

# BISEN: Efficient Boolean Searchable Symmetric Encryption with Verifiability and Minimal Leakage

Bernardo Ferreira<sup>1</sup>, Bernardo Portela<sup>2</sup>, Tiago Oliveira<sup>2</sup>, Guilherme Borges<sup>1</sup>,  
Henrique Domingos<sup>1</sup> and João Leitão<sup>1</sup>

<sup>1</sup>NOVA LINES & DI-FCT-UNL    <sup>2</sup>HASLab INESC TEC & DCC-FC-UP

**Abstract.** The prevalence and availability of cloud infrastructures has made them the *de facto* solution for storing and archiving data, both for organizations and individual users. Nonetheless, the cloud’s wide spread adoption is still hindered by dependability and security concerns, particularly in applications with large data collections where efficient search and retrieval services are also major requirements. This leads to an increased tension between security, efficiency, and search expressiveness, which current state of the art solutions try to balance through complex cryptographic protocols that tradeoff efficiency and expressiveness for near optimal security.

In this paper we tackle this tension by proposing BISEN, a new provably-secure boolean searchable symmetric encryption scheme that improves these three complementary dimensions by exploring the design space of isolation guarantees offered by novel commodity hardware such as Intel SGX, abstracted as Isolated Execution Environments (IEEs). BISEN is the first scheme to enable highly expressive and arbitrarily complex boolean queries, with minimal information leakage regarding performed queries and accessed data, and verifiability regarding fully malicious adversaries. Furthermore, by exploiting trusted hardware and the IEE abstraction, BISEN reduces communication costs between the client and the cloud, boosting query execution performance. Experimental validation and comparison with the state of art shows that BISEN provides better performance with enriched search semantics and security properties.

## 1 Introduction

Cloud computing has had a profound impact on the way that we design and operate systems and applications. In particular, data storage and archiving is now commonly delegated to cloud infrastructures, both by companies and individual users. Companies typically want to archive large volumes of data, such as e-mails or historical documents, overcoming limitations or lowering costs of their on-premise infrastructures [2], while individual users aim at making their documents easily accessible from multiple devices, or simply avoid consuming storage capacity of their mobile devices [15].

However, data being outsourced to the cloud is often sensitive and should be protected regarding all aspects of dependability. Private information incidents are constant reminders of the growing importance of these issues: governmental agencies impose increasing pressure on cloud companies to disclose users’ data and deploy backdoors [23]; cloud providers are responsible, maliciously or accidentally, for critical data disclosures [20]; and even external hackers have gained remote access to users data for limited time windows [28]. Cloud outsourcing services are thus highly incentivized to address these dependability and security requirements. In particular, when storing and updating large volumes of data in the cloud, it is essential to offer efficient, secure, and precise mechanisms to search and retrieve relevant data objects from the archive. This highlights the need for cloud-based systems to balance security, efficiency, and query expressiveness.

To address this tension, Searchable Symmetric Encryption (SSE) [7] has emerged as an important research topic in recent years, allowing one to efficiently search and update an encrypted database within an untrusted cloud server with security guarantees. Efficiency in SSE is achieved by building an encrypted index of the database and storing it in the cloud [17]. At search time, a cryptographic token specific to the query is used to access the index, and retrieved index entries are decrypted and processed. To minimize communication overhead, most SSE schemes delegate the execution of cryptographic computations to the cloud, as multiple index entries would otherwise have to be requested and downloaded to the client. However, performing sensitive operations in the cloud also leads to significant information leakage, including the leakage of document identifiers matching a query, the repetition of queries, and the compromise of forward and backward privacy [34] (respectively, if new update operations match contents with previously issued queries, and if queries return previously deleted documents). These are common, yet severe, flavors of information leakage that pave the way for strong attacks on SSE, including devastating file-injection attacks [38]. Another relevant

limitation of SSE schemes is query expressiveness, as most solutions only support single keyword match [13] or limited boolean queries (e.g., forcing queries to be in Conjunctive Normal Form and not supporting negations) [24]. This hinders system usability and may force users to perform multiple queries in order to retrieve relevant results, leading to extra communication steps and additional information leakage.

In this paper we address these limitations by presenting BISEN (Boolean Isolated Searchable symmetric EN-cryption), a new provably-secure boolean SSE scheme that improves query expressiveness by supporting arbitrarily complex boolean queries with combinations of conjunctions, disjunctions, and negations (e.g., “(cancer  $\vee$  terminal)  $\wedge$   $\neg$  cirrosis”). This is a significant improvement over the current state of the art, since supporting boolean queries is fundamentally more challenging than single-keyword queries and addressing negations is a non trivial task. Furthermore, BISEN also boosts performance by minimizing the number of communication steps and amount of data transferred between clients and cloud servers. A central insight in the design of BISEN is the fact that we can securely delegate critical computations to the cloud by leveraging on a hybrid solution that combines standard symmetric-key cryptographic primitives (e.g., Pseudo-Random Functions and Block-Ciphers [26]) with remote attestation capabilities offered by modern trusted hardware, formally captured by an abstraction called Isolated Execution Environments (IEEs) [4].

An IEE is an environment that allows applications to execute in isolation from all external interference (including co-located software and even a potentially malicious Hypervisor/OS) and that provides a mechanism for the remote attestation of computed outputs. Until recently, such an abstraction could only be built through hardware that was unfeasible to deploy in cloud infrastructures, however recent advances in trusted computing have made IEEs available in commodity hardware. Prominent examples include Intel SGX [16] and ARM TrustZone [1], which are being deployed in current desktop and mobile processors and will soon become available as part of many cloud infrastructures [32].

A main advantage of designing our system to leverage the IEE abstraction lies in its portability, as our solution can be easily instantiated using different existing (or future) IEE-enabling technologies as they become available in cloud platforms, while preserving security guarantees. This is also relevant considering recent attacks on trusted hardware [35, 36]. To further increase this portability, we extend the IEE formalization to support very lightweight hardware technologies (such as Intel SGX, with its limited Enclave Page Cache size of 128MB), complemented with SSE techniques and cryptographically protected accesses to more abundant untrusted resources in the machine hosting the IEE or in other external cloud storage services. This approach has mutual benefits: SSE techniques allow extending IEE trusted resources beyond hardware limitations in a secure way, minimizing assumptions regarding the underlying technology employed; and IEEs allow increasing the performance, scalability, and security of SSE schemes.

In summary, in this paper we provide the following main contributions:

- We propose and formalize an approach for extending trusted hardware resources, integrating it in the IEE abstraction of Barbosa et al. [4]. This approach allows IEEs to support and operate on very large databases, and may be of particular interest for other applications;
- We design BISEN, a Boolean SSE scheme based on the previous approach and that can support all Boolean operands and formulas. By leveraging IEEs as remote trust anchors, BISEN is able to move most client-side computations to the server, providing overall reduced computation, storage, and communication overheads. BISEN provides verifiability against fully malicious adversaries, supports dynamic updates with forward and backward privacy, and queries only reveal which encrypted index entries are accessed;
- We implement a prototype of BISEN based on Intel SGX, which we run on real world datasets to experimentally validate its efficiency properties. Our prototype is open-source and available at <https://github.com/sgtpepperpt/BISEN>.

## 2 Background and Related Work

**Isolated Execution Environments (IEEs)** From a high level, and as defined by Barbosa et al. [4], an IEE is an idealized random access machine, running a fixed program, and whose behaviour can only be influenced by a well-specified interface that allows input/output interactions with the program. Isolation guarantees in IEEs follow from the requirements that: the I/O behaviour of programs running within them can only depend on themselves, on the semantics of their language, and on inputs received; and that the only information revealed about these programs must be contained in their I/O behaviour. This abstraction allows for the formal treatment of remote attestation mechanisms offered by technologies such as SGX and TrustZone, which were shown in [4] to be sufficient for the deployment of an Outsourced Computation protocol.

Building on these definitions, Bahmani et al. [3] demonstrated how to refine the IEE attestation mechanism to enable for the deployment of *general multiparty computation*. Their design follows two main stages. First, clients leverage remote attestation mechanisms to perform a key exchange agreement with the IEE and establish a secure communication channel. Afterwards, clients use these channels to interact with a reactive functionality on the IEE, exchanging encrypted inputs and outputs with confidentiality and integrity guarantees. The usage of sequence numbers in communications made through these channels also prevents a malicious server from repeating requests. In this work we will leverage on the IEE abstraction and this protocol, further extending it by allowing the IEE to interact with untrusted storage resources with privacy, integrity, and verifiability guarantees.

**Searchable Symmetric Encryption (SSE)** SSE deals with the problem of how to efficiently search and update an encrypted database [17]. To achieve this goal, SSE schemes usually build an encrypted index of the database (e.g., an inverted list index [31]), hence reducing the previous problem to the easier one of searching and updating an encrypted index. This approach has additional advantages, as the index allows search performance to be sub-linear on the database size, and the data itself can eventually be stored on a second storage system with different security guarantees. In its most simple version, keys of this encrypted index are message authentication codes (MACs) of keywords, and values are symmetrically encrypted versions of document identifiers.

To search the encrypted index, the client transforms his query into a cryptographic token (usually composed of a pair of cryptographic keys), which is used by the server to find and decrypt relevant index entries. This approach allows the encrypted index to provide zero-leakage when it is on the rest, however some information patterns must be leaked when it is updated or queried, as a necessary tradeoff for achieving practical performance (also providing zero-leakage during computation would require expensive techniques such as Oblivious-RAM [22]). Information leaked by SSE schemes includes search patterns (if a query has occurred in the past and when) and access patterns (which index entries are accessed by the query). Query-recovery attacks have been demonstrated based on both patterns [12, 30], although requiring large a-priori database knowledge (around 90%) or the adversarial ability to inject files [38].

Forward and backward privacy are also important security definitions in SSE [34]. Forward privacy enforces that update operations should not reveal anything regarding updated keywords, even if combined with previously issued query tokens [8], and helps partially mitigating file-injection attacks. Backward privacy enforces that search operations should only reflect the current state of the database and should reveal nothing regarding deleted keywords [9].

SSE schemes usually only support single keyword queries, as supporting boolean multi-keyword queries with similar security guarantees and performance is a fundamentally more difficult problem [14, 24]. Even the most recent Boolean SSE scheme to date [24] still provides: limited usability, as it does not support negations and queries must be in Conjunctive Normal Form (CNF), possibly forcing users to rewrite their queries; limited performance, requiring quadratic server storage in the number of unique keywords in the database and exhibiting quadratic search performance in the query size; and limited security, as queries leak the search and access patterns of some of their individual keywords and of the resulting conjunctions/disjunctions. This is due to the difficulty of managing complex multi-map data-structures required by the authors for supporting boolean queries, which we show in this work to be avoidable by leveraging remote trust anchors expanded with cryptographically secured accesses to large untrusted storage.

**Cryptographic Protocols based on Trusted Hardware** Recent works have demonstrated the benefits that trusted hardware like Intel SGX can bring to the design of cryptographic protocols. Iron [19] used Intel SGX to develop a practical Functional Encryption (FE) scheme [19]. SSE, as most schemes for privacy-preserving computations, can be seen as specialization of FE, meaning that the approach proposed by the authors could also be employed to solve the problems we address in this work. However, our approach is specifically tailored for solving the challenges posed by searching encrypted data, optimizing performance and efficiency as no general purpose approach traditionally can.

ZeroTrace [33] provided a more efficient protocol for Oblivious-RAM based on SGX. Their techniques can be used to complement our approach, as a way to further reduce information leakage and provide further resilience to side-channel attacks [36]. HardIDX [21] used SGX for efficiently supporting range queries in SSE. Their approach has similarities with ours, but its focus is on a fundamentally different problem (range queries). Additionally, it only supported static databases and required the client to build the encrypted index. In contrast, our work supports dynamic updates with minimal leakage and moves most computations to the cloud in a secure way, both considered essential for practical applications.




Fig. 1: Overview of the proposed approach.

### 3 Technical Overview

The main idea of BISEN is for clients to leverage IEEs as remote trust-anchors within the cloud, supporting efficient update and search operations on their cloud-stored encrypted databases. However, it would be unfeasible to maintain a whole database index within a resource-restricted IEE. As such, our proposal is to leverage a highly efficient environment for computations (the IEE) and a virtually infinite source for external storage (the cloud). BISEN combines these tools in its IEE-sided code, processing queries within its isolated memory, and relying on an cloud service for storing encrypted data.

