$ . The procedures  $s\Sigma.\text{Setup}$  and  $s\Sigma.\text{Search}$  are described in Figure 4-2.

$s\Sigma.\text{Setup}$  takes in the database  $\text{DB}$  and a security parameter  $\lambda$  and outputs an encrypted database  $\text{EDB}$ , a key  $k$  and a client state  $\sigma_C$ .  $\text{EDB}$  is uploaded to the server and  $\sigma_C$  and  $k$  are retained with the client.  $s\Sigma.\text{Setup}$  calls the routine  $\text{dbConversion}(\text{DB})$ , described in Figure 4-1 and converts  $\text{DB}$  to  $\widetilde{\text{DB}}$ . Further  $\widetilde{\text{DB}}$  is sent as input to  $\Sigma.\text{Setup}$ , the setup routine of the base static SSE scheme.

Search for a keyword  $w$  is performed by running the  $s\Sigma.\text{Search}$  protocol described in Figure 4-2. In the protocol, first the client's side  $\text{Search}_C$  algorithm corresponding to the base SSE scheme  $\Sigma$  is executed on the inputs  $(k, \sigma_C, w)$ .  $\text{Search}_C(k, \sigma_C, w)$  returns a search token  $\text{stk}_w$  for the keyword  $w$  and the updated client state  $\sigma_C$ . The search token  $\text{stk}_w$  is sent to the server, which subsequently runs the  $\Sigma.\text{Search}_S(\text{stk}_w, \text{EDB})$  and outputs  $\text{res}$ , the search result, which is sent to the client. The rest of the procedure, i.e., lines 15 to 23 of Figure 4-2, runs on the client side. Where the client obtains the set  $\widetilde{\text{db}}(w)$  from the search result sent by the server. Note, as described in Equation 4.1,  $\widetilde{\text{db}}(w)$  is a collection of tuples which encode a cover of a tree of height  $h_w$ . Using the cover reconstruction algorithm, we reconstruct the configuration of the tree representing the set  $\text{db}(w)$  and finally output it.

It is worth noting a few important characteristics of the scheme  $s\Sigma$ .

**Correctness.** The correctness of  $s\Sigma$  directly follows from the correctness of the base static SSE scheme  $\Sigma$  and the correctness of the cover generation scheme  $\Psi$ .

$s\Sigma.\text{Setup}(1^\lambda, \text{DB})$	$s\Sigma.\text{Search}(k, \sigma_C, w)$
01. $\widetilde{\text{DB}} \leftarrow \text{dbConversion}(\text{DB})$ 02. $(k, \sigma_C, \text{EDB}) \leftarrow \Sigma.\text{Setup}_C(1^\lambda, \widetilde{\text{DB}})$ 03. Send EDB to server and retain $(k, \sigma_C)$	01. Generate $(\sigma_C, \text{stk}_w) \leftarrow \Sigma.\text{Search}_C(k, \sigma_C, w)$ 02. Send $\text{stk}_w$ to the server 03. $\text{res} \leftarrow \Sigma.\text{Search}_S(\text{stk}_w, \text{EDB})$ 04. Send $\text{res}$ to the client 05. Decrypt and generate $\widetilde{\text{db}}(w)$ from $\text{res}$ 06. Set $C_w, \widetilde{\text{db}}(w) \leftarrow \emptyset$ 07. <b>for</b> each $\widetilde{\text{id}} \in \widetilde{\text{db}}(w)$ 08.   Parse $\widetilde{\text{id}}$ as $(c, s, h_w)$ 09. $C_w \leftarrow C_w \cup \{(c, s)\}$ 10. Generate $L_w \leftarrow \Psi.\text{ReConstruct}(h_w, C_w)$ 11. <b>for</b> each $(\ell, s) \in L_w$ 12. <b>if</b> $s = +$ 13. $\widetilde{\text{db}}(w) \leftarrow \widetilde{\text{db}}(w) \cup \{\varphi(\ell)\}$ 13. <b>return</b> $\widetilde{\text{db}}(w)$ to client

Figure 4-2: Generic static SSE using tree-cover scheme

**Benefits.** Performance of any SSE schemes is measured by the search and update time and the amount of data communicated during these processes. All these parameters essentially depend on the databases under consideration. For a keyword  $w$ , let  $n_w = |\text{db}(w)|$  be the number of document identifiers containing the keyword  $w$ . In all the state-of-the-art SSE schemes, either static [42] or dynamic, the search complexity is  $\mathcal{O}(n_w)$  [88, 43, 33]. In many dynamic schemes, like in [60, 28, 23, 25, 38], the search complexity is in order of the number of updates for a keyword, which in the worst case may exceed  $\mathcal{O}(n_w)$ . Asymptotically, the worst-case search complexity of our scheme is also  $\mathcal{O}(n_w)$ , but on average the exact number of tuples that need to be communicated is much less than  $n_w$ . To the best of our knowledge, in no existing scheme the search complexity is smaller than  $n_w$ .

Note that, this improvement of search time is obtained by our scheme because we use a compact but lossless representation of  $\text{db}(w)$ . This compact representation results in a smaller index size which further results in more efficient search and communication. The exact savings obtained by pairing our scheme with existing SSEs are further discussed in Section 4.3. Concrete experimental data on real databases is presented in Section 4.5.

**Extra overhead.** Superficially, it may look that  $s\Sigma$  requires more computation than  $\Sigma$ . As  $s\Sigma$  requires the conversion of  $\text{DB}$  to  $\widetilde{\text{DB}}$  in the setup phase and the conversion of the cover to the configuration (lines 15-23 in Figure 4-2). These extra computations are negligible and take place on the client side. Conversion of  $\text{db}$  to  $\widetilde{\text{db}}$  for a database of decently large size only takes a few seconds whereas reconstruction of  $\text{db}$  from cover takes less than a second. More detailed discussion on this can be found in Section 4.5.1. The important savings that are achieved through  $s\Sigma$  is that, on average, the size of  $\widetilde{\text{DB}}$  is much smaller than  $\text{DB}$ , and this will significantly reduce the costs of the routines  $\Sigma.\text{Search}$ . Moreover, this will lead to a smaller size of  $\text{res}$  which leads to a lower communication cost.

**Security.** The scheme  $s\Sigma$  just does some pre-processing of the input to the base scheme  $\Sigma$  and further does some post-processing of the decrypted output. These pre and post-processing take place on the client side and do not require any computation involving the secret key(s). This implies that the security of the base scheme  $\Sigma$  implies the security of  $s\Sigma$ .

BufferUpdate((op, in), B):
01. <b>if</b> (op = add) $\wedge$ ((del, in) $\in$ B)
02. $B \leftarrow B \setminus \{(\text{del}, \text{in})\}$
03. $B \leftarrow B \cup \{(\text{op}, \text{in})\}$
04. <b>else-if</b> (op = del) $\wedge$ ((add, in) $\in$ B)
05. $B \leftarrow B \setminus \{(\text{add}, \text{in})\}$
06. $B \leftarrow B \cup \{(\text{op}, \text{in})\}$
07. <b>else</b>
08. $B \leftarrow B \cup \{(\text{op}, \text{in})\}$
09. <b>return</b> B

Figure 4-3: Buffer update algorithm

We make this intuition more concrete in Section 4.4, where we present a reductionist security proof of our dynamic scheme  $d\Sigma$  (described later in Section 4.2). The security result and the proof are also applicable to the static scheme.

**Additional security advantage.** Informally, an SSE scheme is called *volume hiding* if an adversary cannot guess the number of tuples that are related to a query by

seeing the result of a query which is transmitted by a server. This additional security property in SSE schemes has been recently studied [79, 59, 7]. In general, making an SSE volume hiding makes it inefficient in terms of communication and storage costs. In our scheme, the size of a search result for a keyword  $w$  is related to the size of the cover of the keyword  $w$ , and it does not directly reveal the number of tuples related to the keyword  $w$ . Thus, a constrained passive adversary, who only sees the communication between the server and the client, will not be able to accurately estimate the size of a query result from observing the response sizes. We believe that by using some additional randomness we can make our scheme to be volume hiding for more powerful adversaries.

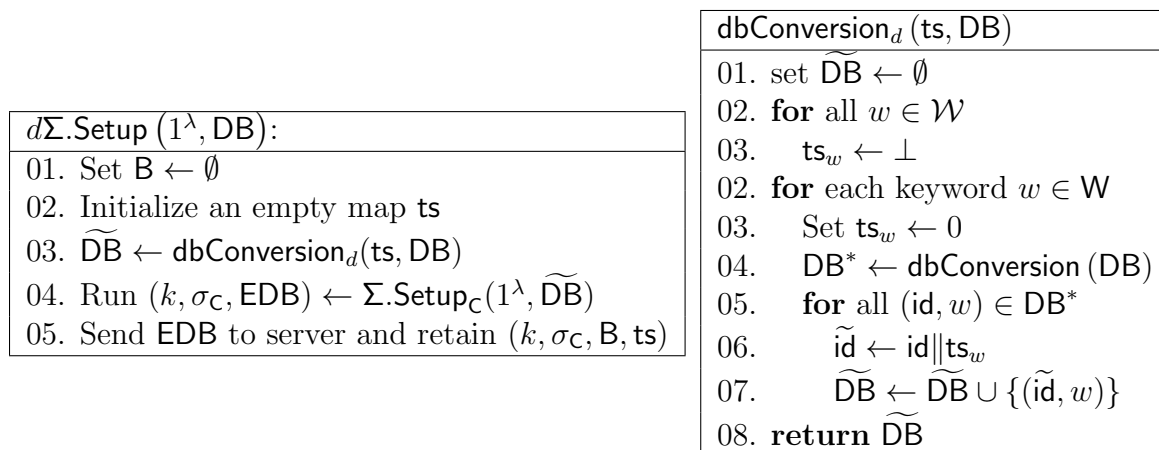


Figure 4-4: A generic dynamic SSE using tree-cover scheme setup phase and DB to DB conversion algorithm

## 4.2 Constructing Dynamic SSE Using Tree Cover

For constructing a dynamic SSE, we need to support modifications in the database. Modification may take place in two ways. One being the addition of new documents in the database and the other being updating an existing document. In the context of SSE, a modification to the database is recorded by adding or removing the corresponding keyword identifier pair involved in the modification. In the context of the tree cover-based SSE scheme, we will still represent  $db(w)$  as a tree, and thus adding a new document would result in adding an extra leaf node to the existing tree, which

in some cases can only be done by increasing the height of the tree. Modification of an existing document would be achieved by assigning or altering the sign of the leaf node associated with the identifiers that were affected by the update operation.

As before we take a secure dynamic SSE scheme  $\Sigma = (\Sigma.\text{Setup}, \Sigma.\text{Search}, \Sigma.\text{Update})$  and convert it into a tree cover based dynamic SSE  $d\Sigma = (d\Sigma.\text{Setup}, d\Sigma.\text{Search}, d\Sigma.\text{Update})$ . The procedures for  $d\Sigma.\text{Setup}$  are shown in Figures 4-4.  $d\Sigma.\text{Update}$  is shown in Figures 4-3 and 4-5.  $d\Sigma.\text{Search}$  is shown in Figure 4-6.

$d\Sigma.\text{Setup}$  described in Figure 4-4 is very similar to  $s\Sigma.\text{Setup}$ . In  $s\Sigma.\text{Setup}$  a given database  $\text{DB}$  is converted to  $\widetilde{\text{DB}}$ , where  $\widetilde{\text{DB}}$  consists of tuples of the form  $(w, c, h_w)$  where  $c$  is an element of the cover corresponding to the keyword  $w$  and  $h_w$  is the height of the corresponding tree representing the set  $\text{db}(w)$ . In the case of  $d\Sigma.\text{Setup}$  the initial database,  $\text{DB}$  is converted into  $\widetilde{\text{DB}}$ , but in this case, the  $\widetilde{\text{DB}}$  consists of tuples of the form  $(w, c, h_w, \text{ts}_w)$ , where  $\text{ts}_w$  is a new variable associated with each keyword, which keeps information about the time at which some identifiers related to  $w$  have been updated. Further, we'll call  $\text{ts}_w$  as the time stamp for the keyword  $w$ .

In the setup phase, for each keyword  $w$ ,  $\text{ts}_w$  is set to zero, signifying that no update has taken place yet. The role of this variable  $\text{ts}_w$  will be more clear from the update operation which we describe next.

Updates, in our case, take place by adding or deleting keyword-identifier pairs. Our main strategy for the update is a *lazy update* model. We assume a buffer memory  $\text{B}$  of restricted size (which may be user defined) at the client's side to store intermediate updates. The client uses this buffer  $\text{B}$  to store a "few" keyword-identifier pairs in un-encrypted form along with the corresponding operation  $\text{op} \in \{\text{add}, \text{del}\}$ , where  $\text{add}$  and  $\text{del}$  represents addition and deletion respectively. Once the buffer is full, the client uploads the contents of the buffer to the server and resets the buffer to empty. The update procedure consists of procedures to update the buffer and procedures to upload the contents of the buffer to the server. The update procedure is summarized in Figure 4-3.

The update procedure shown in Figure 4-5 takes as input  $(\text{op}, \text{in})$ , where  $\text{op} \in \{\text{add}, \text{del}\}$  and  $\text{in}$  is a keyword-identifier pair  $(\text{id}, w)$ . On receiving the input the client



$d\Sigma.\text{Update}(k, \sigma_C, \text{ts}, \text{B}, (\text{op}, \text{in})):$
<p><b>Client Side:</b></p> <ol style="list-style-type: none"> <li>01. <b>if</b> <math>\text{B} \neq \text{FULL}</math></li> <li>02.   <math>\text{B} \leftarrow \text{BufferUpdate}(\text{op}, \text{in}, \text{B})</math></li> <li>03. <b>if</b> <math>\text{B} = \text{FULL}</math></li> <li>04.   set <math>\text{UList} \leftarrow \emptyset</math></li> <li>05.   <b>for</b> each <math>w</math> such that <math>(\text{op}, (\text{id}, w)) \in \text{B}</math></li> <li>06.     set <math>L_w \leftarrow \emptyset</math></li> <li>07.     let <math>m_w</math> be the largest integer such that <math>(\text{op}, (\text{id}_{m_w}, w)) \in \text{B}</math></li> <li>08.     let <math>h_w</math> be the smallest integer such that <math>m_w \leq 2^{h_w}</math></li> <li>09.     initialize an empty full binary tree <math>\mathcal{T}_w</math> of height <math>h_w</math></li> <li>10.     define <math>\mathcal{P} = \{\varphi^{-1}(\text{id}) : (\text{add}, (\text{id}, w)) \in \text{B}\}</math></li> <li>11.     define <math>\mathcal{R} = \{\varphi^{-1}(\text{id}) : (\text{del}, (\text{id}, w)) \in \text{B}\}</math></li> <li>12.     <b>if</b> <math>\text{ts}_w = \perp</math></li> <li>13.       <math>\text{ts}_w \leftarrow 0</math></li> <li>14.       <b>for</b> all <math>p \in \mathcal{P}</math></li> <li>15.         <math>L_w \leftarrow L_w \cup \{(p, +)\}</math></li> <li>16.       <b>for</b> all <math>r \in \text{leaves}(\mathcal{T}_w) \setminus \mathcal{P}</math></li> <li>17.         <math>L_w \leftarrow L_w \cup \{(r, -)\}</math></li> <li>18.       <math>(h_w, C_w) \leftarrow \Psi_d.\text{CoverGen}(L_w)</math></li> <li>19.     <b>else</b></li> <li>20.       <math>\text{ts}_w \leftarrow \text{ts}_w + 1</math></li> <li>21.       <b>for</b> all <math>p \in \mathcal{P}</math></li> <li>22.         <math>L_w \leftarrow L_w \cup \{(p, +)\}</math></li> <li>23.       <b>for</b> all <math>r \in \mathcal{R}</math></li> <li>24.         <math>L_w \leftarrow L_w \cup \{(r, -)\}</math></li> <li>25.       <math>C_w \leftarrow \Psi_d.\text{CoverGen}(h_w, L_w)</math></li> <li>26.       <b>for</b> each <math>(c, s) \in C_w</math></li> <li>27.         <math>\tilde{\text{id}} = (c, s, h_w, \text{ts}_w)</math></li> <li>28.         <math>(\sigma_C, \text{utk}) \leftarrow \Sigma.\text{Update}_C(k, \sigma_C, \text{add}, (\tilde{\text{id}}, w))</math></li> <li>29.         <math>\text{UList} \leftarrow \text{UList} \cup \{\text{utk}\}</math></li> <li>30.   <b>send</b> <math>\text{UList}</math> to server</li> </ol> <p><b>Server Side:</b></p> <ol style="list-style-type: none"> <li>31. <b>for</b> all <math>\text{utk} \in \text{UList}</math></li> <li>32.   <math>\text{EDB} \leftarrow \Sigma.\text{Update}_S(\text{utk}, \text{EDB})</math></li> </ol>

Figure 4-5: A generic dynamic SSE using tree-cover scheme update phase

first checks if  $(\overline{\text{op}}, \text{in}) \in \mathbf{B}$  or not, where  $\overline{\text{op}} = \text{del}$  if  $\text{op} = \text{add}$  and vice versa. If  $(\overline{\text{op}}, \text{in}) \in \mathbf{B}$ , then the client deletes  $(\overline{\text{op}}, \text{in})$  from  $\mathbf{B}$  and add  $(\text{op}, \text{in})$  to the buffer  $\mathbf{B}$ . Otherwise, it only adds  $(\text{op}, \text{in})$  to the buffer  $\mathbf{B}$ . This procedure is summarized in the procedure `BufferUpdate` shown in Figure 4-3. It is important to note that if  $(\overline{\text{op}}, \text{in})$  is present in the buffer, then only deleting  $(\overline{\text{op}}, \text{in})$  from the buffer does not suffice. It is also necessary to add  $(\text{op}, \text{in})$  to the buffer  $\mathbf{B}$  as the client has no knowledge of the current contents of the server, in particular, it is not possible for the client to know during the update process if  $(\overline{\text{op}}, \text{in})$  is currently present in server or not. For example, assume that for an identifier  $\text{id}$  the client wants to delete a keyword  $w$ , i.e.,  $\text{op} = \text{del}$ . Also assume that  $(\text{add}, (\text{id}, w)) \in \mathbf{B}$ . If  $(\text{id}, w)$  currently resides in the server, then deleting  $(\text{add}, (\text{id}, w))$  from buffer and not adding  $(\text{del}, (\text{id}, w))$  to  $\mathbf{B}$  would lead to a wrong configuration for the keyword  $w$ .

After the buffer  $\mathbf{B}$  becomes full the client retrieves all the entries  $(\text{op}, \text{in}) \in \mathbf{B}$  that correspond to each keyword  $w$ . Let  $\text{Udt}(w) = \{(\text{op}, \text{in}) \in \mathbf{B} : \text{in} = (\text{id}, w)\}$ . The set  $\text{Udt}(w)$  gives all identifiers corresponding to  $w$  which were updated in the current phase. The client creates a tree  $\mathcal{T}$  whose leaf nodes based on the identifiers present in  $\text{Udt}(w)$ , the identifiers associated with the operation `add` are labeled  $+$  and the ones associated with `del` are labeled with  $-$ . This labeling gives a configuration  $L_w$  of the tree  $\mathcal{T}$ . This configuration along with the height of the tree is fed to a dynamic cover generation algorithm  $\Psi_d.\text{dCoverGen}$ , which yields a cover  $C_w$  for the configuration  $L_w$ .  $C_w$  consists of pairs  $(i, s)$  where  $i$  is a node of the tree  $\mathcal{T}$  and  $s \in \{+, -\}$ . The update operation uploads a  $\lambda$  bit representation of the tuple  $(i, s, h_w, \text{ts}_w)$  to the server by creating an update token for the string through the client side update procedure of the base SSE, i.e., through  $\Sigma.\text{Update}_C$ . The details are depicted in the Algorithm shown in Figure 4-5.

To perform a search query on  $w$ , the client uses the  $\Sigma.\text{Search}_C$  protocol to search for the keyword  $w$  and obtains the search token  $\text{stk}_w$ . This token is then sent to the server. Upon receiving the search token, the server returns the encrypted search result  $\text{res}_S \leftarrow \Sigma.\text{Search}_S(\text{stk}, \text{EDB})$ , which the client decrypts. Each element of the decrypted search result  $\text{res}$  is an encoding of  $(c, s, h, \text{ts}_w)$ . The client generates a

<p><math>d\Sigma.\text{Search}(k, \sigma_C, \text{ts}, \mathbf{B}, w)</math>:</p> <p><b>Round 1:</b></p> <ol style="list-style-type: none"> <li>01. Generate <math>(\sigma_C, \text{stk}_w) \leftarrow \Sigma.\text{Search}_C(k, \sigma_C, w)</math></li> <li>02. Send <math>\text{stk}_w</math> to server</li> <li>03. Server returns <math>\text{res}_S \leftarrow \Sigma.\text{Search}_S(\text{stk}, \text{EDB})</math></li> <li>04. Parse all <math>r \in \text{res}_S</math> as <math>(c, s, h, i)</math></li> <li>05. <b>for</b> <math>j \in [0, \text{ts}_w - 1]</math></li> <li>06.   Initialize <math>C_j \leftarrow \emptyset</math></li> <li>07.   <math>C_j \leftarrow C_j \cup \{(c, s)\}</math>, and <math>h_j \leftarrow h</math> for all <math>(c, s, h, j) \in \text{res}_S</math></li> <li>08. Generate the final configuration <math>L \leftarrow \Psi_d.\text{dReConstruct}(\{h_j, C_j\}_{j \in [0, \text{ts}_w - 1]})</math></li> <li>09. Generate <math>\text{db}(w)</math> from <math>L</math> using <math>\varphi</math></li> <li>10. Search in <math>\mathbf{B}</math> with keyword <math>w</math> and create the set <math>\text{res}_C</math> from the search result</li> <li>11. <b>for all</b> <math>(\text{op}, (\text{id}, w)) \in \text{res}_C</math></li> <li>12.   <b>if</b> <math>\text{op} = \text{add}</math></li> <li>13.     <math>\text{db}(w) \leftarrow \text{db}(w) \cup \{\text{id}\}</math></li> <li>14.   <b>else</b></li> <li>15.     <math>\text{db}(w) \leftarrow \text{db}(w) \setminus \{\text{id}\}</math></li> </ol> <p><b>Round 2:</b></p> <p><b>Client Side:</b></p> <ol style="list-style-type: none"> <li>16. Set <math>\text{UList} \leftarrow \emptyset</math></li> <li>17. Set <math>\text{ts}_w \leftarrow 0</math></li> <li>18. <math>\widetilde{\text{db}}(w) \leftarrow \text{Conversion}(\text{db}(w))</math></li> <li>19. <b>for all</b> <math>\text{id} \in \widetilde{\text{db}}(w)</math></li> <li>20.   <math>\widetilde{\text{id}} \leftarrow (\text{id}, \text{ts}_w)</math></li> <li>21.   <math>(\sigma_C, \text{utk}) \leftarrow \Sigma.\text{Update}_C(k, \sigma_C, \text{add}, (\widetilde{\text{id}}, w))</math></li> <li>22.   <math>\text{UList} \leftarrow \text{UList} \cup \{\text{utk}\}</math></li> <li>23. Send <math>\text{UList}</math> to server</li> </ol> <p><b>Server Side:</b></p> <ol style="list-style-type: none"> <li>24. <b>for all</b> <math>\text{utk} \in \text{UList}</math></li> <li>25.   <math>\text{EDB} \leftarrow \Sigma.\text{Update}_S(\text{utk}, \text{EDB})</math></li> </ol>
--

Figure 4-6: A generic dynamic SSE using tree-cover scheme search phase

sequence of covers  $\{(h_t, C_t)\}_{t \in [0, \text{ts}_w - 1]}$ , where

$$C_t = \{(c, s) : (c, s, h, \text{ts}_w) \in \text{res} \text{ and } , \text{ts}_w = t\},$$

for all  $t \in [0, \text{ts}_w - 1]$ . It is important to note that for every update (that is when the buffer is full and offloaded to the server), the timestamp and corresponding height  $h$  related to all the updates of a particular keyword are the same. The client then feeds this sequence of covers to the dynamic cover reconstruction algorithm  $\Psi_d.\text{dReConstruct}$

(Section 3.2.4). The final configuration produced by the  $\Psi_d$ .dReConstruct algorithm is used to construct  $\text{db}(w)$ . For computing the final search result the client looks for  $w$  in the buffer  $\mathbf{B}$  and denotes the search result as  $\text{res}_{\mathcal{C}}$ . The entries in  $\mathbf{B}$  are in  $\text{res}_{\mathcal{C}}$  and have the format  $(\text{op}, (\text{id}, w))$ . If  $(\text{del}, (\text{id}, w)) \in \text{res}_{\mathcal{C}}$  then client discards the  $\text{id}$  from  $\text{db}(w)$ . Otherwise, if  $(\text{add}, (\text{id}, w)) \in \text{res}_{\mathcal{C}}$  then the client adds  $\text{id}$  to  $\text{db}(w)$ . Subsequently, it resets  $\text{ts}_w$  to 0. The client then re-uploads the search result for  $w$ , in a manner similar to the update phase discussed earlier. The details are in Figure 4-6.

### 4.3 Discussions

In this section we discuss some existing SSEs and the consequences of pairing our scheme with them.

Consider a dynamic SSE scheme applied to an initially empty database. At a certain instance of time, let  $i_w$  be the number of additions, and  $d_w$  be the number of deletions that have taken place for a keyword  $w$ . Thus, the total number of updates  $u_w$  for the keyword  $w$  is given by  $u_w = i_w + d_w$ , and  $n_w = i_w - d_w$  denote the number of identifier pairs currently matching keyword  $w$ . In addition, let  $N$  be the total number of document identifier pairs in the database at that instance. In Table 4.1 we summarize the characteristics of some widely studied recent SSE schemes. The list does not pretend to be a complete one but is a good representative of the existing SSE schemes. The Table has two major columns named **Computation Cost** and **Communication Cost**. The two sub-columns under Computation Cost report the asymptotic computation cost for search and update, respectively. The three sub-columns under Communication Cost list the size of the result of a single keyword search, the size of an update token for a single update and the number of round-trips (RT), i.e., the number of communication rounds necessary between the server and client for a search.

The schemes reported in Table 4.1 can be naturally grouped into two groups as follows:

*Group-1:* The schemes whose search cost is  $\mathcal{O}(u_w)$ . The first four entries of

Table 4.1, i.e., Fides, Mitra,  $\Pi_{BP}$  and Diana<sub>del</sub> falls under this category.

*Group-2*: The schemes whose search cost is not  $\mathcal{O}(u_w)$  but is dominated by  $n_w$ . The last five entries of Table 4.1 fall under this group.

SSE protocols that achieve a  $\mathcal{O}(n_w)$  time complexity for search are called optimal search protocols. The *Group-2* schemes are near-optimal, as the leading term in the search cost of these schemes is dominated by  $n_w$ .

The *Group-1* schemes fail to achieve the optimal search complexity as they treat deletion also as an insertion with a specific tag and thus the search time depends on the number of updates and not on the number of documents currently in the database that contains  $w$ . But the *Group-2* schemes achieve near-optimal search time at an increased cost for updates. All *Group-2* schemes except Janus and LLSE have an update cost of  $\mathcal{O}(\log N)$  whereas most *Group-1* schemes have a constant update cost.

A similar pattern is observed in the case of communication costs. The size of the search results of all schemes in *Group-2* is dominated by  $n_w$ , but this is achieved with an increased number of communication rounds. Most *Group-1* schemes require a small constant number of communication rounds, but the response size for a search query is  $\mathcal{O}(u_w)$ .

**Effect of our Pre-processing:** Our proposed scheme is just a pre-processing step which can be applied to all existing SSEs. For a concrete understanding, we can consider the effect of our pre-processing step on the schemes listed in Table 4.1. Firstly, if the tree cover scheme is paired with any of the listed schemes the asymptotic complexity of the schemes does not change. For each  $w$ , our scheme deals with  $\widetilde{\mathbf{db}}(w)$  instead of  $\mathbf{db}(w)$ , hence the parameters of interest on which the complexity is measured in all the listed schemes will change. In particular, with our scheme the parameter  $n_w = |\mathbf{db}(w)|$  should be replaced by  $\widetilde{n}_w = |\widetilde{\mathbf{db}}(w)|$  and  $N = |\mathbf{DB}|$  should be replaced by  $\widetilde{N} = |\widetilde{\mathbf{DB}}|$ . We have already amply argued that on average for any database we will have  $\widetilde{n}_w < n_w$  and  $\widetilde{N} < N$ . With this, it is easy to see that on average each of the *Group-2* schemes will have a concrete reduction of both computation and communication costs.

The effect of our scheme in the case of the *Group-1* schemes is similar. In these schemes, the search cost grows linearly with the number of updates  $u_w$ . These schemes consider each update, either insertion or deletion, as a new keyword document pair and these are stored in the database. Thus, the number of keyword document pairs currently in the database is  $u_w$ , and this leads to the linear dependence of the search time with the number of updates. But, if paired with the tree cover scheme, the effective number of updates that are to be stored will get reduced on average as instead of the keyword identifier pairs the cover of the configuration related to those pairs will be stored and this would incur a lesser cost. Based on the same argument, there would be a concrete reduction of the response sizes on average. Moreover, the constant update cost and the constant update token sizes, as achieved by these schemes, would be retained if paired with the tree cover scheme.

However, our pre-processing incurs some additional overhead on the client, as discussed in detail in Section 3.2.3 and experimentally analyzed in Section 4.5.1 (specifically in Table 4.3). These results demonstrate that the pre-processing time remains reasonably low for adequately large databases and provides significant benefits in terms of server storage, search efficiency, and communication overhead.

The use of the buffer memory also adds to the overall performance of the scheme. However, the buffer memory is typically very small as compared to the size of the database. While the buffer does incur additional search time, which is linear in the size of the buffer, it significantly reduces communication overhead. Without the buffer, the update protocol would need to be triggered immediately whenever an update occurs. Additionally, the buffer allows the SSE to take advantage of our proposed tree cover scheme.

## 4.4 Security of $d\Sigma$

As already stated, our scheme  $d\Sigma$  acts as a key-less pre-processing step on a base SSE scheme  $\Lambda$ . Thus, if our scheme is used as a pre-processing over an SSE  $\Lambda$ , the resulting scheme would inherit the security of  $\Lambda$ . In this section, we formalize this

Scheme	Computation Cost		Communication Cost		
	Search	Update	Search	Update	RT
Fides [25]	$\mathcal{O}(u_w)$	$\mathcal{O}(1)$	$\mathcal{O}(u_w)$	$\mathcal{O}(1)$	2
Mitra [49]	$\mathcal{O}(u_w)$	$\mathcal{O}(1)$	$\mathcal{O}(u_w)$	$\mathcal{O}(1)$	1
$\Pi_{BP}$ [38]	$\mathcal{O}(u_w)$	$\mathcal{O}(1)$	$\mathcal{O}(n_w)$	$\mathcal{O}(1)$	2
Diana <sub>del</sub> [25]	$\mathcal{O}(u_w)$	$\mathcal{O}(\log u_w)$	$\mathcal{O}(n_w + d_w \log u_w)$	$\mathcal{O}(1)$	2
Janus [25]	$\mathcal{O}(n_w \cdot d_w)$	$\mathcal{O}(1)$	$\mathcal{O}(n_w)$	$\mathcal{O}(1)$	1
QOS [43]	$\mathcal{O}(n_w \log i_w + \log^2  \mathcal{W} )$	$\mathcal{O}(\log^2 N)$	$\mathcal{O}(n_w \log^2 N)$	$\mathcal{O}(\log^3 N)$	$\mathcal{O}(\log  \mathcal{W} )$
Orion [49]	$\mathcal{O}(n_w \log^2 N)$	$\mathcal{O}(\log^2 N)$	$\mathcal{O}(n_w \log^2 N)$	$\mathcal{O}(\log^2 N)$	$\mathcal{O}(\log N)$
Horus [49]	$\mathcal{O}(n_w \log(d_w) \log N)$	$\mathcal{O}(\log^2 N)$	$\mathcal{O}(n_w \log^2 N)$	$\mathcal{O}(\log^2 N)$	$\mathcal{O}(\log N)$
LLSE [33]	$\mathcal{O}((n_w + \log i_w) \cdot \log \log N)$	$\mathcal{O}(\log^2 N)$	$\mathcal{O}(n_w)$	$\mathcal{O}(\log^2 N)$	1

Table 4.1: Characteristics of some existing SSE schemes:  $N$  is the number of  $(id, w)$  pairs,  $|\mathcal{W}|$  is the number of keywords,  $i_w$  and  $d_w$  are the number of insertions and deletions, and  $u_w = i_w + d_w$  is the total number of updates, and  $n_w = i_w - d_w$  is the number of keyword-identifier pairs currently matching the keyword  $w$ . RT is the number of roundtrips for a search query.

intuition.

The security of a Dynamic SSE  $\Lambda$  is determined by a leakage function  $\mathcal{L} = (\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}})$ .  $\mathcal{L}$  denotes the information that the adversary learns from the setup process and each execution of the search and update protocols. The security of SSE schemes is generally argued by showing that an adversary cannot distinguish between a real-world execution and an ideal-world execution (simulated using the leakage function) of the scheme [42, 60, 28].

We reduce the security of our dynamic scheme  $d\Sigma$  to the security of the base SSE scheme  $\Lambda$ .

**Theorem 4.4.1.** *Let  $d\Sigma = (d\Sigma.\text{Setup}, d\Sigma.\text{Search}, d\Sigma.\text{Update})$  be a Dynamic SSE scheme as described in Figures 4-4, 4-5 and 4-6.  $d\Sigma$  is instantiated with a fixed but arbitrary tree cover scheme  $\Psi = (\Psi.\text{CoverGen}, \Psi.\text{ReConstruct})$  and a base dynamic SSE scheme  $\Lambda$  with leakage profile  $\mathcal{L}$ . If  $\Lambda$  is  $\mathcal{L}$ -adaptive secure, then  $d\Sigma$  is also  $\mathcal{L}$ -adaptive secure.*

**Proof.** We say that  $\text{Sim}_{\mathcal{B}}$  is a *compatible simulator* for an adversary  $\mathcal{B}$  attacking  $\Lambda$  if

$$\left| \Pr [\text{SSEReal}_{\mathcal{B}}^{\Lambda}(\lambda) = 1] - \Pr [\text{SSEIdeal}_{\mathcal{B}, \text{Sim}_{\mathcal{B}}}^{\Lambda}(\lambda) = 1] \right| \leq \text{negl}(\lambda).$$

As  $\Lambda$  is  $\mathcal{L}$ -adaptive secure hence by Definition 2.5.1, a compatible simulator for  $\mathcal{B}$  always exists.