Figure 1 provides an overview of BISEN’s architecture and its components. BISEN starts with a bootstrapping phase, where a client contacts a cloud server to initiate the IEE with BISEN’s code. We call this server the Proxy Server, as it operates the IEE, manages all of its communications, and orders concurrent accesses. When started, the IEE initiates its state and asks a Cloud Storage Service to create BISEN’s (initially empty) encrypted index. This storage service will basically be responsible for large-scale storage, and can even be instantiated through pure storage solutions (e.g. AWS S3), as it only needs to support put/get operations. Hardcoded in BISEN’s code is the public key of a trusted Certificate Authority, meaning that the IEE will only accept communications from clients that present a valid certificate signed by it.

After bootstrap, clients can contact the proxy server to remotely attest it created an IEE with BISEN’s code and to establish secure communication channels with it. Secure channels are established through a key-exchange algorithm based on remote attestation and the clients’ public keys, as in [3]. Through these channels, clients can issue update and search operations to the IEE, which it processes by contacting the storage service and accessing BISEN’s index.

Updates allow both adding or removing keywords to/from documents. In either case, a new encrypted entry is added to BISEN’s index, where its key is composed of a deterministic cryptographic token uniquely combining the keyword and document, and its value is an encrypted message that includes the document id and a flag indicating if the operation is an addition or removal. This approach guarantees that both operations are indistinguishable, a necessary condition for preserving forward and backward privacy.

Search operations take a boolean query as input. The IEE processes this query, retrieves and decrypts relevant index entries from the storage service, calculates the resulting set of document ids, and finally returns this set to the client.

**Adversary Model.** The clients, the IEE, and the Certificate Authority are the only trusted participants in BISEN. The Proxy Server and Storage Service are considered fully malicious, i.e., they may attempt to break data privacy, integrity, or computation correctness. Networking channels are also considered untrusted. Denial of service is considered out of scope for this work, as the cloud controls the whole infrastructure, but may be addressed in future works through cross-cloud replication.

**Application Scenario.** An interesting application scenario for BISEN is that of encrypted archival of email in the cloud. In such a scenario, users would be able to securely outsource the storage and management of their emails to a third-party cloud provider, while still being able to have rich search features that are commonly found in today's unsecured email cloud archival services. As studied by Zheng et al. [38], cloud email is an example scenario that can be easily targeted by file-injection attacks, hence this application enforces the need to improve the security of SSE schemes to withstand fully malicious adversaries. Furthermore, preserving forward privacy is known to help mitigate such attacks [38], and backward privacy may have important implications in future attacks as well [9]. Overall, minimizing information leakage should be a top priority when deploying SSE schemes in practical scenarios.

## 4 BISEN

In this section we present BISEN's full details. We start with some required notations and definitions (§ 4.1), then we present BISEN's protocols (§ 4.2), and finally we analyse its security (§ 4.3).

### 4.1 Notations and Definitions

**General Notations.** In this paper we denote by  $\lambda$  the security parameter and  $\mu(\lambda)$  a negligible function in it. We will use the standard security notions of variable-input-length Pseudo-Random Functions (PRF, instantiated as an HMAC in our implementation) [5] and authenticated encryption schemes ensuring *indistinguishability under chosen-ciphertext attacks* (IND-CCA) [26]. We assume the keys of these primitives to be uniformly sampled from  $\{0, 1\}^\lambda$  by the key generation algorithm. We consider adversaries to be probabilistic algorithms, running in time polynomial on security parameter  $\lambda$ .

**IEE Notations.** IEE behavior and interactions with clients can be abstracted as  $\text{IEE} = (\text{Setup}, \text{Send}, \text{Receive})$ , as follows:

- $\text{IEE.Setup}(1^\lambda, \text{pk}_c)$  corresponds to IEE bootstrapping (if it hasn't been initialized yet) and secure channel establishment. Setup takes security parameter  $1^\lambda$  and client public key  $\text{pk}_c$  as input, and produces state  $\text{st}_{\text{IEE}}$  with the exchanged key, if  $\text{pk}_c$  could be verified.
- $\text{IEE.Send}(\text{st}_{\text{IEE}}, m)$  can be used by the client or IEE, and uses the secure channel established by Setup to encrypt  $m$  with the key in  $\text{st}_{\text{IEE}}$ . This outputs  $c$  and updates state  $\text{st}_{\text{IEE}}$ .
- $\text{IEE.Receive}(\text{st}_{\text{IEE}}, c)$  uses the channel to retrieve encrypted message  $c$  using the key in  $\text{st}_{\text{IEE}}$ . This outputs  $m$  and updates state  $\text{st}_{\text{IEE}}$ .

Additionally, and differently from the IEEs original specification [3, 4], we consider IEEs to rely not only on *trusted* state, which is assumed to be incorruptible by the underlying system, but also on *untrusted* state (represented in BISEN through the Storage Service), which has to be explicitly protected through cryptographic algorithms. To establish interactions with this *untrusted* state, and following a dictionary-like notation, we define three new calls in the IEE abstraction:

- $\text{IEE.ulnit}()$  initializes an empty data-structure  $D$  in untrusted storage outside the IEE. It outputs  $D$ , making it available for future  $\text{uPut}$  and  $\text{uGet}$  operations;
- $\text{IEE.uPut}(D, \{l_i, v_i\}_{i=0}^*)$  accesses untrusted storage and stores a group of entries  $\{l_i, v_i\}_{i=0}^*$  in data-structure  $D$ . It outputs updated structure  $D$ .
- $\text{IEE.uGet}(D, \{l_i\}_{i=0}^*)$  accesses untrusted storage and outputs a group of values  $\{v_i\}_{i=0}^*$ , stored in positions  $\{l_i\}_{i=0}^*$  of data-structure  $D$ .

Formally, we consider  $\text{ulnit}$  and  $\text{uPut}$  to additionally produce an execution trace, containing the operation, its input, and the output. In the security experiment this trace is given directly to the adversary, capturing the notion that all data stored through this mechanism is considered leakage. Since we are considering a fully malicious adversary, all values returned by  $\text{uGet}$  can be set by the adversary.

**SSE Notations.** In SSE, a database  $\text{DB}$  is composed by a collection of  $d$  documents, each with a unique identifier  $\text{id}$  and containing a set of keywords  $W$ . For a keyword  $w$ ,  $\text{DB}(w)$  is the set of documents where it occurs. The total number of document/keyword pairs is denoted by  $n$  and is stored in an encrypted index  $I$ , which is a dictionary structure mapping each unique keyword  $w$  to a list of matching documents  $(\text{id}_0, \dots, \text{id}_{|\text{DB}(w)|-1})$  and allowing queries to be performed in

<p><u>Setup</u>(<math>1^\lambda, pk_c</math>)</p> <p><u>Client:</u></p> <ol style="list-style-type: none"> <li>1: <math>st_{IEE} \leftarrow \\$ IEE.Setup(1^\lambda, pk_c)</math></li> </ol> <p><u>Server:</u></p> <ol style="list-style-type: none"> <li>2: <math>IEE.Setup(1^\lambda, pk_c)</math></li> </ol> <p><u>IEE:</u></p> <ol style="list-style-type: none"> <li>3: <math>st_{IEE} \leftarrow \\$ IEE.Setup(1^\lambda, pk_c)</math></li> <li>4: <math>l \leftarrow IEE.ulnit()</math></li> <li>5: <math>k_F \leftarrow \\$ F.Gen(1^\lambda); k_E \leftarrow \\$ \Theta.Gen(1^\lambda)</math></li> <li>6: <math>W \leftarrow Init(); nDocs \leftarrow 0</math></li> </ol> <hr/> <p><u>Update</u>(<math>op, w, id</math>)</p> <p><u>Client:</u></p> <ol style="list-style-type: none"> <li>1: <math>h_w \leftarrow H(w)</math></li> <li>2: <math>m^* \leftarrow \\$ IEE.Send(st_{IEE}, \{op, id, h_w\})</math></li> <li>3: Send <math>m^*</math> to Server.</li> </ol> <p><u>Server:</u></p> <ol style="list-style-type: none"> <li>4: Send <math>m^*</math> to IEE.</li> </ol> <p><u>IEE:</u></p> <ol style="list-style-type: none"> <li>5: <math>\{op, id, h_w\} \leftarrow IEE.Receive(st_{IEE}, m^*)</math></li> <li>6: <math>c \leftarrow Get(W, h_w)</math></li> <li>7: <b>if</b> <math>c = \perp</math> <b>then</b></li> <li>8:     <math>c \leftarrow 0</math></li> <li>9: <b>else</b></li> <li>10:    <math>c \leftarrow c + 1</math></li> <li>11: <math>W \leftarrow Put(W, h_w, c)</math></li> <li>12: <math>k_w \leftarrow F.Run(k_F, h_w); l \leftarrow F.Run(k_w, c);</math></li> <li>13: <math>id^* \leftarrow \\$ \Theta.Enc(k_E, (l, op, id))</math></li> <li>14: <math>l \leftarrow IEE.uPut(l, id^*)</math></li> <li>15: <b>if</b> <math>id &gt; nDocs</math> <b>then</b></li> <li>16:     <math>nDocs++</math></li> </ol>	<p><u>Search</u>(<math>q</math>)</p> <p><u>Client:</u></p> <ol style="list-style-type: none"> <li>1: <math>\{\bar{w}, \phi\} \leftarrow ProcessBooleanQuery(q); C \leftarrow []</math></li> <li>2: <b>for all</b> <math>w \in \bar{w}</math> <b>do</b></li> <li>3:     <math>h_w \leftarrow H(w); C \leftarrow h_w : C</math></li> <li>4: <math>m^* \leftarrow \\$ IEE.Send(st_{IEE}, \{C, \phi\})</math></li> <li>5: Send <math>m^*</math> to Server.</li> </ol> <p><u>Server:</u></p> <ol style="list-style-type: none"> <li>6: Send <math>m^*</math> to IEE.</li> </ol> <p><u>IEE:</u></p> <ol style="list-style-type: none"> <li>7: <math>\{C, \phi\} \leftarrow IEE.Receive(st_{IEE}, m^*); Q \leftarrow Init()</math></li> <li>8: <b>for all</b> <math>H_w \in C</math> <b>do</b></li> <li>9:     <math>k_w \leftarrow F.Run(k_F, h_w)</math></li> <li>10:    <math>c \leftarrow Get(W, h_w); L \leftarrow []</math></li> <li>11:    <b>for all</b> <math>c_i \leftarrow 0 \dots c</math> <b>do</b></li> <li>12:     <math>l \leftarrow F.Run(k_w, c_i); L \leftarrow l : L</math></li> <li>13:    <math>Q \leftarrow Put(Q, k_w, L)</math></li> <li>14: <math>L' \leftarrow Flatten(Q)</math></li> <li>15: <math>\Pi \leftarrow \\$ RandomPermutation(1^\lambda); L' \leftarrow \Pi(L')</math></li> <li>16: <math>D' \leftarrow IEE.uGet(l, L'); D \leftarrow []</math></li> <li>17: <b>for all</b> <math>id^* \in D'; l' \in L'</math> <b>do</b></li> <li>18:     <math>(l, op, id) \leftarrow \Theta.Dec(k_E, id^*); Verify(l, l')</math></li> <li>19:     <math>D \leftarrow \{op, id\} : D</math></li> <li>20: <math>D \leftarrow \Pi^{-1}(D); Q' \leftarrow Join(Q, D)</math></li> <li>21: <math>R \leftarrow Resolve(\phi, Q', nDocs); r^* \leftarrow \\$ IEE.Send(st_{IEE}, R)</math></li> <li>22: Send <math>r^*</math> to Server.</li> </ol> <p><u>Server:</u></p> <ol style="list-style-type: none"> <li>23: Send <math>r^*</math> to Client.</li> </ol> <p><u>Client:</u></p> <ol style="list-style-type: none"> <li>24: <math>R \leftarrow IEE.Receive(st_{IEE}, r^*)</math></li> </ol>
--	--

Fig. 2: Our BISEN scheme based on IEE = (Setup, Send, Receive, ulnit, uPut, uGet), PRF F = (Gen, Run), authenticated encryption scheme  $\Theta$  = (Gen, Enc, Dec), and hash function H.

time sub-linear in  $n$ .  $\phi(\bar{w})$  is a boolean query composed of a set of keywords  $\bar{w}$  and satisfying a boolean formula  $\phi$ .  $DB(\phi(\bar{w}))$  represents the set of documents satisfying  $\phi(\bar{w})$ .