Let  $\mathcal{A}$  be an arbitrary adversary for  $d\Sigma$  instantiated with a fixed but arbitrary tree cover scheme  $\Psi$  and a base dynamic SSE  $\Lambda$ . Based on the protocols involved in  $d\Sigma$  (see Figures 4-4, 4-5, 4-6) it is important to note down the correct interface between the adversary  $\mathcal{A}$  and its challenger:

1. On a setup query  $(1^\lambda, \text{DB})$  for a database  $\text{DB}$  of  $\mathcal{A}$ 's choice its challenger returns  $\text{EDB}$ .
2. On an update query  $(\text{op}, \text{in})$ , its challenger returns  $\text{UList}$  a set of update tokens (see line 30 of Figure 4-5).
3. For a search query  $w$ , the challenger returns a search token  $\text{stoken}_w$  and a set of update tokens  $\text{UList}$  (See lines 02 and 23 of Figure 4-6).

Given an adversary  $\mathcal{A}$  for  $d\Sigma$ , we construct an adversary  $\mathcal{B}$  for the base protocol  $\Lambda$ , which acts as a challenger for  $\mathcal{A}$ .  $\mathcal{B}$  being an adversary for  $\Lambda$  has a challenger which we denote as  $\mathcal{C}$ .  $\mathcal{B}$  provides responses to the queries of  $\mathcal{A}$  with the help of the responses it receives from its challenger  $\mathcal{C}$ . Note the tree cover scheme  $\Psi$  is key-less and is thus accessible to  $\mathcal{B}$ .

Now we describe the interaction of  $\mathcal{A}$ ,  $\mathcal{B}$  and  $\mathcal{C}$  in a sequence of games  $\text{G0}$ ,  $\text{G1}$ ,  $\text{G2}$ . **Game  $\text{G0}$ :** We assume that  $\mathcal{B}$ 's challenger  $\mathcal{C}$  follows the real protocols  $(\Lambda.\text{Setup}, \Lambda.\text{Search}, \Lambda.\text{Update})$  to answer queries of  $\mathcal{B}$ .  $\mathcal{B}$  acts as a challenger to  $\mathcal{A}$  as follows:

1. When a setup request  $\text{Setup}(1^\lambda, \text{DB})$  is issued by  $\mathcal{A}$  then  $\mathcal{B}$  runs lines 01 to 03 of the procedure  $d\Sigma.\text{Setup}(1^\lambda, \text{DB})$  described in Figure 4-4 and obtains  $\widetilde{\text{DB}}$ .  $\mathcal{B}$  then issues the setup request  $(1^\lambda, \widetilde{\text{DB}})$  to its challenger  $\mathcal{C}$  and receives  $\text{EDB}$  as a response which it transmits to  $\mathcal{A}$ . Note, in this process  $\mathcal{B}$  has created a map  $\text{ts}$  and a buffer  $\text{B}$  which it retains with it and updates during the subsequent queries of  $\mathcal{A}$ .
2. On receiving an update query  $(\text{op}, \text{in})$  from  $\mathcal{A}$ ,  $\mathcal{B}$  runs lines 01 to 30 of the procedure  $d\Sigma.\text{Update}$  described in Figure 4-5. In lieu of line 28, it issues a search



query  $(\tilde{\text{id}}, w)$  to its challenger  $\mathcal{C}$  and receives `utoken` in response. It populates the set `UList` using the responses received from  $\mathcal{C}$ , then it executes lines 31 and 32 of the procedure  $d\Sigma.\text{Update}$  and finally sends `UList` to  $\mathcal{A}$ .

3. On receiving a search query  $w$  from  $\mathcal{A}$ ,  $\mathcal{B}$  queries  $\mathcal{C}$  with  $w$  and receives as a response `stokenw`. Then it executes lines 03 to 25 of the procedure  $d\Sigma.\text{Search}$  as described in Figure 4-6. Instead of executing line 21, it queries  $\mathcal{C}$  with an update query  $(\text{add}, (\tilde{\text{id}}, w))$  and receives `utoken` as response. Finally, it sends `UList` as constructed in line 23 to  $\mathcal{A}$ .

After  $\mathcal{A}$  stops querying,  $\mathcal{A}$  outputs a bit  $b \in \{0, 1\}$ ,  $\mathcal{B}$  and  $\text{G0}$  also outputs  $b$ .

We can view  $\text{G0}$  as an interaction between  $\mathcal{A}$  and its challenger  $\mathcal{B}$ , where  $\mathcal{A}$  gets a perfect interface for the real scheme  $d\Sigma$ , thus

$$\Pr [\text{SSEReal}_{\mathcal{A}}^{d\Sigma}(\lambda) = 1] = \Pr[\text{G0} = 1]. \quad (4.2)$$

We can also view adversaries  $\mathcal{A}$  and  $\mathcal{B}$  together as a single adversary  $\mathcal{B}'$  who interacts with  $\mathcal{C}$  (the challenger of  $\mathcal{B}$ ), which provides the perfect interface for  $\Lambda$ . Thus, we have

$$\Pr [\text{SSEReal}_{\mathcal{B}'}^{\Lambda}(\lambda) = 1] = \Pr[\text{G0} = 1]. \quad (4.3)$$

**Game G1:** We make some small changes in Game  $\text{G0}$  to obtain game  $\text{G1}$ . First, observe that for any adversary  $\mathcal{A}$  for  $d\Sigma$ , Game  $\text{G0}$  gives a concrete description for an adversary  $\mathcal{B}$ , and thus of the combined adversary  $\mathcal{B}'$  which attacks  $\Lambda$ . As  $\Lambda$  is  $\mathcal{L}$ -adaptive secure hence a compatible simulator  $\text{Sim}_{\mathcal{B}'}$  for  $\mathcal{B}'$  exists. In Game  $\text{G1}$  the challenger responds to the queries of  $\mathcal{B}'$  using the simulator  $\text{Sim}_{\mathcal{B}'}$  instead of the real protocols of  $\Lambda$ . Thus,

$$\Pr [\text{SSEIdeal}_{\mathcal{B}', \text{Sim}_{\mathcal{B}'}}^{\Lambda} = 1] = \Pr[\text{G1} = 1]. \quad (4.4)$$

As  $\Lambda$  is  $\mathcal{L}$ -adaptive secure hence, we have

$$|\Pr [\text{SSEReal}_{\mathcal{B}'}^{\Lambda}(\lambda) = 1] - \Pr [\text{SSEIdeal}_{\mathcal{B}', \text{Sim}_{\mathcal{B}'}}^{\Lambda}(\lambda) = 1]| \leq \text{negl}(\lambda). \quad (4.5)$$

Thus using Equations (4.3), (4.4), (4.5), we have

$$|\Pr[\text{G0} = 1] - \Pr[\text{G1} = 1]| \leq \text{negl}(\lambda). \quad (4.6)$$

**Game G2:** In this game we view the game G1 a bit differently. We see the description of  $\mathcal{B}$  along with the simulator  $\text{Sim}'_{\mathcal{B}}$  together and call it as  $\text{Sim}_{\mathcal{A}}$ . Note, with this view, we see  $\mathcal{A}$  interacting with a single piece of code  $\text{Sim}_{\mathcal{A}}$ . Thus, we have

$$\Pr[\text{SSEIdeal}_{\mathcal{A}, \text{Sim}_{\mathcal{A}}}^{\Lambda}(\lambda)] = \Pr[\text{G2} = 1], \quad (4.7)$$

and as G1 and G2 are essentially same we have

$$\Pr[\text{G1} = 1] = \Pr[\text{G2} = 1]. \quad (4.8)$$

Now, using Equations (4.6) and (4.8) we have

$$|\Pr[\text{G0} = 1] - \Pr[\text{G2} = 1]| \leq \text{negl}(\lambda). \quad (4.9)$$

Finally, using Equations (4.2), (4.7) and (4.9), we have

$$|\Pr [\text{SSEReal}_{\mathcal{A}}^{\Lambda}(\lambda) = 1] - \Pr [\text{SSEIdeal}_{\mathcal{A}, \text{Sim}_{\mathcal{A}}}^{\Lambda}(\lambda) = 1]| \leq \text{negl}(\lambda),$$

as desired. □

A few more points regarding the security of  $d\Sigma$  are to be noted:

1. Our pre-processing step does not add to the security of the base scheme. The tree cover scheme is meant to enhance the efficiency of the scheme while restoring the security of the base scheme.

2. Theorem 4.4.1 only asserts a basic security guarantee of our scheme. We assumed the base scheme to be  $\mathcal{L}$ -adaptive secure which is a well accepted model of security. But security of SSE schemes is still an active area of research and is not fully understood. It has been claimed that SSE schemes proven secure in the  $\mathcal{L}$ -adaptive model may still succumb to attacks that the security model fails to incorporate. For example, a class of attacks called file injection attacks [97] or a recent generalization in [6] may still be applicable to provably secure schemes. Such weaknesses of the base SSE scheme  $\Lambda$  may affect the security of  $d\Sigma$ .
3. A similar security Theorem holds for our static scheme  $s\Sigma$ .

## 4.5 Experimental Results

To validate our proposed scheme in practical databases, we conducted experiments using the Enron Email Dataset [45], which has about 500,000 documents and 200,000 keywords. We began by extracting keywords from the data set and construction of DB consisting of the keyword identifier pairs. Subsequently, using the algorithm in Figure 4-1 we converted DB to  $\widetilde{DB}$ . We experimented on different sizes of the database by randomly selecting subsets of the database and for each case we compared the sizes of the original and converted versions. These results are shown in Table 4.2 and for easy visual comparison a pictorial representation of the data in Table 4.2 is shown in Figure 4-7. The last column of Table 4.2 computes  $\frac{(|DB| - |\widetilde{DB}|) \times 100}{|DB|}$ . It is evident from Table 4.2 that our protocol results in substantial reductions in database size, ranging from 60% to 35%, even when dealing with reasonably large databases. This demonstrates the effectiveness of our approach in practical databases.

### 4.5.1 Extra Overheads

Our scheme builds over an existing SSE and the extra overhead over the original SSE is the time required for computing covers and cover reconstruction. We experimentally compute the time required for cover generation and computation. We used the

$ \text{DB} $	$ \widetilde{\text{DB}} $	Advantage (%)
$2^{20}$	5,26,492	52.4
$2^{20.5}$	6,08,347	60.2
$2^{21}$	9,24,455	54.8
$2^{21.5}$	15,50,264	49.8
$2^{22}$	29,74,256	43.8
$2^{22.5}$	44,05,133	41.1
$2^{23}$	64,47,481	36.4
$2^{23.5}$	75,17,379	35.6

Table 4.2: The size of original database  $\text{DB}$  and the converted database  $\widetilde{\text{DB}}$  in case of Enron data.

following configuration for our computation:

**CPU:** Intel Core i5-1135G7 @ 2.40GHz  $\times$  8 processor (3.1GHz).

**RAM:** 16 GB

**OS:** Ubuntu 22.04.3 LTS, 64 bits.

**Programming Language:** Python

Table 4.3 shows the time required for cover generation and reconstruction for databases of different sizes. We considered database sizes of  $n = 2^{17}, 2^{18}, 2^{19}$ . For each of these databases, we considered different sizes of  $|\text{db}(w)|$  as shown in the rows of the Table. For each  $|\text{db}(w)|$  we generated 100 random configurations and the times reported for the cover generation and re-construction are the average time required for generation and reconstruction for these 100 configurations. As expected, Table 4.3 clearly shows that the cover generation times increase with both the increase in  $n$  and  $|\text{db}(w)|$ . The reconstruction times reported for  $|\text{db}(w)| < 10^4$ , are negligible. If  $|\text{db}(w)|$  is small it is expected that a pure cover is generated, and reconstruction of a pure cover is immediate.

**Dynamic Scenario:** Now, we test the performance of our scheme for updates. We consider a database containing  $10^9$  keyword-identifier pairs, and a client buffer of size 0.01% of the actual database size. We consider updates for a single keyword  $w$  following the protocol described below.

We consider an initial set of documents  $\mathcal{D}_0$  and a keyword  $w^*$  where  $w^*$  is initially

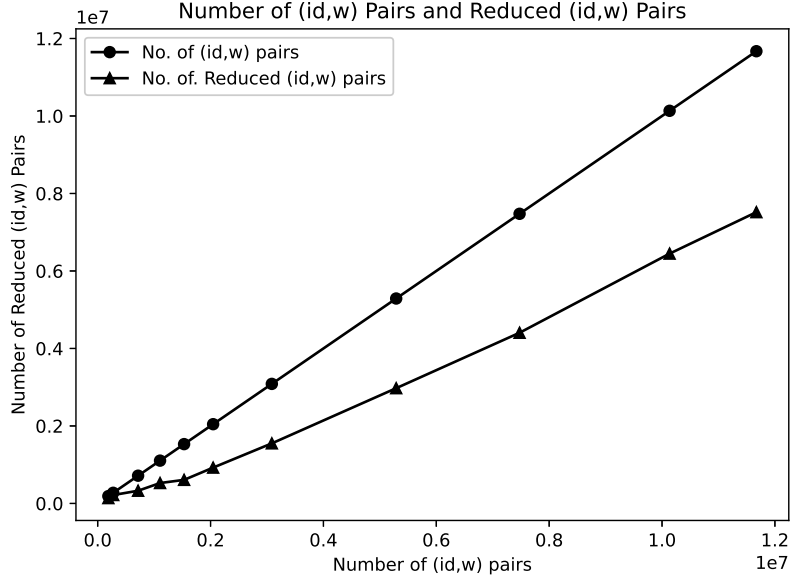


Figure 4-7: A pictorial depiction of the data in Table 4.2 showing the number of tuples in the original database (represented by filled circles) and the number of tuples in the converted database (represented by filled triangles) for different database sizes.

present in 12.5% of the documents in  $\mathcal{D}_0$ . Next, we perform a series of update operations involving  $w^*$ . The updates are made in three phases  $P_1, P_2, P_3$ . The updates in phase  $P_i$  are applied to the documents in  $\mathcal{D}_{i-1}$ , and the updates result in a new set of documents  $\mathcal{D}_i$ . The update operations are designed in such a way that  $w^*$  is present in 25% of the documents in  $\mathcal{D}_1$ , 50% of the documents in  $\mathcal{D}_2$  and 60% documents in  $\mathcal{D}_3$ . In each phase, multiple update operations are performed and 10% of all operations in each phase are delete operations. At the end of each phase of updates, a search operation is performed involving  $w^*$ . Note, that the search operation forces

$ \text{db}(w) $	Time (sec) for $n = 2^{17}$		Time (sec) for $n = 2^{18}$		Time (sec) for $n = 2^{19}$	
	Gen	Re-con	Gen	Re-con	Gen	Re-con
$10^2$	0.22	$10^{-5}$	0.42	$10^{-5}$	0.94	$10^{-5}$
$10^3$	0.27	$10^{-4}$	0.49	$10^{-4}$	1.01	$10^{-4}$
$10^4$	0.51	$10^{-3}$	0.88	$10^{-3}$	1.51	$10^{-3}$
$2^{16}$	0.87	0.17	1.64	0.25	2.97	0.35
$2^{17}$	-	-	1.99	0.37	3.69	0.57
$2^{18}$	-	-	-	-	4.14	0.85

Table 4.3: Cover generation and reconstruction times.

the transmission of all tuples corresponding to  $w^*$  from the buffer to the server.

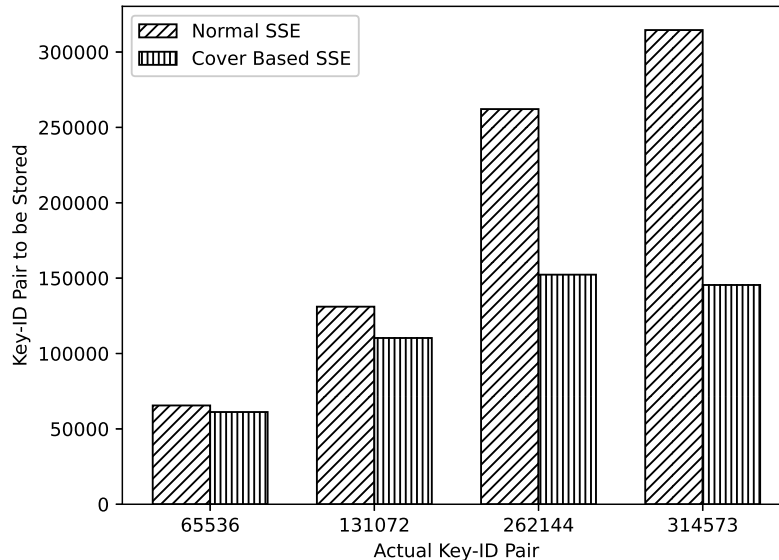


Figure 4-8: Database sizes in dynamic scenario.

We repeated the above protocol 2000 times with random updates following the rules stated above using our dynamic SSE scheme. After each execution, we recorded the size of the database after the search operations in each phase. In Figure 4-8 we report the average size of the database across the 2000 executions in a bar diagram. In Figure 4-8 the database sizes for our scheme and the base SSE are shown. The four pairs of bars correspond to the sizes of the databases after each phase just after the search operation. The first pair correspond to the database for the initial set of documents  $\mathcal{D}_0$ , the second pair is for  $\mathcal{D}_1$  etc. The figure clearly shows that our scheme leads to considerable savings in the dynamic scenario.

## 4.6 Final Remarks

In this chapter we described a novel way to make SSE schemes space efficient. Our scheme converts a given database into a different representation which on average results in a significant amount of space savings. This smaller representation also results in reduced search time and response sizes. In our experiments, we explicitly show that our representation results in smaller index sizes with very low extra overhead. Our

scheme can be used with any secure SSE, and our scheme, being just a pre-processing step, retains the security of the base SSE.





## Conjunctive Dynamic SSE – A Generic Framework

Most SSE schemes to date [42, 60, 28, 23, 87, 25, 49, 89, 38, 90] support single keyword searches per query. Complex searches may involve multiple keywords combined via conjunctions or general boolean expressions. A naive approach to achieve this is by performing individual single keyword searches and then computing set operations on the results. However, this method is inefficient and increases computation and communication between the client and server. This also reveals the number of matching documents for each keyword to the server.

There are a few SSE schemes that support conjunctive queries [29, 47, 2, 93, 58, 66, 80, 102, 78, 100, 95]. The most notable of these is the oblivious cross-tag (OXT) system from Crypto'13 [29]. The OXT scheme is quite efficient with a search complexity of  $\mathcal{O}(n \cdot |\mathbf{db}(w_1)|)$ , where  $n$  is the number of keywords in the conjunction and the set  $\mathbf{db}(w_1)$  contains the ids of documents where the least frequent term in the conjunction  $w_1$  appears. Static schemes like [66, 78] used pre-processing on the data to restrict this leakage.

Our focus in this chapter will be on the construction of forward and backward private dynamic SSE schemes that are adaptively secure and support conjunctive queries. We call such a scheme *conjunctive dynamic searchable symmetric encryption* (CDSSE).

The first CDSSE scheme with claimed forward and backward privacy (Type II) was proposed in 2020 [80]. This was an attempt to extend the OXT scheme to allow dynamic queries while retaining the efficiency of OXT. But, an attack in [100] showed

that the scheme is not forward private.

A few other CDSSE schemes [102, 100, 95] have followed, achieving forward privacy and different flavours of backward privacy with efficiency trade-offs. However, these schemes are far from achieving the efficiency of the OXT scheme.

In this chapter, our goal is the following.

*Construct an adaptively secure forward and backward private CDSSE scheme that has the computational efficiency of the client and server comparable with OXT under standard and well-studied security definitions.*

In this work, we achieve the above goal for a class of databases, which we call as “*non-modifiable databases*”. This chapter has been reported to [3].

## 5.1 Non-modifiable Database

Most of the single-keyword and conjunctive dynamic SSE schemes [60, 28, 23, 25, 87, 89, 49, 101, 38, 80, 102, 100, 95] allow a user to add/delete an identifier-keyword pair  $(id, w)$  to/from the inverted index dynamically at any point of time. So we define “*modifiability*” of a document as

*the ability to add/delete keywords to/from a document at any point.*

However, there are many real-world applications where, once the documents are uploaded to the server, they are not modified. We call such documents “*non-modifiable*”. We assume the inverted index of a database storing non-modifiable documents has the following properties.

- Document identifiers (with associated keywords) can be added to the inverted index and existing identifiers (with associated keywords) can be (marked as) deleted from it *as a whole*.
- Once a document identifier (with associated keywords) has been added to the inverted index, the content of its document may change, but the associated keywords in the inverted index do not change.

We note that in this setting *modifiability can be achieved by deleting the document and then adding the updated version with a different identifier* as is common in SSE literature [28, 25, 51]. Below are some examples of applications handling non-modifiable documents.

1. *Biometric Data.* Stored biometric data (fingerprint, voice, iris scan, facial image, etc.) are typically non-modifiable. It is crucial to secure them for privacy and data protection [61], especially in large databases like Aadhaar [4, 8], for voice data collected in bulk [70], etc.
2. *Media Subscription Services.* In digital media subscription services like Netflix, Spotify, etc., the documents are the DRM-protected media files [75, 14, 5] that once uploaded, do not change.
3. *Digital Libraries and Archives:* Documents in digital libraries like the IEEE Xplore Digital Library and archives are non-modifiable and encrypted for access control [91]. They may also contain sensitive information requiring encryption [83, 82].
4. *Legal Documents:* They are non-modifiable documents containing sensitive information (personal details, financial arrangements, strategic plans, etc.) and should hence be encrypted [9]. Their keywords typically do not change.
5. *Medical Records:* Patient records are non-modifiable documents filled with personally identifiable information (PII), which must be protected under privacy laws<sup>1</sup>, especially for access control and data sharing for collaborative research [21, 84].
6. *Financial Records:* Business and finance documents contain sensitive information about a company's financial status and strategy. Such data do not change and searchability is essential for auditing and making financial decisions. Unauthorised access or revealing access patterns can lead to financial loss, reputational damage, or legal issues.

---

<sup>1</sup>Like HIPAA in the United States: <https://www.ncbi.nlm.nih.gov/books/NBK500019/>

In each of these scenarios, protecting the privacy of the user “by design and by default” is a key responsibility of the service provider according to regulations like GDPR. So while the service providers may use third-party (cloud) storage services for the inverted indices, they must be encrypted to ensure user privacy. In applications like the above, modifiability is either not required (like for biometric data and digital media archives) or should be strictly prohibited (like for legal documents and medical and financial records). The existing definitions of SSE schemes and their security models do not consider this subtlety. That leaves scope for new definitions and scheme designs. As we do away with *modifiability*, several computations and storage requirements from previous SSE schemes are rendered redundant without compromising on security. Hence, we arrive at new efficient schemes that would otherwise be insecure in the general setting of SSE.

## 5.2 Forward and Backward Privacy in CDSSE for Non-modifiable Documents

For backward privacy, the three common notions are called BPIP/Type I, BPUP/Type II, WBP/Type III respectively. Type I is the strongest and Type III is the weakest. In [80], the security notion of the single keyword SSE (SKSSE) scheme was extended to a conjunctive SSE scheme for Type II security only. We have adopted the definition of leakage functions from [23, 25, 38, 80] and appended them as required to define leakage of CDSSE for all three standard notions of backward privacy. Also, [80] missed some leakage components in their definition. We have amended those definitions following the suggestions from [38]. In [38], authors provided another alternate definition of backward privacy called BPLP. However, in a non-modifiable database configuration, BPLP reduces to WBP. Thus we define backward privacy for BPIP, BPUP, and WBP only.

The definition of forward and backward privacy for SSE supporting conjunctive queries is similar to that of single keyword SSE schemes. In [80], the leakage of

SKSSE was extended to define the leakage for CDSSE. However, their extension is for BPUP backward privacy only. We extend the definition of forward and backward privacy of SKSSE to CDSSE for all notions of backward privacy. In [80], the definitions of  $\text{TimeDB}(w)$  and  $\text{Updates}(w)$  for a single keyword  $w$  have been extended to  $\text{TimeDB}(\omega)$  and  $\text{Updates}(\omega)$  for a conjunctive query  $\omega$ . Here, we also extend the other two definitions.

Consider the query sequence  $\text{QS} = \{Q_1, \dots, Q_q\}$ , a conjunctive query  $\omega = w_1 \wedge \dots \wedge w_n$  and its set of included keywords  $L_\omega = \{w_1, \dots, w_n\}$ . We define the following sets of leakages for  $\omega$  as follows.

$\text{sp}(\omega)$ : The search pattern is defined for a pair of keywords  $w_i$  and  $w_j$  in the conjunction  $\omega$ .

$$\text{sp}(w_i, w_j) = \{u \mid (u, \text{srch}, \omega) \in \text{QS}, \text{ and } w_i, w_j \in L_\omega\}$$

Using this definition, the search pattern for the conjunction  $\omega$  is defined as follows.

$$\text{sp}(\omega) = \text{sp}(w_1) \cup \left( \bigcup_{i=2}^n \text{sp}(w_1, w_i) \right)$$

$\text{TimeDB}(\omega)$ : This is the set of all timestamped ids containing keywords in  $L_\omega$  that have been added to but are yet to be deleted from the database. In [80], the definition of  $\text{TimeDB}(w)$  from [25] was extended to  $\text{TimeDB}(\omega)$ . We correct the definition of  $\text{TimeDB}(\omega)$  from [80] with the amendment suggested in [38].

$$\text{TimeDB}(\omega) = \{(\{u_i\}_{i \in [n]}, \text{id}) : w \in L_\omega, ((u_i, \text{add}, (\text{id}, w)) \in \text{QS}) \wedge (\forall u > u_i, (u, \text{del}, (\text{id}, w)) \notin \text{QS})\}.$$

$\text{Updates}(\omega)$ : For a keyword pair  $(w_i, w_j)$  in the conjunction  $\omega$ , we define the set of pairs of timestamps of update operations  $\text{op} \in \{\text{add}, \text{del}\}$  on the database as follows.

$$\text{Updates}(w_i, w_j) = \{(u, u') : ((u, \text{op}, (\text{id}, w_i)) \in \text{QS}) \wedge ((u', \text{op}, (\text{id}, w_j)) \in \text{QS})\}.$$

Using this definition, the update history for the conjunction  $\omega$  is defined as follows.

$$\text{Updates}(\omega) = \text{Updates}(w_1) \cup \left( \bigcup_{i=2}^n \text{Updates}(w_1, w_i) \right)$$

$\text{DelHist}(\omega)$ : We define  $\text{DelHist}(\omega)$  as the set of all pairs of timestamps of addition and deletion for all keywords in a given  $\omega$ .

$$\text{DelHist}(\omega) = \left\{ \{(u_i^{\text{add}}, w_i^{\text{del}})\}_{i \in [n]} : w \in L_\omega, ((u_i^{\text{add}}, \text{add}, (\text{id}, w)) \in \text{QS}) \wedge ((u_i^{\text{del}}, \text{del}, (\text{id}, w)) \in \text{QS}) \right\}.$$

Finally, we define the leakage function  $\mathcal{L}_\Pi = (\mathcal{L}_{\Pi, \text{Setup}}, \mathcal{L}_{\Pi, \text{Update}}, \mathcal{L}_{\Pi, \text{Search}})$ , and correspondingly forward and backward privacy for a CDSSE scheme as follows.

**Definition 5.2.1** (Forward Privacy of CDSSE). *An  $\mathcal{L}_\Pi$ -adaptive secure CDSSE scheme  $\Pi$  is forward private if its leakages due to the Update protocol can be written as the following.*

$$\mathcal{L}_{\Pi, \text{Update}}^{\text{FP}}(\text{op}, \text{in}) = \mathcal{L}_{\Sigma, \text{Update}}^{\text{FP}}.$$

**Definition 5.2.2** (Backward Privacy of CDSSE). *An  $\mathcal{L}_\Pi = (\mathcal{L}_{\Pi, \text{Setup}}, \mathcal{L}_{\Pi, \text{Search}}, \mathcal{L}_{\Pi, \text{Update}})$ -adaptive secure CDSSE scheme  $\Pi$  is backward private as per the notions BPIP, BPUP and WBP if its leakages due to the Setup, Search and, Update protocols can be written as the following.*

$$\begin{aligned} \mathcal{L}_{\Pi, \text{Setup}}^{\text{BPIP}}(\text{DB}) &= \emptyset, & \mathcal{L}_{\Pi, \text{Update}}^{\text{BPIP}}(\text{op}, \text{id}, w) &\preceq \{\text{op}\}, \\ \mathcal{L}_{\Pi, \text{Search}}^{\text{BPIP}}(\omega) &\preceq \{\text{sp}(\omega), \text{TimeDB}(\omega), |\text{Updates}(\omega)|\}, \\ \mathcal{L}_{\Pi, \text{Setup}}^{\text{BPUP}}(\text{DB}) &= \emptyset, & \mathcal{L}_{\Pi, \text{Update}}^{\text{BPUP}}(\text{op}, \text{id}, w) &\preceq \{\text{op}, w\}, \\ \mathcal{L}_{\Pi, \text{Search}}^{\text{BPUP}}(\omega) &\preceq \{\text{sp}(\omega), \text{TimeDB}(\omega), \text{Updates}(\omega)\}, \\ \mathcal{L}_{\Pi, \text{Setup}}^{\text{WBP}}(\text{DB}) &= \emptyset, & \mathcal{L}_{\Pi, \text{Update}}^{\text{WBP}}(\text{op}, \text{id}, w) &\preceq \{\text{op}, w\}, \\ \mathcal{L}_{\Pi, \text{Search}}^{\text{WBP}}(\omega) &\preceq \{\text{sp}(\omega), \text{TimeDB}(\omega), \text{Updates}(\omega), \text{DelHist}(\omega)\}. \end{aligned}$$

### 5.3 Our Construction

Our generic CDSSE construction uses the generic dynamic SKSSE, and the OXT abstraction of [29] to enable conjunctive search queries. The OXT protocol uses two data structures called T-Set and X-Set. The T-Set data structure is used to search all matching documents for every keyword, which is basically a single keyword SSE. The X-Set data structure stores elements that map a keyword to the document identifier which contains it. This X-Set data structure is used to search for the conjunction. The detailed description of OXT protocol is given in Figure 5-2.

In our generic CDSSE construction, we replace T-Set with any generic dynamic SKSSE for the single keyword search on the least frequent keyword  $w_1$ , and the OXT abstraction of [29] remains unaltered. The SKSSE and OXT schemes can be used unaltered, making our CDSSE completely modular. This modular nature allows the flexibility to choose efficient component modules to be used in the generic construction rather than developing a new CDSSE scheme from scratch. It creates the ground for finding better CDSSEs in the future through a more efficient SKSSE, an improved OXT, or a more suitable OXT substitute, instead of a specific construction of a CDSSE. The use of OXT also allows our CDSSE construction to achieve the search complexity of  $\mathcal{O}(nu_1)$ , which was previously attempted by [80].

Next, we describe our generic construction of a forward and backward private CDSSE  $\Pi$  for non-modifiable documents. We assume the existence of a forward and backward private DSSE  $\Sigma$  with support for the **add** operation, and the OXT scheme  $\Gamma$  of [29].  $\Pi$  uses  $\Sigma$  and  $\Gamma$  in a modular fashion as black boxes.

$\Gamma$  assumes that the client knows the *least frequent keyword*  $w_1$  in a conjunction  $\omega = w_1 \wedge \dots \wedge w_n$ . During a search operation in  $\Gamma$ , the server first conducts a single keyword search using the T-Set to find all ids containing the *s-term*  $w_1$ . For all such ids containing  $w_1$ , the X-Set is used to check if the *x-terms*  $w_i, i \geq 2$  are also in id. In  $\Pi$ ,  $\Sigma$  is used to search for the *s-term* and  $\Gamma$  to search for the conjunction.

**A Generic DSSE Scheme.** A generic DSSE scheme  $\Sigma = (\Sigma.Setup, \Sigma.Search, \Sigma.Update)$  is described in Figure 5-1.

$\Sigma.\text{Setup}(1^\lambda, \text{DB}):$ $-\ (\mathbf{K}_s, \sigma_C, \text{TSet}) \leftarrow \Sigma.\text{Setup}_C(1^\lambda, \text{DB})$ $-\ \perp \leftarrow \Sigma.\text{Setup}_S(\text{TSet})$ $\Sigma.\text{Update}(k, \sigma_C, \text{op}, (\text{id}, w); \text{utk}, \text{TSet}):$ $-\ (\sigma_C, \Sigma.\text{utk}) \leftarrow \Sigma.\text{Update}_C(\mathbf{K}_s, \sigma_C, \text{op}, (\text{id}, w))$ $-\ \text{TSet} \leftarrow \Sigma.\text{Update}_S(\Sigma.\text{utk}, \text{TSet})$ $\Sigma.\text{Search}(k, \sigma_C, w; \text{stk}, \text{TSet}):$ $-\ (\sigma_C, \Sigma.\text{stk}) \leftarrow \Sigma.\text{Search}_C(\mathbf{K}_s, \sigma_C, w)$ $-\ (\mathbf{V}_w, \text{TSet}) \leftarrow \Sigma.\text{Search}_S(\Sigma.\text{stk}, \text{TSet})$
---

Figure 5-1: A generic dynamic single-keyword SSE (DSSE) construction  $\Sigma$

**The Oblivious Cross Tag Scheme.** A modular description of the OXT scheme  $\Gamma = (\Gamma.\text{Setup}, \Gamma.\text{Add}, \Gamma.\text{Search}_C, \text{Search}_S)$  is in Figure 5-2. It assumes the existence of three PRF families  $F_p^{(I)} : \mathcal{K} \times \mathcal{I} \times \{0, 1\} \rightarrow \mathbb{Z}_p^*$ ,  $F_p^{(X)} : \mathcal{K} \times \mathcal{W} \rightarrow \mathbb{Z}_p^*$  and  $F_p^{(Z)} : \mathcal{K} \times \mathcal{W} \times \mathbb{N} \rightarrow \mathbb{Z}_p^*$ . It omits the SKSSE part of the original OXT scheme [29], to be substituted by the functionalities of the DSSE  $\Sigma$ . The original OXT scheme [29] only supports static databases where all  $(\text{id}, w)$  pairs are added during setup. To adapt it to our dynamic setting, we have split the setup of [29] into  $\Gamma.\text{Setup}$  and  $\Gamma.\text{Add}$  protocols, without changing their functionalities.  $\Gamma.\text{Search}_S$  also assumes an input  $\mathbf{V}_w$  which has  $\mathbf{V}_w[y_j] = y_j$  for  $j \in [\text{cnt}[w]]$  generated within  $\Gamma$  for a keyword  $w$ . The  $\text{res}$  however accumulates  $\mathbf{V}_w[\text{val}_j]$  from  $\Sigma$ . This step integrates  $\Sigma$  and  $\Gamma$  in a modular fashion within our CDSSE  $\Pi$ .