A *multi-client dynamic boolean searchable symmetric encryption scheme*  $\Pi = (\text{Setup}, \text{Search}, \text{Update})$  consists of three protocols between a client and a server (in our case, the IEE):

- $\Pi.Setup(1^\lambda, pk_c)$  starts the scheme, with inputs security parameter  $1^\lambda$  and client public key  $pk_c$ . At the end of the protocol the client has secret parameter  $K$  and, if the scheme hadn't been initialized yet by another client, the server has the (initially empty) encrypted index  $l$ .
- $\Pi.Update(K, op, w, id)$  updates the database with inputs secret parameter  $K$ , operation  $op = \{\text{add}, \text{del}\}$  (i.e., an addition or deletion of a keyword), keyword  $w$ , and document identifier  $id$ . In the end, the server outputs updated  $l$ .
- $\Pi.Search(K, \phi(\bar{w}))$  queries the database with inputs  $K$  and boolean query  $\phi(\bar{w})$ . In the end, the client outputs a set of document identifiers. If for any possible inputs this output is  $DB(\phi(\bar{w}))$ , we say that  $\Pi$  is *correct*.

## 4.2 The Scheme

Figure 2 details BISEN's protocols. In the Setup algorithm, the client starts by contacting the proxy server and invoking the IEE.Setup protocol. If the IEE hasn't been initiated yet, the proxy server starts it, and then redirects the client's setup message to it. The IEE starts by performing a key-exchange algorithm with the client to create the secure channel between the two. From this algorithm results state  $st_{IEE}$ , which contains the corresponding communication key. Then, if this is its first execution, the IEE initiates its state: a key  $k_F$ , to be used with PRF F; key  $k_E$ , to be used with encryption scheme  $\Theta$ ; dictionary of counters  $W$ , mapping each unique keyword  $w$  in the database to an integer counter  $c$  initialized at zero (each increment in  $c$  represents a new document containing  $w$ ); and counter  $nDocs$ , which counts the number of unique database documents and will be used in the Search protocol to help resolve Boolean queries with negations. Finally, the IEE also asks the storage service to initiate index  $l$ , through the IEE.ulnit call.

The Update protocol can be used both for adding or deleting keywords to/from documents, depending only on the value of input  $op$ . Moreover, the protocol follows the same specification for both, and for adverseries, they are indistinguishable. When performing an update, the client starts by sending  $(op, id, h_w)$  to the IEE through their secure

channel (IEE. {Send, Receive} calls), where  $h_w$  is a hash of keyword  $w$ , performed to normalize keyword lengths and make updates for different keywords indistinguishable. Upon receiving this message, the IEE builds  $l$ , the label (or index key) for this update, by first applying  $F$  on  $h_w$  and  $k_F$  to produce  $k_w$ , and then on  $k_w$  and  $c$  (the keyword’s counter, retrieved from  $W$ ). This double application of  $F$  allows merging both  $h_w$  and  $c$  with a secret parameter ( $k_F$ ), producing a secure label that is unique for this update.  $l$  determines the position in index  $I$  where the update will be stored. As index value, the IEE encrypts  $(l, op, id)$  with  $\Theta$ , an authenticated encryption scheme.  $\Theta$  ensures the preservation of both privacy and integrity of encrypted index values. Furthermore, by including  $l$  in the encrypted index value, the IEE can validate during Search operations that the server is returning correct responses when it requests index values from untrusted storage. Finally, the IEE sends this new index entry to the server for storage, through the IEE.uPut call, and increments  $nDocs$  if this is a new document.

When searching with a boolean query, the client also sends  $h_w$  to the IEE, for each keyword  $w$  in the query. Additionally the client sends  $\phi$ , the boolean formula of the query that the IEE needs for computing search results. Given this message, the IEE recalculates all labels for the inputted keywords (as in Update), requesting the respective index positions from the server through IEE.uGet. To hide any possible patterns in the query structure, the IEE randomly permutes label order and requests all at once (or alternatively in single, but successive requests, if the storage service does not provide such semantics). The IEE proceeds to decrypt these entries with  $\Theta$ , verifies if the server returned correct index values for each label (aborting the protocol otherwise), and resolves the boolean query by applying  $\phi$  to the results.

In this setting, the process of resolving a boolean query can be described in light of set operations. Searching for a keyword results in a set of document identifiers. When two or more keywords are queried, their sets can be unionized or intersected, depending if  $\phi$  specifies disjunctions or conjunctions between them, respectively. For queries of three or more keywords, parentheses can also be used to specify precedence between boolean operands. Performing negations is somewhat more complicated however, since inverting sets implies having knowledge of the range of all possible values (in this case, all document ids). To circumvent this issue we define that documents are identified by the incremental values of counter  $nDocs$ , starting at zero. Additionally, correctness of document identifiers is assured by enforcing that the ids inputted on Update belong to the range  $[0..nDocs + 1]$ . Using this approach, the system can easily filter results for all existing documents, and thus efficiently support negations by searching for a keyword and inverting its document set<sup>1</sup>.

**Optimizations and Extensions.** An important goal in BISEN is being able to support lightweight IEE technologies, such as Intel SGX with its restricted EPC size of 128MB. The proposal to extend IEE storage with cryptographically secured accesses to untrusted storage partially supports this goal. However, when performing a search in very large databases, intermediary metadata that the IEE needs to process may still be too large for such hardware restrictions. In these scenarios, incremental computing principles can be applied to ensure scalability: the IEE can dynamically calculate how many index entries will fit in its limited trusted storage, request that many entries through the IEE.uGet call, process and discard them, preserving only partial search results and merging them with the results of previous iterations of this algorithm.

Increasing the number of IEEs is also a useful extension, allowing BISEN to scale regarding both database size and client concurrency. The best way to achieve this is to make IEEs stateless: clients could generate keys  $k_E$  and  $k_F$  and share them with new IEEs (which would also help with possible IEE termination issues by the proxy server), while remaining state ( $W$  and  $nDocs$ ) could be stored in the storage service in a secure way and with a concurrency control mechanism. Encrypted index  $I$  is already monotonically crescent, hence avoiding concurrency issues.

Another design choice is how to perform delete operations. In the original approach by Cash et al. [13], delete tokens should be stored in a separate set  $D$ , indexed by a PRF over both the deleted keyword and document id. This approach has the advantage of only storing one index entry per keyword/document deletion, while BISEN stores a new entry for each deletion submitted by the client (even if repeated). However it makes keyword deletions distinguishable from additions, and requires contacting the server twice, first for accessing index  $I$  and then for accessing  $D$ .

### 4.3 Security Analysis

In BISEN our goal is for Update operations to have no leakage and Search operations to only reveal message lengths and which encrypted index entries are accessed, which we abstract as labels. This is similar to the access pattern of

<sup>1</sup> We assume that ids are never effectively removed, i.e., even if a document has all of its keywords deleted, its id will still exist and will represent an empty document. This approach has other benefits as well, including the possibility of recycling document ids.

previous SSE schemes, however it captures a stronger security notion since document identifiers are protected at all times. Moreover, executing Search for two distinct queries can leak the same label set, thus reducing the adversarial ability to distinguish between queries. For instance, boolean formulas  $\phi_1 = w_1 \vee w_2$  and  $\phi_2 = w_1 \wedge w_2$ , although representing different queries, access the same label set.

Formally, BISEN’s security is parametrized by three leakage functions ( $\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Update}}, \mathcal{L}_{\text{Search}}$ ). From these, only  $\mathcal{L}_{\text{Search}}$  produces leakage, detailed as follows:

$$\mathcal{L}_{\text{Search}}(q) = ( (|\phi| + N), |\text{Resolve}(\phi, D, n\text{Docs})|, L)$$

where the first part corresponds to the length of the input message ( $|\phi|$  is the length of the boolean formula of the query and  $N$  is the number of distinct keywords in it), the second part is the length of the query response, and  $L$  is the set of labels relevant for the resolution of the query.

The Update protocol has no leakage as all input and output messages are of equal length and cryptographic operations are performed inside the IEE. Having updates with zero leakage also ensures forward privacy [8]. Backward privacy (specifically, backward privacy with update pattern [9]) is ensured by storing additions and deletions in an indistinguishable fashion and filtering results in the IEE.

Moreover, the employed cryptographic mechanisms ensure security as follows: prf-security and uniqueness in keywords and counters ensures indistinguishability of labels from outputs from a random function applied to a unique counter; unforgeability of  $\Theta$  ensures that adversaries cannot produce a ciphertext that does not exactly match the stored data for said keyword/counter pair on the corresponding update request; the security of the IEE channel and the sequence numbers used prevent adversaries from emulating a fake BISEN execution, forging client requests, altering the order of messages exchanged, or performing replay attacks. In the Appendix Section we provide a formal security proof of these statements in the real/ideal standard cryptographic model [26].

## 5 Implementation

We implemented a prototype of BISEN in C/C++, with around 6200 lines of code. Our prototype is based on Intel SGX [16], using its remote-attestation and enclave management primitives to provide the IEE functionalities required by BISEN. To bootstrap the IEE and establish secure Client-IEE channels, we leveraged the SGX-based open-source implementation of Bahmani et al. [3], adapting it for BISEN. Their original implementation employed the NaCl cryptographic library [6] for elliptic curve algorithms and other cryptographic primitives. Our adapted implementation relies on LibSodium [29] instead, which is a more complete and up-to-date constant-time cryptographic library, partly based on AES-NI. Constant-time cryptographic algorithms based on hardware implementations and oblivious primitives allow us to prevent side-channel leakage of the SGX enclave, including page and cache level leakage.<sup>2</sup>

We instantiate PRF  $F$  with LibSodium’s SHA256-HMAC implementation, and  $\Theta$  with its authenticated encryption algorithm, XSalsa20 stream cipher with Poly1305 MACs. Since LibSodium is not ready for SGX deployment, we prepared an SGX-compatible version by (among other steps) removing all unsupported functions in SGX and replacing randomness functions with their equivalents from Intel’s RNG library.

Regarding attestation, the employed mechanism follows the design originally proposed in [4], where each program running on an IEE must produce a signature of its code and I/O trace thus far. For Intel SGX, this relies on the Quoting enclave, which uses the EPID group signature scheme [11] to produce a signature (quote) binding the enclave execution trace with the code that produced such trace. Verification of quotes is performed by the client through Intel’s Attestation Service.

For the IEE to interact with the encrypted index  $I$ , we leveraged on SGX `ocalls`. Additionally, concurrent accessed are managed by leveraging enclave multi-threading and by using concurrent data structures that only block on writes. Our implementation is open-source and available at: <https://github.com/sgtpepper/BISEN>.

## 6 Experimental Evaluation

We now experimentally evaluate BISEN, using the prototype implementation described in the previous section.

<sup>2</sup> Note that these countermeasures are not sufficient to protect from devastating attacks such as [35]. This is where BISEN benefits from relying on the IEE abstraction, as one can instead implement it on hardware resilient to speculative execution attacks if this is a realistic concern, e.g. MI6 [10].






Fig. 3: Performance of the Update protocol.

**Experimental Test-Bench.** We present performance results for BISEN and its Search and Update protocols. As IEE and proxy server, we used an Intel NUC i3-7100U with built-in SGX support, 2.4GHz of CPU frequency, 8GB of RAM, 256GB of SSD storage, running Ubuntu Server 18.04.1. As storage service we used a server with an AMD Opteron 6272 CPU with 64GB of RAM. Both machines were deployed on a one gigabit ethernet network. To evaluate the impact of remote communications, and since we already had the previous hardware available, we leveraged the cloud to deploy the client instead, using an AWS EC3 t3.large instance. The round-trip time between client and proxy server was 41.377ms and the max transmission rate was 50Mb/s. As dataset, we used a the english wikipedia dump of August 2018 [37] with around 60GB of uncompressed text data, 5.5 million documents, and 464 million keyword/document pairs. Measurements are based on an average of 50 independent executions.

**Experimental Evaluation Roadmap.** The goal of our experimental work is to answer the following questions: *i.*) what are the storage costs of BISEN; *ii.*) what is the performance cost (i.e., total time consumed) to process and store a whole dataset through a batch of Update protocol invocations, and how does this performance evolve as we scale the dataset’s size; *iii.*) what is the performance cost of executing different types of Search queries, including queries with multiple conjunctions, disjunctions, and negations, considering different database sizes, the selectivity of queried keywords (i.e., the size of returned results), and the query size; and *iv.*) how does BISEN’s performance compare with the state of art in Boolean SSE, namely the recent IEX-2LEV scheme [24].

## 6.1 Storage Costs

In BISEN, clients only store one cryptographic key (32 bytes), which is used for secure communication with the IEE. The IEE also stores this key, plus  $k_F$  and  $k_E$  ( $3 * 32 = 96$  bytes). Additionally it stores a dictionary of counters  $W$ , which keeps a counter (4 bytes) and a hash (32 bytes) per entry, with one entry per unique keyword in the database. For the English Oxford dictionary containing 616500 unique word-forms, this results in an upper bound of around 20MB IEE storage, while e-mail date searches for individual days over 200 years could correspond to around 3MB IEE storage. The storage service stores index  $I$ , which can grow due to the security guarantees provided (83 bytes per entry), nonetheless with cloud storage this can be more seamlessly scaled.

## 6.2 Update Performance

Figure 3 reports the performance results for the Update protocol of BISEN. The y-axis represents time elapsed (in seconds), while the x-axis represents the update size in terms of keyword-document pairs (i.e., how many entries are being added to index  $I$  at once, with a single batch of multiple Update protocol invocations). Results were measured at different batch update sizes (up to 464 million pairs) and are reported for networking and for the three main protocol executors in separate, namely the client, IEE, and storage service. Proxy performance is omitted for simplicity, as it only forwards messages and its execution is highly efficient. Total results are also reported for convenience of the reader.

Analysing the obtained results, one can conclude that BISEN’s performance scales linearly with the size of the batch update (Total line in Figure 3). An update for a single document with 640 keywords takes 29ms, while a batch update of multiple documents totaling 464 million keywords takes 11.239 seconds (around 3 hours). This means that




Fig. 4: Performance of each participant in the Search protocol, for an example conjunctive query of five keywords.

performance of Update invocations is mostly unaffected by the current database size. This is a natural observation, since this protocol does not depend on previous operations. These results also reflect the good performance properties of modern trusted hardware technologies, namely Intel SGX. The results for network performance basically show the cost of uploading data to the cloud, as BISEN adds very little cryptographic expansion: communications are encrypted with standard symmetric-key cryptography, and keywords are only hashed.

Regarding the performance of each protocol participant in separate, we can observe that time spent in the IEE and Storage Service is roughly similar, with a tendency for the Storage to become a bottleneck for larger operations. While we consider a single storage server, distributing this service across multiple machines might mitigate its weight in the operation. In turn, the IEE is responsible for simple cryptographic computations, entering enclave mode in SGX, and exiting this mode to store data through SGX ocalls, which is reflected in the latency for the maximum update (1300 seconds for 464 million keyword/document pairs). The largest slice of processing is on the client, which seems contradictory as from BISEN’s specification (Figure 2), the client performs very few computations. From our analysis, we argue that these results are due to necessary pre-processing: the client has to process the whole dataset from disk, parsing its keywords, stemming them and filtering stop-words [31]. The introduction of parallel processing on our client prototype can help improve these results. Additionally, in applications where documents are created and edited online for instance, this overhead would be greatly reduced.

### 6.3 Search Performance

To analyse the performance of the Search protocol, we conducted experiments with different types of queries, measuring in all cases how performance scaled with the increase in database size. For transparency in evaluation, in the following experiments we used the most popular keywords in the english language, i.e., the keywords that appear in more documents (also known as having high selectivity). From first to twelve, these are: *time, person, year, way, day, thing, man, world, life, hand, part, and child*

**Performance of each Participant.** We start by analysing the performance of networking and of each protocol participant in separate when executing the Search protocol. For this analysis we used an example conjunctive query with the five most popular keywords in the database, measuring performance at increasing database sizes. Figure 4 presents the results. In contrast with the previous results for Update, client processing in Search is very efficient. This performance cost is mostly dependent on the query size, nonetheless even for a query of five keywords it is almost close to zero (an average of  $80\mu s$ ). Networking also exhibits similar results.

The remaining performance cost is divided between the storage service and the IEE, with the IEE being the least efficient of the three components. This is due to most computations in Search being performed by the IEE. This aspect can potentially be improved by exploring parallelism in our prototype implementation based on SGX.

**Boolean Formulas and Query Size.** With this test (Figure 5) we wanted to assess the impact of both the type of operators and the length of the query on overall latency. We used queries in both Conjunctive (CNF) and Disjunctive (DNF) Normal Forms, with one, three and five conjunctions and disjunctions. These correspond, for example, to queries of the form  $(A \vee B) \wedge (C \vee D)$  (one conjunction) or  $(A \vee B) \wedge (C \vee D) \wedge (E \vee F) \wedge (G \vee H)$  (three conjunctions) for CNF; the same logic applies to the DNF.




Fig. 5: Impact of the boolean formula and query size on the performance of the Search protocol.




Fig. 6: Impact of query selectivity on the performance of the Search protocol.

Analysing the results, we can conclude that BISEN supports queries in any boolean formula with equal performance. For this experiment, the determining factors in performance were the database and query sizes. Increasing the database size leads to a linear increase in the time required for resolving queries, as was already noted in the previous experiment. Moreover, increasing the query size (from one to three and five conjunctions/disjunctions) also increases search latency, but by a smaller fraction. This means performance costs tend to amortize when increasing query sizes.

**Query Selectivity.** Next we study the impact of query selectivity (i.e., the size of search results) on Search performance. In these experiments, we performed single-keyword queries with different selectivity levels, by choosing query keywords based on their database popularity. Figure 6 shows the results for queries returning from 0.2% to 25% of the database. As expected, query selectivity has a high impact on Search performance. Just by searching a different, more popular keyword, Search performance can go from 1 to 16 seconds. This is not surprising, as more popular keywords appear in more documents, and hence the IEE will have to request, decrypt, and verify additional index entries. Nonetheless, results seem to amortize towards larger databases. These results are also consistent with the performance measurements of Figure 5, whose keyword searches have very high selectivity.

**Negations.** In Table 1 we present the impact of negations for queries of fixed size (10 keywords), varying the number of negated keywords – one, five and ten; then a fully negated query – of the form  $\neg(A \wedge B)$ , and finally the equivalent version of the latter using De Morgan’s laws. Our objective was to assess the impact performance of the negation operation across different types of queries and numbers of negations. Results show that the number of negation operations performed has minimal impact, even for larger database sizes, which can be explained by the low overhead of Boolean processing. Since all queries require the same number of entries to be fetched from Storage Service, which is where the main bottleneck lies, their latency is therefore similar.

DB Size	1 Neg.	5 Neg.	10 Neg.	Fully Neg.
35 996 207	4.286	4.498	3.052	4.319
76 672 004	9.335	9.241	9.610	7.185
156 143 147	18.653	18.092	21.095	16.589
333 784 724	52.265	58.227	50.850	51.996
464 054 543	86.057	82.289	85.041	86.938

Table 1: Performance (in seconds) of negations in the Search protocol.

Database Size (Nr of pairs w/id)	Update		Search CNF		Search DNF	
	BISEN	IEX-2LEV	BISEN	IEX-2LEV	BISEN	IEX-2LEV
9 793	0.151	5143	0.004	12	0.004	15
27 446	0.423	15568	0.021	173	0.012	249
56 238	0.862	29274	0.061	216	0.034	427

Table 2: Performance comparison between BISEN and IEX-2LEV [24]. All times are in seconds. Queries composed of eight keywords.

#### 6.4 Comparison with IEX-2LEV

We now compare the performance of BISEN with the state of the art in Boolean SSE, in particular the recent IEX-2LEV scheme [24]. To this end, we used the author’s open-source implementation [18] (with a filtering parameter of 0.2, as reported in their evaluation [24]), and conducted experiments with the Enron database [27], an email archive with 2.6GB of text data used by the authors.

Since IEX-2LEV requires large volatile storage and was originally evaluated on a machine with 60 GB of RAM and a 60-core CPU, we followed a similar test-bench and deployed IEX-2LEV in our AMD Opteron 6272 CPU with 64 cores and 64GB of RAM. For experimental comparison we deployed BISEN on the same machine, executing IEE computations in SGX simulated mode. Table 2 presents the results obtained for BISEN and IEX-2LEV, considering increasing database sizes (up to 56238 keyword-document pairs, as we were unable to execute IEX-2LEV with higher database sizes), and different operations: Update (performed as Setup in IEX-2LEV), and Search with queries with eight keywords (selected at random from the Enron database) in both CNF and DNF.

Analysing the results we can conclude that BISEN is much more efficient than the state of the art in Boolean SSE. This phenomenon can be observed both for the Update operation, where IEX-2LEV requires eight hours to index a database with 56 238 pairs while BISEN only requires 0.151 seconds; and the Search operation, where IEX-2LEV is more efficient but still requires 216 seconds to search the largest database with a CNF boolean query while BISEN performs the same query in 0.061 seconds. Furthermore, the improvement in storage performance is also evident from these results, since BISEN could process and index large databases with 10 million pairs in a machine with only 8 GB of RAM and IEX-2LEV could only support little more than 56 thousand pairs in a machine with 64 GB. These results can be explained by the difficulty of managing complex multi-map data-structures that IEX-2LEV needs to employ in order to achieve its security guarantees. In BISEN, by leveraging the natural synergy between standard cryptographic primitives and IEEs deployed as remote trust anchors, we are able to improve performance and scalability by a large fraction, while further improving security and minimizing leakage.

## 7 Conclusions

In this paper, we have identified and addressed one of the fundamental security issues in Searchable Symmetric Encryption (SSE) schemes, which is the outsourcing of critical cryptographic computations to the untrusted server. This was achieved by proposing a new hybrid approach to SSE that combines standard symmetric-key cryptographic primitives with modern attestation-based trusted hardware. In our approach we minimize assumptions and requirements on the employed hardware technology, in particular regarding its trusted storage capacity. Instead, trusted hardware is used as a limited-capacity Isolated Execution Environment abstraction, extending its resources through standard cryptographic primitives over more abundant (local, or even remote) untrusted resources. Based on this hybrid approach, we proposed BISEN, a new dynamic boolean SSE scheme with both forward and backward privacy, minimal leakage, and optimal computation, storage, and communication overheads. BISEN is shown to be provably secure against

active adversaries under the standard security model. Experimental results obtained through real-world datasets and an open-source implementation of BISEN demonstrate its optimal performance and efficiency properties.

## Acknowledgments

This work was supported by FCT/MCTES through project HADES (PTDC/CCI-INF/31698/2017) and NOVA LINCIS (UID/CEC/04516/2013), and the EU through project LightKone (grant agreement  $n^{\circ}$  732505).

## References

1. T. Alves and D. Felton. TrustZone: Integrated hardware and software security. *ARM white paper*, 3(4):18–24, 2004.
2. M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia. A view of cloud computing. *Communications of the ACM*, 53(4):50–58, 2010.
3. R. Bahmani, M. Barbosa, F. Brasser, B. Portela, A.-R. Sadeghi, G. Scerri, and B. Warinschi. Secure multiparty computation from SGX. In *Financial Cryptography and Data Security - FC'17*, 2017.
4. M. Barbosa, B. Portela, G. Scerri, and B. Warinschi. Foundations of hardware-based attested computation and application to SGX. In *EURO S&P'16*, pages 245–260, 2016.
5. M. Bellare and P. Rogaway. Introduction to modern cryptography. *Ucsd Cse*, 207:207, 2005.
6. D. Bernstein, T. Lange, and P. Schwabe. The security impact of a new cryptographic library. In *LATINCRYPT'12*. Springer, 2012.
7. C. Bösch, P. Hartel, W. Jonker, and A. Peter. A Survey of Provably Secure Searchable Encryption. *ACM CSUR*, 47(2):18:1—18:51, 2015.
8. R. Bost. Sophos - Forward Secure Searchable Encryption. In *CCS'16*. ACM, 2016.
9. R. Bost, B. Minaud, and O. Ohrimenko. Forward and Backward Private Searchable Encryption from Constrained Cryptographic Primitives. In *CCS'17*. ACM, 2017.
10. T. Bourgeat, I. Lebedev, A. Wright, S. Zhang, S. Devadas, et al. Mi6: Secure enclaves in a speculative out-of-order processor. *arXiv preprint arXiv:1812.09822*, 2018.
11. E. Brickell and J. Li. Enhanced privacy id from bilinear pairing for hardware authentication and attestation. *International Journal of Information Privacy, Security and Integrity* 2, 1(1):3–33, 2011.
12. D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. Leakage-Abuse Attacks Against Searchable Encryption. In *CCS'15*, pages 668–679. ACM, 2015.
13. D. Cash, J. Jaeger, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Dynamic searchable encryption in very-large databases: Data structures and implementation. In *NDSS'14*, volume 14, 2014.
14. D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Rosu, and M. Steiner. Highly-Scalable Searchable Symmetric Encryption with Support for Boolean Queries. In *CRYPTO'13*, pages 353–373. Springer, 2013.
15. ComScore. The 2017 U.S. Mobile App Report. <http://tinyurl.com/ya8kkxan>, 2017.
16. V. Costan and S. Devadas. Intel sgx explained. Technical report, Cryptology ePrint Archive, Report 2016/086, 2016. <https://eprint.iacr.org/2016/086>, 2016.
17. R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky. Searchable Symmetric Encryption: Improved Definitions and Efficient Constructions. In *CCS'06*, pages 79–88, 2006.
18. Encrypted Systems Lab, Brown University. The Clusion library. <https://github.com/encryptedsystems/Clusion>, 2018.
19. B. A. Fisch, D. Vinayagamurthy, D. Boneh, and S. Gorbunov. Iron: Functional encryption using intel sgx. In *CCS'17*. ACM, 2017.
20. T. Frieden. VA will pay \$20 million to settle lawsuit over stolen laptop's data. CNN. <http://tinyurl.com/lg4os9m>, 2009.
21. B. Fuhry, R. Bahmani, F. Brasser, F. Hahn, F. Kerschbaum, and A.-R. Sadeghi. Hardidx: practical and secure index with sgx. In *IFIP DBSec*, pages 386–408. Springer, 2017.
22. S. Garg, P. Mohassel, and C. Papamanthou. TWORAM: efficient oblivious RAM in two rounds with applications to searchable encryption. In *Crypto'16*, pages 563–592. Springer, 2016.
23. G. Greenwald and E. MacAskill. NSA Prism program taps in to user data of Apple, Google and others. The Guardian. <http://tinyurl.com/oea3g8t>, 2013.
24. S. Kamara and T. Moataz. Boolean Searchable Symmetric Encryption with Worst-Case Sub-Linear Complexity. In *EURO-CRYPT'17*. IACR, 2017.
25. S. Kamara, C. Papamanthou, and T. Roeder. Dynamic searchable symmetric encryption. In *CCS'12*, pages 965–976. ACM, 2012.
26. J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. CRC PRESS, 2007.
27. B. Klimt and Y. Yang. Introducing the Enron Corpus. In *CEAS*, 2004.
28. D. Lewis. iCloud Data Breach: Hacking And Celebrity Photos. Forbes. <https://tinyurl.com/nohznmr>, 2014.