### 5.3.1 Our Generic CDSSE Scheme

Our generic CDSSE scheme  $\Pi = (\Pi.\text{Setup}, \Pi.\text{Search}, \Pi.\text{Update})$  is described using  $\Sigma$  and  $\Gamma$  in Figure 5-3. Note the arguments passed to  $\Gamma.\text{Search}_S$  in step 3 of  $\Pi.\text{Search}_S$  Figure 5-3. The input  $\text{cnt}$  to  $\Gamma.\text{Search}_S$  is  $\text{st.cnt}[w_1]$  from  $\Pi.\text{Search}_S$ . The input  $\Pi.\mathbf{V}_{w_1}$  is from the single keyword search  $\Sigma.\text{Search}_S$  on the  $s$ -term  $w_1$ , containing pairs  $(\text{val}_j, y_j)$  for  $j \in [\text{st.cnt}[w_1]]$ .

On receiving the result  $\text{res}$  from the server in step 4 of  $\Pi.\text{Search}_S$ , the client finds the search result for the conjunction. The identifier in the search result is of the form  $\text{id}||\text{op}$ . To get the final search result  $\text{db}(\omega)$ , the client does the following. Initialize



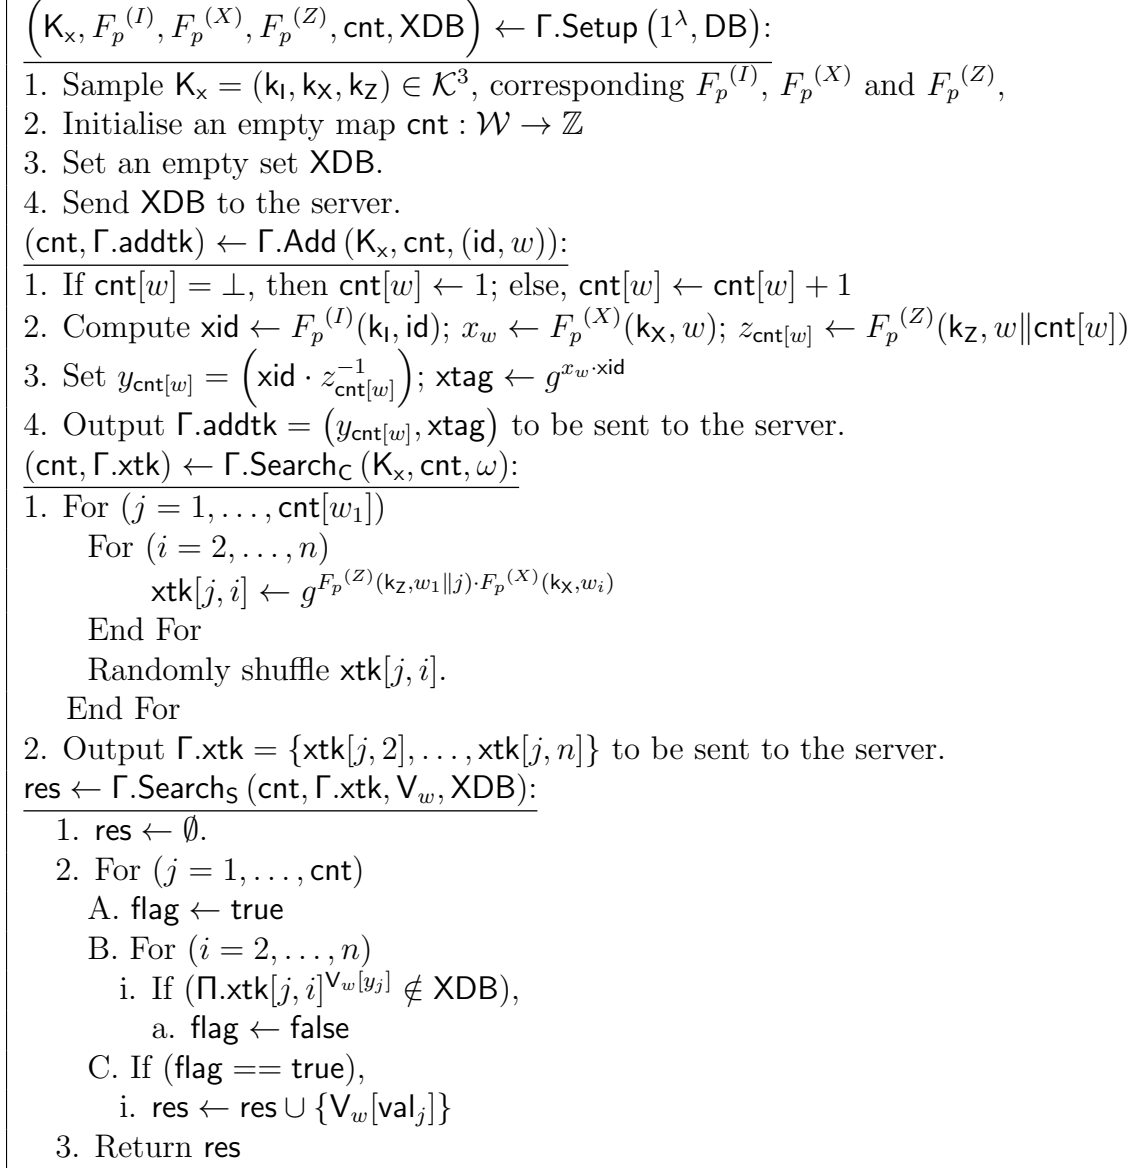


Figure 5-2: A modular description of  $\Gamma$  without the SKSSE part of [29]

an empty set  $\mathcal{I}_\omega \leftarrow \emptyset$ . For each  $r \in \text{res}$ , parse  $r = \text{id} \parallel \text{op}$ . Now if ( $\text{op} = \text{add}$ ),  $\mathcal{I}_\omega \leftarrow \mathcal{I}_\omega \cup \{\text{id}\}$ , and if ( $\text{op} = \text{del}$ ).  $\mathcal{I}_\omega \leftarrow \mathcal{I}_\omega \setminus \{\text{id}\}$ . Finally output,  $\mathcal{I}_\omega$ .

### 5.3.2 Design Rational

With our generalisation in the previous section, the scheme of [80] is now just an instantiation of our construction. The work of [80] essentially club Mitra [49] and their proposed ODXT. The scheme of [80] has a large search token size due to the

<p><math>\Pi</math>.Setup (<math>1^\lambda, \text{DB}</math>):</p> <p><b>Client:</b></p> <p><math>(\text{K}, \text{st}, \Pi.\text{EDB}) \leftarrow \Pi.\text{Setup}_C (1^\lambda, \text{DB})</math>:</p> <p><math>(\text{K}_s, \sigma_C, \text{TSet}) \leftarrow \Sigma.\text{Setup}_C (\text{DB}, 1^\lambda)</math>; assign <math>\text{st}.\sigma_C \leftarrow \sigma_C</math> and <math>\Pi.\text{TSet} \leftarrow \text{TSet}</math></p> <p><math>(\text{K}_x, F_p^{(I)}, F_p^{(X)}, F_p^{(Z)}, \text{cnt}, \text{XDB}) \leftarrow \Gamma.\text{Setup}(\text{DB}, 1^\lambda)</math></p> <p>Assign <math>\text{st}.\text{cnt} \leftarrow \text{cnt}</math> and <math>\Pi.\text{XDB} \leftarrow \text{XDB}</math></p> <p>Set <math>\text{K} \leftarrow (\text{K}_s, \text{K}_x)</math>, and <math>\text{st} \leftarrow (\text{st}.\sigma_C, \text{st}.\text{cnt})</math></p> <p>Send, <math>\Pi.\text{EDB} \leftarrow (\Pi.\text{TSet}, \Pi.\text{XDB})</math> to the server</p> <p><b>Server:</b></p> <p><math>\perp \leftarrow \Pi.\text{Setup}_S (\Pi.\text{EDB})</math>: Server stores <math>\Pi.\text{EDB}</math></p> <p><math>\Pi</math>.Update (<math>k, \text{st}, \text{op}, (\text{id}, w)</math>; <math>\Pi.\text{utk}, \Pi.\text{EDB}</math>):</p> <p><b>Client:</b></p> <p><math>(\text{st}, \Pi.\text{utk}) \leftarrow \Pi.\text{Update}_C (\text{K}, \text{st}, \text{op}, (\text{id}, w))</math>:</p> <ol style="list-style-type: none"> <li>1. Call <math>\Sigma.\text{Update}_C (\text{K}_s, \text{st}.\sigma_C, \text{add}, (w, (\text{id} \parallel \text{op})))</math> to get <math>\Sigma.\text{utk}</math></li> <li>2. Call <math>\Gamma.\text{add} (\text{K}_x, \text{st}.\text{cnt}, w, \text{id} \parallel \text{op})</math> to get <math>\Gamma.\text{addtk} = (y_{\text{cnt}[w]}, \text{xtag})</math></li> <li>3. Set <math>\Pi.\text{utk}_s \leftarrow (\Sigma.\text{utk}, y_{\text{cnt}[w]})</math> and <math>\Pi.\text{utk}_x \leftarrow \{\text{xtag}\}</math>;</li> <li>4. Send <math>\Pi.\text{utk} \leftarrow (\Pi.\text{utk}_s, \Pi.\text{utk}_x)</math> to the server</li> </ol> <p><b>Server:</b></p> <p><math>\Pi.\text{EDB} \leftarrow \Pi.\text{Update}_S (\Pi.\text{utk}, \Pi.\text{EDB})</math>:</p> <ol style="list-style-type: none"> <li>1. Call <math>\Sigma.\text{Update}_S (\Pi.\text{utk}_s, \Pi.\text{TSet})</math> and to get the output <math>\text{TSet}</math> as following <ol style="list-style-type: none"> <li>a. Parse <math>(\Sigma.\text{utk}, y_{\text{cnt}[w]}) \leftarrow \Pi.\text{utk}_s</math></li> <li>b. Further parse <math>(\text{loc}, \text{val}) \leftarrow \Sigma.\text{utk}</math></li> <li>c. Update <math>\Pi.\text{TSet}[\text{loc}] \leftarrow (\text{val}, y_{\text{cnt}[w]})</math></li> </ol> </li> <li>2. Set <math>\Pi.\text{XDB}</math> as <math>\Pi.\text{XDB} \leftarrow \Pi.\text{XDB} \cup \Pi.\text{utk}_x</math></li> <li>3. Output <math>\Pi.\text{EDB} \leftarrow (\Pi.\text{TSet}, \Pi.\text{XDB})</math></li> </ol> <p><math>\Pi</math>.Search (<math>k, \text{st}, \omega</math>; <math>\Pi.\text{stk}, \Pi.\text{EDB}</math>):</p> <p><b>Client:</b></p> <p><math>(\Pi.\text{stk}, \Pi.\text{xtk}) \leftarrow \Pi.\text{Search}_C (\text{K}, \text{st}, \omega)</math>:</p> <ol style="list-style-type: none"> <li>1. Call <math>\Sigma.\text{Search}_C (\text{K}_s, \text{st}.\sigma_C, w_1)</math> to get <math>\Sigma.\text{stk}</math>; assign <math>\Pi.\text{stk} \leftarrow \Sigma.\text{stk}</math></li> <li>2. Call <math>\Gamma.\text{Search} (\text{K}_x, \text{st}.\text{cnt}, \omega)</math> to get <math>\Gamma.\text{xtk}</math>; assign <math>\Pi.\text{xtk} \leftarrow \Gamma.\text{xtk}</math></li> <li>3. Send <math>(\Pi.\text{stk}, \Pi.\text{xtk})</math> to the server</li> </ol> <p><b>Server:</b></p> <p><math>\text{res} \leftarrow \Pi.\text{Search}_S ((\Pi.\text{stk}, \Pi.\text{xtk}), \Pi.\text{EDB})</math>:</p> <ol style="list-style-type: none"> <li>1. <math>\text{res} \leftarrow \phi</math>.</li> <li>2. Call <math>\Sigma.\text{Search}_S (\Pi.\text{stk}, \Pi.\text{TSet})</math> to get <math>\Pi.\text{V}_{w_1} = \{(\text{val}_j, y_j) : j \in [\text{st}.\text{cnt}[w_1]]\}</math></li> <li>3. <math>\text{res} \leftarrow \Gamma.\text{Search}_S (\text{st}.\text{cnt}, \Pi.\text{xtk}, \Pi.\text{V}_{w_1}, \Pi.\text{XDB})</math></li> <li>4. Send <math>\text{res}</math> to the client</li> </ol>
---

Figure 5-3: Our generic CDSSE construction  $\Pi$ .

use of Mitra. With our generalisation, we can do away with the large search token size of [80] via our instantiations  $\Pi_{\text{BP-OXT}}$  which performances better as can be seen

from our experimental results.

Using a generic SKSSE and the OXT scheme in a black box fashion we achieved the computation and communication complexity for search operation as  $\mathcal{O}(n \cdot \text{db}(w_1))$ , and update operation of  $\mathcal{O}(1)$  for our construction. In [95], the computation and communication complexities of the search operation are both  $\mathcal{O}(u_1 + n \cdot \text{db}(w_1))$  and involves two rounds, where  $u_1$  is the number of updates related to  $w_1$ ; the complexities of the update operation are  $\mathcal{O}(|W|/|W_d|)$  for adding a document,  $\mathcal{O}(|D|)$  for editing and  $\mathcal{O}(1)$  for deleting, where  $|W|$  is the number of keywords in the database and  $|W_d|$  denotes the number of keywords contained in the updated document. The storage requirement of [95] is  $\mathcal{O}(|W||D|)$ , as compared to  $\mathcal{O}(N)$  for our scheme, where  $N$  is the number of identifier-keyword pairs in the database. Typically,  $N$  is much smaller than  $|W||D|$  in practice. There is no additional storage requirement due to the keyword updates in [95]. Compared to these, our scheme enjoys the search and update complexities as  $\mathcal{O}(n \cdot \text{db}(w_1))$ , and  $\mathcal{O}(1)$  respectively.

The scheme of [102] performs  $\mathcal{O}(|D|)$  computation for search, where  $|D|$  is the number of documents in the database, making it prohibitively slow. The scheme of [100] uses a single keyword SSE to store and fetch the elements of both TSet and XSet which makes it very slow as well. The scheme of [95] runs in sub-linear time but is still far from the practical performance of OXT. See Table 5.1 for a detailed comparison.

Scheme	Search			Update			Storage		FP	BP	KPRP
	Computation	Communication	RT	Add	Edit	Delete	Client	Server			
VBTree [93]	$\mathcal{O}(u \text{db}(w_1)  \log  D )$	$\mathcal{O}(n +  \text{db}(\omega) )$	1	$\mathcal{O}(W_d)$	$\mathcal{O}(v)$	$\mathcal{O}(1)$	$\mathcal{O}( W  \log  D )$	$\mathcal{O}( W h)$	Yes	No	No
BDXT [80]	$\mathcal{O}(u_1 + n \text{db}(w_1) )$	$\mathcal{O}(u_1 + n \text{db}(w_1) )$	2	$\mathcal{O}(1)$	-	$\mathcal{O}(1)$	$\mathcal{O}( W  \log  D )$	$\mathcal{O}(N)$	No	Type-II	No
ODXT [80]	$\mathcal{O}(nu_1)$	$\mathcal{O}(nu_1)$	1	$\mathcal{O}(1)$	-	$\mathcal{O}(1)$	$\mathcal{O}( W  \log  D )$	$\mathcal{O}(N)$	No	Type-II	No
FBSSE-CQ [102]	$\mathcal{O}(u D )$	$\mathcal{O}(n +  D )$	1	$\mathcal{O}( D )$	$\mathcal{O}( D )$	$\mathcal{O}( D )$	$\mathcal{O}( W  \log  D  + \lambda)$	$\mathcal{O}(N D )$	Yes	Type-II	Yes
HDXT [95]	$\mathcal{O}(u_1 + n \text{db}(w_1) )$	$\mathcal{O}(u_1 + n \text{db}(w_1) )$	2	$\mathcal{O}( W /W_d)$	$\mathcal{O}( D )$	$\mathcal{O}(1)$	$\mathcal{O}( W  \log  D  + \lambda)$	$\mathcal{O}( W  D )$	Yes	Type-II	Yes
Our construction (Sec 5.3)	$\mathcal{O}(nu_1)$	$\mathcal{O}(nu_1)$	1	$\mathcal{O}(1)$	$\mathcal{O}(W_d)$	$\mathcal{O}(1)$	$\mathcal{O}( W  \log  D )$	$\mathcal{O}(N)$	NMD	Type-II	No

Table 5.1: Comparison of our construction with previous schemes.  $W$ : the set of keywords.  $D$ : Set of documents.  $W_d$ : Number of keywords in the document to be updated.  $N$ : Number of identifier-keyword pairs in the database.  $n$ : Number of keywords in a conjunction.  $u_i$ : Updates performed for keyword  $w_i$ .  $u = \sum_{i=1}^n u_i$ : Total updates for a conjunction of  $n$  keywords.  $v$ : Keywords involved in the edit operation.  $h$ : Average number of documents matched by any two keywords.  $\text{db}(\omega)$ : Set of matching document identifiers for query  $\omega = w_1 \wedge \dots \wedge w_n$ , where  $n \geq 1$ . RT: Round-Trip. FP: Forward Private. BP: Backward Private. KPRP: Keyword-Pair Result Pattern Hiding. NMD: Non-Modifiable Documents.

**Leakage Suppression.** The scheme of [95] pays with extra storage, computation and an additional round-trip to prevent *keyword-pair result pattern* (KPRP) leakage. In a conjunctive query for  $w_1 \wedge w_2 \wedge \dots \wedge w_n$ , KPRP leakage contains  $\text{db}(w_i) \cap \text{db}(w_j)$  for all  $1 \leq i < j \leq n$ . Although our scheme does not prevent KPRP leakage completely, we have taken some countermeasures prescribed in [80] to restrict the leakage of our scheme that makes the attack of [97] using KPRP leakage very difficult to mount. We have discussed these countermeasures later while discussing the security of our scheme. According to [80] using those countermeasures, the best-known attack described in [97] which uses KPRP leakage of SSE, has a success probability of less than 5% even when the fraction of file injection is more than 60% in the database. By the same argument, despite not hiding the KPRP leakage, our construction is suitable for practical applications.

## 5.4 Security of Our Construction

The security of SSE schemes is generally argued based on leakage to the server that is upper bounded by some well-defined and “*sensible*” leakage functions that have been studied over many years. The practical consequences of such leakage functions especially through resulting attacks, are well studied and understood. So it is always good to have an SSE scheme that is secure against a well-studied leakage function. Our generic construction enjoys this advantage, as the security of our scheme is inherited from the security of the individual components that are secure under well-studied leakage functions. We note here that previous forward and backward private CDSSEs [102] and [100] allowing “modifiable” documents have proven their security under non-standard security notions that warrant more investigations.

We provide a rigorous security analysis of our generic scheme, assuming the standard and maximum possible leakage of the individual components. The reason to do so is to keep room for possible future improvements in efficiency or the adaptability of the generic construction in other ways.

We recollect that  $\Sigma$  is an adaptively secure forward and backward private SKSSE

scheme as per Definitions 2.5.1, 2.5.5 and 2.5.6. With this assumption, we prove our generic CDSSE scheme  $\Pi$  to be adaptively secure as well. The leakage functions in the backward privacy definition of CDSSE (see Definition 5.2.1 and Definition 5.2.2) subsume the actual leakages of our generic scheme. As discussed in Section 5.2, the leakage of WBP is a superset of the leakage of the other two notions of backward privacy – BPUP and BPIP – for non-modifiable documents, and that of forward privacy as well. We prove  $\Pi$  to be secure assuming the leakage for WBP. So our scheme would be secure when instantiated with a BPIP and BPUP secure SKSSE scheme as well. We follow the common practice of defining and subsequently proving the security of SSE, assuming more than the actual leakages. We later provide the exact leakages of our scheme as well. We assume the leakage of our conjunctive dynamic SSE  $\Pi$  to be the following.

$$\mathcal{L}_{\Pi, \text{Update}}(\text{op}, \text{id}, w) = \text{op}, \quad \mathcal{L}_{\Pi, \text{Search}}(\omega) = \{\text{sp}(\omega), \text{TimeDB}(\omega), \text{Updates}(\omega)\}.$$

Note that our protocol never used `del` as an operation, so  $\text{DelHist}(\omega)$  is always empty, in our case. To the best of our knowledge, all forward and backward private schemes known in the literature have, at most, this leakage. Also, it is essential to point out that the leakage of our construction becomes exactly  $\mathcal{L}_{\Sigma}$  if  $\omega = w$ .

**Theorem 5.4.1.** *Assuming that (1)  $\Sigma$  is an  $\mathcal{L}_{\Sigma}$ -adaptively secure dynamic DSSE, (2)  $F_p$  is a secure PRF, and (3) the eddh problem is hard in the group output by  $\mathcal{G}$ , our scheme  $\Pi$  is  $\mathcal{L}_{\Pi}^{\text{BPUP}}$ -adaptively secure conjunctive dynamic SSE.*

**Proof Idea.** We first describe six games  $G_0, \dots, G_6$  where we replace all the cryptographic primitives used in our construction one by one with ideal primitives and show that no PPT adversary can detect the change in any of the steps. Finally, we show that there exists a simulator that takes the leakages  $\mathcal{L}_{\Pi} = (\mathcal{L}_{\Pi, \text{Setup}}, \mathcal{L}_{\Pi, \text{Update}}, \mathcal{L}_{\Pi, \text{Search}})$  as input and simulates the protocol’s transcript that is indistinguishable from the actual execution of the protocol with ideal primitives. If all leakages have been considered, then the simulator can simulate all the server’s inputs. Since the server has no other

source of input in the protocol, hence there is no other leakage. This proves that our construction  $\Pi$  does not leak more than what is defined by the leakage  $\mathcal{L}_\Pi$ .

In our security proof, Game  $G_0$  is the original execution of the protocol  $\Pi$ . In Game  $G_1$  replace all cryptographic primitives used in the  $\Sigma$  protocol with ideal primitives, and show that the adversary has an advantage at most as an SKSSE adversary against the secure  $\Sigma$  protocol. In Game  $G_2$  we replace all the PRFs  $F_p^{(I)}$ ,  $F_p^{(X)}$ , and  $F_p^{(Z)}$  one by one with randomly sampled elements from  $\mathbb{Z}_p^*$ , and show that an adversary has negligible advantage in Game  $G_2$  over Game  $G_1$ . In Game  $G_3$ , we simulate the part of the search token of  $\Pi$  in the same fashion as proof of the underlying  $\Sigma$  and show that an adversary has a negligible advantage in Game  $G_2$  over Game  $G_3$ , which is upper bounded by the SKSSE adversary. In Game  $G_4$  and Game  $G_5$  we change the way  $\text{xtk}$  and  $y_{\text{cnt}[w]}$  are generated respectively and show Game  $G_3$ ,  $G_4$  and Game  $G_5$  are all equivalent. In the final Game  $G_6$  we change the way  $\text{xtags}$  are generated and show that for all adversary the distinguishing advantage of  $G_6$  over  $G_5$  is upper bounded by the `eddh` advantage of any PPT adversary. Finally, we demonstrate a poly-time simulator that can simulate a computationally indistinguishable transcript from what the adversary got in Game  $G_6$ . Thus, we conclude our theorem.

*Proof.* We now describe the six games  $G_0, G_1, \dots, G_6$  in the following.

GAME  $G_0$ . The experiment starts with the real execution of the protocol  $\Pi$  as described in Section 5.3.1. Thus we have,

$$\Pr [\text{SSEReal}_A^\Pi(1^\lambda, q) = 1] = \Pr[G_0 = 1].$$

As we have assumed  $\Sigma$  to be adaptively secure, so it can be simulated exactly as in its security proof. So in the following steps, we do not explicitly argue about the simulation of those elements generated from  $\Sigma$ .

GAME  $G_1$ . In game  $G_1$ , we replace all cryptographic primitives used in the construction of  $\Sigma.\text{utk}$ , same as the done in the security proof of  $\Sigma$ . We can do so as  $\Sigma$  is  $\mathcal{L}_\Sigma$ -adaptive secure. Thus the advantage of any adversary distinguishing game  $G_0$

and  $G_1$  is bounded by the SSE advantage of  $\Sigma$ . We write it as follows.

$$|\Pr[G_0 = 1] - \Pr[G_1 = 1]| \leq \text{Adv}_{\mathcal{A}, \mathcal{S}, \mathcal{L}_\Sigma}^{\Sigma, \text{SSE}} \leq \text{negl}(\lambda).$$

**GAME  $G_2$ .** Next the challenger replaces all executions of  $F_p^{(I)}$ ,  $F_p^{(X)}$  and  $F_p^{(Z)}$  one by one with randomly sampled strings from  $\mathbb{Z}_p^*$ . The experiment maintains three lists  $I$ ,  $X$  and  $Z$  to store the values sampled for  $F_p^{(I)}$ ,  $F_p^{(X)}$  and  $F_p^{(Z)}$ , thus can be reused later if needed. The entries of  $I$ ,  $X$  and  $Z$  are indexed by  $(\text{id}, \text{op})^2$ ,  $w$  and  $(w, \text{cnt}[w])$ . For an update operation of  $(\text{op}, (\text{id}, w))$ , if  $I[\text{op}, \text{id}] = \perp$  then the challenger picks a random value from  $\mathbb{Z}_p^*$  and uses it for  $F_p^{(I)}(k_1, \text{id}||\text{op})$ , and also stores the random value in  $I[\text{op}, \text{id}]$ . If  $I[\text{op}, \text{id}] \neq \perp$ , then the challenger uses the value stored in  $I[\text{op}, \text{id}]$ . The other two PRFs  $F_p^{(X)}$  and  $F_p^{(Z)}$  are simulated similarly. Now if there exist PPT adversaries  $\mathcal{A}_I$ ,  $\mathcal{A}_X$  and  $\mathcal{A}_Z$  who can distinguish between the games  $G_1$  and  $G_2$  based on the use of  $F_p^{(I)}$ ,  $F_p^{(X)}$  and  $F_p^{(Z)}$  respectively, we can use these adversaries to break the security of  $F_p$ . Thus,

$$\begin{aligned} |\Pr[G_1 = 1] - \Pr[G_2 = 1]| &\leq \text{Adv}_{F_p^{(I)}, \mathcal{A}_I}^{\text{prf}}(1^\lambda) + \text{Adv}_{F_p^{(X)}, \mathcal{A}_X}^{\text{prf}}(1^\lambda) + \text{Adv}_{F_p^{(Z)}, \mathcal{A}_Z}^{\text{prf}}(1^\lambda) \\ &\leq \text{negl}(\lambda). \end{aligned}$$

**GAME  $G_3$ .** For a search query  $\omega = w_1 \wedge w_2 \wedge \dots \wedge w_n$  in step 1 of  $\Pi.\text{Search}_c$  the client generates  $\Pi.\text{stk}$  using  $\Sigma.\text{stk}$ . In the game  $G_3$ , the challenger proceeds exactly as the security proof of  $\Sigma$  to show the generation of  $\Sigma.\text{stk}$  is indistinguishable from its simulated form. The same argument will hold in this proof as well. Since the difference between  $G_2$  and  $G_3$  is only in the generation of  $\Sigma.\text{stk}$  (using ideal primitives) and the advantage of any such adversary is at most  $\text{Adv}_{\Sigma, \mathcal{A}, \mathcal{S}}^{\text{sse}}$ , hence we have

$$|\Pr[G_2 = 1] - \Pr[G_3 = 1]| \leq \text{Adv}_{\mathcal{A}, \mathcal{S}, \mathcal{L}_\Sigma}^{\Sigma, \text{SSE}} \leq \text{negl}(\lambda).$$

**GAME  $G_4$ .** For a search query  $\omega = w_1 \wedge w_2 \wedge \dots \wedge w_n$ , the challenger next changes the

---

<sup>2</sup>Note that the input to  $F_p^{(I)}$  is  $\text{id}$  along with the key  $k_1$ . However, when called,  $\text{id}||\text{op}$  is passed in place of  $\text{id}$  and hence the simulator will index the list by  $(\text{id}, \text{op})$ .

manner in which  $\text{xtk}_{i,j}$  is generated. Recall,  $\text{xtk}_{i,j} = g^{F_p^{(Z)}(\mathbf{k}_Z, w_1 \| j) \cdot F_p^{(X)}(\mathbf{k}_X, w_i)}$ , where  $j \in [\text{cnt}[w]]$  and  $i \in [2, n]$ . The challenger knows all  $\text{ids}$  that were updated corresponding to the  $s$ -term  $w_1$ . The challenger computes the  $\text{xtag}$  (as computed in steps 2 and 3 of  $\Pi.\text{Update}_C$ ) values corresponding to all the  $\text{ids}$  that were updated for  $w_1$  and all the keywords in the  $x$ -terms. Notice that the challenger can do this using its lists  $I$  and  $X$ . Then the challenger computes the values for  $y_{i,j} = F_p^{(I)}(\mathbf{k}_I, \text{id} \| \text{op}) \cdot F_p^{(Z)}(\mathbf{k}_Z, w \| \text{cnt}[w])$  for all the  $\text{ids}$  updated for  $w_1$  and for all keywords in the  $x$ -terms using the lists  $I$  and  $Z$ . Finally the challenger computes  $\text{xtk}_{i,j} = \text{xtag}_{i,j}^{y_{i,j}^{-1}}$ . Now, the view of the adversary in games  $\mathbf{G}_3$  and  $\mathbf{G}_4$  are identical, and no adversary has any advantage of distinguishing between these two games. Thus,

$$\Pr[\mathbf{G}_3 = 1] = \Pr[\mathbf{G}_4 = 1].$$

**GAME  $\mathbf{G}_5$ .** Next, the challenger samples random values  $y \xleftarrow{\$} \mathbb{Z}_p^*$  and uses them in place of  $y_{\text{cnt}[w]}$ 's in the update operations. We recollect that  $y_{\text{cnt}[w]} = \text{xid} \cdot z_{\text{cnt}[w]}^{-1} \in \mathbb{Z}_p^*$ . Now if  $a, b, c \xleftarrow{\$} \mathbb{Z}_p^*$ , then  $c$  is statistically indistinguishable from  $(a \cdot b) \in \mathbb{Z}_p^*$  for any adversary. The secure PRF  $F_p$  is used to generate  $\text{xid}$  and  $z_{\text{cnt}[w]}$ . So  $\text{xtag}$  and  $z_{\text{cnt}[w]}^{-1}$  are also indistinguishable for any random element from  $\mathbb{Z}_p^*$ . In the Games  $\mathbf{G}_2$ , we have already replaced all uses of  $F_p$  with randomly sampled values from  $\mathbb{Z}_p^*$ . Replacing the value of  $y_{\text{cnt}[w]}$  with a randomly sampled element from  $\mathbb{Z}_p^*$  does not change the advantage of any adversary. Thus,

$$\Pr[\mathbf{G}_4 = 1] = \Pr[\mathbf{G}_5 = 1].$$

**GAME  $\mathbf{G}_6$ .** In game  $\mathbf{G}_6$ , for each update operation, the challenger samples a random value  $\gamma \xleftarrow{\$} \mathbb{Z}_p^*$  and replaces the value of  $\text{xtag}$  in every update operation with the value  $g^\gamma$ . Recall that  $\text{xtag} = g^{F_p^{(I)}(\mathbf{k}_I, \text{id}) \cdot F_p^{(X)}(\mathbf{k}_X, w)}$  where  $F_p$  is a secure PRF. We have already replaced the values of every execution of  $F_p$  with random values from  $\mathbb{Z}_p^*$ . In particular, we already have  $\text{xtag} = g^{\alpha\beta}$  for random values  $\alpha, \beta \in \mathbb{Z}_p^*$ . Upon replacing the values of  $\text{xtag}$  with  $g^\gamma$  if any adversary  $\mathcal{A}$  can detect the change between the games  $\mathbf{G}_5$  and  $\mathbf{G}_6$ , we can construct another adversary  $\mathcal{B}$  which uses  $\mathcal{A}$



as a subroutine to break the `eddh` problem. Due to [80], breaking the `eddh` problem is equivalent to breaking the `ddh` assumption if the number of queries is bounded by  $q = \text{poly}(\lambda)$ . Thus,

$$|\Pr[\mathbf{G}_5 = 1] - \Pr[\mathbf{G}_6 = 1]| \leq \text{Adv}_B^{\text{eddh}} \leq \text{negl}(\lambda).$$

**SIMULATION.** Finally, we demonstrate a simulator that perfectly simulates a transcript, which would be indistinguishable from the transcript in game  $\mathbf{G}_6$ . The simulator does not have the actual queries in  $\Pi$ . Instead, the simulator only has access to the leakage functions  $\mathcal{L}_\Pi$ . The simulation has two parts, one is to simulate the components generated in the execution of  $\Sigma$ , and another is to simulate the components of  $\Gamma$ , which together construct  $\Pi$ . Now the components generated using  $\Sigma$  can be simulated perfectly as  $\Sigma$  is  $\mathcal{L}_\Sigma$ -adaptively secure SSE. We note that  $\mathcal{L}_\Sigma \subseteq \mathcal{L}_\Pi$ .

We demonstrate how to simulate the components generated in the execution of  $\Gamma$ . To simulate update queries, the simulator has access to leakage function  $\mathcal{L}_{\Pi, \text{Update}}^{\text{BPUP}}(\text{op}, w, \text{id}) \preceq \{\text{op}, w\}$ . The simulator knows all the time stamps of update queries. For every update query, the simulator simulates  $y_{\text{cnt}[w]}$  and  $\text{xtag}$  by sampling two uniform random values  $\alpha, \beta$  from  $\mathbb{Z}_p^*$  and provides  $\alpha$  and  $g^\beta$  in place of  $y_{\text{cnt}[w]}$  and  $\text{xtag}$ . This is the same as what is done in game  $\mathbf{G}_6$ . So the simulation is indistinguishable. The simulator also stores the values of  $y_{\text{cnt}[w]}$  and  $\text{xtag}$  sampled for each update in a map `Map` indexed by the timestamps of the queries.

Now to simulate a conjunctive search query  $\omega = w_1 \wedge w_2 \wedge \dots \wedge w_n$ , the simulator simulates  $\text{xtk}_{i,j}$  as follows. From the leakage functions `Updates`( $w$ ) and `TimeDB`( $w$ ), the simulator knows all timestamps of update queries for the  $s$ -term and their corresponding `ids`. From `TimeDB`( $\omega$ ) and `Updates`( $\omega$ ), the simulator knows the timestamps where a keyword in the  $x$ -term matches these `ids`. The simulator uses these timestamps to determine the corresponding  $\text{xtag}_{i,j}$  and  $y_{i,j}$  values to construct

$$\text{xtk}_{i,j} = \text{xtag}_{i,j}^{y_{\text{cnt}_{i,j}}^{-1}}.$$

This step is similar to what is done in game  $\mathbf{G}_6$ . So the simulation is indistinguishable.