<pre> Setup(<math>1^\lambda</math>)   <u>Client:</u>   1: <math>k_E \leftarrow \text{\\$} \Theta.\text{Gen}(1^\lambda)</math>   2: <math>k_F \leftarrow \text{\\$} F.\text{Gen}(1^\lambda)</math>   3: <math>W \leftarrow \text{Init}()</math>   4: <math>n\text{Docs} \leftarrow 0</math>   <u>Server:</u>   5: <math>I \leftarrow \text{Init}()</math> </pre> <hr style="border: 0.5px solid black;"/> <pre> Update(<math>op, w, id</math>)   <u>Client:</u>   1: <b>if</b> <math>id &gt; n\text{Docs}</math> <b>then</b>   2:   <math>n\text{Docs}++</math>   3:   <math>c \leftarrow \text{Get}(W, w)</math>   4:   <b>if</b> <math>c = \perp</math> <b>then</b>   5:     <math>c \leftarrow 0</math>   6:   <b>else</b>   7:     <math>c \leftarrow c + 1</math>   8:   <math>k_w \leftarrow F.\text{Run}(k_F, w)</math>   9:   <math>l \leftarrow F.\text{Run}(k_w, c)</math>   10: <math>id^* \leftarrow \text{\\$} \Theta.\text{Enc}(k_E, \{l, op, id\})</math>   11: Send <math>l, id^*</math> to Server.   12: <math>W \leftarrow \text{Put}(W, w, c)</math>   <u>Server:</u>   13: <math>I \leftarrow \text{Put}(I, l, id^*)</math> </pre>	<pre> Search(<math>q</math>)   <u>Client:</u>   1: <math>\{\bar{w}, \phi\} \leftarrow \text{ProcessBooleanQuery}(q)</math>   2: <math>Q \leftarrow \text{Init}()</math>   3: <b>for all</b> <math>w \in \bar{w}</math> <b>do</b>   4:   <math>k_w \leftarrow F.\text{Run}(k_F, w); c \leftarrow \text{Get}(W, w)</math>   5:   <math>L \leftarrow []</math>   6:   <b>for all</b> <math>c_i \leftarrow 0 \dots c</math> <b>do</b>   7:     <math>l \leftarrow F.\text{Run}(k_w, c_i); L \leftarrow l : L</math>   8:   <math>Q \leftarrow \text{Put}(Q, w, L)</math>   9: <math>L' \leftarrow \text{Flatten}(Q); \Pi \leftarrow \text{\\$} \text{RandomPermutation}(1^\lambda)</math>   10: <math>L' \leftarrow \Pi(L')</math>   11: Send <math>L'</math> to Server.   <u>Server:</u>   12: <math>D' \leftarrow []</math>   13: <b>for all</b> <math>l \in L'</math> <b>do</b>   14:   <math>id^* \leftarrow \text{Get}(I, l); D' \leftarrow id^* : D'</math>   15: Send <math>D'</math> to Client.   <u>Client:</u>   16: <math>D' \leftarrow \Pi^{-1}(D'); D \leftarrow []</math>   17: <b>for all</b> <math>id^* \in D'; l' \in L'</math> <b>do</b>   18:   <math>\{l, op, id\} \leftarrow \Theta.\text{Dec}(k_E, id^*); \text{Verify}(l, l')</math>   19:   <math>D \leftarrow \{op, id\} : D</math>   20: <math>Q' \leftarrow \text{Join}(Q, D)</math>   21: <math>R \leftarrow \text{Resolve}(\phi, Q', n\text{Docs})</math> </pre>
---	---

Fig. 7: A variant of BISEN without IEEs, based solely on PRF  $F = (\text{Gen}, \text{Run})$  and authenticated encryption scheme  $\Theta = (\text{Gen}, \text{Enc}, \text{Dec})$ .

29. libsodium Development Team. The sodium crypto library (libsodium). <https://libsodium.org>, 2018.
30. C. Liu, L. Zhu, M. Wang, and Y.-A. Tan. Search pattern leakage in searchable encryption: Attacks and new construction. *Information Sciences*, 265:176–188, 2014.
31. C. D. Manning, P. Raghavan, and H. Schütze. *An Introduction to Information Retrieval*, volume 1. Cambridge University Press, 2009.
32. M. Russinovich. Introducing Azure confidential computing. <https://tinyurl.com/y3qqwguk>, 2017.
33. S. Sasy, S. Gorbunov, and C. W. Fletcher. Zerotracer: Oblivious memory primitives from intel sgx. In *NDSS'18*, 2018.
34. E. Stefanov, C. Papamanthou, and E. Shi. Practical Dynamic Searchable Encryption with Small Leakage. In *NDSS'14*, 2014.
35. J. Van Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx. Foreshadow: Extracting the keys to the intel sgx kingdom with transient out-of-order execution. In *Security'18*. Usenix, 2018.
36. W. Wang, G. Chen, X. Pan, Y. Zhang, X. Wang, V. Bindschaedler, H. Tang, and C. A. Gunter. Leaky cauldron on the dark land: Understanding memory side-channel hazards in sgx. In *CCS'17*, 2017.
37. I. Wikimedia Foundation. Wikipedia:Database download. [https://en.wikipedia.org/wiki/Wikipedia:Database\\_download](https://en.wikipedia.org/wiki/Wikipedia:Database_download), 2018.
38. Y. Zhang, J. Katz, and C. Papamanthou. All Your Queries Are Belong to Us: The Power of File-Injection Attacks on Searchable Encryption. In *Security'16*. USENIX Association, 2016.

## Building BISEN without IEEs

Figure 7 of this Appendix Section details a variant of BISEN that can be used in scenarios where trusted hardware and IEEs are unavailable. In a nutshell, computations previously performed by the IEE are now performed by the client instead. This approach achieves the same security guarantees as the original BISEN scheme, however it imposes a higher communication overhead.

## Formal Security model

Semantic security of an SSE scheme is defined with respect to a well defined leakage function  $\mathcal{L} = (\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}})$  [17, 25]. This definition of security follows the simulation-based real/ideal paradigm that is standard for security definitions in cryptography [26]. Leakage function  $\mathcal{L}$  describes precisely what information each protocol in the scheme is allowed to reveal. Definition 1 captures the security model of SSE.

**Definition 1.** Let  $\Pi = (\text{Setup}, \text{Search}, \text{Update})$  be a dynamic boolean SSE scheme and  $\mathcal{L} = (\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Search}}, \mathcal{L}_{\text{Update}})$  a leakage function. For algorithms  $\mathcal{A}$  (the adversary) and  $\mathcal{S}$  (a simulator), we define security games  $\mathbf{Real}_{\Pi, \mathcal{A}}(1^\lambda)$  and  $\mathbf{Ideal}_{\mathcal{L}, \mathcal{S}, \mathcal{A}}(1^\lambda)$  as follows:

$\mathbf{Real}_{\Pi, \mathcal{A}}(1^\lambda)$ : run  $(\text{EDB}, \text{K}) \leftarrow_{\$} \text{Setup}(1^\lambda)$  and give EDB to  $\mathcal{A}$ .  $\mathcal{A}$  then adaptively requests executions of Search and Update, selecting client inputs  $\text{inp}$ . The game responds by executing the requested protocols with input  $(\text{K}, \text{inp})$ , allowing  $\mathcal{A}$  to select the server input EDB and arbitrarily respond to uGet requests. The execution transcripts are then provided to  $\mathcal{A}$ . Eventually,  $\mathcal{A}$  returns a bit, which is the output of the game.

$\mathbf{Ideal}_{\mathcal{L}, \mathcal{S}, \mathcal{A}}(1^\lambda)$ : run  $(\text{EDB}, \text{st}) \leftarrow_{\$} \mathcal{S}(\mathcal{L}^{\text{Setup}}(1^\lambda))$  and give EDB to  $\mathcal{A}$ .  $\mathcal{A}$  then repeatedly requests protocols Search and Update, selecting client inputs  $\text{inp}$ , server input EDB and arbitrarily respond to uGet requests. The game responds by performing  $\mathcal{S}(\text{st}, \mathcal{L}^{\text{Search}}(\text{inp}))$  and  $\mathcal{S}(\text{st}, \mathcal{L}^{\text{Update}}(\text{inp}))$ , respectively, and returning the simulated transcript back to  $\mathcal{A}$ . Eventually,  $\mathcal{A}$  returns a bit, which is the output of the game.

$\Pi$  is  $\mathcal{L}$ -secure against adaptive attacks if, for any active adversary  $\mathcal{A}$ , there exists a simulator  $\mathcal{S}$  such that:

$$|\Pr[\mathbf{Real}_{\Pi, \mathcal{A}}(1^\lambda) = 1] - \Pr[\mathbf{Ideal}_{\mathcal{L}, \mathcal{S}, \mathcal{A}}(1^\lambda) = 1]| \leq \mu(\lambda)$$

Given the previous definition, Figure 8 extends it with a game-based approach. In  $\mathbf{Real}_{\Pi, \mathcal{A}}$ , we first initialize the IEE-enabling machine (IEE.Init), and then allow  $\mathcal{A}$  to interact with the BISEN protocol, detailed in a similar fashion in Figure 9. Since  $\mathcal{A}$  is an active adversary, we allow him full control over the untrusted storage component.  $\mathcal{A}$  is given feedback whenever a ulnit and uPut is executed, and is allowed to freely specify the output of any uGet request (thus the split of Search<sub>1</sub>, Search<sub>2</sub>). In  $\mathbf{Ideal}_{\mathcal{L}, \mathcal{S}}$ , simulator  $\mathcal{S}$  first initializes the public parameters, and then  $\mathcal{A}$  is allowed to interact with  $\mathcal{S}$ , as described in Figure 11, receiving the leakage associated with each operation via  $\mathcal{L} = (\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Update}}, \mathcal{L}_{\text{Search}})$ , detailed in Figure 10.

For clarity of presentation, our security model prevents the adversary from creating arbitrary IEEs, from attempting to forge requests from the IEE to the client and vice-versa, and from changing the order of requests. Indeed, our proof can be extended to an active adversary in the IEE model of [3] with all these capabilities. However, since this cannot be done in a black-box manner, the technical details orthogonal to our contribution (e.g. bookkeeping of arbitrary inputs to arbitrary IEEs) would make the model and proof significantly more convoluted. We thus prefer to present a model that highlights the necessary mechanisms for BISEN to be secure against active adversaries, assuming the security of the underlying mechanisms. Briefly, the intuition as to why these behaviours are not an issue is threefold. The security of attested key exchange ensures that no more than a single instance of an IEE loaded with the exact code of BISEN successfully performs the setup (key exchange) with the client, so no advantage is gained from launching additional IEEs; the secure channel prevents an external adversary from producing any valid inputs to either the client or the IEE; and the usage of sequence numbers allows for the rejection of any request that is presented in an incorrect order.