Lastly, if two search queries  $\omega_1$  and  $\omega_2$  have different *s-terms* but the same *x-term* and share some common *ids*, then the simulator should simulate the *xtk* such that the same *xtag* is generated for those queries. The simulator knows this by considering  $\text{sp}(\omega_1)$ ,  $\text{Updates}(\omega_1)$ ,  $\text{sp}(\omega_2)$  and  $\text{Updates}(\omega_2)$ . Hence this simulation is also indistinguishable, and we have

$$\Pr [\text{SSEIdeal}_{\mathcal{A}, \mathcal{S}, \mathcal{L}_n}^{\Pi}(1^\lambda, q) = 1] = \Pr[\mathbf{G}_6 = 1].$$

Thus,

$$|\Pr [\text{SSEReal}_{\mathcal{A}}^{\Pi}(1^\lambda, q) = 1] - \Pr [\text{SSEIdeal}_{\mathcal{A}, \mathcal{S}, \mathcal{L}_n}^{\Pi}(1^\lambda, q) = 1]| \leq \text{negl}(\lambda).$$

□

## 5.5 Leakage Analysis

Although we have argued the security of our scheme assuming the standard and maximum possible leakage functions defined in the literature, we also argue about the precise leakage of our scheme. It has been found historically that without a proper leakage analysis SSE has been attacked by exploiting those leakages missed in the analysis. Like leakage miss in dynamic SSE was exploited in [97]. Also, the update leakage of [80] was assumed to be empty. Whereas, in some cases, it leaks the keyword, as shown in [102]. Thus, a thorough leakage analysis is essential for any SSE. In order to achieve this, we analyze the leakage of our scheme through all possible combinations of queries. This novel technique gives confidence in the actual leakage of any SSE scheme. This discussion helps in understanding the consequences of the effect of combining different secure components and the exact leakage of the combination.

The actual leakage of our CDSSE construction is much less than those assumed above in Section 5.4. The following are public at the beginning of  $\Pi$ : the security parameter  $\lambda$ , the description of  $\Pi$ , the set of all possible keywords  $\mathcal{W}$ , the set  $\text{db}$  from

which each document is assigned a unique identifier, the key space  $\mathcal{K}$ , the descriptions of the data-structures  $\Pi.\text{TSet}$  and  $\Pi.\text{XDB}$  to be stored in the server, the descriptions of the PRF families  $F_p^{(I)}(\cdot, \cdot)$ ,  $F_p^{(X)}(\cdot, \cdot)$ ,  $F_p^{(Z)}(\cdot, \cdot)$ .

### Setup Leakage

In the setup phase, we assume that the initial database  $\text{DB}$  that is input to  $\Sigma.\text{Setup}_C$  is empty. If the initial database is not empty, we would outsource an empty database in the setup phase and then update the database using the update protocol. Hence, the initial structures of  $\Pi.\text{TSet}$  and  $\Pi.\text{XDB}$  are already known to the adversary. The server receives as input  $\Pi.\text{EDB} = (\Pi.\text{TSet}, \Pi.\text{XDB})$  and stores it. Since  $\Pi.\text{EDB}$  does not have any data from  $\text{db}$  or  $\mathcal{W}$ , there is no leakage during the setup of  $\Pi$ . In other words,  $\mathcal{L}_{\Pi, \text{Setup}} = \emptyset$ .

After the setup, the server receives additional information with each update or search query. To find the leakage due to a query  $q_i \in \text{QS}$ , we should consider all information that the server had before  $q_i$  and those that it attains due to and after processing  $q_i$ .

### Update Leakage

**Single Update Leakage.** Before an update query for  $(\text{op}, (\text{id}, w))$ , let the state of the client be  $\text{st} = (\text{st}.\sigma_C, \text{st}.\text{cnt})$  while the server has the encrypted database  $\Pi.\text{EDB} = (\Pi.\text{TSet}, \Pi.\text{XDB})$ . During an update operation, the client calls  $\Pi.\text{Update}_C$  and the output  $\Pi.\text{utk} = ((\Sigma.\text{utk}, y_{\text{cnt}[w]}), \text{xtag})$  is sent to the server. So the leakage due to the update operation must be from  $\Pi.\text{utk}$  only. Note that  $\Sigma.\text{utk}$  is an output of  $\Sigma.\text{Update}_C$  and  $(y_{\text{cnt}[w]}, \text{xtag})$  is an output of  $\Gamma.\text{Add}$ . We argue that these two outputs together do not leak any more than  $\mathcal{L}_\Pi$ .

1.  $\Sigma.\text{utk}$  *at most* leaks  $\emptyset$ . We first note that the protocol  $\Pi$  calls  $\Sigma.\text{Update}_C$  only with  $\text{op} = \text{add}$  and never uses  $\text{op} = \text{del}$ . By Definition 2.5.5 and Definition 2.5.6 of forward and backward privacy of  $\Sigma$  respectively, the update token  $\Sigma.\text{utk}$  output by  $\Sigma.\text{Update}_C(\mathcal{K}_s, \sigma_C, \text{op}, (\text{id}, w))$  leaks at most the operation  $\text{op}$  which is

always **add** in  $\Pi$ . However, as for both cases  $\text{op} = \text{add}$  as well as  $\text{op} = \text{del}$  the algorithm only issues an **add** token, so **add** and **del** operations are indistinguishable. Hence, there is no leakage from  $\Sigma.\text{utk}$ .

2. **xtag** *at most leaks*  $\emptyset$ . Recall that,  $x_w = F_p^{(X)}(\mathbf{k}_X, w)$  is fixed for every  $w$ , and  $\text{xid} = F_p^{(I)}(\mathbf{k}_I, \text{id} \parallel \text{op})$  is fixed for every  $(\text{id} \parallel \text{op})$ . In our setting for non-modifiable documents, the tuple  $(\text{op}, w, \text{id})$  is always unique for every update operation (as  $(\text{id}, w)$  is added or deleted only once). The pseudo-randomness of  $F_p$  ensures that  $x_w$  and  $\text{xid}$  for a single update are two random elements from the group  $G$ . Thus **ddh** hardness of  $G$  (Definition 2.2.2) ensures that no efficient adversary can distinguish the **xtag** which is  $g^{x_w \cdot \text{xid}}$  from a random element of  $G$ .
3.  $y_{\text{cnt}[w]}$  *at most leaks*  $\emptyset$ . The value  $\text{cnt}[w]$  changes (incremented by 1) for every update operation corresponding to every keyword  $w$  and never repeats. As a result,  $(w \parallel \text{cnt}[w])$  is always unique. The pseudo-randomness of  $F_p$  ensures that  $z_{\text{cnt}[w]}$  which is  $F_p^{(Z)}(\mathbf{k}_Z, w \parallel \text{cnt}[w])$  and consequently  $y_{\text{cnt}[w]}$  which is  $\text{xid} \cdot z_{\text{cnt}[w]}^{-1}$  are random elements from  $G^3$ . Thus,  $y_{\text{cnt}[w]}$  also leaks nothing.
4. *Together,  $((\Sigma.\text{utk}, y_{\text{cnt}[w]}), \text{xtag})$  at most leak*  $\emptyset$ .  $\Sigma.\text{utk}$  is generated from an independent protocol from  $(y_{\text{cnt}[w]}, \text{xtag})$ . Thus,  $\Sigma.\text{utk}$  together with  $(y_{\text{cnt}[w]}, \text{xtag})$  leaks nothing. Additionally,  $y_{\text{cnt}[w]}$  and  $\text{xtag}$  are two random elements of the group  $G$ . So there is no leakage altogether.

This concludes that the leakage of a single update operation is  $\emptyset$ .

### Update-Update Leakage.

1.  $\Sigma.\text{utk}$  *of many updates leaks*  $\emptyset$ . The leakage from an update of SKSSE is  $\emptyset$  even when the server sees many updates. This ensures that  $\Sigma.\text{utk}$  of our scheme for two or more updates is also  $\emptyset$ .

---

<sup>3</sup>A formal argument provided in Game  $G_5$  of Theorem 5.4.1

2.  $y_{\text{cnt}[w]}$  at most leaks  $\emptyset$ . For every update operation, the input  $w \parallel \text{cnt}[w]$  to  $F_p^{(Z)}$  is always unique. The pseudo-randomness of  $F_p$  ensures that  $y_{\text{cnt}[w]}$  in several updates are indistinguishable from a random element of  $G$ .
3.  $\text{xtag}$  at most leaks  $\emptyset$ . From the above discussion, we see that for any two update operations, both  $x_w$  and  $\text{xid}$  cannot be the same. Moreover, the server never gets to see  $x_w$  or  $\text{xid}$  in clear in any operation. When the server sees many update operations, it might happen that two updates use the same  $x_w$  (when a keyword is the same) or the same  $\text{xid}$  (when  $\text{id}$  and  $\text{op}$  are the same). However, both  $x_w$  and  $\text{xid}$  would never be the same for two update operations. Hence,  $\text{eddh}$  hardness of  $G$  (Definition 2.2.3) ensures that no efficient adversary can distinguish the  $\text{xtag}$  which is  $g^{x_w \cdot \text{xid}}$  from a random element of  $G$  even when the server sees many  $\text{xtags}$  either with same  $x_w$  or with same  $\text{xid}$ <sup>4</sup>.
4. Together,  $((\Sigma.\text{utk}, y_{\text{cnt}[w]}), \text{xtag})$  at most leak  $\emptyset$ . From the games  $G_1$ ,  $G_5$  and  $G_6$  in the Proof of Theorem 5.4.1, we know that each of these values can be simulated independently of each other. Since the adversary cannot distinguish between any of these games (except with negligible probability), the server does not get any additional information from this tuple, other than what it does from each of its elements individually.

**Search-Update Leakage.** Consider the following sequence of two queries.

1. The current update query  $q_\beta$  at time instance  $\beta$  is,  $q_\beta = (\beta, \text{op}_\beta, (\text{id}_\beta, w_\beta))$ , that is input to  $\Pi.\text{Update}_C$  and outputs  $((\Sigma.\text{utk}_\beta, y_\beta), \text{xtag}_\beta)$ , where  $\text{op}_\beta \in \{\text{add}, \text{del}\}$ .
2. For  $\alpha < \beta$ , a previous search query  $q_\alpha = (\alpha, \text{srch}_\alpha, \omega_\alpha)$ , where  $\omega_\alpha = w_1 \wedge w_2 \wedge \dots \wedge w_n$ , that is input to  $\Pi.\text{Search}_C$  and outputs  $(\Sigma.\text{stk}_\alpha, \Gamma.\text{xtk}_\alpha)$ .

We argue that the leakage due to the outputs of  $q_\alpha$  and  $q_\beta$  in each case is no more when combined. Both  $\Sigma.\text{utk}_\beta$  and  $\Sigma.\text{stk}_\alpha$  are outputs to the server for the underlying single keyword search protocol  $\Sigma$ . Hence the leakage due to the combination of these

---

<sup>4</sup>The argument is similar to the argument in Game  $G_6$  of Theorem 5.4.1.

two parameters is no more than has been accounted for in  $\mathcal{L}_{\Sigma, \text{Update}}$ . Now, as we only invoke  $\Sigma$  with the **add** operation, there is no leakage.

In the proof of Theorem 5.4.1, we have argued in Game  $G_5$  that the values  $y_j \leftarrow F_p^{(I)}(\mathbf{k}_I, \text{id} \parallel \text{op}) \cdot (F_p^{(Z)}(\mathbf{k}_Z, w \parallel j))^{-1}$ , are essentially random elements of  $G$  to the adversary. Due to our setting, the tuple  $(\text{op}, w, \text{id})$  received in  $q_\beta$  has never appeared in any previous update. This ensures that  $y_\beta$  received in the current update operation when combined with  $\text{xtk}_\alpha$  received during the search operations at time instance  $\alpha$ , does not generate an  $\text{xtag} = \text{xtk}_\alpha[i, j]^{y_\beta}$  that is present in XDB. Similarly,  $\text{xtag}_\beta$  received in  $q_\beta$  is also not present in XDB. Thus,  $\text{xtag}_\beta$  and  $y_\beta$  received in the current update combined with any previous search query  $q_\alpha$  do not leak any more than the leakage after the current update. All of the above together show that the update leakage for our scheme is  $\mathcal{L}_{\Pi, \text{Update}} = \emptyset$ .

## Search Leakage

**Single Search Leakage.** For a conjunctive search query  $\omega = w_1 \wedge w_2 \wedge \dots \wedge w_n$ ,  $w_1$  is the *s-term* and the search token is  $(\Pi.\text{stk}, \Pi.\text{xtk})$ .

1. From  $\Pi.\text{stk} = \Sigma.\text{stk}$ , the server will know  $\text{sp}(w_1)$ ,  $\text{TimeDB}(w_1)$  and  $\text{Updates}(w_1)$ , as these are leakages of  $\Sigma$  (see Definition 2.5.6). Note that  $\text{DelHist}(w_1)$  shall not be leaked as we call  $\Sigma.\text{Update}_c$  with **op** = **add**. The actual update operation is appended to the **id**. So even if the **id** in two updates are the same,  $(\text{id} \parallel \text{op})$  is always different for the same keyword. Thus, if a keyword is added and deleted from the same identifier, the **id** in both cases would be different, and the server cannot link the add and delete operations.
2. Next, the server computes

$$\text{xtk}[j, i]^{y_j} = g^{F_p^{(Z)}(\mathbf{k}_Z, w_1 \parallel j) \cdot x_{w_i} \cdot (\text{id} \cdot F_p^{(Z)}(\mathbf{k}_Z, w_1 \parallel j))^{-1}},$$

for all  $j \in [\text{cnt}[w_1]]$  and  $i \in [2, n]$ , using  $\Pi.\text{xtk}$  output by  $\Gamma.\text{Search}$ , where  $x_{w_i} = F_p^{(X)}(\mathbf{k}_X, w_i)$ . The server then checks if the generated  $\text{xtag}_{i,j} = \text{xtk}[j, i]^{y_j}$  is in XDB or not.

- (a) If  $\text{xtag}_{i,j} \in \text{XDB}$ , the server knows that the  $\text{xid} = F_p(\text{k}_X, \text{id} \parallel \text{op})$  that was added for the  $s$ -term must have also been added for the  $x$ -term. Thus, the server comes to know those timestamps when the pair  $(\text{op}, \text{id})$  in the  $s$ -term matches other keywords in the conjunction.
- (b) Otherwise,  $\text{xtag}_{i,j} \notin \text{XDB}$ . Our non-modifiability setting ensures that this event will not occur in future updates. So the  $\text{xtag}_{i,j}$  computed as  $\text{xtk}[j, i]^{y_j}$  which is not in  $\text{XDB}$ , does not convey any information other than the result of the conjunction to the server.

Thus, the whole leakage is subsumed by the leakage through  $\text{Updates}(\omega)$  and  $\text{TimeDB}(\omega)$ .

Here we note that  $\text{DelHist}(\omega)$  is never leaked as we only add an  $(\text{id}, w)$  pair using  $\Sigma$ , even in case of deletion for the protocol  $\Pi$ . So  $\Sigma$  only works with  $\text{op} = \text{add}$ , and the actual operation remains hidden from  $\Sigma$  and hence from the server.

We note that  $\text{TimeDB}(\omega)$  is a huge overestimation of the actual leakage of  $\Pi$ . The actual leakage of  $\text{TimeDB}(\omega)$  in our scheme is the following.

$$\overline{\text{TimeDB}}(\omega) = \text{TimeDB}(w_1) \cup \{(u, \text{id}) : (i \geq 2) \wedge ((\text{id}, *) \in \text{TimeDB}(w_1)) \wedge ((u, \text{add}, (\text{id}, w_i)) \in \text{QS})\}.$$

In other words, the timestamps of those identifiers are leaked for  $x$ -terms which are only present in the  $s$ -term.

**Search-Search Leakage.** For  $\alpha < \beta$ , consider two search queries  $q_\alpha = (\alpha, \text{srch}_\alpha, \omega_\alpha)$  and  $q_\beta = (\beta, \text{srch}_\beta, \omega_\beta)$ . Any search token  $\Pi.\text{xtk}$  contains  $\text{xtk}[j, i] = g^{F_p^{(Z)}(\text{k}_Z, w_1 \parallel j) \cdot F_p^{(X)}(\text{k}_X, w_i)}$  for  $j \in [\text{cnt}[w_1]]$  and  $i \in [2, n]$ . Now say,  $\omega_\alpha$  and  $\omega_\beta$  have the same  $s$ -terms. So the  $\text{xtks}$  are the same whenever  $w_1$  and  $w_i$  are repeated in the two searches. Hence, the server learns if a pair  $(w_1, w_i)$  has appeared in multiple conjunctive search queries through the equality of  $\text{xtks}$ . This is true for any  $X = w_{i_1} \wedge w_{i_2} \wedge \dots \wedge w_{i_\ell}$ , where  $w_{i_j} \in \tilde{\mathcal{L}}_\omega$  for  $j \in [\ell]$ . That is, for the least frequent keyword  $w_1$  and an arbitrary conjunction

of set of keywords in  $\tilde{\mathcal{L}}(\omega)$ , the server comes to know the joint-search pattern

$$\text{sp}(w_1, X) = \{u \mid (u, \text{srch}, \omega) \in \text{QS}, \omega = w_1 \wedge \dots \wedge X \wedge \dots\}.$$

Hence, for a conjunctive query  $\omega = w_1 \wedge \dots \wedge w_n$ , the total search pattern leakage of our scheme is

$$\text{sp}(\omega) = \text{sp}(w_1) \cup \left( \bigcup_{i=2}^n \text{sp}(w_1, w_i) \right).$$

**Update-Search Leakage.** Consider a search query  $q_\beta = (\beta, \text{srch}_\beta, \omega_\beta)$ . For  $\alpha < \beta$ , consider a previous update query  $q_\alpha = (\alpha, \text{op}_\alpha, (\text{id}_\alpha, w_\alpha))$ , where  $\text{op}_\alpha \in \{\text{add}, \text{del}\}$  and  $\omega_\beta = w_1 \wedge w_2 \wedge \dots \wedge w_n$ . If  $w_\alpha \notin \mathcal{L}(\omega_\beta)$ , that is, if a keyword that was updated previously is not present in the current search query, then such update and search queries are unrelated and hence there is no additional leakage. If  $w_\alpha \in \mathcal{L}(\omega_\beta)$ , the leakage is precisely that captured by the single search leakage described above, as an update operation does not leak anything.

All of the above together show that the search leakage for our scheme is subsumed by the following leakage we have assumed for the proof of Theorem 5.4.1.

$$\mathcal{L}_{\Pi, \text{Search}} = (\text{sp}(\omega), \overline{\text{TimeDB}(\omega)}, \text{Updates}(\omega)).$$

**Leakage Profile of  $\Pi$ .** We finally summarise the overall leakage of our CDSSE scheme. By combining the leakage of  $\Sigma$  and  $\Gamma$ , we indeed get the leakage of  $\Pi$  that is subsumed by the following leakage.

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

**Remark 1.** As mentioned in [95], for a file injection attack [97] to work properly, the adversary must know the results of all the sub-queries of the form  $w_i \wedge w_j$  for all  $w_i, w_j$  in the conjunction and  $1 \leq i < j \leq n$ , that is the KPRP leakage. Although our scheme does not prevent such leakage, the precise leakage of our scheme is  $w_1 \wedge w_i$



for all  $w_i$  in the conjunction. Moreover, as shown by [80], random shuffling of  $x$ tags reduces the success probability of the file injection attack like [97] to almost 5% even when the amount of injected files is 60% of the entire database.

**Remark 2.** *It is important to note that our  $x$ tag is deterministic. Consider an update operation  $(\text{op}, (\text{id}, w_1))$  while  $(\text{op}, (\text{id}, w_i))$  has not occurred yet. The  $x$ tag that would be generated for the search  $\omega = w_1 \wedge \dots \wedge w_i \wedge \dots \wedge w_n$  would not be present in the XDB. Now, if at a later stage, the update  $(\text{op}, (\text{id}, w_i))$  indeed happens, then the  $x$ tag for that update would be exactly the same  $x$ tag that server has computed at the time of the search operation  $\omega = w_1 \wedge \dots \wedge w_i \wedge \dots \wedge w_n$ . Thus, the server would be able to link the update operation  $(\text{op}, (\text{id}, w_i))$  with the search operation  $\omega = w_1 \wedge \dots \wedge w_i \wedge \dots \wedge w_n$ , breaking the forward privacy of  $\Pi$ . The above event resembles that at the time of adding the document identifier  $\text{id}$ ,  $w_1$  was present in that document and at a later stage  $w_i$  was added to the document. However, if an update  $(\text{op}, (\text{id}, w_1))$  has occurred and,  $(\text{op}, (\text{id}, w_i))$  has not, in that case, our setting of non-modifiable documents ensures that this operation would never happen in any future update, preventing such kind of forward privacy breach. The above is the reason why  $\Pi$  is not secure in the general case of SSE. This was missed by [80].*

**Remark 3.** *From the concrete leakage analysis it is quite evident that, even if our construction can be instantiated with a BPIP secure SKSSE, the resulting CDSSE is not BPIP secure. The reason is that even if we store the update for every keyword with a BPIP secure scheme, the  $x$ tags are stored in a set membership query data structure. Now, at the time of a search operation, when an  $x$ tag generated for any keyword in the conjunction is found in the set, it reveals the time when the element was added to the set, degrading the whole security to BPUP. However, if our scheme is instantiated using a WBP secure single keyword construction, the security is enhanced to BPUP. The reason is, in the update phase we always call our single keyword SSE with  $\text{op} = \text{add}$ , and the actual operation is embedded with the  $\text{id}$  (as we replace  $\text{id}$  with  $\text{id}||\text{op}$ ). So the server is unaware of the actual operation as the server never sees the original identifiers. Thus the server cannot link the add operation to the delete*

operation as it cannot distinguish between these two operations. Achieving BPIP secure CDSSE is still an open problem.

## 5.6 Implementation and Experimental Results

We implemented three instantiations of our generic construction with three different DSSEs [87, 49, 38]. Our code is available at [77]. We conducted experiments on the Enron Email database[45] to evaluate the performances of these three implementations. All three schemes performed very well. Our implementation proved to be very practical and scalable.

**Implementation Details.** We have instantiated our generic scheme with three different SKSSE schemes – one of which is only forward private [87], and two are state-of-the-art forward and backward private schemes [49] and [38]. The schemes were implemented in C++. The code is available at [77]. The basic codes for these experiments were written by Theo Henault.

All the experiments were performed on a desktop computer with Intel Core i7-6700 CPU, 3.40 GHz with 8 Cores and 16GB memory. The database was stored using RocksDB [44], and all cryptographic primitives were implemented using Crypto++<sup>5</sup>. We report our results in a similar fashion as reported in [29, 80, 100]. We have used the Enron Email Dataset[45] for our experiments. In [87, 38], when two consecutive queries have the same keyword, the server stores the previous search result to optimise the search. However, for a fair comparison of the time taken by the three schemes, we have not implemented this optimisation – all times reported are for fresh search of the *s-terms*. We do not report our search results with different probabilities of search operation. However, the modification suggested in [87, 38] can also be applied for single keyword searches, which will improve the search performance of the corresponding implementations.

---

<sup>5</sup>Available at: <https://www.cryptopp.com/>

**EDB Creation and Update Operation.** The EDB was created using the Enron Email Dataset. We wrote a C++ program to extract IDs and keywords from the dataset. We got 517,401 IDs and 72,252,961 ID-keyword pairs. Table 5.2 shows the sizes of the encrypted databases and clients’ state sizes for different schemes. We also report the total time and the time per pair, required to create the EDB. The EDB creation essentially calls the update process of our protocol (without the network delay), so the average update time per pair was calculated while generating the EDB and is reported in Table 5.2. The update time for an id-keyword pair is almost constant in all three instantiations. We have performed an independent experiment by inserting  $10^4$  pairs and taking the average time. The experiment also matches the result reported in Table 5.2.

Schemes	Time (Hr.)	Time/Pair (ms)	Server Storage (GB)	Client Storage (MB)
FAST-OXT	12.6	0.63	15.5	20
Mitra-OXT	10.4	0.52	9.1	2
$\Pi_{BP}$ -OXT	11.1	0.55	13.3	3.1

Table 5.2: Storage cost and time required for each update.

**Search Operation.** To evaluate the performance of search operations, we follow the same setting as in [80], and [100]. We report the time for search queries that are conjunctions of two keywords ( $w_1 \wedge w_2$ ). We have conducted two types of search operations,

1. In the first experiment, we keep the size of the number of updates performed on the second keyword to a constant  $|\text{Updates}(w_2)| = 10^5$  and vary  $|\text{Updates}(w_1)|$  from 100 to  $10^5$  (Figure 5-4). The exact values are in Table 5.3
2. In the second case, we keep the size of the number of updates performed on the first keyword to a constant, that is,  $|\text{Updates}(w_1)| = 10$  and vary  $|\text{Updates}(w_2)|$  from 10 to  $10^5$  (Figure 5-5).

In both cases, we report the search time. The time required to search two keywords both having order of  $10^5$  updates takes  $\sim 80$  sec for FAST-OXT,  $\sim 60$  sec for Mitra-OXT and  $\sim 52$  sec for  $\Pi_{BP}$ -OXT (Figure 5-4). Our results concur with the results




Figure 5-4: Search time comparison between the three instantiations of our generic CDSSE for fixed  $|\text{Updates}(w_2)| = 10^5$  with varying  $|\text{Updates}(w_1)|$  from  $10^2$  to  $10^5$ .




Figure 5-5: Search time comparison between the three instantiations of our generic CDSSE for fixed  $|\text{Updates}(w_1)| = 10$  with varying  $|\text{Updates}(w_2)|$  from 10 to  $10^5$ .

from [38], as Mitra-OXT outperforms  $\Pi_{BP}$ -OXT when the search output size is small and  $\Pi_{BP}$ -OXT outperforms Mitra-OXT for large search output sizes. The search time is high for FAST-OXT as decryption is done at the server.

We did not perform searches with more than 2 keywords in conjunction as the search for elements on XDB can be *parallelised* and hence will essentially take similar time as a conjunctive query with 2 keywords.

The work of [100] uses an SSE to store the entries of *s-term* and *x-term* as well. The implementation of [100] is not available. However, as reported in their paper, they did their experiment on a Ubuntu 20.04.3 LTS workstation with Intel @Xeon(R) W-2123 CPU 3.60GHz with 8 cores and 32GB RAM. Their setup uses more process-

ing power with a double-sized RAM. As reported in their paper, their protocol has a higher running time when the number of updates for both  $w_1$  and  $w_2$  is high. This is understandably so, as for each  $s$ -term and  $x$ -term they have to access an SSE. However, their running time overshoots our scheme by a huge margin when the number of updates for the  $s$ -term is small but the number of updates for the  $x$ -term is high. This is because we only fetch very few  $x$ -terms (precisely the number of times the  $s$ -term has been updated) whereas they have to fetch all the  $x$ -terms related to the keyword. Similarly, as reported in [95], they only test their scheme for a database with 23,643 documents, 60,879 keywords, and 8,373,977 keyword-document pairs, which is almost a magnitude smaller than our test case. Their server machines were running Ubuntu 18.04 LTS and had 16 cores (Intel Core i9-9900 CPU 3.10 GHz), 31GB RAM, and 483 GB SSD disk space. However, their running time is also similar to our results when the matching search output size is small and their search time grows rapidly when the matching search output size increases, making their scheme much less suitable for large databases compared to ours.

$ \text{Updates}(w_1) $	Mitra-OXT (s)	FAST-OXT (s)	$\Pi_{\text{BP-OXT}}$ (s)
$10^2$	0.06	0.06	0.13
$10^3$	0.57	0.54	0.53
$10^4$	6.01	6.15	5.31
$10^5$	59.79	80.41	52.49

Table 5.3: Search time comparison for fixed  $|\text{Updates}(w_2)| = 10^5$  with varying  $|\text{Updates}(w_1)|$  from  $10^2$  to  $10^5$

## 5.7 Final Remarks

We proposed a generic dynamic searchable symmetric encryption scheme that allows conjunctive queries while being both forward and backward private. Our scheme is specifically crafted to protect *non-modifiable* documents that, once uploaded, their keywords remain unaltered. This is a common characteristic in a variety of applications, including biometric databases, digital archives, legal documents, medical records, etc. Our construction allows the keyword set associated with a document

to be modified with some extra overhead. With efficiency closest to OXT (among known dynamic schemes) and by ensuring both forward and backward privacy, our proposed construction offers a viable solution for preserving the privacy of dynamic datasets while facilitating conjunctive queries.

The OXT framework has been extended [29] to support general boolean queries and to the multi-client setting [57] in a modular black-box fashion. Both these extensions facilitate a lot of practical use cases. Our construction uses OXT in a black-box fashion as well, where the functionality of OXT has not been altered. Hence, the extension of our generic scheme to support general boolean queries and multi-clients seems quite plausible for future work.

## Updatable Message Authentication Codes

Our next goal is to construct a fault-tolerant, verifiable DSSE. The main tool that we use for our construction is a message authentication code with some added functionality which we call an updatable message authentication code (UdMAC). In this Chapter, we systematically develop UdMACs, and in the next chapter, we use them to construct a fault-tolerant verifiable DSSE.

Message authentication codes (MACs) are algorithms that have been traditionally used to provide integrity of a message. A MAC scheme is a pair of algorithms (*MAC generation algorithm*, *Verification algorithm*). On a given message and a key *MAC generation algorithm* outputs a *tag*. The sender sends this *tag* along with the message. Given a message, the key and a *tag*, the *Verification algorithm* outputs either 0 or 1, where 1 indicates that the message is authentic and 0 indicates the message is not authentic. There are several paradigms for constructing MACs. For our purpose, we need to extend the basic functionality of a MAC. We consider a scenario where a client outsources its message  $M$  to be stored in a server. To ensure the integrity of the message, the client computes a MAC  $t$  for the message and transmits  $M$  along with  $t$  to be stored in the server. Over time, if the client needs to update  $M$ , then (s)he has to download the message  $M$  update it, re-compute the MAC and then send it back to the server. If the message is large and frequent updates are required, this naive solution may be inefficient. We propose a new type of MAC, which we call *updatable MAC* (UdMAC), which has the property that for updates, the client does not need to download the message and recompute the MAC. Suppose the message

$M$  is to be updated by  $m$  and we denote the updated message by  $M\Delta m$ ; the client should be able to compute the update MAC  $t_u$  for  $M\Delta m$  using only  $m$  and  $t$ . Thus, the client only sends the update  $m$  along with  $t_u$  to the server, and the server stores the updated message  $M\Delta m$  and the updated tag  $t_u$ .

This specialized functionality provided by UdMACs is not available in traditional MACs. We define the syntax and security of UdMACs, and provide two constructions, namely, **ConCatU** and **XoRU**. The two constructions work for two different types of updates, concatenation and xor difference, respectively. We analyze in detail both the constructions and prove their security.

There is a long history of cryptographic schemes with updating capabilities starting from the work reported in [12, 13], where the concepts of incremental collision resistant hash functions and signatures were introduced. These ideas were further refined and formalized in [10]. Another line of work focuses on key updatable encryption. A key updatable encryption scheme is a symmetric encryption scheme that allows the key holder to update keys and to compute an update token, which can be given to a party storing ciphertexts, and can be used to update existing ciphertexts to ones under the new key [19, 18, 46]. Studies on key updatable signatures and message authentication codes have been reported in [39]. In [39] the key updatable MACs are called updatable MACs (UMAC). We do not consider key updates, but updates in the message. We also call our MACs updatable MACs but use a different acronym, UdMAC. Our work aligns more with the scenarios considered in [12, 13, 10] but our formulations and constructions are different.

## 6.1 Message Authentication Codes

A message authentication code (MAC) is a map  $F : \mathcal{K} \times \mathcal{M} \rightarrow \{0, 1\}^\tau$ , where  $\mathcal{K}$  is the key space and  $\mathcal{M}$  the message space. We often write  $F_K(\cdot)$  to denote  $F(K, \cdot)$ . The output of a MAC is called the tag, and  $\tau$  is called the tag length. The security of a MAC  $F$  is defined using an interaction of  $F$  with an adversary  $\mathcal{A}$ . It is assumed that  $\mathcal{A}$  has an oracle access to  $F_K(\cdot)$ , where  $K \xleftarrow{\$} \mathcal{K}$ . For a query  $x \in \mathcal{M}$  of  $\mathcal{A}$  the oracle



responds by sending  $y = F_K(x)$ . Let,  $\mathcal{A}$  queries  $x_1, x_2, \dots, x_q$  and gets  $y_1, y_2, \dots, y_q$  as responses from the oracle. These queries are performed adaptively. Finally,  $\mathcal{A}$  outputs a pair  $(x^*, y^*)$ , where  $x^* \notin \{x_1, x_2, \dots, x_q\}$ . This pair is called a forgery and it is said that  $\mathcal{A}$  has successfully *forged*  $F$  if  $F_K(x^*) = y^*$ . The **auth**-advantage of  $\mathcal{A}$  is defined as

$$\text{Adv}_F^{\text{auth}}(\mathcal{A}) = \Pr[K \xleftarrow{\$} \mathcal{K} : \mathcal{A} \text{ forges}].$$

We say that  $F$  is  $(\epsilon, t)$  secure if for every adversary  $\mathcal{A}$ , which runs for time at most  $t$ ,  $\text{Adv}_F^{\text{auth}}(\mathcal{A}) \leq \epsilon$ . It is well known [20] that for any arbitrary adversary  $\mathcal{A}$  for the MAC  $F$  there exists a PRF adversary  $\mathcal{B}$  for  $F$  such that

$$\text{Adv}_F^{\text{auth}}(\mathcal{A}) \leq \text{Adv}_F^{\text{prf}}(\mathcal{B}) + \frac{1}{2^\tau}, \quad (6.1)$$

where  $\mathcal{B}$  and  $\mathcal{A}$  both run almost for the same time and ask almost the same number of queries.

## 6.2 Polynomial MACs

Let  $\mathbb{F}_{2^n}$  denote the field with  $2^n$  elements,  $n$ -bit strings can be seen as elements in  $\mathbb{F}_{2^n}$ . Thus additions and multiplication of  $n$ -bit strings are addition and multiplications in  $\mathbb{F}_{2^n}$ . For  $x, y \in \{0, 1\}^n$ , we denote their addition and multiplication by  $x \oplus y$  and  $xy$  respectively.

For  $x \in \{0, 1\}^*$  and  $h \in \{0, 1\}^n$ , let  $(x_1, x_2, \dots, x_m) = \text{parse}(x)$ . We define the function

$$\text{Poly}_h(x) = x_1 h \oplus x_2 h^2 \oplus \dots \oplus \text{Pad}_n(x_m) h^m$$

where the additions and multiplications are in  $\mathbb{F}_{2^n}$  and for  $y \in \{0, 1\}^*$ ,  $|y| \leq n$ ,  $\text{Pad}_n(y) = y \parallel 0^{n-|y|}$ .

A related function is

$$\text{PPoly}_h(x) = \text{Poly}_h(x) \oplus |x| h^{m+1}.$$

The following result is well known:

**Proposition 6.2.1.** *Let  $x_1, x_2 \in \{0, 1\}^*$ ,  $x_1 \neq x_2$  and let  $\ell_1 = \lceil |x_1|/n \rceil$  and  $\ell_2 = \lceil |x_2|/n \rceil$  then for any  $c \in \{0, 1\}^n$*

$$\Pr[h \xleftarrow{\$} \{0, 1\}^n : \text{PPoly}_h(x_1) \oplus \text{PPoly}_h(x_2) = c] \leq \frac{\max\{\ell_1, \ell_2\} + 1}{2^n}.$$

The nonce based Wegman Carter MAC [26, 92] is a widely used MAC algorithm. A popular variant of this MAC takes in a message  $x \in \{0, 1\}^*$ , two keys  $K, h \in \{0, 1\}^n$  and a nonce  $N \in \{0, 1\}^n$  and produces an  $n$ -bit tag as

$$\text{WC}_{K,h}(x; N) = \text{PPoly}_h(x) \oplus E_K(N), \quad (6.2)$$

where  $E_K()$  is a pseudo-random function family.

The Theorem below regarding the security of Wegman Carter MAC is well known (for example see [20]).

**Theorem 6.2.1.** *Let Wegman Carter MAC be as defined in Equation (6.2) and  $\mathcal{A}$  be a nonce respecting<sup>1</sup> forgery adversary who asks at most  $q$  queries. Among the  $q$  queries, let the message queried of the largest length be of length  $\ell$ . Then there exists a PRF adversary  $\mathcal{B}$  for  $E_K()$ , such that*

$$\text{Adv}_{\text{WC}}^{\text{auth}}(\mathcal{A}) \leq \text{Adv}_E^{\text{prf}}(\mathcal{B}) + \frac{\ell + 1}{2^n}.$$

We will need the following result regarding the function  $\text{Poly}_h()$ .

**Proposition 6.2.2.** *Let  $x \in \{0, 1\}^{n\ell_1}$  and  $y \in \{0, 1\}^{n\ell_2}$ , and let  $(x_1, x_2, \dots, x_{\ell_1}) = \text{parse}_n(x)$  and  $(y_1, y_2, \dots, y_{\ell_2}) = \text{parse}_n(y)$  and  $x_{\ell_1} \neq 0^n$  and  $y_{\ell_2} \neq 0^n$  and  $x \neq y$ . Then, for any  $c \in \{0, 1\}^n$ ,*

$$\Pr[h \xleftarrow{\$} \{0, 1\}^n : \text{Poly}_h(x) \oplus \text{Poly}_h(y) = c] \leq \frac{\max\{\ell_1, \ell_2\}}{2^n}.$$

---

<sup>1</sup>A nonce respecting adversary is an adversary who never repeats a nonce in its queries.

*Proof.* Without loss of generality, let  $\ell_1 \geq \ell_2$ . Thus, we have

$$\begin{aligned}\text{Poly}_h(x) &= x_1h \oplus x_2h^2 \oplus \cdots \oplus x_{\ell_2}h^{\ell_2} \oplus \cdots \oplus x_{\ell_1}h^{\ell_1} \\ \text{Poly}_h(y) &= y_1h \oplus y_2h^2 \oplus \cdots \oplus y_{\ell_2}h^{\ell_2}.\end{aligned}$$

If  $\ell_1 = \ell_2$  we have

$$\text{Poly}_h(x) \oplus \text{Poly}_h(y) = (x_1 \oplus y_1)h \oplus (x_2 \oplus y_2)h^2 \oplus \cdots \oplus (x_{\ell_1} \oplus y_{\ell_1})h^{\ell_1},$$

which is a non zero polynomial of degree at most  $\ell_1$  as  $x \neq y$ . If  $\ell_1 > \ell_2$  we have

$$\text{Poly}_h(x) \oplus \text{Poly}_h(y) = (x_1 \oplus y_1)h \oplus (x_2 \oplus y_2)h^2 \oplus \cdots \oplus (x_{\ell_2} \oplus y_{\ell_2})h^{\ell_2} \oplus \cdots \oplus x_{\ell_1}h^{\ell_1},$$

which is a non zero polynomial of degree at most  $\ell_1$  as  $x_{\ell_1} \neq 0^n$ .

Thus, in both cases we have  $\text{Poly}_h(x) \oplus \text{Poly}_h(y) \oplus c = 0$  have at most  $\ell_1$  roots in  $\mathbb{F}_{2^n} = \{0, 1\}^n$ .

Similar argument holds if  $\ell_2 > \ell_1$  and in that case we have  $\text{Poly}_h(x) \oplus \text{Poly}_h(y) \oplus c = 0$  have at most  $\ell_2$  roots in  $\mathbb{F}_{2^n} = \{0, 1\}^n$ . Thus we have

$$\Pr[h \stackrel{\$}{\leftarrow} \{0, 1\}^n : \text{Poly}_h(x) \oplus \text{Poly}_h(y) = c] \leq \frac{\max\{\ell_1, \ell_2\}}{2^n}.$$

□

## 6.3 Updatable MACs

An updatable MAC works on four associated sets  $\mathcal{K}$ ,  $\mathcal{M}$ ,  $\Sigma$  and  $\mathcal{T}$  called the key space, message space, state space and tag space, respectively. We will consider  $\mathcal{M} \subset \{0, 1\}^*$  and an associative binary operation  $\Delta : \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$ , which we will further call as an update function. An updatable MAC  $\Psi$  is a triple of algorithms  $\Psi = (\text{KeyGen}, \text{MacUpdate}, \text{MacVerify})$  described as follows:

**KeyGen**( $1^\lambda$ ): takes as input the security parameter  $\lambda$  outputs a key  $K$ , selected

uniformly at random from the key space  $\mathcal{K}$ .

**MacUpdate**( $K, m, \sigma$ ): takes as input  $K \in \mathcal{K}$ ,  $m \in \mathcal{M}$  and  $\sigma \in \Sigma \cup \{\perp\}$  and outputs a  $t \in \mathcal{T}$  and an updated state  $\tilde{\sigma}$ .

**MacVerify**( $K, m, \sigma, t$ ): takes in the key  $K$ , a  $m \in \mathcal{M}$ ,  $\sigma \in \Sigma$  and  $t \in \mathcal{T}$  and outputs a bit  $b \in \{0, 1\}$ .

**Correctness:** Let  $m_1, m_2, \dots, m_s \in \mathcal{M}$  be a sequence of messages and  $\Psi = (\text{KeyGen}, \text{MacUpdate}, \text{MacVerify})$ . For an updatable MAC  $\Psi$ , we define a procedure  $\Psi.\text{SeqUpdt}$  in Figure 6-1 which takes in a key  $K \in \mathcal{K}$  generated by  $\Psi.\text{KeyGen}(1^\lambda)$ , a sequence of messages  $(m_1, m_2, \dots, m_s)$  and outputs  $(\sigma, t) \in \Sigma \times \mathcal{T}$ . We say that  $\Psi$  is correct, if for any message sequence  $(m_1, \dots, m_s)$  and any  $K$  generated by  $\Psi.\text{KeyGen}$ ,

$$\Psi.\text{MacVerify}(K, M, \Psi.\text{SeqUpdt}(K, (m_1, m_2, \dots, m_s))) = 1,$$

where  $M = m_1 \Delta m_2 \Delta \dots \Delta m_s$ .

```

Ψ.SeqUpdt( $K, (m_1, m_2, \dots, m_s)$ )
01.  ( $t_1, \sigma_1$ )  $\leftarrow$   $\Psi.\text{MacUpdate}(K, m_1, \perp)$ ;
02.  for  $i = 2$  to  $s$ ,
03.    ( $\sigma_i, t_i$ )  $\leftarrow$   $\Psi.\text{MacUpdate}(K, m_i, \sigma_{i-1})$ ;
04.  end for
05.  return ( $\sigma_s, t_s$ );

```

Figure 6-1: The procedure **SeqUpdt** used to define correctness of an updatable MAC scheme  $\Psi$ .

The definition of correctness specifies the usefulness of updatable MAC. Consider the **MacUpdate** procedure is applied on the messages  $m_1, m_2, \dots, m_s$ , one after the other in the same sequence as in the procedure **SeqUpdt** of Figure 6-1. Thus, we start with the message  $m_1$  and an empty state and in the subsequent messages the input to **MacUpdate** is the current message and the state obtained for the previous

message. The final state-tag pair  $(\sigma_s, t_s)$  obtained should be a valid MAC for the message  $M = m_1 \Delta m_2 \Delta \dots \Delta m_s$ .

### 6.3.1 Security of Updatable MACs

For defining security of an updatable MAC  $\Psi = (\text{KeyGen}, \text{MacUpdate}, \text{MacVerify})$ , we give  $\mathcal{A}$  oracle access of  $\Psi.\text{MacUpdate}(K, \cdot, \cdot)$ , instantiated with a key  $K$  generated by  $\Psi.\text{KeyGen}(1^\lambda)$ .  $\mathcal{A}$  can query its oracle with  $(m, \sigma) \in \mathcal{M} \times \Sigma$  and get as response  $(\tilde{\sigma}, t) \in \Sigma \times \mathcal{T}$ . Let  $\mathbb{Q} = \{(m_1, \sigma_1), (m_2, \sigma_2), \dots, (m_q, \sigma_q)\}$  be the set of queries made by  $\mathcal{A}$  and for the  $i^{\text{th}}$  query  $(m_i, \sigma_i)$ , let the response of the oracle be  $(\tilde{\sigma}_i, t_i)$ . We denote the set of messages queried by  $\mathcal{A}$  by  $\mathbb{M} = \{m : (m, \sigma) \in \mathbb{Q}\}$  and we call the set

$$\mathbb{P}_{\mathcal{A}} = \{(\sigma_i, m_i, \tilde{\sigma}_i, t_i) : 1 \leq i \leq q\}$$

as the query profile of the adversary  $\mathcal{A}$ . Finally, after  $\mathcal{A}$  stops querying, it outputs a forgery  $(M^*, \sigma^*, t^*)$ .

We call a forgery  $(M^*, \sigma^*, t^*)$  of an adversary  $\mathcal{A}$  with query profile  $\mathbb{P}_{\mathcal{A}}$  as “invalid” if there exists a sequence of query-responses  $(\sigma^{(1)}, m^{(1)}, \tilde{\sigma}^{(1)}, t^{(1)})$ ,  $(\sigma^{(2)}, m^{(2)}, \tilde{\sigma}^{(2)}, t^{(2)})$ ,  $\dots$ ,  $(\sigma^{(\ell)}, m^{(\ell)}, \tilde{\sigma}^{(\ell)}, t^{(\ell)})$ , such that the following are true simultaneously

1.  $\{(\sigma^{(1)}, m^{(1)}, \tilde{\sigma}^{(1)}, t^{(1)}), (\sigma^{(2)}, m^{(2)}, \tilde{\sigma}^{(2)}, t^{(2)}), \dots, (\sigma^{(\ell)}, m^{(\ell)}, \tilde{\sigma}^{(\ell)}, t^{(\ell)})\} \subseteq \mathbb{P}_{\mathcal{A}}$ .
2.  $m^{(1)} \Delta m^{(2)} \Delta \dots \Delta m^{(\ell)} = M^*$
3.  $\sigma^{(1)} = \perp$ ,  $\tilde{\sigma}^{(\ell)} = \sigma^*$  and  $\sigma^{(i)} = \tilde{\sigma}^{(i-1)}$  for  $1 \leq i \leq \ell$ .

A forgery  $(M^*, \sigma^*, t^*)$  is valid if it is not invalid. An adversary  $\mathcal{A}$  is successful if its forgery  $(M^*, \sigma^*, t^*)$  is valid and  $\Psi.\text{MacVerify}(M^*, \sigma^*, t^*) = 1$ . Let  $\text{Succ}(\mathcal{A})$  be the event that  $\mathcal{A}$  is successful. We define the forging advantage of an adversary  $\mathcal{A}$  for a  $\Delta$ -updatable MAC as

$$\text{Adv}_F^{\text{Uauth}}(\mathcal{A}) = \Pr[\text{Succ}(\mathcal{A})]. \quad (6.3)$$

The probability is taken over the selection of the key  $K$  with which the oracle is instantiated, the (possible) randomness in  $\Psi$  and the (possible) randomness in  $\mathcal{A}$ .

### 6.3.2 Some New Notations

In the subsequent sections we describe two constructions MACs which support two different update functions. The first one is called **ConCatU** which converts any given PRF  $f : \mathcal{K} \times \mathcal{M} \rightarrow \{0, 1\}^n$  to an  $\parallel$ -updatable MAC, i.e., where the update function is concatenation. The second construction converts a Wegman Carter MAC [26] to an  $\oplus$ -updatable MAC, i.e., a UdMAC where the update function is xor differences.

We introduce a bit of new notation for convenience. Let  $(\mathbb{N})^*$  denote the set of all lists whose elements are natural numbers including the empty list. Thus a finite list  $L \in (\mathbb{N})^*$  can be seen as an array, and its  $i^{\text{th}}$  element will be denoted by  $L[i]$ , and its length, i.e., the number of elements present in it, will be denoted by  $\|L\|$ . An empty list is denoted by  $\langle \rangle$ . For an  $a \in \mathbb{N}$ , **append**( $L, a$ ) appends  $a$  at the end of  $L$ . By **sum**( $L$ ) we will mean the sum of all the elements in  $L$ . Let,  $L \in (\mathbb{N})^*$  be a finite length list of length  $\|L\| = \ell$  where  $L = \langle x_1, x_2, \dots, x_\ell \rangle$  and let  $M \in \{0, 1\}^*$  be a string of length **sum**( $L$ ). Then, **Parse**( $M, L$ ) returns  $(m_1, m_2, \dots, m_\ell)$  such that  $|m_i| = L[i]$  and  $m_1 \parallel m_2 \parallel \dots \parallel m_\ell = M$ . For a list  $L = \langle x_1, x_2, \dots, x_\ell \rangle$ , by **str**( $L$ ) we mean the  $n\ell$  bit string  $\text{bin}_n(x_1) \parallel \text{bin}_n(x_2) \parallel \dots \parallel \text{bin}_n(x_\ell)$ . If  $L = \langle \rangle$ ,  $\|L\| = 0$  and **str**( $L$ ) is the empty string.

## 6.4 ConCatU: An Updatable MAC for Concatenation

Let  $M \subset \{0, 1\}^*$ ,  $\mathcal{K} = \{0, 1\}^k \times \{0, 1\}^n$ ,  $\Sigma \subset (\mathbb{N})^* \times \{0, 1\}^n \cup \{\perp\}$ ,  $\mathcal{T} = \{0, 1\}^n$  be the message space, key space, state space and tag space of **ConCatU** respectively. Let  $f : \mathcal{K} \times \mathcal{M} \rightarrow \{0, 1\}^n$  be a *PRF* family. We use  $f$  to construct a  $\Delta$ -updatable MAC **ConCatU**, where  $\Delta$  is a the concatenation function  $\parallel$ .

A state  $\sigma$  of **ConCatU** is either  $\perp$  or it contains two distinct fields denoted as  $\sigma = (L_\sigma, \text{pTag}_\sigma)$ , where  $L_\sigma$  is a finite list of positive integers, and  $\text{pTag}_\sigma$  is a  $n$ -bit binary string.

**ConCatU.KeyGen**() outputs a pair  $(K, h)$ , where  $K \xleftarrow{\$} \{0, 1\}^k$  and  $h \xleftarrow{\$} \{0, 1\}^n$ .

The procedures `ConCatU.MacUpdate` and `ConCatU.MacVerify` are shown in Figure 6-2.

To understand the workings of the procedures `ConCatU.MacUpdate` and `ConCatU.MacVerify` it is important see the role of the state  $\sigma$ .

Consider an initial message  $m_1$ , and let subsequent updates on it be  $m_2, \dots, m_k$ . As the update function is a concatenation function, hence the message after the  $i^{\text{th}}$  update,  $1 \leq i \leq k$ , will be  $M_i = m_1 \| m_2 \| \dots \| m_i$ . We see the initial message also as an update on an empty message. The intended usage of `ConCatU.MacUpdate` is to obtain a correct authentication tag of the updated message after each update without knowing the whole message. To accomplish this, each update  $m$  that is given as an input to the `ConCatU.MacUpdate` procedure is associated with a state  $\sigma$ . The state  $\sigma$  carries some information regarding the message on which the current update  $m$  is to be applied. In particular,  $\sigma$  keeps a record of the history of updates on the message. The procedure produces an authentication tag for the updated message and also updates the state with the information of the current update.

For an input  $(m, \sigma)$ , if  $\sigma = \perp$ , then it signifies  $m$  is the initial message, i.e., an update  $m$  on an empty message is being sought. The `ConCatU.MacUpdate` procedure on such an input should produce a MAC tag for the message  $m$ , which in this case is the updated message. Further, the procedure should populate the initially empty list  $L$  with the length of  $m$  and set the current state to  $(L, \text{tag})$  and output it along with the tag.

When an input  $(m, \sigma)$  is received with  $\sigma \neq \perp$ , then a MAC is being sought for the message  $M \| m$ , where  $M$  is some message which is not being explicitly provided to the algorithm. But, the information regarding  $M$  is encoded in the state  $\sigma = (L, \text{pTag})$ , which carries the following information:

1. The message  $M$ , over which the update is being sought was obtained after  $k = \|L\|$  updates.
2. The initial update on the empty message, i.e., the initial message was of length  $L[1]$  bits.
3. For  $2 \leq i \leq k$ , the  $i^{\text{th}}$  update increased the length of the message by  $L[i]$  bits.

Thus, for a correct usage,  $|M| = \text{sum}(L)$ .

4. **pTag** was the tag computed by `ConCatU.MacUpdate` for the message  $M$ , i.e., the last call to `ConCatU.MacUpdate` for this specific message returned **pTag**.

<pre> ConCatU.MacUpdate((K, h), m, σ) 01.  if σ = ⊥ 02.    ℓ ← 0; 03.    pTag ← 0<sup>n</sup>; 04.    L = []; 05.  else 06.    (L, pTag) ← σ; 07.    ℓ ←   L  ; 08.    X ← 1<sup>n</sup>; 09.    for i = 1 to ℓ, 10.      X ← X    bin<sub>n</sub>(L[i]); 11.    p ← Poly<sub>h</sub>(X); 12.    tag ← f<sub>K</sub>(p    pTag    m); 13.    L ← append(L,  m ); 14.    σ ← (L, tag); 15.    return (σ, tag); </pre>	<pre> ConCatU.MacVerify((K, h), M, σ, t) 21.  if σ = ⊥ 22.    return 0; 23.  else 24.    (L, pTag) ← σ; 25.    if sum(L) ≠  M  or pTag ≠ t, 26.      return 0 27.    ℓ ←   L  ; 28.    (m<sub>1</sub>, m<sub>2</sub>. . . . m<sub>ℓ</sub>) ← Parse(L, m); 29.    X ← 1<sup>n</sup>; tag ← 0<sup>n</sup> 30.    for i = 0 to ℓ 31.      p ← Poly<sub>h</sub>(X); 32.      tag ← f<sub>K</sub>(p    tag    m<sub>i</sub>) ; 33.      X ← X    bin<sub>n</sub>(L[i]); 34.    if tag = t; 35.      return 1; 36.    else return 0; </pre>
--	---

Figure 6-2: Specification of `ConCatU` using a prf  $f_K : \{0, 1\}^* \rightarrow \{0, 1\}^n$ .

As specified in Figure 6-2, the update procedure `ConCatU.MacUpdate` on an input  $(m, \sigma)$ , parses  $\sigma = (L, \text{pTag})$  if  $\sigma \neq \perp$  and if  $\sigma = \perp$  it sets  $L$  to an empty list and **pTag** to  $0^n$ . Further, it computes a string  $X$  as the concatenation of the  $n$ -bit binary encodings of the entries in  $L$ . In particular, if  $\sigma = \perp$ , then  $X$  is set to  $1^n$  and otherwise it is set as  $X = 1^n || \text{bin}_n(L[1]) || \dots || \text{bin}_n(L[\ell]) = 1^n || \text{str}(L) ||$ , where  $\ell = ||L||$ . Then,  $p$  is computed as  $\text{Poly}_h(X) \oplus |X|h^{\ell+2}$  and finally the **tag** is produced as  $f_K(p || \text{pTag} || m)$ . The new state  $\sigma$  is constructed as  $(\text{append}(L, |m|), \text{tag})$ .

The verification process takes a message  $M$ , a state  $\sigma$  and a tag  $t$ . First, it does some checks to verify if the state  $\sigma$  and the message  $M$  are compatible. In particular, if the following are true, then the verification procedure rejects by returning 0.



1.  $\sigma = \perp$ . This is an incompatible state for any message for verification, as  $\sigma = \perp$  signifies that the message has never gone through the `ConCatU.MacUpdate` procedure.
2.  $\sigma \neq \perp$  and  $\sigma = (L, \text{pTag})$  but  $\text{sum}(L) \neq |M|$ . This also represents an incompatible state for the message  $M$ . As the list  $L$  keeps track of the updates performed on a message. The size of the list  $\|L\|$  represents the number of updates that have taken place and  $L[i]$  denotes the number of bits concatenated to the message in the  $i$ -th update. Hence, if  $\text{sum}(L) \neq |M|$ , then  $\sigma$  represents an incompatible state for  $M$ .
3.  $\sigma \neq \perp$  and  $\sigma = (L, \text{pTag})$  but  $\text{pTag} \neq t$ . This is incompatible, as the `ConCatU.MacUpdate` always writes the current computed tag in the `pTag` field of  $\sigma$ .

After these checks of compatibility, it sets  $\ell = \|L\|$ , and parses the message  $M$  as  $(m_1, m_2, \dots, m_\ell)$  such that  $M = m_1 \| m_2 \| \dots \| m_\ell$  and  $|m_i| = L[i]$ . Then, in lines 29 to 33, it basically runs the procedure `ConCatU.SeqUpdt((K, h), (m_1, m_2, \dots, m_\ell))` as described in Figure 6-1, i.e., it sequentially computes updates of  $m_1, m_2, \dots, m_\ell$  starting from an empty message. Finally, if the computed tag matches the input tag  $t$ , then it accepts by returning 1, else it rejects by returning 0. Based on the verification procedure described in Figure 6-2, it is easy to verify that `ConCatU` is a correct UdmAC.

### 6.4.1 Instantiations and Efficiency

We have motivated UdmACs to be used in client-server scenarios, where a client stores a message  $M$  along with a tag  $t$  in an un-trusted server and it requires updating the message. In the context of `ConCatU` the allowed update is concatenation. The client wants to compute the authentication tag for the updated message without downloading the previous message from the server. We will view the procedures of `ConCatU` in this context as this would help us to evaluate the efficiency of the procedure.

A motivating use case may be the following. Suppose the client records real time video footage through a surveillance camera and every hour uploads the raw video data to a server. Thus, every hour, the message in the server is updated by concatenation and by using **ConCatU**, one can compute the tag using just the data of one hour without needing any access to the whole data.

In such a client server scenario. If a single message is being updated multiple times, it is best for the client to store the state of the message with itself. In the case of **ConCatU**, the state contains the tag and the list of updates. The size of the state for a message which has undergone  $k$  updates would be  $n(k + 1)$  bits,  $n$ -bits for the tag and  $nk$  bits for the list  $L$ . For our example, the size of the state grows just  $n$ -bits per hour, whereas the size of the message grows by about a GB per hour (this would depend on the resolution of the video etc). Note,  $n$  being the tag length, an acceptable value for it would be just 128.

Storing the state, along with some extra information with the client would have efficiency implications while computing the tag. As per the procedure depicted in Figure 6-2, an update with a message  $m$  and state  $(L, \mathbf{pTag})$  requires computing  $p = \text{Poly}_h(X)$ , and one computation of  $f_k(p||\mathbf{pTag}||m)$ . The other computations required are not significant. Note,  $X = 1^n||\text{str}_n(L)$  and thus  $|X| = (|L| + 1)n$ , hence  $\text{Poly}_h(X)$  is of degree at most  $(|L| + 1)$ , and computing  $\text{Poly}_h(X)$  will require  $(|L| + 1)$  additions and multiplications in the field  $\mathbb{F}_{2^n}$ . But, the client may keep a bit more information to speed up this computation. Along with the state, it also stores the current value of  $p$  (see line 11 of Figure 6-2) for each update. When a message  $m_1$  is first uploaded then  $p$  is computed as  $(1^n)h$  and  $L$  is initialized as a list containing a single element, where  $L[1] = |m_1|$ . The client stores both  $p$  and  $L$ . When  $m_1$  is updated with  $m_2$ , then instead of recomputing  $p$  the client just updates it as  $p = (p + |m_2|)h$ , and uses this value in 12. This update incurs just one multiplication and one addition. This can be continued, and each update would require just one multiplication and one addition for computing  $p$  irrespective of the current state of the message.

In addition to  $p$ , the client needs to compute  $f_k(p||\mathbf{pTag}||m)$ . The only require-

ment for  $f_k()$  is that it has to be a variable input length PRF. Thus,  $f_k()$  can be instantiated with any block-cipher based secure deterministic message authentication code like PMAC [16, 81], PAuth [32], secure variants of CBC-MAC [15, 56, 74, 62] etc. The parallelizable MACs like PMAC and PAuth are particularly suited as they are efficient across several platforms and require just  $\lceil |M|/n \rceil$  many block-cipher calls to authenticate a message  $M$ . Thus, a single update call for a message  $m$  will require  $\lceil (|m| + 2)/n \rceil \approx \lceil |m|/n \rceil + 2$  block-cipher calls and one multiplication in  $\mathbb{F}_{2^n}$ .

### 6.4.2 Security of ConCatU

First, let us consider a small modification in the procedure `ConCatU.MacUpdate` where the procedure only returns the `tag`, i.e., it does not output the updated state, though it computes it. We call this modified construction as `ConCatU.MU1`. Our first observation is that if  $f$  is a PRF, then `ConCatU.MU1` is also a PRF. Before we formalize this fixing a notation would be useful.

**Definition 6.4.1.** *A  $(q, r, \ell)$ -adversary is an adversary who asks  $q$  queries in the query phase where the queries are  $(\sigma_i, m_i)$ ,  $1 \leq i \leq q$ , each  $\sigma_i = (L_i, \mathbf{pTag}_i) \in (\mathbb{N})^* \times \{0, 1\}^n$  and  $r = \max_i\{|L_i|\}$  and  $\ell = \max_i\{|m_i|\}$ . Note if  $\sigma = \perp$  we consider  $\sigma = (\langle \rangle, 0^n)$ .*

**Proposition 6.4.1.** *Let  $\mathcal{A}$  be an arbitrary  $(q, r, \ell)$ -PRF adversary for `ConCatU.MU1`. Then there exists a PRF adversary  $\mathcal{B}$  for  $f_K$  such that*

$$\text{Adv}_{\text{ConCatU.MU1}}^{\text{prf}}(\mathcal{A}) \leq \text{Adv}_f^{\text{prf}}(\mathcal{B}) + \binom{q}{2} \frac{r}{2^n}.$$

*Proof.* We assume that  $\mathcal{A}$  does not repeat queries. As it is a  $(q, r, \ell)$ -PRF adversary, it asks  $q$  distinct queries.

Consider the  $i$ -th query  $(\sigma_i, m_i)$  of  $\mathcal{A}$ . If  $\sigma_i = \perp$  we set  $L_i = \langle \rangle$ ,  $X_i = 1^n$  and  $\mathbf{pTag}_i = 0^n$ . If  $\sigma_i \neq \perp$  then  $\sigma_i = (L_i, \mathbf{pTag}_i)$ , and we set  $X_i = 1^n \parallel \text{str}_n(L_i)$ . For the  $i$ -th query of  $\mathcal{A}$  we define the string  $Z_i = p_i \parallel \mathbf{pTag}_i \parallel m_i$ , where

$$p_i = \text{Poly}_h(X_i).$$

Note that on the  $i$ -th query, the function  $f_K()$  is called with input  $Z_i$ . Let the set of all inputs to  $f_K()$  during the queries of the adversary be

$$\mathbb{S} = \{Z_i : 1 \leq i \leq q\}.$$

Let us define the following events:

**G0:**  $\mathcal{A}$  interacting with `ConCatU.MU1` instantiated with  $K \xleftarrow{\$} \{0, 1\}^k$  and  $h \xleftarrow{\$} \{0, 1\}^n$  and outputting a 1.

**G1:**  $\mathcal{A}$  interacting with `ConCatU.MU1`, where  $f_K()$  is replaced by a uniform random function  $\rho$  drawn from the set of all functions mapping  $\{0, 1\}^*$  to  $\{0, 1\}^n$ , and outputting 1.

**G2:**  $\mathcal{A}$  interacting with `ConCatU.MU1`, where  $f_K()$  is replaced by  $\$(\cdot)$ , which on any input returns a uniform random  $n$  bit string, and outputting a 1. Note that `ConCatU.MU1` where  $f_K()$  is replaced by  $\$(\cdot)$  essentially outputs a  $n$ -bit random string irrespective of its input.

**COLL:** While interacting with `ConCatU.MU1`, where  $f_K()$  is replaced by a uniform random function  $\rho$ ,  $\mathcal{A}$  asks two distinct queries  $i, j \in [q]$ ,  $i \neq j$  such that  $Z_i = Z_j$ .

By definition of PRF advantage, we have

$$\text{Adv}_{\text{ConCatU.MU1}}^{\text{prf}}(\mathcal{A}) = |\Pr[\text{G0}] - \Pr[\text{G2}]|. \quad (6.4)$$

By a standard reduction, we can construct an adversary  $\mathcal{B}$  such that

$$|\Pr[\text{G0}] - \Pr[\text{G1}]| \leq \text{Adv}_f^{\text{prf}}(\mathcal{B}), \quad (6.5)$$

Also, the events **G1** and **G2** are same unless the event **COLL** occurs, hence by the difference lemma, we have

$$|\Pr[\text{G1}] - \Pr[\text{G2}]| \leq \Pr[\text{COLL}]. \quad (6.6)$$

Using the above equations, we have the following sequence of inequalities.

$$\begin{aligned}
\text{Adv}_{\text{ConCatU.MU1}}^{\text{prf}}(\mathcal{A}) &= |\Pr[\text{G0}] - \Pr[\text{G2}]| \\
&= |(\Pr[\text{G0}] - \Pr[\text{G1}]) + (\Pr[\text{G1}] - \Pr[\text{G2}])| \\
&\leq |(\Pr[\text{G0}] - \Pr[\text{G1}])| + |(\Pr[\text{G1}] - \Pr[\text{G2}])| \\
&\leq \text{Adv}_f^{\text{prf}}(\mathcal{B}) + \Pr[\text{COLL}].
\end{aligned}$$

We are only left to bound  $\Pr[\text{COLL}]$ . For which we consider two distinct queries  $(\sigma_i, m_i)$ ,  $(\sigma_j, m_j)$  for which  $\rho()$  gets called with  $Z_i$  and  $Z_j$  respectively. We consider the following cases:

**Case 1.**  $m_i \neq m_j$ . In this case  $\Pr[Z_i = Z_j] = 0$ .

**Case 2.**  $m_i = m_j$ . As  $(\sigma_i, m_i) \neq (\sigma_j, m_j)$ , thus  $\sigma_i \neq \sigma_j$ . We have to consider two sub cases here. If  $\text{pTag}_i \neq \text{pTag}_j$ , then  $\Pr[Z_i = Z_j] = 0$ . If  $\text{pTag}_i = \text{pTag}_j$ , then  $L_i \neq L_j$ . Which implies  $X_i \neq X_j$  and by Proposition 6.2.2 we have  $\Pr[Z_i = Z_j] = \Pr[p_i = p_j] \leq (\max(\|L_i\|, \|L_j\|))/2^n$ .

Thus, by the above two cases, the fact that  $\mathbb{S}$  contains  $q$  elements, and the union bound, we have

$$\Pr[\text{COLL}] \leq \binom{q}{2} \frac{r}{2^n}. \quad (6.7)$$

This completes the proof.  $\square$

The following Theorem claims security of **ConCatU** against a forgery adversary.

**Theorem 6.4.1.** *Let  $\mathcal{A}$  be an arbitrary  $(q, r, \ell)$ -forgery adversary attacking **ConCatU**, also let the forgery produced by  $\mathcal{A}$  be  $(\sigma, m, t)$ , where  $\sigma = (L, \text{ptag})$ , where  $\|L\| = r^*$ . Then there exists a PRF adversary  $\mathcal{B}$  for  $f$  such that*

$$\text{Adv}_{\text{ConCatU}}^{\text{Uauth}}(\mathcal{A}) \leq \text{Adv}_f^{\text{prf}}(\mathcal{B}) + \frac{1 + (q+1)r^*}{2^n} + \frac{q \max\{r, r^*\}}{2^n} + \binom{q}{2} \frac{r}{2^n}.$$

*Proof.* We denote the real scheme instantiated with  $(K, h) \xleftarrow{\$} \{0, 1\}^k \times \{0, 1\}^n$  by  $\text{ConCatU}(f_K, h)$ . By  $\text{ConCatU}(\rho, h)$  we denote the update and verification algorithms

of  $\text{ConCatU}$ , where the function  $f_K()$  is replaced by a uniform random function  $\rho$  drawn from the set of all functions mapping  $\{0, 1\}^*$  to  $\{0, 1\}^n$ .

We call the event of  $\mathcal{A}$  attacking  $\text{ConCatU}(f_K, h)$  and producing a valid forgery which passes verification by  $\text{Succ}^{(f_K, h)}(\mathcal{A})$ . Similarly, we call the event of  $\mathcal{A}$  attacking  $\text{ConCatU}(\rho, h)$  and producing a valid forgery which passes verification by  $\text{Succ}^{(\rho, h)}(\mathcal{A})$ .

By an easy reduction, we can construct a PRF adversary  $\mathcal{B}$  such that

$$\text{Adv}_{\text{ConCatU}}^{\text{Uauth}}(\mathcal{A}) = \Pr[\text{Succ}^{(f_K, h)}(\mathcal{A})] \leq \text{Adv}_f^{\text{prf}}(\mathcal{B}) + \Pr[\text{Succ}^{(\rho, h)}(\mathcal{A})]. \quad (6.8)$$

Now we try to bound  $\Pr[\text{Succ}_{\mathcal{A}}^{(\rho, h)}]$ .