Given the previous definitions, Theorem 1 states BISEN's security:

**Theorem 1.** *If encryption scheme  $\Theta$  ensures IND-CCA security,  $F$  is a pseudorandom function and IEE provides a secure channel with IND-CCA security, then the BISEN scheme presented in Figure 2 is  $(\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Update}}, \mathcal{L}_{\text{Search}})$ -secure according to Definition 1.*

## Proof of Theorem 1

Let  $\Theta = (\text{Gen}, \text{Enc}, \text{Dec})$  be an IND-CCA encryption scheme following the security definitions of [26]. Let  $F = (\text{Gen}, \text{Run})$  be a pseudo-random function of domain  $D$  and output range  $R$  ensuring prf-security for generator  $\text{Func}(D, R)$  following the definitions of [5]. Let  $\Gamma = (\text{New}, \text{Put}, \text{Get})$  and  $\Gamma = (\text{ulnit}, \text{uPut}, \text{uGet})$  be structures for safe/unsafe storage (respectively) described in Section 4.1. Let IEE = (Setup, Send, Receive) be the secure channel protocol for IEEs described in Section 4.1. Let PBQ refer to ProcessBooleanQuery detailed in Section 4. For clarity in presentation, we simplify the process of lines 15 – 23 where entries are inserted using a  $\Gamma$  structure to use lists, and thus we denote Sort as the probabilistic algorithm that sorts a list and produces the employed permutation, which can afterwards be recovered using Reorder. Let  $\text{id}_s$  be the fixed length of document identifiers,  $\text{op}_s$  be the fixed length of operations,  $f_s$  be the fixed length of  $F$  output, and let  $U_{\text{size}}$  denote the fixed size of updates, such that

$$U_{\text{size}} = \text{op}_s + \text{id}_s + f_s$$

<p><b>Game Real</b><math>\Pi, \mathcal{A}(1^\lambda)</math>:</p> <pre> prms <math>\leftarrow</math> IEE.Init(<math>1^\lambda</math>) (st, t) <math>\leftarrow</math> <math>\Pi</math>.Setup(<math>1^\lambda</math>, prms) st<math>_A</math> <math>\leftarrow</math> t Return <math>\mathcal{A}_1^{\text{Update, Search}}(1^\lambda, t)</math> </pre> <p><b>Oracle Update</b>(op, w, id):</p> <pre> (st, t) <math>\leftarrow</math> <math>\Pi</math>.Update(st, op, w, id) st<math>_A</math> <math>\leftarrow</math> <math>\mathcal{A}_2(\text{st}_A, t)</math> Return st<math>_A</math> </pre>	<p><b>Oracle Search</b>(q):</p> <pre> (st, t<math>_1</math>) <math>\leftarrow</math> <math>\Pi</math>.Search<math>_1</math>(st, q) (st<math>_A</math>, m) <math>\leftarrow</math> <math>\mathcal{A}_3(\text{st}_A, t_1)</math> (st, r, t<math>_2</math>) <math>\leftarrow</math> <math>\Pi</math>.Search<math>_2</math>(st, m) (st<math>_A</math>) <math>\leftarrow</math> <math>\mathcal{A}_4(\text{st}_A, t_2)</math> Return (r, st<math>_A</math>) </pre>
<p><b>Game Ideal</b><math>\mathcal{L}, \mathcal{A}, \mathcal{S}(1^\lambda)</math>:</p> <pre> (st<math>_L</math>, l) <math>\leftarrow</math> <math>\mathcal{L}</math>Setup(<math>1^\lambda</math>) (st<math>_S</math>, t) <math>\leftarrow</math> <math>\mathcal{S}_1(1^\lambda, l)</math> st<math>_A</math> <math>\leftarrow</math> t Return <math>\mathcal{A}_1^{\text{Update, Search}}(1^\lambda, t)</math> </pre> <p><b>Oracle Update</b>(op, w, id):</p> <pre> (st<math>_L</math>, l) <math>\leftarrow</math> <math>\mathcal{L}</math>Update(st<math>_L</math>, op, w, id) (st<math>_S</math>, t) <math>\leftarrow</math> <math>\mathcal{S}_2(\text{st}_S, l)</math> st<math>_A</math> <math>\leftarrow</math> <math>\mathcal{A}_2(\text{st}_A, t)</math> Return st<math>_A</math> </pre>	<p><b>Oracle Search</b>(q):</p> <pre> (st<math>_L</math>, l) <math>\leftarrow</math> <math>\mathcal{L}</math>Search(st<math>_L</math>, q) (st<math>_S</math>, t<math>_1</math>) <math>\leftarrow</math> <math>\mathcal{S}_3(\text{st}_S, l)</math> (st<math>_A</math>, m) <math>\leftarrow</math> <math>\mathcal{A}_3(\text{st}_A, t_1)</math> (st<math>_S</math>, r, t<math>_2</math>) <math>\leftarrow</math> <math>\mathcal{S}_4(\text{st}_S, m)</math> (st<math>_A</math>) <math>\leftarrow</math> <math>\mathcal{A}_4(\text{st}_A, t_2)</math> Return (r, st<math>_A</math>) </pre>

Fig. 8: Security experiment

<p><b>Algorithm Setup</b>(<math>1^\lambda</math>, prms):</p> <pre> k<math>_f</math> <math>\leftarrow</math> F.Gen(<math>1^\lambda</math>) (k<math>_c</math>, t) <math>\leftarrow</math> IEE.Setup(<math>1^\lambda</math>, prms) W <math>\leftarrow</math> <math>\Gamma</math>.New() k<math>_e</math> <math>\leftarrow</math> <math>\Theta</math>.Gen(<math>1^\lambda</math>) l <math>\leftarrow</math> <math>\Gamma</math>.ulnit() Return ((k<math>_c</math>, k<math>_e</math>, k<math>_f</math>, 0, W), (t, l)) </pre> <p><b>Algorithm Update</b>(st, op, id, w):</p> <pre> (k<math>_c</math>, k<math>_e</math>, k<math>_f</math>, nDocs, W) <math>\leftarrow</math> st k<math>_w</math> <math>\leftarrow</math> F.Run(k<math>_f</math>, w) c* <math>\leftarrow</math> IEE.Send(k<math>_c</math>, (op, id, k<math>_w</math>)) If (c <math>\leftarrow</math> <math>\Gamma</math>.Get(W, w)):   c = c + 1 Else: c <math>\leftarrow</math> 0 W <math>\leftarrow</math> <math>\Gamma</math>.Put(W, w, c) If id &gt; nDocs:   nDocs <math>\leftarrow</math> nDocs + 1 l <math>\leftarrow</math> F.Run(k<math>_w</math>, c) id* <math>\leftarrow</math> <math>\Theta</math>.Enc(k<math>_e</math>, (l, (op, id))) (l, l<math>_t</math>) <math>\leftarrow</math> <math>\Gamma</math>.uPut(l, l, id*) Return ((k<math>_c</math>, k<math>_e</math>, k<math>_f</math>, nDocs, W), (c*, l<math>_t</math>)) </pre>	<p><b>Algorithm Search<math>_1</math></b>(st, q):</p> <pre> (k<math>_c</math>, k<math>_e</math>, k<math>_f</math>, nDocs, W) <math>\leftarrow</math> st L <math>\leftarrow</math> []; n <math>\leftarrow</math> 0 q* <math>\leftarrow</math> IEE.Send(k<math>_c</math>, q) (W<math>_q</math>, <math>\phi</math>) <math>\leftarrow</math> PBQ(q) For w <math>\in</math> W<math>_q</math>:   k<math>_w</math> <math>\leftarrow</math> F.Run(k<math>_f</math>, w)   c <math>\leftarrow</math> <math>\Gamma</math>.Get(W, w)   For k <math>\in</math> [0 .. c]:     l <math>\leftarrow</math> F.Run(k<math>_w</math>, k)     L <math>\leftarrow</math> l : L; n <math>\leftarrow</math> n + 1 (<math>\Delta</math>, L<math>_p</math>) <math>\leftarrow</math> Sort(L) Return ((st, <math>\Delta</math>, L<math>_p</math>, n), (L<math>_p</math>, q*)) </pre> <p><b>Algorithm Search<math>_2</math></b>((st, <math>\Delta</math>, L<math>_p</math>, n), m):</p> <pre> D <math>\leftarrow</math> [] For i <math>\in</math> [0..n]:   (l, (op, id)) <math>\leftarrow</math> <math>\Theta</math>.Dec(k<math>_e</math>, m[i])   If L<math>_p</math>[i] <math>\neq</math> l: abort   D <math>\leftarrow</math> (op, id) : D D' <math>\leftarrow</math> Reorder(<math>\Delta</math>, D) r <math>\leftarrow</math> Resolve(<math>\phi</math>, D', nDocs) r* <math>\leftarrow</math> IEE.Send(k<math>_c</math>, r) Return (st, r, r*) </pre>
--	---

Fig. 9: Boolean SSE protocol

<p><b>Algorithm <math>\mathcal{L}</math>Setup</b>(<math>1^\lambda</math>):</p> <pre> Return (([], <math>\Gamma</math>.New, 0, 0), <math>\perp</math>) </pre> <p><b>Algorithm <math>\mathcal{L}</math>Update</b>(st, op, id, w):</p> <pre> (W, A, c, nDocs) <math>\leftarrow</math> st If w <math>\notin</math> W: W <math>\leftarrow</math> w : W A <math>\leftarrow</math> <math>\Gamma</math>.Put(A, c, (op, w, id, c)) If id &gt; nDocs: nDocs <math>\leftarrow</math> nDocs + 1 st <math>\leftarrow</math> (W, A, c + 1, nDocs) Return (st, <math>\perp</math>) </pre>	<p><b>Algorithm <math>\mathcal{L}</math>Search</b>(st, q):</p> <pre> (W, A, n, nDocs) <math>\leftarrow</math> st C <math>\leftarrow</math> []; D <math>\leftarrow</math> []; N <math>\leftarrow</math> 0 (W<math>_q</math>, <math>\phi</math>) <math>\leftarrow</math> PBQ(q) For w <math>\in</math> W<math>_q</math>: If w <math>\in</math> W: N <math>\leftarrow</math> N + 1 qi<math>_{\text{size}}</math> <math>\leftarrow</math> <math> \phi  + N * f_s</math> For (op, w, id, c) <math>\in</math> st <math>\wedge</math> w <math>\in</math> W<math>_q</math>:   C <math>\leftarrow</math> (c : C); D <math>\leftarrow</math> (id : D) (<math>\cdot</math>, C') <math>\leftarrow</math> Sort(C) r <math>\leftarrow</math> Resolve(<math>\phi</math>, D, nDocs) qo<math>_{\text{size}}</math> <math>\leftarrow</math>  r  Return (st, r, (qi<math>_{\text{size}}</math>, qo<math>_{\text{size}}</math>, C')) </pre>
--	---

Fig. 10: Leakage functions



<p><b>Algorithm <math>\mathcal{S}_1(1^\lambda, l)</math>:</b>  <math>g \leftarrow \text{Func}(D, R)</math>  <math>(st_S, t) \leftarrow \mathcal{S}_{IEE}(1^\lambda)</math>  <math>k_e \leftarrow \Theta.\text{Gen}(1^\lambda)</math>  <math>l \leftarrow \Gamma.\text{ulnit}()</math>  Return <math>((st_S, k_e, g, 0), (t, l))</math></p> <p><b>Algorithm <math>\mathcal{S}_2(st, \perp)</math>:</b>  <math>(st_S, k_e, g, c) \leftarrow st</math>  <math>c^* \leftarrow \mathcal{S}_{IEE}(st_S, (\{0\}^{U_S}))</math>  <math>l \leftarrow g(c)</math>  <math>id^* \leftarrow \Theta.\text{Enc}(k_e, (l, \{0\}^{ids+ops}))</math>  <math>(l, l_t) \leftarrow \Gamma.\text{uPut}(l, l, id^*)</math>  Return <math>((st_S, k_e, g, c+1), (c^*, l_t))</math></p>	<p><b>Algorithm <math>\mathcal{S}_3(st, q_{i\_size}, q_{o\_size}, C)</math>:</b>  <math>(st_S, k_e, g, c) \leftarrow st</math>  <math>L \leftarrow []; n \leftarrow 0</math>  <math>q_{in} \leftarrow \{0\}^{q_{i\_size}}</math>  <math>q_{in}^* \leftarrow \mathcal{S}_{IEE}(st_S, q_{in})</math>  For <math>k \in C</math>:  <math>l \leftarrow g(k)</math>  <math>L \leftarrow l : L; n \leftarrow n + 1</math>  Return <math>((st, L, n, q_{o\_size}), (L, q_{in}))</math></p> <p><b>Algorithm <math>\mathcal{S}_4((st, L, n, q_{o\_size}), m)</math>:</b>  <math>(st_S, k_e, g, c) \leftarrow st</math>  <math>q_{out} \leftarrow \{0\}^{q_{o\_size}}</math>  <math>q_{out}^* \leftarrow \mathcal{S}_{IEE}(st_S, q_{out})</math>  For <math>i \in [0..n]</math>:  <math>(l, *) \leftarrow \Theta.\text{Dec}(k_e, m[i])</math>  If <math>L[i] \neq l</math>: abort  Return <math>((st_S, k_e, g, c), q_{out}^*)</math></p>
--	--