We assume that  $\mathcal{A}$  does not repeat any query and makes  $q$  distinct queries. Let its  $i^{\text{th}}$  query be  $(\sigma_i, m_i)$ , and for such a query it gets a response  $(\tilde{\sigma}_i, t_i)$ . Thus, the query profile of  $\mathcal{A}$  is

$$\mathbb{P}_{\mathcal{A}} = \{(\sigma_i, m_i, \tilde{\sigma}_i, t_i) : 1 \leq i \leq q\}.$$

Also, we denote

$$\mathbb{Q}_{\mathcal{A}} = \{(\sigma_i, m_i) : 1 \leq i \leq q\},$$

the set of distinct queries made by  $\mathcal{A}$ .

As in the proof of Proposition 6.4.1, consider the  $i$ -th query  $(\sigma_i, m_i)$ . If  $\sigma_i = \perp$  we set  $L_i = \langle \rangle$ ,  $X_i = 1^n$  and  $\text{pTag}_i = 0^n$ . If  $\sigma_i \neq \perp$  then  $\sigma_i = (L_i, \text{pTag}_i)$ , and we set  $X_i = 1^n \parallel \text{str}_n(L_i)$ . For the  $i$ -th update query of  $\mathcal{A}$  we define the string  $Z_i = p_i \parallel \text{pTag}_i \parallel m_i$ , where

$$p_i = \text{Poly}_h(X_i).$$

Note that on the  $i$ -th query, the random function  $\rho$  is called on the string  $Z_i$ . Let the set of all inputs to  $\rho$  during the queries of the adversary be

$$\mathbb{S} = \{Z_i : 1 \leq i \leq q\}.$$

**Claim 6.4.1.** *Let  $E_0$ , be the event that there is a collision in the set  $\mathbb{S}$ . Then,*

$$\Pr[E_0] \leq \binom{q}{2} \frac{r}{2^n}. \quad (6.9)$$

As the event  $E_0$  is exactly the event COLL defined in the proof of Proposition 6.4.1, hence Equation (6.7) proves the above claim.

Let the valid forgery produced by  $\mathcal{A}$  be  $(\sigma^*, m^*, t^*)$ . It is required that  $\sigma^* \neq \perp$ , as otherwise the success probability of  $\mathcal{A}$  will be zero. Let  $\sigma^* = (L^*, \mathbf{pTag}^*)$ , where  $L^* = \langle \ell_1^*, \ell_2^*, \dots, \ell_{r^*}^* \rangle$ , and

$$\ell_1^* + \ell_2^* + \dots + \ell_{r^*}^* = \ell^* = |m^*|.$$

Also, let  $(m_1^*, m_2^*, \dots, m_{r^*}^*) = \text{Parse}(m^*, L^*)$ . Also define  $L_0^* = \langle \rangle$  and for  $1 \leq i \leq r^*$ ,  $L_i^* = \text{append}(L_{i-1}^*, \ell_i^*)$ .

For ease of exposition, we define some variables. For  $0 \leq i \leq r^*$  we define

$$X^{(i)} = \begin{cases} 1^n & \text{if } i = 1 \\ X^{(i-1)} \parallel \text{bin}_n(\ell_{i-1}) & \text{otherwise,} \end{cases} \quad (6.10)$$

and for  $1 \leq j \leq r^*$

$$p^{(j)} = \text{Poly}_h(X^{(j-1)}).$$

For  $0 \leq i \leq r^*$

$$\text{tag}^{(i)} = \begin{cases} 0^n & \text{if } i = 0 \\ f_K(p^{(i)} \parallel \text{tag}^{(i-1)} \parallel m_i^*) & \text{otherwise,} \end{cases} \quad (6.11)$$

and for  $1 \leq i \leq r^*$ ,

$$Y^{(i)} = p^{(i)} \parallel \text{tag}^{(i-1)} \parallel m_i^*.$$

and let  $\mathbb{Y} = \{Y^{(i)} : 1 \leq i \leq r^*\}$ . Note that while verification of the forgery produced by  $\mathcal{A}$ ,  $\rho$  gets evaluated on the strings in  $\mathbb{Y}$ .

Now, we define a set of queries  $\mathbb{Q} = \{Q_i : 1 \leq i \leq r^*\}$ , where

$$Q_i = \begin{cases} (\perp, m_i^*) & \text{if } i = 1 \\ ((L_{i-1}^*, \mathbf{tag}^{(i-1)}), m_i^*) & \text{if } i > 1. \end{cases}$$

Note that, if  $Q_i$  is given as an update query, then the function  $\rho$  gets evaluated on the string  $Y^{(i)}$ . Also this query set  $\mathbb{Q}$  is defined by the forgery attempt of  $\mathcal{A}$ .

As the forgery produced by  $\mathcal{A}$  is a valid forgery, hence according to the definition of a valid forgery,  $\mathcal{A}$  has not asked the queries  $Q_1, Q_2, \dots, Q_{r^*}$  in the query phase in the same order.

**Definition 6.4.2.** *Given the forgery of  $\mathcal{A}$  the query set  $\mathbb{Q}_{\mathcal{A}}$  is called out of order, if there exists a  $i \in [r^* - 1]$ , such that query  $Q_{i+1}$  was asked before query  $Q_i$ .*

**Definition 6.4.3.** *Given the forgery of  $\mathcal{A}$  the query set  $\mathbb{Q}_{\mathcal{A}}$  is called incomplete, if there exists a  $i \in [r^*]$ , such that query  $Q_i \notin \mathbb{Q}_{\mathcal{A}}$ .*

Thus for the forgery produced by  $\mathcal{A}$  to be valid its query set  $\mathbb{Q}_{\mathcal{A}}$  has to be either out of order or incomplete.

**Claim 6.4.2.** *Let E1 be the event that the query set  $\mathbb{Q}_{\mathcal{A}}$  is out of order. Then,*

$$\Pr[\mathbf{E1}] \leq \frac{r^*}{2^n} + \binom{q}{2} \frac{r}{2^n}. \quad (6.12)$$

**Proof of Claim:** Let for  $i \in [q - 1]$  the query  $Q_{i+1}$  was asked before  $Q_i$ . For these queries, the random function  $\rho$  was evaluated on  $Z_{i+1}$  and  $Z_i$  respectively in the same order. Assuming that there was no collision on  $\mathbb{S}$ , we have  $Z_i \neq Z_{i+1}$ . Also,

$$\begin{aligned} Z_{i+1} &= p_i + 1 \parallel \mathbf{pTag}_{i+1} \parallel m_{i+1}, \\ Z_i &= p_i \parallel \mathbf{pTag}_i \parallel m_i, \end{aligned}$$

and  $\mathbf{pTag}_{i+1} = \rho(Z_i)$ . As all  $Z_i$ 's are distinct thus  $\mathcal{A}$  has not seen the output of  $\rho(Z_i)$  before it asked the query  $Q_i$ , as  $\rho$  is a random function. Thus the probability that  $\mathcal{A}$  could have asked the query  $Q_{i+1} = ((L_i, t), m_i)$  where  $t = \rho(Z_i)$  is at most  $1/2^n$ .



Hence using the union bound over all possible  $i \in [r^*]$  which are out of order, we have  $\Pr[\mathbf{E1}|\mathbf{E0}] \leq r^*/2^n$ . Hence,

$$\begin{aligned} \Pr[\mathbf{E1}] &\leq \Pr[\mathbf{E1}|\mathbf{E0}] + \Pr[\mathbf{E0}] \\ &\leq \frac{r^*}{2^n} + \binom{q}{2} \frac{r}{2^n}. \end{aligned} \tag{6.13}$$

This completes the proof of the claim.

Now, we assume that the query set  $\mathbb{Q}_{\mathcal{A}}$  is not out of order, then if the forgery of  $\mathcal{A}$  is valid, then  $\mathbb{Q}_{\mathcal{A}}$  must be incomplete, i.e., there exists a  $j^* \in [r^*]$  such that  $Q_{j^*}$  is not in  $\mathbb{Q}_{\mathcal{A}}$ . Thus we have the following claim.

**Claim 6.4.3.** *For  $j^* \leq k \leq r^*$ , let  $\mathbf{E2}^{(k)}$  be the event that  $Y^{(k)} \in \mathbb{S}$ . Then*

$$\Pr[\mathbf{E2}^{(j^*)}] \leq \frac{q \max\{r, j^*\}}{2^n} + \frac{r^*}{2^n} + \binom{q}{2} \frac{r}{2^n}. \tag{6.14}$$

and for  $j^* + 1 \leq k \leq r^*$ ,

$$\Pr[\mathbf{E2}^{(k)} | \neg \mathbf{E2}^{(k-1)}] \leq \frac{q}{2^n}. \tag{6.15}$$

For a proof of Equation (6.14), notice that, if  $Q_{j^*} = ((L_{j^*}^*, \mathbf{pTag}_{j^*}), m_{j^*})$  was never asked in the query phase then  $\Pr[Y^{(j^*)} = Z_i]$  for some  $i \in [q]$  is at most  $\max\{r, j^*\}/2^n$ . This follows from Proposition 6.2.2. Thus as there are  $q$  many  $Z_i$ 's in  $\mathbb{S}$ , by the union bound we have,  $\Pr[\mathbf{E2}^{(j^*)} | \neg \mathbf{E1}]$  is at most  $q \max\{r, j^*\}/2^n$ . Hence,

$$\begin{aligned} \Pr[\mathbf{E2}^{(j^*)}] &\leq \Pr[\mathbf{E2}^{(j^*)} | \neg \mathbf{E1}] + \Pr[\mathbf{E1}] \\ &\leq \frac{q \max\{r, j^*\}}{2^n} + \frac{r^*}{2^n} + \binom{q}{2} \frac{r}{2^n}. \end{aligned}$$

For Equation (6.15), we observe that if  $Y^{(k-1)}$ , not in  $\mathbb{S}$ , then  $\mathbf{tag}^{(k-1)}$  is a  $n$ -bit uniform random string, and thus  $Y^{(k)} = p^{(k)} \parallel \mathbf{tag}^{(k-1)} \parallel m_k^*$  is equal to any  $Z_i \in \mathbb{S}$  with probability at most  $1/2^n$ , and hence

$$\Pr[\mathbf{E2}^{(k)} | \neg \mathbf{E2}^{(k-1)}] \leq \frac{q}{2^n}.$$

as claimed.

Finally, we would like to bound the probability of success of  $\mathcal{A}$ , i.e., we need to bound  $\Pr[\text{Succ}_{\mathcal{A}}^{(\rho,h)}]$ .

$$\begin{aligned}
\Pr[\text{Succ}_{\mathcal{A}}^{(\rho,h)}] &= \Pr[t^* = \text{tag}^{(r^*)}] \\
&\leq \Pr[t^* = \text{tag}^{(r^*)} | \neg \text{E2}^{(r^*)}] + \Pr[\text{E2}^{(r^*)}] \\
&\leq \frac{1}{2^n} + \Pr[\text{E2}^{(r^*)}].
\end{aligned} \tag{6.16}$$

The last inequality is due to the fact that if  $Y^{(r^*)} \notin \mathbb{S}$ , then  $\text{tag}^{(r^*)}$  is a uniform random  $n$ -bit string and  $t^*$  is a fixed string fixed by  $\mathcal{A}$  before  $\rho()$  was evaluated on  $Y^{(r^*)}$ .

Now, by using repeated conditioning, we have

$$\begin{aligned}
\Pr[\text{E2}^{(r^*)}] &= \sum_{k=0}^{r^*-j^*} \Pr[\text{E2}^{(r^*-k)} | \neg \text{E2}^{(r^*-(k-1))}] + \Pr[\text{E2}^{(j^*)}] \\
&\leq \frac{(r^* - j^*)q}{2^n} + \frac{q \max\{r, j^*\}}{2^n} + \frac{r^*}{2^n} + \binom{q}{2} \frac{r}{2^n} \\
&< \frac{r^*q}{2^n} + \frac{q \max\{r, r^*\}}{2^n} + \frac{r^*}{2^n} + \binom{q}{2} \frac{r}{2^n}.
\end{aligned} \tag{6.17}$$

Finally using Equations (6.8),(6.16) and (6.17), we get

$$\begin{aligned}
\Pr[\text{Succ}_{\mathcal{A}}^{(f\kappa,h)}] &\leq \text{Adv}_f^{\text{prf}}(\mathcal{B}) + \frac{1}{2^n} + \frac{r^*q}{2^n} + \frac{q \max\{r, r^*\}}{2^n} + \frac{r^*}{2^n} + \binom{q}{2} \frac{r}{2^n} \\
&= \text{Adv}_f^{\text{prf}}(\mathcal{B}) + \frac{1 + (q+1)r^*}{2^n} + \frac{q \max\{r, r^*\}}{2^n} + \binom{q}{2} \frac{r}{2^n}.
\end{aligned}$$

□

## 6.5 XoRU: An Updatable MAC for $\oplus$ Difference

Here, we discuss a scheme for UdMAC where the update function is  $\oplus$ , i.e., on a message  $M$ , if an update  $m$  is applied then it results in the message  $M \oplus m$ . We always assume that the length of the updated message is at least the size of the

original message. In particular, it is important to note the following points:

1. If the current message is  $M$  and the intended updated message is  $M'$ , where  $|M| = |M'|$ , then the update request should be with the string  $m = M \oplus M'$ , and here  $|m| = |M|$ .
2. If the current message is  $M$  and the intended updated message is  $M'$  where  $|M| < |M'|$ , then an update request should be issued with  $m = M \parallel 0^{|M'| - |M|}$ . In this case  $|m| = |M'|$ .
3. If the current message is  $M$  and the intended updated message is  $M'$  where  $|M| > |M'|$ , such an update is not allowed. It is possible to modify the construction of XoRU that we present next to allow such kinds of updates, but it will require some extra information to be stored in the state. We do not further explore this here, as such updates will not be required in the context of the verifiable SSE that we design in the next Chapter.

Let  $x \in \{0, 1\}^*$ , and  $x_1, x_2, \dots, x_m = \text{parse}_n(x)$ . For  $h \in \{0, 1\}^n$ , define

$$\text{Poly}_h(x) = x_1 h \oplus x_2 h^2 \oplus \dots \oplus \text{Pad}_n(x_m) h^m$$

where the multiplications are in  $\mathbb{F}_{2^n}$ .

For XoRU we consider the key space to be  $\mathcal{K} = \{0, 1\}^k \times \{0, 1\}^n$ , the state space to be  $\Sigma \cup \{\perp\}$ , where  $\sigma$  is either  $\perp$  or can be parsed as  $\sigma = (\text{pNonce}, \text{pLen}, \text{pTag})$ , where  $\text{pNonce}, \text{pLen}, \text{pTag} \in \{0, 1\}^n$ , and  $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  is a fixed input length PRF.

The MacUpdate and Verify procedures for XoRU are described in Figure 6-3. Next we note down some elementary observations regarding the construction.

**Proposition 6.5.1.** *Let  $m_1, m_2, \dots, m_s \in \mathcal{M}$  and  $\text{XoRU.SeqUpdt}((K, h), m_1, m_2, \dots, m_s) = (\sigma_s, t_s)$ . Let  $\sigma_s = (\mathbf{N}, \mathbf{L}, \mathbf{T})$ . Then*

1.  $\mathbf{L} = \max(|m_1|, |m_2|, \dots, |m_s|)$
2.  $t_s = \mathbf{T} = \text{Poly}_h(m_1 \oplus m_2 \oplus \dots \oplus m_s) \oplus \mathbf{L}h^{\lceil \mathbf{L}/n \rceil} \oplus E_K(\mathbf{N})$ .

<u>XoRU.MacUpdate(<math>K, \sigma, N, m</math>)</u>	<u>XoRU.MacVerify(<math>K, \sigma, m, t</math>)</u>
01. <b>if</b> $\sigma = \perp$ <b>then</b>	21. <b>if</b> $\sigma = \perp$
02. $\text{tag} \leftarrow \text{Poly}_h(m) \oplus  m h^{\lceil  m /n \rceil} \oplus E_K(N)$ ;	22. <b>return</b> 0;
03. <b>else</b>	23. <b>end if</b>
04. $(\text{pNonce}, \text{pLen}, \text{pTag}) \leftarrow \sigma$ ;	24. $(\text{pNonce}, \text{pLen}, \text{pTag}) \leftarrow \sigma$ ;
05. $\ell_1 = \lceil \frac{\text{pLen}}{n} \rceil$ ;	25. <b>if</b> $\text{pLen} \neq  m $ ,
06. $\ell_2 = \lceil \frac{ m }{n} \rceil$ ;	26. <b>return</b> 0;
07. $\text{tag} \leftarrow \text{pTag} \oplus h^{\ell_1} \text{pLen} \oplus E_K(\text{pNonce})$ ;	27. $\ell = \lceil \frac{\text{pLen}}{n} \rceil$ ;
08. <b>if</b> $\text{pLen} =  m $ ,	28. $\text{tag} \leftarrow \text{Poly}_h(m) \oplus h^\ell \text{pLen} \oplus E_K(\text{pNonce})$ ;
09. $\text{tag} \leftarrow \text{tag} \oplus \text{Poly}_h(m) \oplus h^{\ell_1} \text{pLen} \oplus E_K(N)$ ;	29. <b>if</b> $\text{tag} \neq t$ ,
10. <b>else</b>	30. <b>return</b> 0;
11. $\text{tag} \leftarrow \text{tag} \oplus \text{Poly}_h(m) \oplus h^{\ell_2}  m  \oplus E_K(N)$ ;	31. <b>else</b>
12. <b>end if</b>	32. <b>return</b> 1;
13. <b>endif</b>	33. <b>end if</b>
14. $\sigma \leftarrow (N, \max\{\text{pLen},  m \}, \text{tag})$ ;	
15. <b>return</b> $(\sigma, \text{tag})$ ;	

Figure 6-3: Specification of XoRU.MacUpdate and XoRU.MacVerify.

Note that, point (2) in Proposition 6.5.1 essentially gives the correctness of XoRU.

### 6.5.1 Efficiency of XoRU

The bulk computation for update takes place in lines 07 and line 09 or line 11 of the procedure described in Figure 6-3. We first give a general estimate of the computation required.

First, we observe that if the condition in line number 08 of Figure 6-3 is satisfied then the computation of  $h^{\ell_1} \text{pLen}$  in both lines 07 and 10 can be avoided. As the updated tag would be computed as

$$\text{tag} = \text{pTag} \oplus \text{Poly}_h(m) \oplus E_K(\text{pNonce}) \oplus E_K(N).$$

Thus, an update satisfying the condition in Line 08 will require the computation of  $\text{Poly}_h(m)$ , which is just an evaluation of a degree  $\ell_2$  polynomial with no constant term. If the Horner's rule is used for evaluating the polynomial, it will take  $\ell_2$  multiplications and  $\ell_2 - 1$  additions in  $\mathbb{F}_{2^n}$ . In addition to the computation of  $\text{Poly}_h(m)$  three more additions in  $\mathbb{F}_{2^n}$  and two calls to the PRF  $E_K(\cdot)$  (which can be instantiated with a block-cipher like AES). would be required. If we assume the cost of one block-cipher call to be  $T_{\text{AES}}$  and the cost of one multiplication and one addition in  $\mathbb{F}_{2^n}$  to be  $T_{\text{add}}$

and  $T_{\text{mult}}$  respectively, then the total cost if  $\text{pLen} > |m|$  (the condition in line 08) will be

$$T_{\text{add}}(\ell_2 + 2) + T_{\text{mult}}\ell_2 + 2T_{\text{AES}}. \quad (6.18)$$

Now, let us analyze the case where the condition in line 08 is not satisfied, i.e., we have  $\text{pLen} < |m|$ . Note, in this case computation in line 07 will require two additions in  $\mathbb{F}_{2^n}$ , a computation of  $h^{\ell_1}$  and a multiplication in  $\mathbb{F}_{2^n}$ . If the square and multiply algorithm for exponentiation is used then computation of  $h^{\ell_1}$  will require at most  $\log(\ell_2)$  squarings and  $\log(\ell_2)$  multiplications in  $\mathbb{F}_{2^n}$ . Though squares in fields of characteristic 2 can be computed much more efficiently than multiplications in various platforms, for simplicity, we consider the cost of squaring to be the same as that of multiplication. With such an assumption line 07 can be computed with  $2\log(\ell_2)$  multiplications. For the computation in line 11,  $\text{Poly}_h(m) \oplus h^{\ell_2}|m|$  can be computed with  $\ell_2$  additions and  $\ell_2 + 1$  multiplications in  $\mathbb{F}_{2^n}$ . Thus summarizing, the total cost will be

$$(\ell_2 + 4)T_{\text{add}} + (2\log(\ell_2) + \ell_2 + 1)T_{\text{mult}} + 2T_{\text{AES}}. \quad (6.19)$$

**A Special Case:** XoRU can be used to support the concatenation of messages by proper update requests. For example, if the current message is  $M$  and the intended message is  $M||m$ , then  $M$  can be updated with an update request of  $0^{|M|}||m$ . The SSE scheme which we discuss in the next chapter will require such kinds of updates and in that case both ConCatU and XoRU can be used.

Note, for each update, if  $|m| = kn$  where  $k$  is a small constant, then the client can store  $h^{\ell_1}$  where  $\ell_1 = \lceil \text{pLen}/n \rceil$  for each updated message, and also store  $\text{Poly}(M)$  for the current message  $M$ . This will require just  $(2k + 2)$  multiplications for computing each updated tag and also for updating  $h^{\ell_1}$  (as  $\ell_1$  increases by  $k$  in each update), and the polynomial of the current message.

## 6.5.2 Security of XoRU

We claim security of XoRU against a *nonce respecting* adversary, i.e., for the mac-update queries, an adversary never repeats a nonce.

For technical convenience, we consider a little change in the construction of XoRU.MacUpdate, where the procedure returns only the tag instead of  $(\sigma, \text{tag})$ . We call this construction as XoRU.MacUpdate<sup>(1)</sup>. This change has no effect on security as the updated state  $\sigma$  can be constructed by the adversary.

We define the *rnd*-advantage for a nonce respecting adversary  $\mathcal{A}$  attacking XoRU as the following:

$$\text{Adv}_{\text{XoRU}}^{\text{rnd}}(\mathcal{A}) = \left| \Pr[K, h \xleftarrow{\$} \{0, 1\}^n : \mathcal{A}^{\text{XoRU.MacUpdt}^{(1)}}(K, h, \cdot, \cdot, \cdot) = 1] - \Pr[\mathcal{A}^{\$(\cdot, \cdot, \cdot)} = 1] \right|,$$

where the  $\$(\cdot, \cdot, \cdot)$  oracle on invocation with any  $(\sigma, N, m) \in \Sigma \cup \{\perp\} \times \{0, 1\}^n \times \mathcal{M}$  returns a uniform random  $n$ -bit string.

We claim that an efficient nonce-respecting adversary interacting with the XoRU.MacUpdate oracle instantiated with keys  $(K, h) \xleftarrow{\$} \{0, 1\}^n \times \{0, 1\}^n$  cannot distinguish it from the oracle  $\$$  if  $E_K : \{0, 1\}^n \leftarrow \{0, 1\}^n$  is a pseudorandom function family. In other words,

**Proposition 6.5.2.** *For any efficient nonce respecting rnd adversary  $\mathcal{A}$  attacking XoRU, there exists an (almost) equally efficient PRF adversary  $\mathcal{B}$  for the function family  $E_K : \{0, 1\}^n \rightarrow \{0, 1\}^n$ , such that*

$$\text{Adv}_{\text{XoRU.MacUpdate}}^{\text{rnd}}(\mathcal{A}) = \text{Adv}_E^{\text{prf}}(\mathcal{B}).$$

The proof of the above proposition is an easy reduction. The main observation that leads to the claim is that if  $E_K()$  in the construction of XoRU.MacUpdate<sup>(1)</sup> is replaced by a uniform random function then the procedure essentially outputs uniform random strings.

Now, we are in a position to state the security of XoRU.

**Theorem 6.5.1.** *Let  $\mathcal{A}$  be any efficient forging adversary for XoRU who asks  $q$  many*

update queries and finally outputs a valid forgery  $(M^*, \sigma^*, t^*)$ , then there exists an (almost) equally efficient PRF adversary  $\mathcal{B}$  attacking  $E_K$  such that

$$\text{Adv}_{\text{XoRU}}^{\text{Uauth}}(\mathcal{A}) \leq \text{Adv}^{\text{prf}}(\mathcal{B}) + \left( \frac{|M^*| + 1}{n} \right) \frac{1}{2^n} + \frac{1}{2^n}.$$

*Proof.* Let  $\text{XoRU}(K, h)$  be the real update and verification algorithms described in Figure 6-3 instantiated with uniform random  $n$ -bit keys  $K, h$ . Let  $\text{XoRU}(\rho, h)$  denote the update and verification algorithms where  $E_K()$  in Figure 6-3 is replaced by a function  $\rho$  sampled uniformly at random from the set of all functions from  $\{0, 1\}^n$  to  $\{0, 1\}^n$ .

We call the event of  $\mathcal{A}$  attacking  $\text{XoRU}(K, h)$  and producing a valid forgery which passes verification by  $\text{Succ}_{\mathcal{A}}^{(K, h)}$ . Similarly, we call the event of  $\mathcal{A}$  attacking  $\text{XoRU}(\rho, h)$  and producing a valid forgery which passes verification by  $\text{Succ}_{\mathcal{A}}^{(\rho, h)}$ .

By an easy reduction, we can construct a PRF adversary  $\mathcal{B}$  such that

$$\text{Adv}_{\text{XoRU}}^{\text{Uauth}}(\mathcal{A}) = \Pr[\text{Succ}_{\mathcal{A}}^{(K, h)}] \leq \text{Adv}^{\text{prf}}(\mathcal{B}) + \Pr[\text{Succ}_{\mathcal{A}}^{(\rho, h)}]. \quad (6.20)$$

Now we try to bound  $\Pr[\text{Succ}_{\mathcal{A}}^{(\rho, h)}]$ .

First, note that according to Proposition 6.5.2, when  $\mathcal{A}$  asks update queries to  $\text{XoRU}(\rho, h)$  (s)he gets as response random strings in place of the real tags. Thus, in the query phase, irrespective of the queries it asks,  $\mathcal{A}$  gets no information regarding the real scheme.

Finally  $\mathcal{A}$  outputs a valid forgery  $(M^*, \sigma^*, t^*)$ . We consider  $\sigma^* = (N^*, L^*, \text{tag}^*)$ . We have the following cases to consider:

**Case I: There was no query with the nonce  $N^*$ :** Consider the verification procedure described in Figure 6-3 where  $E_K()$  is replaced by  $\rho()$ . Thus, in this case, in the verification procedure,  $\rho(N^*)$  would be a uniform random string. Thus, the **tag** computed in line 27 of the procedure described in Figure 6-3 will in turn be a uniform random string. Thus, as  $t^*$  is a fixed sting,  $\Pr[\text{tag} = t^*] =$

$1/2^n$ , Thus,

$$\Pr[\text{Succ}_{\mathcal{A}}^{(\rho,h)}] = \frac{1}{2^n}. \quad (6.21)$$

**Case II: There was a query with the nonce  $N^*$ :** In this case,  $\rho(N^*)$  was fixed, say to  $t$  when a query with nonce  $N^*$  was asked by  $\mathcal{A}$ . Thus the tag computed by the verification algorithm is

$$\text{tag} = \text{Poly}_h(M^*) \oplus h^{\lceil L^*/n \rceil} L^* \oplus t.$$

For  $\mathcal{A}$  to be successful it is required that  $\text{tag} = t^*$ . Thus, in this case, we have

$$\Pr[\text{Succ}_{\mathcal{A}}^{(\rho,h)}] = \Pr[\text{tag} = t^*] \quad (6.22)$$

$$= \Pr[\text{Poly}_h(M^*) \oplus h^{\lceil L^*/n \rceil} L^* \oplus t \oplus t^* = 0]. \quad (6.23)$$

The above probability is computed on the randomness of the choice of  $h$ . As,

$$p(h) = \text{Poly}_h(M^*) \oplus h^{\lceil L^*/n \rceil} L^* \oplus t \oplus t^*,$$

is a nonzero polynomial of degree  $\lceil L^*/n \rceil$ , hence there are at most  $\lceil L^*/n \rceil$  many values of  $h$  which makes  $p(h) = 0$ . Thus, we have

$$\Pr[\text{Succ}_{\mathcal{A}}^{(\rho,h)}] = \Pr[\text{tag} = t^*] \quad (6.24)$$

$$\leq \frac{1}{2^n} \left\lceil \frac{L^*}{n} \right\rceil \quad (6.25)$$

$$= \frac{1}{2^n} \left\lceil \frac{|M^*|}{n} \right\rceil. \quad (6.26)$$

Now, combining cases I and II and using the union bound, we have

$$\Pr[\text{Succ}_{\mathcal{A}}^{(\rho,h)}] \leq \frac{1}{2^n} \left( \frac{|M^*|}{n} + 1 \right) + \frac{1}{2^n} \quad (6.27)$$



Finally, using Equations (6.20) and (6.27), we get

$$\text{Adv}_{\text{XoRU}}^{\text{Uauth}}(\mathcal{A}) \leq \text{Adv}^{\text{prf}}(\mathcal{B}) + \left( \frac{|M^*| + 1}{n} \right) \frac{1}{2^n} + \frac{1}{2^n},$$

as desired. □

## 6.6 Final Remarks

We proposed a new type of MAC UdMAC which can be used to update the authentication tag of a message and compute the correct tag for the updated message only with access to the update. We defined the security of UdMACs and finally proposed two constructions. This class of MACs will be used in the next chapter to construct verifiable SSE schemes.



## Fault-tolerant Verifiable DSSE

So far in this thesis, all the adversaries we've considered for an SSE are honest-but-curious. However, the server, which is typically viewed as the adversary in an SSE system, can potentially act maliciously by improperly storing user data or by providing incorrect responses to user queries in order to conserve storage and CPU resources. In addressing such adversarial behavior, SSE systems must incorporate a mechanism by which search results may be verified for correctness. An SSE equipped with a mechanism to verify the correctness of search queries is called a verifiable SSE.

It is not feasible for an honest client to keep track of the current state of the database stored on the server. As a result, a client may issue faulty updates, such as adding duplicate entries to the database or attempting to delete non-existent entries (see Definition 2.5.4). Since the database is encrypted, the server cannot assist the client in identifying such incorrect behavior. Informally, an SSE which inherently corrects such faulty updates is called a fault tolerant SSE.

Both verifiability and fault tolerance are essential characteristics required to safeguard the interests of clients. Recent work by [94] proposed the first fault-tolerant verifiable DSSE (FVDSSE) that achieves forward privacy.

In this work, we present the first generic construction for a DSSE scheme that is both forward and backward private, as well as verifiable and fault-tolerant. The construction by [94] provided the verification for the search results using an *authenticated encryption* (AE) scheme. For each update, in addition to the update token, they also outsource an additional AE tag, which they call an AE *proof*. When retrieving the

results, the server sends the AE proofs along with the results, and the client verifies the results using these AE proofs.

In an SSE system, the search result for a keyword  $w$  is the set  $\text{db}(w)$ . An update to a keyword essentially modifies the set  $\text{db}(w)$ , either by adding new identifiers or deleting some from the set. The previous schemes which ensure verifiability [24, 98, 48, 99] generally use an incremental multi-set, which is unable to provide proof in the presence of faulty updates. The work of [94] generates a proof for every update using an AE scheme and sends the update token and the proof/tag to the server. During the search, the server finds all the updates and corresponding tags and sends them to the client. The client verifies each update by invoking the AE verify algorithm.

In our proposed scheme, we see all the updates related to a keyword  $w$  as a message related to the keyword  $w$  which is getting updated with each update operation. Thus, using any updatable MAC, we can generate a tag which acts as proof for every update corresponding to a  $w$ . The specific syntax of a UdMAC allows us then to have a single tag of constant length for all the updates made for the keyword. This tag can be stored with the client, as most SSE schemes [23, 25, 24, 87, 49, 90, 38] allows a  $\mathcal{O}(|W|)^1$  storage at the client side. Please note that keeping the tag safe on the client side is not necessary for the purpose of security. We can also outsource the tag with the update query and ask the server to save the tag. In the proof, we'll provide the adversary with the tag and the state of the UdMAC algorithm for every keyword.

## 7.1 Generic Fault-tolerant Verifiable DSSE

In this section, we explain how to use an UdMAC scheme  $\Psi$  to transform any dynamic secure SSE  $\Sigma$  (see Section 2.3) which is also forward and backward private, into an equivalently secure dynamic forward and backward private verifiable SSE  $\mathbf{v}\Sigma$  that supports faulty updates. We assume the existence of a forward and backward private secure DSSE scheme  $\Sigma$ , with leakage  $\mathcal{L}_\Sigma$ .

An updatable MAC is best described in a client-server framework. Where the

---

<sup>1</sup>Recall  $W$  is the set of distinct keyword present in the database.

client with the messages wants to store it in the server and then update the messages. While querying the client needs the guarantee that the server is sending the current state of the message and has not tampered with the message.

Recall the primary motivation behind a UdMAC is to avoid recomputing the MAC of an entire message when updates are made. This is particularly beneficial when we outsource the message and its tag (potentially to a third-party server). Without an updatable MAC, updating the message would necessitate downloading the entire message and its tag, verifying it, making updates, generating a new tag, and then outsourcing the updated message and tag together, which can be costly.

Notice that the MAC-update algorithm of UdMAC  $\Psi$  in Section 6.3 has three inputs. The secret key  $K$ , the message  $m$ , and a state  $\sigma$ , and outputs a tag  $t$  and updated state  $\sigma$ . Now, the state  $\sigma$  is specific to every message. Thus, the client has to maintain a state for every new message. Now, on query for a message, the server replies with the latest updated message along with its current tag.

A MAC algorithm dictates that only the key should be secret, and the adversary should have control over every other parameter in the forgery. We'll prove the security of our scheme in a model where the adversary controls every parameter of the MAC algorithm except the key. However, the client keeping the state of the MAC algorithm is a requirement of the model (for convenience), and has nothing to do with the security. This does not imply that we can outsource the state of the SSE from client to the server side, and still have correct verification. We believe that to outsource the state of the SSE we need to have some sort of verification on that as well. This needs further investigation.