Fig. 11: Simulator behavior

<p><b>Game <math>G_0^A(1^\lambda)</math>:</b>  <math>prms \leftarrow \text{IEE.Init}(1^\lambda)</math>  <math>k_f \leftarrow F.\text{Gen}(1^\lambda)</math>  <math>W \leftarrow \Gamma.\text{New}()</math>  <math>(k_c, t) \leftarrow \text{IEE.Setup}(1^\lambda, prms)</math>  <math>k_e \leftarrow \Theta.\text{Gen}(1^\lambda)</math>  <math>l \leftarrow \Gamma.\text{ulnit}()</math>  <math>st \leftarrow (k_c, k_e, k_f, 0, W)</math>  <math>st_A \leftarrow (t, l)</math>  Return <math>\mathcal{A}_1^{\text{Update, Search}}(1^\lambda, st_A)</math></p> <p><b>Oracle Update(<math>st, id, w</math>):</b>  <math>(k_c, k_e, k_f, nDocs, W) \leftarrow st</math>  <math>k_w \leftarrow F.\text{Run}(k_f, w)</math>  <math>c^* \leftarrow \text{IEE.Send}(k_c, (op, id, k_w))</math>  If <math>(c \leftarrow \Gamma.\text{Get}(W, w))</math>:  <math>c = c + 1</math>  Else: <math>c \leftarrow 0</math>  <math>W \leftarrow \Gamma.\text{Put}(W, w, c)</math>  If <math>id &gt; nDocs</math>: <math>nDocs \leftarrow nDocs + 1</math>  <math>l \leftarrow F.\text{Run}(k_w, c)</math>  <math>id^* \leftarrow \Theta.\text{Enc}(k_e, (l, (op, id)))</math>  <math>(l, l_t) \leftarrow \Gamma.\text{uPut}(l, l, id^*)</math>  <math>st \leftarrow (k_c, k_e, k_f, nDocs, W)</math>  <math>st_A \leftarrow \mathcal{A}_2(st_A, (c^*, l_t))</math>  Return <math>st_A</math></p>	<p><b>Oracle Search(<math>st, q</math>):</b>  <math>(k_c, k_e, k_f, nDocs, W) \leftarrow st</math>  <math>L \leftarrow []; n \leftarrow 0; D \leftarrow []</math>  <math>q^* \leftarrow \text{IEE.Send}(k_c, q)</math>  <math>(W_q, \phi) \leftarrow \text{PBQ}(q)</math>  For <math>w \in W_q</math>:  <math>k_w \leftarrow F.\text{Run}(k_f, w)</math>  <math>c \leftarrow \Gamma.\text{Get}(W, w)</math>  For <math>k \in [0..c]</math>:  <math>l \leftarrow F.\text{Run}(k_w, k)</math>  <math>L \leftarrow l : L; n \leftarrow n + 1</math>  <math>(\Delta, L_p) \leftarrow \text{Sort}(L)</math>  <math>(st_A, m) \leftarrow \mathcal{A}_3(st_A, (L_p, q^*))</math>  For <math>i \in [0..n]</math>:  <math>(l, (op, id)) \leftarrow \Theta.\text{Dec}(k_e, m[i])</math>  If <math>L_p[i] \neq l</math>: abort  <math>D \leftarrow (op, id) : D</math>  <math>D' \leftarrow \text{Reorder}(\Delta, D)</math>  <math>r \leftarrow \text{Resolve}(\phi, D', nDocs)</math>  <math>r^* \leftarrow \text{IEE.Send}(k_c, r)</math>  <math>(st_A) \leftarrow \mathcal{A}_4(st_A, r^*)</math>  Return <math>(r, st_A)</math></p>
---	---

Fig. 12: Extended real world.

*Proof.* Our proof is a sequence of eight games, presented in Figures 12 to 19.

Game  $G_0^A$  is the real world of Figure 8, extended with the protocol of Figure 9. In game  $G_1^A$ , a uPut is always accompanied by an idealised storage Put, and Search instead uses the ideal storage to process the query. We upper bound the distance between these two games, by constructing an adversary  $\mathcal{B}$  against the existential unforgeability of  $\Theta$ , such that

$$|\Pr[G_0^A(1^\lambda) \Rightarrow T] - \Pr[G_1^A(1^\lambda) \Rightarrow T]| \leq \frac{\text{Adv}_{\Theta, \mathcal{B}}^{\text{uf}}(\lambda)}{s * i}$$

Adversary  $\mathcal{B}$  simulates the environment of  $G_1^A$  as follows. At the beginning of the game,  $\mathcal{B}$  has to try and guess which uGet to Search will be distinguishable. As such, it samples uniformly from  $[1..s]$  a request  $s$ , and from a maximum size of searched labels  $[1..i]$  a document  $i$ . Whenever Update is required to perform an encryption,  $\mathcal{B}$  requests the ciphertext to the corresponding oracle in  $\text{IND-CCA}_{\Theta, \mathcal{B}}$ . Upon the  $s$ -th call to Search, and upon the  $i$ -th decrypted document,  $\mathcal{B}$  presents to the  $\text{IND-CCA}_{\Theta, \mathcal{B}}$  experiment the ciphertext  $m[i]$ . Observe that, if

$$\begin{aligned} (\text{op}, \text{id}) &\leftarrow \Gamma.\text{Get}(l', L'_p[i]) \\ (l, (\text{op}', \text{id}')) &\leftarrow \Theta.\text{Dec}(k_e, m[i]) \\ (\text{op}, \text{id}) &= (\text{op}', \text{id}') \end{aligned}$$

then  $\Pr[G_0^A(1^\lambda) \Rightarrow T] = \Pr[G_1^A(1^\lambda) \Rightarrow T]$ . It remains to show that, whenever  $(\text{op}, \text{id}) \neq (\text{op}', \text{id}')$ ,  $m[i]$  is a valid forgery.

To see this, observe that this is a valid ciphertext, as the decryption of  $m[i]$  is performed previously, and the result was not  $\perp$ . It suffices to establish that  $m[i]$  could not have been constructed by the encryption oracle of  $\text{IND-CCA}_{\Theta, \mathcal{B}}$ . From the construction of Update and the behaviour of  $\mathcal{B}$ , one can infer that the encryption oracle in  $\text{IND-CCA}_{\Theta, \mathcal{B}}$  is only called once for every  $l$ , as each  $w$  has a unique counter that is incremented on Update. That exact call  $(\text{op}, \text{id})$  is stored in  $l'$ . Since we know that  $l = L'_p[i]$ , and since we have the precondition of  $(\text{op}, \text{id}) \neq (\text{op}', \text{id}')$  for label  $l$ , then  $m[i]$  could not have been produced by the encryption oracle in  $\text{IND-CCA}_{\Theta, \mathcal{B}}$ , and is thus a forgery.

<p><b>Game <math>G_1^A(1^\lambda)</math>:</b>  <math>\text{prms} \leftarrow \text{IEE.Init}(1^\lambda)</math>  <math>k_f \leftarrow \text{F.Gen}(1^\lambda)</math>  <math>W \leftarrow \Gamma.\text{New}()</math>  <math>(k_c, t) \leftarrow \text{IEE.Setup}(1^\lambda, \text{prms})</math>  <math>k_e \leftarrow \Theta.\text{Gen}(1^\lambda)</math>  <math>l \leftarrow \Gamma.\text{ulnit}()</math>  <math>l' \leftarrow \Gamma.\text{New}()</math>  <math>\text{st} \leftarrow (k_c, k_e, k_f, 0, W, l')</math>  <math>\text{st}_A \leftarrow (t, l)</math>  <math>\text{Return } \mathcal{A}_1^{\text{Update, Search}}(1^\lambda, \text{st}_A)</math></p> <p><b>Oracle Update(st, id, w):</b>  <math>(k_c, k_e, k_f, \text{nDocs}, W, l') \leftarrow \text{st}</math>  <math>k_w \leftarrow \text{F.Run}(k_f, w)</math>  <math>c^* \leftarrow \text{IEE.Send}(k_c, (\text{op}, \text{id}, k_w))</math>  <b>If</b> <math>(c \leftarrow \Gamma.\text{Get}(W, w))</math>:  <math>c = c + 1</math>  <b>Else:</b> <math>c \leftarrow 0</math>  <math>W \leftarrow \Gamma.\text{Put}(W, w, c)</math>  <b>If</b> <math>\text{id} &gt; \text{nDocs}</math>: <math>\text{nDocs} \leftarrow \text{nDocs} + 1</math>  <math>l \leftarrow \text{F.Run}(k_w, c)</math>  <math>\text{id}^* \leftarrow \Theta.\text{Enc}(k_e, (l, (\text{op}, \text{id})))</math>  <math>(l, l_t) \leftarrow \Gamma.\text{uPut}(l, l, \text{id}^*)</math>  <math>l' \leftarrow \Gamma.\text{Put}(l', l, (\text{op}, \text{id}))</math>  <math>\text{st} \leftarrow (k_c, k_e, k_f, \text{nDocs}, W)</math>  <math>\text{st}_A \leftarrow \mathcal{A}_2(\text{st}_A, (c^*, l_t))</math>  <math>\text{Return } \text{st}_A</math></p>	<p><b>Oracle Search(st, q):</b>  <math>(k_c, k_e, k_f, \text{nDocs}, W, l') \leftarrow \text{st}</math>  <math>L \leftarrow []; n \leftarrow 0; D \leftarrow []</math>  <math>q^* \leftarrow \text{IEE.Send}(k_c, q)</math>  <math>(W_q, \phi) \leftarrow \text{PBQ}(q)</math>  <b>For</b> <math>w \in W_q</math>:  <math>k_w \leftarrow \text{F.Run}(k_f, w)</math>  <math>c \leftarrow \Gamma.\text{Get}(W, w)</math>  <b>For</b> <math>k \in [0 \dots c]</math>:  <math>l \leftarrow \text{F.Run}(k_w, k)</math>  <math>L \leftarrow l : L; n \leftarrow n + 1</math>  <math>(\Delta, L_p) \leftarrow \text{Sort}(L)</math>  <math>(\text{st}_A, m) \leftarrow \mathcal{A}_3(\text{st}_A, (L_p, q^*))</math>  <b>For</b> <math>i \in [0..n]</math>:  <math>(\text{op}, \text{id}) \leftarrow \Gamma.\text{Get}(l', L_p[i])</math>  <math>D \leftarrow (\text{op}, \text{id}) : D</math>  <math>(l, *) \leftarrow \Theta.\text{Dec}(k_e, m[i])</math>  <b>If</b> <math>L_p[i] \neq l</math>: <b>abort</b>  <math>D' \leftarrow \text{Reorder}(\Delta, D)</math>  <math>r \leftarrow \text{Resolve}(\phi, D', \text{nDocs})</math>  <math>r^* \leftarrow \text{IEE.Send}(k_c, r)</math>  <math>(\text{st}_A) \leftarrow \mathcal{A}_4(\text{st}_A, r^*)</math>  <math>\text{Return } (r, \text{st}_A)</math></p>
--	---

Fig. 13: Game 1.