### 7.1.1 A Generic FVDSSE Scheme $v\Sigma$ Using $\Psi$ and $\Sigma$

In the beginning, we fix an updatable-MAC scheme  $\Psi$  and a secure DSSE scheme  $\Sigma$  and demonstrate how to construct an FVDSSE scheme  $v\Sigma$ . The setup process summarised in Figure 7-1, starts by initializing a key for the  $\Psi$ -MAC. A general practice for designing SSE is to assume an empty database at the setup and update it accordingly. We call the setup phase of  $\Sigma$  to get back a key  $K_\Sigma$  a state  $\sigma_C$  and

<b>Setup(<math>1^\lambda, \text{DB}</math>):</b> 01. $K_{\text{mac}} \leftarrow \Psi.\text{KeyGen}(1^\lambda)$ 02. $(K_\Sigma, \sigma_C, \text{EDB}) \leftarrow \Sigma.\text{Setup}(1^\lambda, \perp)$ 03. $\text{UT}, \text{ST} \leftarrow [] \quad \backslash \backslash$ empty map 04. $\text{K} \leftarrow (K_{\text{mac}}, K_\Sigma)$ 05. $\text{State} \leftarrow (\sigma_C, \text{ST}, \text{UT})$ 06. Keep $(\text{K}, \text{State})$ to client 07. Send EDB to server
---

Figure 7-1: Setup phase of  $\mathbf{v\Sigma}$ .

an empty encrypted database EDB. We also initialize two empty maps UT and ST. These two maps are used to store the number of updates for a keyword and the state of the  $\Psi$  for each keyword, respectively.

<b>Update(<math>\text{K}, \text{State}, \text{op}, (\text{id}, w)</math>):</b> <b>Client Side:</b> 01. <b>if</b> $\text{UT}[w] = \perp$ <b>set</b> $\text{UT}[w] \leftarrow 1$ , <b>else</b> $\text{UT}[w] \leftarrow \text{UT}[w] + 1$ 02. $\text{ut}_w \leftarrow \text{UT}[w]$ 03. $\tilde{\text{id}} \leftarrow (\text{ut}_w, \text{op}, \text{id})$ 04. $(\sigma_C, \text{utk}) \leftarrow \Sigma.\text{Update}_C(K_\Sigma, \sigma_C, \text{add}, (\tilde{\text{id}}, w))$ 05. <b>if</b> $\text{ST}[w] = \perp$ 06. $(\sigma_{\text{mac}}, \text{tag}) \leftarrow \Psi.\text{MacUpdate}(K_{\text{mac}}, \text{ST}[w], w)$ 07. $\text{ST}[w] \leftarrow (\sigma_{\text{mac}}, \text{tag})$ 08. $(\sigma_{\text{mac}}, \text{tag}) \leftarrow \Psi.\text{MacUpdate}(K_{\text{mac}}, \text{ST}[w], \tilde{\text{id}})$ 09. $\text{ST}[w] \leftarrow (\sigma_{\text{mac}}, \text{tag})$ 10. Send $\text{utk}$ to server <b>Server Side:</b> 20. Server updates EDB using $\text{utk}$
---

Figure 7-2: Update protocol of  $\mathbf{v\Sigma}$ .

In the update phase, for every update operation  $(\text{op}, \text{id}, w)$ , we fetch the update number  $\text{ut}_w$  from the map UT. If the keyword is being updated for the first time, we set the update number to 1. Otherwise, we increment the update number by 1. Then we construct a new  $\tilde{\text{id}}$  as an encoding of  $(\text{ut}, \text{op}, \text{id})$ . With this new  $\tilde{\text{id}}$  we now call the base SSE scheme with update query  $(\text{add}, \tilde{\text{id}}, w)$ .

If the keyword is being updated for the first time, we call the  $\Psi$  with an empty state and a  $|\tilde{\text{id}}|$  bit encoding of the keyword  $w$  and save the state. That is, the first message for every keyword is the keyword itself. This is done to have a signature

of the keyword on every update on the keyword. We then update the message for the keyword with the current update  $\tilde{id}$ , and save the current state and tag for the keyword. Next, we call the base SSE scheme  $\Sigma$  with  $(\text{add}, \tilde{id}, w)$  to get back an update token  $\text{stk}$ . Finally, we send the update token (from SSE) to the server. The server updates the EDB with the  $\text{stk}$ . The process is summarised in Figure 7-2.

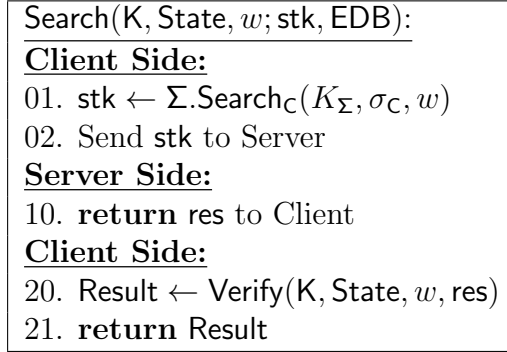


Figure 7-3: Search protocol of  $v\Sigma$ .

During the search operation, the client calls the base SSE scheme with the search keyword  $w$ . The client provides the server with the search token it got from the base SSE scheme. The correctness of the base SSE dictates that the server is able to find all the updates related to the search keyword  $w$  with the search token. Upon receiving the search result the client runs the verify algorithm with the search result as input. The entire process of search is described in Figure 7-3.

The verify algorithm described in Figure 7-4 checks the search result output by the server and returns the correct search result upon verification else returns “reject”. The algorithm first finds all  $\tilde{id}_i$  of the form  $i||\text{op}_i||\text{id}_i$  from the search result. Now from  $i = 1$  to  $\text{UT}[w]$  in sequence, if  $\text{op}_i = \text{add}$ , then the algorithm adds the  $\text{id}_i$  to the result set, else if  $\text{op}_i = \text{del}$  then the algorithm discards  $\text{id}_i$  from the result set. The algorithm also constructs a message  $m$  for verification. It initializes the message with a  $|\tilde{id}|$  bit encoding of the keyword  $w$ , and append every  $\tilde{id}_i$  for  $1 \leq i \leq \text{UT}[w]$  in sequence. Finally, the algorithm fetches the state of the keyword  $\text{ST}[w]$  and sends the state and the  $m$  it constructed to the  $\Psi.\text{MacVerify}$  algorithm. If the output of  $\Psi.\text{MacVerify}$  is “accept”, then the algorithm returns the result set. Else, it returns “reject”.

Verify(K, State, $w$ , res, tag):
00. Result, $V \leftarrow \emptyset$
01. $m \leftarrow w$
02. <b>for</b> $r \in \text{res}$ <b>do</b>
03.   extract $\tilde{\text{id}}$ from $r$ and add it to $V$
04. $m \leftarrow m \parallel \tilde{\text{id}}$
05. <b>for</b> $i = 1$ to $\text{ut}_w$ and $(i, \text{op}, \text{id}) \in V$
06. <b>if</b> $\text{op} = \text{add}$
07.     Result $\leftarrow \text{Result} \cup \{\text{id}\}$
08. <b>else</b>
09.     Result $\leftarrow \text{Result} \setminus \{\text{id}\}$
10. $\beta \leftarrow \Psi.\text{MacVerify}(K_{\text{mac}}, m, \text{ST}[w], \text{tag})$
11. <b>if</b> $\beta = 1$
12. <b>return</b> Result
13. <b>else</b>
14. <b>return</b> Reject

Figure 7-4: Verify algorithm of  $v\Sigma$ .

### 7.1.2 Fault-tolerance and Correctness of the Search Result

Notice that, upon verification, the Verify algorithm of  $v\Sigma$  returns a set of document identifiers matching the searched keyword  $w$ . Now, if  $v\Sigma$  is sound, that is, the server could not forge the search result, then the result returned by the server is the set of all ids updated for  $w$ . Recall, for an update query  $(\text{op}_i, \text{id}_i, w)$  on keyword  $w$ , all updates to the  $v\Sigma$  is of the form  $(\text{add}, \tilde{\text{id}}_i, w)$ , where  $\tilde{\text{id}}_i$  is a  $\lambda$  bit encoding of the tuple  $(i, \text{op}_i, \text{id}_i)$ , and  $i$  is the number of updates performed on  $w$  till that point. This value  $i$  is kept at a map UT indexed by every keyword  $w$  in the database. After every update  $\text{UT}[w]$ , which was initialized to 1 during the first update call on  $w$ , is incremented by 1. Thus, if the  $v\Sigma$  is sound, we get back all  $\text{UT}[w]$  updates on the keyword  $w$ . Now, the Verify algorithm of  $v\Sigma$  scans each  $\tilde{\text{id}} = (i, \text{op}_i, \text{id}_i)$  in the result set res returned by the server in order of  $i$ . To compute the final result, the Verify algorithm initialises an empty set Result. For every  $\text{op}_i = \text{add}$ , it adds id to the result set (i.e.,  $\text{Result} \leftarrow \text{Result} \cup \{\text{id}\}$ ). And for every  $\text{op}_i = \text{del}$ , it deletes id to the result set (i.e.,  $\text{Result} \leftarrow \text{Result} \setminus \{\text{id}\}$ ). Thus, in the presence of a faulty update, that is add id while id is already in the database, or delete id where id is not in the database, the set operations performed by the algorithm Verify will ensure the correctness of the



search result. Details of the `Verify` algorithm can be found in Figure 7-4.

## 7.2 Comparisons and Actual Overheads

In this section, we compare the overheads of our scheme with some of the existing works in the literature. Table 7.1 provides a summary of selected studies focused on forward and backward private verifiable fault-tolerant DSSE. The table clearly shows that none of the existing works achieves verifiability along with backward privacy and fault tolerance. In this context, our construction is the first to accomplish all of these features.

Next, we compare the overheads of our scheme with those of existing constructions. The incremental multi-set hash proposed in [41] and adopted in [24, 98] involves multiplication over a DDH-hard group, which is a computationally expensive operation. Additionally, these schemes require storing the incremental tags on the server. The most efficient scheme presented in [94] uses an AE to generate and store proofs for every update. The construction proposed in [94] also introduces a generic method to integrate their approach with any forward-secure single-keyword SSE. However, it fails to achieve backward privacy. The instantiation in [94] utilizes the FAST protocol proposed in [87] as the single-keyword SSE. In this scheme, the AE proof generated for each update is outsourced and stored along with the update token of the FAST protocol on the server. During a query, the server retrieves all updates along with their proofs and returns them to the client, who then runs AE verification on each result to ensure correctness. The overhead of the most efficient FVDSSE scheme [94] is twice that of the base single-keyword SSE scheme, FAST, due to the storage and communication of the tags. Additionally, the scheme requires verification to be performed on each value returned by the server. The client-side storage is also doubled compared to what is required by the FAST protocol. In contrast, our protocol incurs no additional server-side storage and communication overhead as we save a fixed length tag for every update on the client side. This amount of storage is allowed in the DSSE scheme, achieving forward and backward privacy and sub-linear

Construction	Computation		Communication		Client Storage	FP	FT	BP
	Search	Update	Search	Update				
[64]	$\mathcal{O}( D )$	$\mathcal{O}(u D )$	$\mathcal{O}(m)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	N	N	N
GVS-Hash [24]	$\mathcal{O}(m + \log W)$	$\mathcal{O}(u \cdot \log W)$	$\mathcal{O}(m + \log W)$	$\mathcal{O}(\log W)$	$\mathcal{O}(1)$	Y	N	N
GVS-Acc-Pairing [24]	$\mathcal{O}(m)$	$\mathcal{O}(uW^\epsilon)$	$\mathcal{O}(m + \log W)$	$\mathcal{O}(\log W)$	$\mathcal{O}(1)$	Y	N	N
GVS-Acc-RSA [24]	$\mathcal{O}(m + W^\epsilon)$	$\mathcal{O}(u)$	$\mathcal{O}(m + \log W)$	$\mathcal{O}(\log W)$	$\mathcal{O}(1)$	Y	N	N
[98]	$\mathcal{O}(u)$	$\mathcal{O}(1)$	$\mathcal{O}(m)$	$\mathcal{O}(1)$	$\mathcal{O}(\lambda W)$	Y	N	N
Verifiable Linear SPS [24]	$\mathcal{O}(\alpha + \log N)$	$\mathcal{O}(u \cdot \log^2 N)$	$\mathcal{O}(m + \log N)$	$\mathcal{O}(\log N)$	$\mathcal{O}(\lambda \log N)$	Y	Y	N
Verifiable Sublinear SPS [24]	$\mathcal{O}(m \cdot \log^3 N)$	$\mathcal{O}(u \cdot \log^2 N)$	$\mathcal{O}(m + \log N)$	$\mathcal{O}(\log N)$	$\mathcal{O}(\lambda \log N)$	Y	Y	N
[94]	$\mathcal{O}(u)$	$\mathcal{O}(1)$	$\mathcal{O}(u)$	$\mathcal{O}(1)$	$\mathcal{O}(W \log  D )$	Y	Y	N
$\mathbf{v}\Sigma$ [Section 7.1]	$\mathcal{O}(u)$	$\mathcal{O}(1)$	$\mathcal{O}(u)$	$\mathcal{O}(1)$	$\mathcal{O}(\lambda W)$	Y	Y	Y

Table 7.1: Comparison of our schemes with existing VDSSE schemes. FP, BP and FT stand for forward privacy, backward privacy, and fault-tolerant, respectively. ‘Computation’ and ‘Communication’ refer to the computational and communication complexity, respectively.  $W$  is the number of distinct keywords in the database, and  $|D|$  denotes the number of documents in the database.  $N$  is the number of keyword-identifier pairs.  $0 < \epsilon < 1$  is a fixed constant, and  $\alpha$  denotes the number of times the queried keyword was historically added to the database.  $m$  denotes the number of matching entries for the searched keyword.

search [23, 25, 49, 38, 90]. Furthermore, storing the tag on the client side does not compromise security, as demonstrated in the security proof (Theorem 7.3.2) of our scheme, where the adversary is required to store and submit the tags at the time of query and forgery. We have added a row in Table 7.1 which reflects the overhead of our scheme in comparison to others. As our scheme can be instantiated with any DSSE scheme, to make this type of comparison we have to assume a particular construction. In this case, we assume the construction of [38] or [49].

It is clear from the description of  $\mathbf{v}\Sigma$  any updatable MAC scheme satisfying the Definition 6.3 can be used as  $\Psi$  in our construction  $\mathbf{v}\Sigma$ . So, we can use both our proposals ConCatU and XoRU as our  $\Psi$  in the construction of  $\mathbf{v}\Sigma$ . Below we describe each extra overhead of  $\mathbf{v}\Sigma$  while clubbed with each of the ConCatU or XoRU.

**Overhead due to ConCatU.** From the description of the efficiency of ConCatU in Section 6.4.1, it is clear that a single update call for a message  $m$  will require  $\lceil (|m| + 2)/n \rceil \approx \lceil |m|/n \rceil + 2$  block-cipher calls and one multiplication in  $\mathbb{F}_p$ , where  $|m|$  is the size of the message and  $n$  is the block size. This also requires the ConCatU algorithm to save an intermediate value  $p$ , which is of size  $n$  at the client side. However, one might argue that ConCatU requires storing the list  $L$  of the size of every update made

to the message. The size of the list  $L$  is the order of the number of updates made for the keyword. So storing the list  $L$  will defeat the purpose of SSE.

We like to mention that, in the context of SSE, every update for a keyword  $w$  has a fixed size, that is, the size  $\widetilde{\text{id}}$ . So the information of  $L$  is implicit in the context of  $\mathbf{v}\Sigma$ , and the algorithm **ConCatU** in the context of  $\mathbf{v}\Sigma$  does not need to explicitly store the list  $L$  as its state. So the extra overhead of  $\mathbf{v}\Sigma$  over the base SSE scheme  $\Sigma$  is the tag value for every keyword and the value of  $p$ , and nothing else. At this point, we would like to mention that most SSE schemes, which are forward and backward private [23, 25, 87, 38, 90], store the number update for every keyword, that is, the map **UT** at the client end. This can also be verified from Table 7.1. However, if the base SSE does not require to store **UT** then another extra overhead would be of order  $\log |D|$ .

**Overhead due to XoRU.** As discussed in the special case of Section 6.5.1, where **XoRU** is used only for concatenation, the message is always updated with a string of the form  $0^{|M|}||m$ , where  $M$  is the previous message and  $m$  is the update which is to be concatenated with the current message. In scenarios where,  $|m| = kn$ , with  $k$  being a small constant, the update would require only  $(2k + 2)$  multiplications in  $\mathbb{F}_{2^n}$  and 2 AES computations, provided that the value of  $h^{\ell_1}$  is stored for every keyword. Please refer to Section 6.5.1 for a detailed analysis.

In the case of an SSE update, we update with only one  $(\text{id}, w)$  pair at a time. As a result, for each update the message length increases by only one  $n$ -bit block, i.e., the message  $m$  which is to be concatenated, is always  $n$ -bits long. Consequently, the Equation (6.19) boils down to the following:

$$(\ell_2 + 4)T_{\text{add}} + 4T_{\text{mult}} + 2T_{\text{AES}},$$

where  $T_{\text{add}}$  and  $T_{\text{mult}}$  are the times taken for a multiplication and addition in  $\mathbb{F}_{2^n}$  respectively and  $T_{\text{AES}}$  is the time taken for one call to AES. This amount of computation can be easily handled by a client in any reasonable computing platform.

## 7.3 Security

In the construction of the FVDSSE scheme  $v\Sigma$ , we keep the tag generated for each update to the client itself. This is done to reduce the overhead of our scheme. However, this is not necessary, and the client may choose to outsource the tag generated for every update. Keeping or outsourcing the tag does not affect the security of the SSE. We prove both of our security requirements, i.e., adaptive security and soundness, by allowing the adversary to see the tag along with the update token. We consider that the SSE stores the tag in a similar fashion as it stores the `id`.

### 7.3.1 Adaptive Security of $v\Sigma$

The adaptive security of SSE is established by capturing the maximum permissible leakage of a scheme and designing a simulator that, using these leakage functions, can generate a transcript indistinguishable from the one produced by the real-world execution of the protocol [60, 23, 38]. We assume that the leakage of the base SSE scheme  $\Sigma$  is  $\mathcal{L}_\Sigma = (\mathcal{L}_{\Sigma, \text{Setup}}, \mathcal{L}_{\Sigma, \text{Search}}, \mathcal{L}_{\Sigma, \text{Update}})$  as defined in Definition 5.2.2.

Our design,  $v\Sigma$ , mirrors the setup phase of  $\Sigma$  exactly, resulting in identical setup leakage for both. In the update phase, along with the update token of  $\Sigma$ , we also send a `tag` for the update. The `tag` can be modelled as the output of a PRF following Proposition 6.4.1 and 6.5.2 in the real-world execution of our protocol. Therefore, our scheme has the same update leakage as the base SSE scheme  $\Sigma$ , as the `tag` value does not convey any information about the input of the update function. Finally, in the search protocol, we exactly follow the search protocol of the base SSE scheme. Thus, the leakage of protocol  $v\Sigma$  is  $\mathcal{L}_{v\Sigma} = \mathcal{L}_\Sigma$ . Below, we formally prove the confidentiality and soundness of our scheme.

**Theorem 7.3.1.** *If  $\Sigma$  is a  $\mathcal{L}_\Sigma$ -adaptive secure forward and backward private DSSE scheme, then our protocol  $v\Sigma$  is also  $\mathcal{L}_\Sigma$ -adaptive secure forward and backward private DSSE.*

*Proof.* As  $\Sigma$  is an  $\mathcal{L}_\Sigma$ -adaptive secure forward and backward private DSSE scheme,

with the leakage function  $\mathcal{L}_\Sigma$ , it implies the existence of a simulator  $\mathcal{S}$  that can, using the leakage function  $\mathcal{L}_\Sigma$ , simulate a transcript  $T_{\text{sim}}^\Sigma$  indistinguishable from the real-world transcript  $T_{\text{real}}^\Sigma$  of the protocol  $\Sigma$ .

In this proof, we will demonstrate that a simulator  $\tilde{\mathcal{S}}$  can efficiently generate a transcript  $T_{\text{sim}}^{\mathbf{v}\Sigma}$  of our protocol using the transcript  $T_{\text{sim}}^\Sigma$  produced by the simulator  $\mathcal{S}$ . Furthermore, we will show that any efficient distinguisher that can distinguish between the transcript  $T_{\text{sim}}^{\mathbf{v}\Sigma}$  generated by  $\tilde{\mathcal{S}}$  and the real-world transcript of our protocol  $T_{\text{real}}^{\mathbf{v}\Sigma}$  can always be used to construct a distinguisher for  $\Sigma$ , or a PRF used to model the UdMAC  $\Psi$  used in our construction.

Hence, if  $\Sigma$  is  $\mathcal{L}_\Sigma$ -adaptive secure, then our scheme is also  $\mathcal{L}_\Sigma$ -adaptive secure, as the leakage of  $\Sigma$  is exactly the same as the leakage of our scheme  $\mathbf{v}\Sigma$ . Therefore, the theorem. We demonstrate our proof using a sequence of games.

**Game  $G_0$ :** At the beginning of the protocol, the setup and every search query we follow the protocol  $\mathbf{v}\Sigma$ . For each update query  $(\text{op}, \text{id}, w)$  made by the adversary  $\mathcal{A}$ , the challenger maintains a list **UT** that tracks the number of updates performed on the keyword  $w$ . For every update query  $(\text{op}, \text{id}, w)$  to  $\Sigma$ , the challenger replaces  $\text{id}$  with  $\tilde{\text{id}} = (\text{UT}[w], \text{op}, \text{id})$  and sets the operation  $\text{op}$  to **add**, as shown in Figure 7-2. The  $\Sigma$  outputs an update token for  $(\text{add}, \tilde{\text{id}}, w)$ .

Also, for every update query, the challenger generates a **tag** using the UdMAC  $\Psi$ , as described in lines 5 to 8 of Figure 7-2. Finally, the challenger sends the update token generated by  $\Sigma$  for input  $(\text{add}, \tilde{\text{id}}, w)$  and the tag generated by  $\Psi$  protocol. Thus,

$$\Pr[G_0 = 1] = \Pr[\text{SSEReal}_{\mathcal{A}}^{\mathbf{v}\Sigma}(\lambda, q) = 1].$$

**Game  $G_1$ :** In the next game, for every update query, we randomly sample a value from  $\{0, 1\}^\tau$  and replace the value of **tag** with this newly sampled value. Now for any adversary  $\mathcal{A}$  that can distinguish between game  $G_0$  and game  $G_1$ , we can construct an adversary  $\mathcal{B}_1$  which can distinguish the output  $\Psi$  from a PRF, which contradicts the Proposition 6.4.1 and 6.5.2. Thus,

$$|\Pr[G_1 = 1] - \Pr[G_2 = 1]| \leq \text{Adv}_{\Psi, \mathcal{B}_1}^{\text{prf}} \leq \text{negl}(\lambda).$$

Recall that,  $\Sigma$  is a  $\mathcal{L}_\Sigma$ -adaptively secure SSE. Thus, for the real-world transcript produced by executing the protocol  $\Sigma$ , (denoted by  $T_{\text{real}}^\Sigma$ ), there exists a simulator  $\mathcal{S}$  that, with the assistance of  $\mathcal{L}_\Sigma$ , can generate a transcript  $T_{\text{sim}}^\Sigma$  that is indistinguishable from  $T_{\text{real}}^\Sigma$ . Now, when all the updates of the transcript  $T_{\text{real}}^\Sigma$  are appended with a randomly sampled value from  $\{0, 1\}^\tau$ , the newly formed transcript is exactly the real-world execution of the protocol  $\mathbf{v}\Sigma$ . We denote this transcript as  $T_{\text{real}}^{\mathbf{v}\Sigma}$ .

**Construction of  $\tilde{\mathcal{S}}$ :** Finally, we construct a simulator  $\tilde{\mathcal{S}}$  which can simulate an indistinguishable transcript from  $T_{\text{real}}^{\mathbf{v}\Sigma}$  with the help of the indistinguishable transcript  $T_{\text{sim}}^\Sigma$  output by the simulator  $\mathcal{S}$ . The construction of  $\tilde{\mathcal{S}}$  is the following.

In the setup phase, the challenger sends an empty database to the adversary, which is exactly how the simulator  $\mathcal{S}$  will simulate it.

The search queries of  $\mathbf{v}\Sigma$  are simulated exactly the same way as the search queries of  $T_{\text{sim}}^\Sigma$ , are simulated. This simulation is correct as in  $\mathbf{v}\Sigma$  we use the search protocol of  $\Sigma$  unaltered.

For the update query, the simulator  $\tilde{\mathcal{S}}$  replace all the update tokens of  $T_{\text{real}}^{\mathbf{v}\Sigma}$  (except the tag part) with the update token as simulated by  $\mathcal{S}$ . The tag is simulated with randomly sampled strings from  $\{0, 1\}^\tau$ . The final transcript is exactly the transcript of  $T_{\text{sim}}^{\mathbf{v}\Sigma}$ . Formally,

$$\left| \Pr[G_2 = 1] - \Pr \left[ \text{SSEIdeal}_{\mathcal{A}, \tilde{\mathcal{S}}}^{\mathbf{v}\Sigma}(\lambda) = 1 \right] \right| \leq \text{Adv}_{\mathcal{B}_2, \mathcal{S}}^\Sigma \leq \text{negl}(\lambda).$$

Now any efficient distinguisher, distinguishing between the real-world transcript  $T_{\text{real}}^{\mathbf{v}\Sigma}$  and the simulate transcript  $T_{\text{sim}}^{\mathbf{v}\Sigma}$  of  $\mathbf{v}\Sigma$  output by  $\tilde{\mathcal{S}}$ , we can construct a distinguisher that distinguishes transcript simulated by  $\mathcal{S}$  from the real-world execution of  $\Sigma$ . Thus,

$$\left| \Pr \left[ \text{SSEReal}_{\mathcal{A}, \tilde{\mathcal{S}}}^{\mathbf{v}\Sigma}(\lambda) = 1 \right] - \Pr \left[ \text{SSEIdeal}_{\mathcal{A}, \tilde{\mathcal{S}}}^{\mathbf{v}\Sigma}(\lambda) = 1 \right] \right| \leq \text{negl}(\lambda).$$

□

### 7.3.2 Forward and Backward Privacy of $v\Sigma$

A DSSE scheme is said to be forward and backward private if the leakage of that scheme is bounded by some well defined leakage functions defined in Section 2.5.5. Now, in Theorem 7.3.1 we proved that if the base SSE scheme  $\Sigma$  used in our  $v\Sigma$  construction has leakage  $\mathcal{L}_\Sigma$ , then our FVDSSE scheme  $v\Sigma$  also has the same leakage  $\mathcal{L}_\Sigma$ . As the base SSE scheme  $\Sigma$  is forward and backward private, the  $\mathcal{L}_\Sigma$  is subsumed by the leakage defined in the forward and backward privacy Definition 2.5.5 and 2.5.6 respectively. Thus leakage of  $v\Sigma$  is also subsumed by the leakage defined in the Definition 2.5.5 and 2.5.6. So, by Definition 2.5.5 and 2.5.6, our construction  $v\Sigma$  also achieves forward privacy and the same level of backward privacy as achieved by the base SSE scheme  $\Sigma$ .

### 7.3.3 Soundness of $v\Sigma$

While arguing about soundness of  $v\Sigma$  we let the adversary play the soundness game of DSSE as defined in Definition 2.5.3. We prove in the following theorem that if  $v\Sigma$  is instantiated with a correct SSE scheme  $\Sigma$  as in Definition 2.5.2, then given any adversary  $\mathcal{A}$  breaking the soundness of  $v\Sigma$  as in Definition 2.5.3, we can construct another adversary  $\mathcal{B}$  which can either break the authenticity of UdMAC  $\Psi$  or can break the correctness of  $\Sigma$ .

The adversary  $\mathcal{A}$  initially submits a database and receives the encrypted database. In the subsequent phase,  $\mathcal{A}$  makes a search and update query of its choice. Adversary  $\mathcal{B}$  acts as a challenger for  $\mathcal{A}$  and responds to the query of  $\mathcal{A}$  by running an SSE instance on its own. For the tag, it sends along with the update query,  $\mathcal{B}$  uses its challenger to compute the tag. Finally, any successful forgery by  $\mathcal{A}$  according to Definition 2.5.3 we show that  $\mathcal{B}$  can construct a successful forgery for  $\Psi$ .

**Theorem 7.3.2.** *If  $\Sigma$  is a correct DSSE scheme according to Definition 2.5.2, then for any ppt adversary  $\mathcal{A}$  which can break the soundness of  $v\Sigma$  as in Definition 2.5.3, there exists another ppt adversary  $\mathcal{B}$  which will break the security of UdMAC  $\Psi$  as defined in Section 6.3.1.*

*Proof.* We establish the soundness of our scheme  $\mathbf{v}\Sigma$  under the assumption that  $\Sigma$  is always correct. This implies that given the latest search token  $\mathbf{stk}$  for a keyword  $w$ , the server can always accurately retrieve all updates made to EDB for keyword  $w$  up to that point. Failure to satisfy this condition would result in the soundness of  $\mathbf{v}\Sigma$ , reducing to the correctness of  $\Sigma$ .

Let's assume a ppt adversary  $\mathcal{A}$ , which makes polynomial many queries to the SSE  $\mathbf{v}\Sigma$ , in the soundness game of Definition 2.5.3 and wins. We then construct another UDMAC adversary  $\mathcal{B}$  that uses  $\mathcal{A}$  as a subroutine to break the security of  $\Psi$  as described in Section 6.3.1. We follow the security definition of 2.5.3 and allow the adversary to manipulate the search results arbitrarily. Adversary  $\mathcal{B}$  acts as the challenger of  $\mathcal{A}$  and plays the soundness game defined in Definition 2.5.3.  $\mathcal{B}$  runs an SSE instance on its own and only uses its challenger for  $\Psi$  to compute the tags.  $\mathcal{B}$  runs  $\Sigma$ .Setup to obtain a key, a state, and an encrypted database EDB for the SSE.  $\mathcal{B}$  then sends EDB to  $\mathcal{A}$ . Here, we assume the common practice of outsourcing an empty database during setup and then updating it. In cases where the initial database is not empty, we can still outsource an empty database and update it subsequently.

Now, for every update query  $(\mathbf{op}, \mathbf{id}, w)$  that adversary  $\mathcal{A}$  makes, adversary  $\mathcal{B}$  runs  $\Sigma$ .Update<sub>c</sub> and returns the  $\mathbf{utk}$  to adversary  $\mathcal{A}$  following soundness game of Figure 2-2. Along with the  $\mathbf{utk}$ , adversary  $\mathcal{B}$ , following steps 5 to 8 of the update algorithm of  $\mathbf{v}\Sigma$  in Figure 7-2, generates the tag for the update query. To generate the tag, adversary  $\mathcal{B}$ , queries its challenger for the UDMAC  $\Psi$ . To do this, adversary  $\mathcal{B}$  maintains the state ST and UT described in Figure 7-2. The adversary  $\mathcal{B}$  sends the updated tag received from its challenger along with the update token  $\mathbf{utk}$ . For a search query on keyword  $w$  from adversary  $\mathcal{A}$ , adversary  $\mathcal{B}$  runs the  $\Sigma$ .Search protocol and returns the search token  $\mathbf{stk}$  to  $\mathcal{A}$ .

Finally, the adversary stops by outputting a forgery for any search query of its choice. However, in response to any search query, it must return the correct number of search results corresponding to the number of updates made for the keyword. This can easily be checked by the adversary  $\mathcal{B}$  using the map UT that it stores. Sending an equal number of search results is necessary as the correctness of  $\Psi$  can verify any



partial search result. In this security game, the adversary submits a search result along with a tag as its forgery.

For every update query, adversary  $\mathcal{B}$  also maintains its own state  $\mathbf{D}$ , initialized to empty. The state  $\mathbf{D}$  is a list of sets for each keyword  $w$ . For every update operation  $(\text{op}, \text{id}, w)$ , adversary  $\mathcal{B}$  performs the following steps:

- If  $\text{op} = \text{add}$ ,  $\mathbf{D}[w] \leftarrow \mathbf{D}[w] \cup \{\text{id}\}$ .
- If  $\text{op} = \text{del}$ ,  $\mathbf{D}[w] \leftarrow \mathbf{D}[w] \setminus \{\text{id}\}$ .

Finally, for a forgery attempt by  $\mathcal{A}$  for a search query on a keyword  $w$ ,  $\mathcal{B}$  constructs its forgery as follows. Let the forgery of  $\mathcal{A}$  be  $\text{res} = (L, t)$  consists of two components: (i) the list of all updates  $L$  (or an encrypted list, depending on the description of  $\Sigma$ ) made for the keyword  $w$  up to that point, and (ii) a tag  $t$ . If  $|L| \neq \text{UT}[w]$ , then  $\text{res}$  is not a valid forgery by  $\mathcal{A}$ , and  $\mathcal{B}$  aborts. Otherwise, the entries of  $L$  (or the decrypted version) are of the form  $(i, \text{op}, \text{id})$ , where  $i$  is the update sequence number,  $\text{op}$  is the operation corresponding to that update, and  $\text{id}$  is the identifier involved in that update.

Now the adversary  $\mathcal{B}$  constructs its forgery as in Figure 7-5. Notice, if  $\tilde{\mathbf{D}} = \mathbf{D}[w]$  in

<p><b>Forgery of <math>\mathcal{B}</math>:</b></p> <pre> 01. <math>m \leftarrow w</math> 02. <math>\tilde{\mathbf{D}} \leftarrow \emptyset</math> 03. <b>for</b> <math>i = 1</math> to <math> L </math>, <b>and</b> <math>\ell = i \parallel \text{op} \parallel \text{id} \in L</math> <b>do</b> 04.   <math>m \leftarrow m \parallel \ell</math> 05.   <b>if</b> <math>\text{op} = \text{add}</math> 06.     <math>\tilde{\mathbf{D}} \leftarrow \tilde{\mathbf{D}} \cup \{\text{id}\}</math> 07.   <b>else if</b> <math>\text{op} = \text{del}</math> 08.     <math>\tilde{\mathbf{D}} \leftarrow \tilde{\mathbf{D}} \setminus \{\text{id}\}</math> 09.   <b>end if</b> 10. <b>end for</b> 11. <b>if</b> <math>\tilde{\mathbf{D}} = \mathbf{D}[w]</math> 12.   <b>abort</b> 13. <b>return</b> <math>(m, \text{ST}[w], t)</math> </pre>
---

Figure 7-5: Forger game of FVDSSE.

Line No. 11 of Figure 7-5, then  $\text{res}$  is not a valid forgery attempt by  $\mathcal{A}$ , and  $\mathcal{B}$  aborts.

If the forgery submitted by  $\mathcal{A}$  is a valid forgery then adversary  $\mathcal{B}$  constructs its forgery as  $(m, \text{ST}[w], t)$  and submits to its challenger of  $\Psi$ . Notice that, if  $\mathcal{A}$  submits a valid forgery, that is  $\tilde{D} \neq D[w]$  then the message constructed in Line Number 03 to 10 of Figure 7-5 has never been queried by  $\mathcal{B}$  to its challenger. Thus,  $(m, \text{ST}[w], t)$  is a valid forgery attempt by  $\mathcal{B}$ . So,  $\mathcal{A}$  breaking the soundness of  $v\Sigma$  implies,  $\mathcal{B}$  certainly breaks the  $\Psi$ -MAC.

At this point, we allow the underlying SSE to be correct except with negligible probability. By  $\Sigma_{\text{corr}}$  we denote the event that the  $\Sigma$  is correct. So,

$$\begin{aligned} \Pr[\mathcal{A} \text{ success}] &= \Pr[\mathcal{B} \text{ wins} \mid \Sigma_{\text{corr}}] \Pr[\Sigma_{\text{corr}}] + \Pr[\mathcal{B} \text{ wins} \mid \overline{\Sigma_{\text{corr}}}] \Pr[\overline{\Sigma_{\text{corr}}}] \\ &\leq \Pr[\mathcal{B} \text{ wins} \mid \Sigma_{\text{corr}}] + \Pr[\overline{\Sigma_{\text{corr}}}] \\ \text{Adv}_{\mathcal{A}}^{\text{Sound}, \Sigma} &\leq \text{Adv}_{\mathcal{B}}^{\Psi} + \text{Adv}_{\mathcal{A}}^{\text{Corr}, \Sigma}. \end{aligned}$$

□

## 7.4 Final Remark

In this chapter, we addressed the open problem of designing a fault-tolerant verifiable SSE, which is both forward and backward private. We provided the first generic construction of FVDSSE using a generic DSSE which is only secure against an honest-but-curious adversary and our newly proposed updatable-MAC. Our construction is generic and requires no extra communication and storage on the server side, with the expense of one tag per keyword stored on the client side. Compared to other state-of-the-art SSE which store and communicate one tag per update on the server side. A storage of  $\mathcal{O}(|W|)$  is considered acceptable in the SSE literature. In fact, almost all DSSE schemes keep a client state of  $\mathcal{O}(|W|)$ .

## Conclusion

In this thesis, we explored the security and efficiency of SSE under different adversarial settings and for various classes of queries.

To recall, we started this thesis with a four-way goal of (i) storing the inverted index in a succinct manner to reduce overheads of SSE, (ii) possibly supporting richer queries while maintaining efficiency, and (iii) & (iv) provide correct search result in presence of a malicious server and a client producing faulty update. To this end, we conclude that we have addressed all of the above objectives in a satisfactory way.

**To address the first problem,** we describe a novel way to make SSE schemes space efficient. Our scheme converts a given database into a different representation, which, on average, results in significant space savings. The smaller representation also results in reduced search time and response sizes. Our scheme depends on representing the set  $\text{db}(w)$  using binary trees where the leaves are labeled. This representation has interesting combinatorial properties, which we explore in detail. Our experiments show that our representation results in smaller index sizes with minimal extra overhead. Our scheme can be used with any secure single keyword SSE, and our scheme, being just a pre-processing step, retains the security of the base SSE.

**Regarding our second goal,** we have proposed a generic DSSE scheme that allows conjunctive queries while being both forward and backward private. Our scheme is specifically crafted to protect *non-modifiable* documents that, once uploaded, their keywords remain unaltered. This is a common characteristic in various applications, including biometric databases, digital archives, legal documents, medical records,

etc. Our construction does allow the keyword set associated with a document to be modified, but with some extra overhead. With efficiency closest to OXT (among known dynamic schemes) and by ensuring both forward and backward privacy, our proposed construction offers an ideal solution for preserving the privacy of dynamic datasets while facilitating conjunctive queries.

**To achieve our third and fourth objective**, we first introduce a novel message authentication scheme called the updatable MAC. Unlike traditional MAC methods, where updating a message necessitates recomputing the tag for the entire updated message, our updatable MAC system maintains a small state for each message. When a message is modified, our scheme recalculates the tag using only the stored state and the previous tag, eliminating the need to recompute the tag for the entire updated message. We establish specific security criteria for these MACs and propose two distinct MAC designs tailored for two different types of update functions. Our constructions support, updating the message through concatenation and XOR-difference.

**Finally**, we focus on designing a secure SSE scheme that is resilient against a malicious server, starting from an SSE scheme that is only secure against an honest-but-curious server. By using the updatable MAC introduced in this thesis, we convert a generic single keyword SSE scheme—originally forward and backward private against honest-but-curious adversaries—into a scheme that also offers security against malicious servers while maintaining forward and backward privacy. This approach is the first of its kind. We also achieve it without increasing communication or storage on the server’s end from the base SSE construction. We only store a constant size state for every keyword at the client’s end, which is well-accepted in SSE literature. On the contrary, state-of-the-art SSE only provides forward privacy and needs to store and communicate one tag per update.

Throughout our thesis, we provide generic constructions that use an SSE scheme in a black-box manner to create new SSE systems with desired security or performance requirements. This increases the acceptability of these works as these works can be integrated into existing systems without needing to modify the whole structure.

## Future Directions

We list some aspects of our work which we plan to explore in the near future:

1. Our current study on cover based SSE involves only single keyword select queries. It seems that with some small modifications, it may be possible to equip the scheme to handle range queries.
2. We compute an estimate of the average size of a cover in the case of pure covers. Our estimate is quite accurate, as demonstrated by the experiments, but we failed to obtain a semantically useful analytical expression of the estimate. It would be nice to try to estimate the average size in some alternative way to obtain a more meaningful expression.
3. We would like to see the performance of our cover-based scheme in a real environment, which would require careful implementation of our scheme paired with a base SSE in a real cloud environment.
4. An attractive property of the cover based SSE scheme is that given the response of a query produced by our scheme, it is not possible to know the number of documents that match the query directly. This hints that our scheme is volume hiding to some extent. But, a thorough investigation of this property is required. We think that by using some additional randomization, it may be possible to convert our scheme into a volume hiding scheme against a large class of powerful adversaries.
5. The recent C-SSE scheme of [78] is similar to [29]. However, [78] pre-computes all possible 2-conjunctions to facilitate conjunctive search. This is difficult to achieve in a dynamic setting where the database is not known beforehand. It will be interesting to use [78] instead of OXT for  $\Gamma$  without a significant loss in efficiency.
6. The OXT framework has been extended to support multi-client settings in a modular black-box fashion [57]. Our construction also uses OXT in a black-box

manner without altering its functionality. Therefore, our generic scheme may be extended to support multi-clients. We leave this as another future work.

7. We provide constructions for updatable macs for two kinds of updates, namely, concatenation and xor difference. An interesting theoretical question could be to ask which kinds of updates can be handled by updatable MACs. Also it would be interesting to design schemes for other update functions.
8. In the context of our work on verifiable SSE, a natural question is: Is it possible to use updatable MAC to construct a verifiable SSE, which supports more complex queries than a single keyword search? We wish to take this up in near future.

## Bibliography

- [1] Mohamed Ahmed Abdelraheem, Tobias Andersson, Christian Gehrman, and Cornelius Glackin. Practical attacks on relational databases protected via searchable encryption. In Liqun Chen, Mark Manulis, and Steve A. Schneider, editors, *Information Security - 21st International Conference, ISC 2018, Guildford, UK, September 9-12, 2018, Proceedings*, volume 11060 of *Lecture Notes in Computer Science*, pages 171–191. Springer, 2018.
- [2] Mohamed Ahmed Abdelraheem, Christian Gehrman, Malin Lindström, and Christian Nordahl. Executing boolean queries on an encrypted bitmap index. In *Proceedings of the 2016 ACM on Cloud Computing Security Workshop, CCSW '16*, page 11–22. Association for Computing Machinery, 2016.
- [3] Mohamed Ahmed Abdelraheem, Theo Henault Sanjay Bhattacharjee, and Avishek Majumder. Conjunctive dynamic searchable symmetric encryption – a generic framework. *The IACR Communications in Cryptology (CiC) (submitted)*.
- [4] Shweta Agrawal, Subhashis Banerjee, and Subodh Sharma. Privacy and security of aadhaar: A computer science perspective. *Economic and Political Weekly*, 52(37):93–102, 2017.
- [5] Shweta Agrawal, Sanjay Bhattacharjee, Duong Hieu Phan, Damien Stehlé, and Shota Yamada. Efficient public trace and revoke from standard assumptions: Extended abstract. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, page 2277–2293. Association for Computing Machinery, 2017.
- [6] Ghous Amjad, Seny Kamara, and Tarik Moataz. Injection-secure structured and searchable symmetric encryption. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology - ASIACRYPT 2023 - 29th International Conference on the Theory and Application of Cryptology and Information Security*,

Guangzhou, China, December 4-8, 2023, *Proceedings, Part VI*, volume 14443 of *Lecture Notes in Computer Science*, pages 232–262. Springer, 2023.

- [7] Ghous Amjad, Sarvar Patel, Giuseppe Persiano, Kevin Yeo, and Moti Yung. Dynamic volume-hiding encrypted multi-maps with applications to searchable encryption. *Proc. Priv. Enhancing Technol.*, 2023(1):417–436, 2023.
- [8] Nishant Anand. New principles for governing aadhaar: Improving access and inclusion, privacy, security, and identity management. *Journal of Science Policy & Governance*, 18(01):1–14, 2021.
- [9] John Christopher Anderson. Transmitting legal documents over the internet: How to protect your client and yourself. *Rutgers Computer & Tech. LJ*, 27:1, 2001.
- [10] Vivek Arte, Mihir Bellare, and Louiza Khati. Incremental cryptography revisited: Prfs, nonces and modular design. In Karthikeyan Bhargavan, Elisabeth Oswald, and Manoj Prabhakaran, editors, *Progress in Cryptology - INDOCRYPT 2020 - 21st International Conference on Cryptology in India, Bangalore, India, December 13-16, 2020, Proceedings*, volume 12578 of *Lecture Notes in Computer Science*, pages 576–598. Springer, 2020.
- [11] Arnab Bag, Debadrita Talapatra, Ayushi Rastogi, Sikhar Patranabis, and Debdeep Mukhopadhyay. Two-in-one-sse: Fast, scalable and storage-efficient searchable symmetric encryption for conjunctive and disjunctive boolean queries. *Proc. Priv. Enhancing Technol.*, 2023(1):115–139, 2023.
- [12] Mihir Bellare, Oded Goldreich, and Shafi Goldwasser. Incremental cryptography: The case of hashing and signing. In Yvo Desmedt, editor, *Advances in Cryptology - CRYPTO '94, 14th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1994, Proceedings*, volume 839 of *Lecture Notes in Computer Science*, pages 216–233. Springer, 1994.
- [13] Mihir Bellare, Oded Goldreich, and Shafi Goldwasser. Incremental cryptography and application to virus protection. In Frank Thomson Leighton and Allan Borodin, editors, *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing, 29 May-1 June 1995, Las Vegas, Nevada, USA*, pages 45–56. ACM, 1995.
- [14] Sanjay Bhattacharjee and Palash Sarkar. Complete tree subset difference broadcast encryption scheme and its analysis. *Des. Codes Cryptogr.*, 66(1-3):335–362, 2013.
- [15] John Black and Phillip Rogaway. CBC MACs for arbitrary-length messages: The three-key constructions. In Mihir Bellare, editor, *Advances in Cryptology - CRYPTO 2000, 20th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2000, Proceedings*, volume 1880 of *Lecture Notes in Computer Science*, pages 197–215. Springer, 2000.



- [16] John Black and Phillip Rogaway. A block-cipher mode of operation for parallelizable message authentication. In Lars R. Knudsen, editor, *Advances in Cryptology - EUROCRYPT 2002, International Conference on the Theory and Applications of Cryptographic Techniques, Amsterdam, The Netherlands, April 28 - May 2, 2002, Proceedings*, volume 2332 of *Lecture Notes in Computer Science*, pages 384–397. Springer, 2002.
- [17] Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7):422–426, 1970.
- [18] Dan Boneh, Saba Eskandarian, Sam Kim, and Maurice Shih. Improving speed and security in updatable encryption schemes. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part III*, volume 12493 of *Lecture Notes in Computer Science*, pages 559–589. Springer, 2020.
- [19] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 410–428. Springer, 2013.
- [20] Dan Boneh and Victor Shoup. A graduate course in applied cryptography. (this reference is still in the form of an incomplete draft.). 2023.
- [21] Joppe W. Bos, Kristin Lauter, and Michael Naehrig. Private predictive analysis on encrypted medical data. *Journal of Biomedical Informatics*, 50:234–243, 2014. Special Issue on Informatics Methods in Medical Privacy.
- [22] Christoph Bösch, Richard Brinkman, Pieter H. Hartel, and Willem Jonker. Conjunctive wildcard search over encrypted data. In Willem Jonker and Milan Petkovic, editors, *Secure Data Management - 8th VLDB Workshop, SDM 2011, Seattle, WA, USA, September 2, 2011, Proceedings*, volume 6933 of *Lecture Notes in Computer Science*, pages 114–127. Springer, 2011.
- [23] Raphael Bost.  $\sum\text{o}\phi\text{o}\varsigma$ : Forward secure searchable encryption. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, page 1143–1154. Association for Computing Machinery, 2016.
- [24] Raphael Bost, Pierre-Alain Fouque, and David Pointcheval. Verifiable dynamic symmetric searchable encryption: Optimality and forward security. *IACR Cryptol. ePrint Arch.*, 2016:62, 2016.
- [25] Raphaël Bost, Brice Minaud, and Olga Ohrimenko. Forward and backward private searchable encryption from constrained cryptographic primitives. In

- Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, page 1465–1482. Association for Computing Machinery, 2017.
- [26] J Lawrence Carter and Mark N Wegman. Universal classes of hash functions. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 106–112, 1977.
- [27] David Cash, Paul Grubbs, Jason Perry, and Thomas Ristenpart. Leakage-abuse attacks against searchable encryption. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, page 668–679. Association for Computing Machinery, 2015.
- [28] David Cash, Joseph Jaeger, Stanislaw Jarecki, Charanjit S. Jutla, Hugo Krawczyk, Marcel-Catalin Rosu, and Michael Steiner. Dynamic searchable encryption in very-large databases: Data structures and implementation. In *21st Annual Network and Distributed System Security Symposium, NDSS 2014, San Diego, California, USA, February 23-26, 2014*. The Internet Society, 2014.
- [29] David Cash, Stanislaw Jarecki, Charanjit Jutla, Hugo Krawczyk, Marcel-Cătălin Roşu, and Michael Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *Advances in Cryptology – CRYPTO 2013*, pages 353–373. Springer Berlin Heidelberg, 2013.
- [30] Qi Chai and Guang Gong. Verifiable symmetric searchable encryption for semi-honest-but-curious cloud servers. In *2012 IEEE International Conference on Communications (ICC)*, pages 917–922, 2012.
- [31] Debrup Chakraborty, Avishek Majumder, and Subhabrata Samajder. Making searchable symmetric encryption schemes smaller and faster. *Int. J. Inf. Sec.*, 24(1):10, 2025.
- [32] Debrup Chakraborty and Palash Sarkar. On modes of operations of a block cipher for authentication and authenticated encryption. *Cryptogr. Commun.*, 8(4):455–511, 2016.
- [33] Javad Ghareh Chamani, Dimitrios Papadopoulos, Mohammadamin Karbasforushan, and Ioannis Demertzis. Dynamic searchable encryption with optimal search in the presence of deletions. In Kevin R. B. Butler and Kurt Thomas, editors, *31st USENIX Security Symposium, USENIX Security 2022, Boston, MA, USA, August 10-12, 2022*, pages 2425–2442. USENIX Association, 2022.
- [34] Yan-Cheng Chang and Michael Mitzenmacher. Privacy preserving keyword searches on remote encrypted data. In John Ioannidis, Angelos D. Keromytis, and Moti Yung, editors, *Applied Cryptography and Network Security, Third International Conference, ACNS 2005, New York, NY, USA, June 7-10, 2005, Proceedings*, volume 3531 of *Lecture Notes in Computer Science*, pages 442–455, 2005.

- [35] Melissa Chase and Seny Kamara. Structured encryption and controlled disclosure. In *Advances in Cryptology - ASIACRYPT 2010*, pages 577–594, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [36] Melissa Chase and Emily Shen. Substring-searchable symmetric encryption. *Proc. Priv. Enhancing Technol.*, 2015(2):263–281, 2015.
- [37] Sanjit Chatterjee, Manish Kesarwani, Jayam Modi, Sayantan Mukherjee, Shra-  
van Kumar Parshuram Puria, and Akash Shah. Secure and efficient wildcard  
search over encrypted data. *Int. J. Inf. Sec.*, 20(2):199–244, 2021.
- [38] Sanjit Chatterjee, Shra-  
van Kumar Parshuram Puria, and Akash Shah. Efficient  
backward private searchable encryption. *J. Comput. Secur.*, 28(2):229–267,  
2020.
- [39] Valerio Cini, Sebastian Ramacher, Daniel Slamanig, Christoph Striecks, and  
Erkan Tairi. Updatable signatures and message authentication codes. In  
Juan A. Garay, editor, *Public-Key Cryptography - PKC 2021 - 24th IACR In-  
ternational Conference on Practice and Theory of Public Key Cryptography,  
Virtual Event, May 10-13, 2021, Proceedings, Part I*, volume 12710 of *Lecture  
Notes in Computer Science*, pages 691–723. Springer, 2021.
- [40] Dwaine Clarke, Srinivas Devadas, Marten van Dijk, Blaise Gassend, and G. Ed-  
ward Suh. Incremental multiset hash functions and their application to mem-  
ory integrity checking. In Chi-Sung Laih, editor, *Advances in Cryptology -  
ASIACRYPT 2003*, pages 188–207, Berlin, Heidelberg, 2003. Springer Berlin  
Heidelberg.
- [41] Dwaine E. Clarke, Srinivas Devadas, Marten van Dijk, Blaise Gassend, and  
G. Edward Suh. Incremental multiset hash functions and their application to  
memory integrity checking. In Chi-Sung Laih, editor, *Advances in Cryptology  
- ASIACRYPT 2003, 9th International Conference on the Theory and Appli-  
cation of Cryptology and Information Security, Taipei, Taiwan, November 30  
- December 4, 2003, Proceedings*, volume 2894 of *Lecture Notes in Computer  
Science*, pages 188–207. Springer, 2003.
- [42] Reza Curtmola, Juan Garay, Seny Kamara, and Rafail Ostrovsky. Search-  
able symmetric encryption: Improved definitions and efficient constructions. In  
*Proceedings of the 13th ACM Conference on Computer and Communications  
Security, CCS '06*, page 79–88, New York, NY, USA, 2006. Association for  
Computing Machinery.
- [43] Ioannis Demertzis, Javad Ghareh Chamani, Dimitrios Papadopoulos, and Char-  
alampos Papamanthou. Dynamic searchable encryption with small client stor-  
age. In *27th Annual Network and Distributed System Security Symposium,  
NDSS 2020, San Diego, California, USA, February 23-26, 2020*. The Internet  
Society, 2020.

- [44] Siying Dong, Andrew Kryczka, Yanqin Jin, and Michael Stumm. Rocksdb: Evolution of development priorities in a key-value store serving large-scale applications. 2021.
- [45] Enron Corp and Cohen, W. W. Enron email dataset. United States Federal Energy Regulatory Commission, comp, 2015.
- [46] Adam Everspaugh, Kenneth G. Paterson, Thomas Ristenpart, and Samuel Scott. Key rotation for authenticated encryption. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part III*, volume 10403 of *Lecture Notes in Computer Science*, pages 98–129. Springer, 2017.
- [47] Sky Faber, Stanislaw Jarecki, Hugo Krawczyk, Quan Nguyen, Marcel-Catalin Rosu, and Michael Steiner. Rich queries on encrypted data: Beyond exact matches. In *Computer Security - ESORICS 2015 - 20th European Symposium on Research in Computer Security, Vienna, Austria, September 21-25, 2015, Proceedings, Part II*, volume 9327 of *Lecture Notes in Computer Science*, pages 123–145. Springer, 2015.
- [48] Xinrui Ge, Jia Yu, Hanlin Zhang, Chengyu Hu, Zengpeng Li, Zhan Qin, and Rong Hao. Towards achieving keyword search over dynamic encrypted cloud data with symmetric-key based verification. *IEEE Transactions on Dependable and Secure Computing*, 18(1):490–504, 2021.
- [49] Javad Ghareh Chamani, Dimitrios Papadopoulos, Charalampos Papamanthou, and Rasool Jalili. New constructions for forward and backward private symmetric searchable encryption. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, page 1038–1055. Association for Computing Machinery, 2018.
- [50] Eu-Jin Goh. Secure indexes. *IACR Cryptol. ePrint Arch.*, page 216, 2003.
- [51] Cheng Guo, Wenfeng Li, Xinyu Tang, Kim-Kwang Raymond Choo, and Yinling Liu. Forward private verifiable dynamic searchable symmetric encryption with efficient conjunctive query. *IEEE Transactions on Dependable and Secure Computing*, 2023.
- [52] Yu Guo, Chen Zhang, and Xiaohua Jia. Verifiable and forward-secure encrypted search using blockchain techniques. In *ICC 2020 - 2020 IEEE International Conference on Communications (ICC)*, pages 1–7, 2020.
- [53] Florian Hahn, Nicolas Loza, and Florian Kerschbaum. Practical and secure substring search. In Gautam Das, Christopher M. Jermaine, and Philip A. Bernstein, editors, *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, pages 163–176. ACM, 2018.

- [54] Changhui Hu and Lidong Han. Efficient wildcard search over encrypted data. *Int. J. Inf. Sec.*, 15(5):539–547, 2016.
- [55] Mohammad Saiful Islam, Mehmet Kuzu, and Murat Kantarcioglu. Access pattern disclosure on searchable encryption: ramification, attack and mitigation. In *NDSS*, volume 20, page 12. Citeseer, 2012.
- [56] Tetsu Iwata and Kaoru Kurosawa. OMAC: one-key CBC MAC. In Thomas Johansson, editor, *Fast Software Encryption, 10th International Workshop, FSE 2003, Lund, Sweden, February 24-26, 2003, Revised Papers*, volume 2887 of *Lecture Notes in Computer Science*, pages 129–153. Springer, 2003.
- [57] Stanislaw Jarecki, Charanjit Jutla, Hugo Krawczyk, Marcel Rosu, and Michael Steiner. Outsourced symmetric private information retrieval. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security, CCS '13*, page 875–888, New York, NY, USA, 2013. Association for Computing Machinery.
- [58] Seny Kamara and Tarik Moataz. Boolean searchable symmetric encryption with worst-case sub-linear complexity. In *Advances in Cryptology – EUROCRYPT 2017*, pages 94–124. Springer International Publishing, 2017.
- [59] Seny Kamara and Tarik Moataz. Computationally volume-hiding structured encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 183–213. Springer, 2019.
- [60] Seny Kamara, Charalampos Papamanthou, and Tom Roeder. Dynamic searchable symmetric encryption. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security, CCS '12*, page 965–976, New York, NY, USA, 2012. Association for Computing Machinery.
- [61] Els J Kindt. *Privacy and data protection issues of biometric applications*, volume 1. Springer, 2016.
- [62] Kaoru Kurosawa and Tetsu Iwata. TMAC: two-key CBC MAC. In Marc Joye, editor, *Topics in Cryptology - CT-RSA 2003, The Cryptographers' Track at the RSA Conference 2003, San Francisco, CA, USA, April 13-17, 2003, Proceedings*, volume 2612 of *Lecture Notes in Computer Science*, pages 33–49. Springer, 2003.
- [63] Kaoru Kurosawa and Yasuhiro Ohtaki. Uc-secure searchable symmetric encryption. In Angelos D. Keromytis, editor, *Financial Cryptography and Data Security*, pages 285–298, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [64] Kaoru Kurosawa and Yasuhiro Ohtaki. How to update documents verifiably in searchable symmetric encryption. In Michel Abdalla, Cristina Nita-Rotaru, and Ricardo Dahab, editors, *Cryptology and Network Security*, pages 309–328, Cham, 2013. Springer International Publishing.

- [65] Kaoru Kurosawa, Keisuke Sasaki, Kiyohiko Ohta, and Kazuki Yoneyama. Uc-secure dynamic searchable symmetric encryption scheme. In Kazuto Ogawa and Katsunari Yoshioka, editors, *Advances in Information and Computer Security*, pages 73–90, Cham, 2016. Springer International Publishing.
- [66] Shangqi Lai, Sikhar Patranabis, Amin Sakzad, Joseph K. Liu, Debdeep Mukhopadhyay, Ron Steinfeld, Shi-Feng Sun, Dongxi Liu, and Cong Zuo. Result pattern hiding searchable encryption for conjunctive queries. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18, page 745–762. Association for Computing Machinery, 2018.
- [67] Feng Li, Jianfeng Ma, Yinbin Miao, Pengfei Wu, and Xiangfu Song. Beyond volume pattern: Storage-efficient boolean searchable symmetric encryption with suppressed leakage. In Gene Tsudik, Mauro Conti, Kaitai Liang, and Georgios Smaragdakis, editors, *Computer Security - ESORICS 2023 - 28th European Symposium on Research in Computer Security, The Hague, The Netherlands, September 25-29, 2023, Proceedings, Part I*, volume 14344 of *Lecture Notes in Computer Science*, pages 126–146. Springer, 2023.
- [68] Chang Liu, Liehuang Zhu, Mingzhong Wang, and Yu an Tan. Search pattern leakage in searchable encryption: Attacks and new construction. *Information Sciences*, 265:176–188, 2014.
- [69] Jinlu Liu, Bo Zhao, Jing Qin, Xi Zhang, and Jixin Ma. Multi-keyword ranked searchable encryption with the wildcard keyword for data sharing in cloud computing. *Comput. J.*, 66(1):184–196, 2023.
- [70] Judith A Markowitz. Voice biometrics. *Communications of the ACM*, 43(9):66–73, 2000.
- [71] Meixia Miao, Yunling Wang, Jianfeng Wang, and Xinyi Huang. Verifiable database supporting keyword searches with forward security. volume 77, page 103491, 2021.
- [72] Tarik Moataz, Indrajit Ray, Indrakshi Ray, Abdullatif Shikfa, Frédéric Cuppens, and Nora Cuppens. Substring search over encrypted data. *J. Comput. Secur.*, 26(1):1–30, 2018.
- [73] Tarik Moataz and Abdullatif Shikfa. Boolean symmetric searchable encryption. In Kefei Chen, Qi Xie, Weidong Qiu, Ninghui Li, and Wen-Guey Tzeng, editors, *8th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '13, Hangzhou, China - May 08 - 10, 2013*, pages 265–276. ACM, 2013.
- [74] Mridul Nandi. Fast and secure CBC-Type MAC algorithms. In Orr Dunkelman, editor, *Fast Software Encryption, 16th International Workshop, FSE 2009, Leuven, Belgium, February 22-25, 2009, Revised Selected Papers*, volume 5665 of *Lecture Notes in Computer Science*, pages 375–393. Springer, 2009.

- [75] Dalit Naor, Moni Naor, and Jeff Lotspiech. Revocation and tracing schemes for stateless receivers. In *Advances in Cryptology — CRYPTO 2001*, pages 41–62, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- [76] Alexandros Ntoulas and Junghoo Cho. Pruning policies for two-tiered inverted index with correctness guarantee. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '07, page 191–198, New York, NY, USA, 2007. Association for Computing Machinery.
- [77] Our Code. <https://github.com/anonCod/CD-SSE>.
- [78] Sarvar Patel, Giuseppe Persiano, Joon Young Seo, and Kevin Yeo. Efficient boolean search over encrypted data with reduced leakage. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2021 - 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6-10, 2021, Proceedings, Part III*, volume 13092 of *Lecture Notes in Computer Science*, pages 577–607. Springer, 2021.
- [79] Sarvar Patel, Giuseppe Persiano, Kevin Yeo, and Moti Yung. Mitigating leakage in secure cloud-hosted data structures: Volume-hiding for multi-maps via hashing. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 79–93, 2019.
- [80] Sikhar Patranabis and Debdeep Mukhopadhyay. Forward and backward private conjunctive searchable symmetric encryption. In *28th Annual Network and Distributed System Security Symposium, NDSS 2021, virtually, February 21-25, 2021*. The Internet Society, 2021.
- [81] Phillip Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In *Advances in Cryptology - ASIACRYPT 2004, 10th International Conference on the Theory and Application of Cryptology and Information Security, Jeju Island, Korea, December 5-9, 2004, Proceedings*, pages 16–31, 2004.
- [82] Moe Sabry and Reza Samavi. ArchiveSafe LT: Secure long-term archiving system. In *Proceedings of the 38th Annual Computer Security Applications Conference*, ACSAC '22, page 936–948, New York, NY, USA, 2022. Association for Computing Machinery.
- [83] Moe Sabry, Reza Samavi, and Douglas Stebila. Archivesafe: Mass-leakage-resistant storage from proof-of-work. In Joaquin Garcia-Alfaro, Guillermo Navarro-Arribas, and Jordi Herrera-Joancomarti, editors, *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, pages 89–107, Cham, 2020. Springer International Publishing.

- [84] James Scheibner, Jean Louis Raisaro, Juan Ramón Troncoso-Pastoriza, Marcello Ienca, Jacques Fellay, Effy Vayena, and Jean-Pierre Hubaux. Revolutionizing medical data sharing using advanced privacy-enhancing technologies: Technical, legal, and ethical synthesis. *J Med Internet Res*, 23(2):e25120, Feb 2021.
- [85] Saeed Sedghi, Peter van Liesdonk, Svetla Nikova, Pieter H. Hartel, and Willem Jonker. Searching keywords with wildcards on encrypted data. In Juan A. Garay and Roberto De Prisco, editors, *Security and Cryptography for Networks, 7th International Conference, SCN 2010, Amalfi, Italy, September 13-15, 2010. Proceedings*, volume 6280 of *Lecture Notes in Computer Science*, pages 138–153. Springer, 2010.
- [86] Dawn Xiaoding Song, D. Wagner, and A. Perrig. Practical techniques for searches on encrypted data. In *Proceeding 2000 IEEE Symposium on Security and Privacy. S&P 2000*, pages 44–55, 2000.
- [87] Xiangfu Song, Changyu Dong, Dandan Yuan, Qiuliang Xu, and Minghao Zhao. Forward private searchable symmetric encryption with optimized i/o efficiency. *IEEE Transactions on Dependable and Secure Computing*, 17(5):912–927, 2020.
- [88] Emil Stefanov, Charalampos Papamanthou, and Elaine Shi. Practical dynamic searchable encryption with small leakage. In *21st Annual Network and Distributed System Security Symposium, NDSS 2014, San Diego, California, USA, February 23-26, 2014*. Internet Society, 2014.
- [89] Shi-Feng Sun, Xingliang Yuan, Joseph K. Liu, Ron Steinfeld, Amin Sakzad, Viet Vo, and Surya Nepal. Practical backward-secure searchable encryption from symmetric puncturable encryption. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, page 763–780. Association for Computing Machinery, 2018.
- [90] Shifeng Sun, Ron Steinfeld, Shangqi Lai, Xingliang Yuan, Amin Sakzad, Joseph K. Liu, Surya Nepal, and Dawu Gu. Practical non-interactive searchable encryption with forward and backward privacy. In *28th Annual Network and Distributed System Security Symposium, NDSS 2021, virtually, February 21-25, 2021*. The Internet Society, 2021.
- [91] Na Wang, Junsong Fu, Bharat K. Bhargava, and Jiwen Zeng. Efficient retrieval over documents encrypted by attributes in cloud computing. *IEEE Transactions on Information Forensics and Security*, 13(10):2653–2667, 2018.
- [92] Mark N Wegman and J Lawrence Carter. New classes and applications of hash functions. In *20th Annual Symposium on Foundations of Computer Science (SFCS 1979)*, pages 175–182. IEEE, 1979.
- [93] Zhiqiang Wu and Kenli Li. Vbtree: forward secure conjunctive queries over encrypted data for cloud computing. *VLDB J.*, 28(1):25–46, 2019.



- [94] Dandan Yuan, Shujie Cui, and Giovanni Russello. We can make mistakes: Fault-tolerant forward private verifiable dynamic searchable symmetric encryption. In *7th IEEE European Symposium on Security and Privacy, EuroS&P 2022, Genoa, Italy, June 6-10, 2022*, pages 587–605. IEEE, 2022.
- [95] Dandan Yuan, Cong Zuo, Shujie Cui, and Giovanni Russello. Result-pattern-hiding conjunctive searchable symmetric encryption with forward and backward privacy. *Proc. Priv. Enhancing Technol.*, 2023(2):40–58, 2023.
- [96] Xi Zhang, Ye Su, Zhongkai Wei, Wenting Shen, and Jing Qin. Efficient wildcard searchable symmetric encryption with forward and backward security. In Jaideep Vaidya, Moncef Gabbouj, and Jin Li, editors, *Artificial Intelligence Security and Privacy - First International Conference on Artificial Intelligence Security and Privacy, AIS&P 2023, Guangzhou, China, December 3-5, 2023, Proceedings, Part I*, volume 14509 of *Lecture Notes in Computer Science*, pages 342–357. Springer, 2023.
- [97] Yupeng Zhang, Jonathan Katz, and Charalampos Papamanthou. All your queries are belong to us: The power of file-injection attacks on searchable encryption. In *Proceedings of the 25th USENIX Conference on Security Symposium, SEC’16*, page 707–720. USENIX Association, 2016.
- [98] Zhongjun Zhang, Jianfeng Wang, Yunling Wang, Yaping Su, and Xiaofeng Chen. Towards efficient verifiable forward secure searchable symmetric encryption. In Kazue Sako, Steve Schneider, and Peter Y. A. Ryan, editors, *Computer Security – ESORICS 2019*, pages 304–321, Cham, 2019. Springer International Publishing.
- [99] Jie Zhu, Qi Li, Cong Wang, Xingliang Yuan, Qian Wang, and Kui Ren. Enabling generic, verifiable, and secure data search in cloud services. *IEEE Transactions on Parallel and Distributed Systems*, 29(8):1721–1735, 2018.
- [100] Cong Zuo, Shangqi Lai, Xingliang Yuan, Joseph K. Liu, Jun Shao, and Huaxiong Wang. Searchable encryption for conjunctive queries with extended forward and backward privacy. *IACR Cryptol. ePrint Arch.*, page 1585, 2021.
- [101] Cong Zuo, Shi-Feng Sun, Joseph K. Liu, Jun Shao, and Josef Pieprzyk. Dynamic searchable symmetric encryption with forward and stronger backward privacy. In *Computer Security – ESORICS 2019*, pages 283–303. Springer International Publishing, 2019.
- [102] Cong Zuo, Shifeng Sun, Joseph K. Liu, Jun Shao, Josef Pieprzyk, and Guiyi Wei. Forward and backward private dynamic searchable symmetric encryption for conjunctive queries. *IACR Cryptol. ePrint Arch.*, page 1357, 2020.