In game  $G_2^A$ , the secure channel with the IEE is now handled by a simulator executing  $S_{\text{IEE}}$ . Intuitively, any adversary that is able to distinguish these two games can actively break the secrecy of the secure channel with the IEE. We are directly using the key exchange scheme proposed in [3], we can apply the same Utility theorem. Since we are only establishing a secure channel with a single party, this can be applied only once, and thus

$$|\Pr[G_1^A(1^\lambda) \Rightarrow T] - \Pr[G_2^A(1^\lambda) \Rightarrow T]| \leq \text{Adv}_{\text{AttKE}, \mathcal{A}}^{\text{UT}}(\lambda)$$

<p><b>Game <math>G_2^A(1^\lambda)</math>:</b></p> <p><math>(k_c, t) \leftarrow \mathcal{S}_{IEE}(1^\lambda)</math>  <math>k_f \leftarrow \mathcal{F}.Gen(1^\lambda)</math>  <math>W \leftarrow \Gamma.New()</math>  <math>k_e \leftarrow \mathcal{E}.Gen(1^\lambda)</math>  <math>l \leftarrow \Gamma.ulnit()</math>  <math>l' \leftarrow \Gamma.New()</math>  <math>st \leftarrow (k_c, k_e, k_f, 0, W, l')</math>  <math>st_{\mathcal{A}} \leftarrow (t, l)</math>  Return <math>\mathcal{A}_1^{Update, Search}(1^\lambda, st_{\mathcal{A}})</math></p> <p><b>Oracle Update(st, id, w):</b></p> <p><math>(k_c, k_e, k_f, nDocs, W, l') \leftarrow st</math>  <math>k_w \leftarrow \mathcal{F}.Run(k_f, w)</math>  <math>c^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{U_s})</math>  If <math>(c \leftarrow \Gamma.Get(W, w))</math>:  <math>c = c + 1</math>  Else: <math>c \leftarrow 0</math>  <math>W \leftarrow \Gamma.Put(W, w, c)</math>  If <math>id &gt; nDocs</math>: <math>nDocs \leftarrow nDocs + 1</math>  <math>l \leftarrow \mathcal{F}.Run(k_w, c)</math>  <math>id^* \leftarrow \mathcal{E}.Enc(k_e, (l, (op, id)))</math>  <math>(l, l_t) \leftarrow \Gamma.uPut(l, l, id^*)</math>  <math>l' \leftarrow \Gamma.Put(l', l, (op, id))</math>  <math>st \leftarrow (k_c, k_e, k_f, nDocs, W)</math>  <math>st_{\mathcal{A}} \leftarrow \mathcal{A}_2(st_{\mathcal{A}}, (c^*, l_t))</math>  Return <math>st_{\mathcal{A}}</math></p>	<p><b>Oracle Search(st, q):</b></p> <p><math>(k_c, k_e, k_f, nDocs, W, l') \leftarrow st</math>  <math>L \leftarrow []; n \leftarrow 0; D \leftarrow []</math>  <math>q^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{lq})</math>  <math>(W_q, \phi) \leftarrow PBQ(q)</math>  For <math>w \in W_q</math>:  <math>k_w \leftarrow \mathcal{F}.Run(k_f, w)</math>  <math>c \leftarrow \Gamma.Get(W, w)</math>  For <math>k \in [0 \dots c]</math>:  <math>l \leftarrow \mathcal{F}.Run(k_w, k)</math>  <math>L \leftarrow l : L; n \leftarrow n + 1</math>  <math>(\Delta, L_p) \leftarrow \text{Sort}(L)</math>  <math>(st_{\mathcal{A}}, m) \leftarrow \mathcal{A}_3(st_{\mathcal{A}}, (L_p, q^*))</math>  For <math>i \in [0..n]</math>:  <math>(op, id) \leftarrow \Gamma.Get(l', L_p[i])</math>  <math>D \leftarrow (op, id) : D</math>  <math>(l, \star) \leftarrow \mathcal{E}.Dec(k_e, m[i])</math>  If <math>L_p[i] \neq l</math>: abort  <math>D' \leftarrow \text{Reorder}(\Delta, D)</math>  <math>r \leftarrow \text{Resolve}(\phi, D', nDocs)</math>  <math>r^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{l r })</math>  <math>(st_{\mathcal{A}}) \leftarrow \mathcal{A}_4(st_{\mathcal{A}}, r^*)</math>  Return <math>(r, st_{\mathcal{A}})</math></p>
--	---

Fig. 14: Game 2.

In game  $G_3^A$ , the resolution of the query is no longer performed by the originally described Search. Instead, all requests for Update are stored in the ideal structure  $l'$ . Search becomes a full table scan for all entries of  $l'$  for which the identifiers are relevant, and Resolve executes upon that structure. This hop is derived from the correctness property of BISEN.

$$\text{Search}(K, \phi(\bar{w}), DB) = DB(\phi(\bar{w}))$$

which ensures that the output of Search according to the original BISEN description (left side), is exactly the same as that of simply executing the query on the clean database (right side). Since a plaintext database  $l'$  is managed on Update, and the set of  $\phi(D)$  is selected from  $l'$  according to  $(W_q, \phi) \leftarrow PBQ(q)$ , such that  $\forall w \in W_q \cap l' : w \in D$  it follows that

$$\Pr[G_2^A(1^\lambda) \Rightarrow T] = \Pr[G_3^A(1^\lambda) \Rightarrow T]$$

In game  $G_4^A$ , we replace the encryption of document identifiers by dummy messages of the same length. Observe that, since query resolution is being performed over  $l'$ , these identifiers are no longer necessary for Search. Let  $u$  be the number of calls to Oracle Update. Since  $\mathcal{A}$  does not have access to  $k_e$ , we upper bound the distance between these two games, by constructing an adversary  $\mathcal{C}$  against the IND-CCA security of  $\mathcal{E}$  such that

$$|\Pr[G_3^A(1^\lambda) \Rightarrow T] - \Pr[G_4^A(1^\lambda) \Rightarrow T]| \leq \frac{\text{Adv}_{\mathcal{E}, \mathcal{C}}^{\text{IND-CCA}}(\lambda)}{u}$$

Adversary  $\mathcal{C}$  simulates the environment of  $G_4^A$  as follows. At the beginning of the game,  $\mathcal{C}$  as to try and guess which call to Update will be distinguishable. As such, it samples uniformly from  $[1..u]$  a request  $u$ . Upon the  $u$ -th call to Update,  $\mathcal{C}$  presents to the  $\text{IND-CCA}_{\mathcal{E}, \mathcal{C}}$  experiment the message pair  $((op, id), \{0\}^{U_s})$  and proceeds  $G_4^A$  with the received ciphertext.  $\mathcal{C}$  presents the result of  $G_4^A$  as the guessing bit of  $\text{IND-CPA}_{\mathcal{E}, \mathcal{B}}$ . Given that the difference between the two games is exactly that of presenting either the encryption of  $(op, id)$  or  $\{0\}^{U_s}$ , the advantage of  $\mathcal{A}$  distinguishing between  $G_4^A$  and  $G_3^A$  is exactly that of breaking the IND-CCA security of  $\mathcal{E}$  for the  $u$ -th instance of Update.

In game  $G_5^A$ , instead of maintaining a counter  $c$  for each unique word, we maintain a global counter  $c'$ , as well as a structure for counting unique keywords  $W$ . This means that, for label generation, instead of running  $\mathcal{F}$  over  $w$ , and then  $\mathcal{F}$  over that specific  $c$ , we run  $\mathcal{F}$  over a unique  $c'$ . We want to show that

$$\Pr[G_4^A(1^\lambda) \Rightarrow T] = \Pr[G_5^A(1^\lambda) \Rightarrow T]$$

This has two major implications. It changes how labels are computed in Update, and changes how they are recovered in Search. We argue that this is indistinguishable for an  $\mathcal{A}$  without access to  $k_f$  by analysing each process individually.

<p><b>Game <math>G_3^A(1^\lambda)</math>:</b></p> <p><math>(k_c, t) \leftarrow \mathcal{S}_{IEE}(1^\lambda)</math>  <math>k_f \leftarrow \mathcal{F}.Gen(1^\lambda)</math>  <math>W \leftarrow \Gamma.New()</math>  <math>k_e \leftarrow \mathcal{O}.Gen(1^\lambda)</math>  <math>l \leftarrow \Gamma.ulnit()</math>  <math>l' \leftarrow \Gamma.New()</math>  <math>st \leftarrow (k_c, k_e, k_f, 0, 0, W, l')</math>  <math>st_A \leftarrow (t, l)</math>  Return <math>\mathcal{A}_1^{Update, Search}(1^\lambda, st_A)</math></p> <p><b>Oracle Update(st, id, w):</b></p> <p><math>(k_c, k_e, k_f, nDocs, c', W, l') \leftarrow st</math>  <math>k_w \leftarrow \mathcal{F}.Run(k_f, w)</math>  <math>c^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{U_s})</math>  If <math>(c \leftarrow \Gamma.Get(W, w))</math>:  <math>c = c + 1</math>  Else: <math>c \leftarrow 0</math>  <math>W \leftarrow \Gamma.Put(W, w, c)</math>  If <math>id &gt; nDocs</math>: <math>nDocs \leftarrow nDocs + 1</math>  <math>l \leftarrow \mathcal{F}.Run(k_w, c)</math>  <math>id^* \leftarrow \mathcal{O}.Enc(k_e, (l, (op, id)))</math>  <math>(l, l_t) \leftarrow \Gamma.uPut(l, l, id^*)</math>  <math>l' \leftarrow \Gamma.Put(l', c', (op, w, id, c'))</math>  <math>c' \leftarrow c' + 1</math>  <math>st \leftarrow (k_c, k_e, k_f, nDocs, W)</math>  <math>st_A \leftarrow \mathcal{A}_2(st_A, (c^*, l_t))</math>  Return <math>st_A</math></p>	<p><b>Oracle Search(st, q):</b></p> <p><math>(k_c, k_e, k_f, nDocs, c', W, l') \leftarrow st</math>  <math>L \leftarrow []; n \leftarrow 0; D \leftarrow []</math>  <math>q^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{lq})</math>  <math>(W_q, \phi) \leftarrow PBQ(q)</math>  For <math>(op, w, id, c') \in A \wedge w \in W_q</math>:  <math>D \leftarrow (id : D)</math>  For <math>w \in W_q</math>:  <math>k_w \leftarrow \mathcal{F}.Run(k_f, w)</math>  <math>c \leftarrow \Gamma.Get(W, w)</math>  For <math>k \in [0 \dots c]</math>:  <math>l \leftarrow \mathcal{F}.Run(k_w, k)</math>  <math>L \leftarrow l : L; n \leftarrow n + 1</math>  <math>(\Delta, L_p) \leftarrow \text{Sort}(L)</math>  <math>(st_A, m) \leftarrow \mathcal{A}_3(st_A, (L_p, q^*))</math>  For <math>i \in [0 \dots n]</math>:  <math>(l, \star) \leftarrow \mathcal{O}.Dec(k_e, m[i])</math>  If <math>L_p[i] \neq l</math>: abort  <math>r \leftarrow \text{Resolve}(\phi, D, nDocs)</math>  <math>r^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{l r })</math>  <math>(st_A) \leftarrow \mathcal{A}_4(st_A, r^*)</math>  Return <math>(r, st_A)</math></p>
--	--

Fig. 15: Game 3.

<p><b>Game <math>G_4^A(1^\lambda)</math>:</b></p> <p><math>(k_c, t) \leftarrow \mathcal{S}_{IEE}(1^\lambda)</math>  <math>k_f \leftarrow \mathcal{F}.Gen(1^\lambda)</math>  <math>W \leftarrow \Gamma.New()</math>  <math>k_e \leftarrow \mathcal{O}.Gen(1^\lambda)</math>  <math>l \leftarrow \Gamma.ulnit()</math>  <math>l' \leftarrow \Gamma.New()</math>  <math>st \leftarrow (k_c, k_e, k_f, 0, 0, W, l')</math>  <math>st_A \leftarrow (t, l)</math>  Return <math>\mathcal{A}_1^{Update, Search}(1^\lambda, st_A)</math></p> <p><b>Oracle Update(st, id, w):</b></p> <p><math>(k_c, k_e, k_f, nDocs, c', W, l') \leftarrow st</math>  <math>k_w \leftarrow \mathcal{F}.Run(k_f, w)</math>  <math>c^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{U_s})</math>  If <math>(c \leftarrow \Gamma.Get(W, w))</math>:  <math>c = c + 1</math>  Else: <math>c \leftarrow 0</math>  <math>W \leftarrow \Gamma.Put(W, w, c)</math>  If <math>id &gt; nDocs</math>: <math>nDocs \leftarrow nDocs + 1</math>  <math>l \leftarrow \mathcal{F}.Run(k_w, c)</math>  <math>id^* \leftarrow \mathcal{O}.Enc(k_e, (l, \{0\}^{ids+ops}))</math>  <math>(l, l_t) \leftarrow \Gamma.uPut(l, l, id^*)</math>  <math>l' \leftarrow \Gamma.Put(l', c', (op, w, id, c'))</math>  <math>c' \leftarrow c' + 1</math>  <math>st \leftarrow (k_c, k_e, k_f, nDocs, W)</math>  <math>st_A \leftarrow \mathcal{A}_2(st_A, (c^*, l_t))</math>  Return <math>st_A</math></p>	<p><b>Oracle Search(st, q):</b></p> <p><math>(k_c, k_e, k_f, nDocs, c', W, l') \leftarrow st</math>  <math>L \leftarrow []; n \leftarrow 0; D \leftarrow []</math>  <math>q^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{lq})</math>  <math>(W_q, \phi) \leftarrow PBQ(q)</math>  For <math>(op, w, id, c') \in A \wedge w \in W_q</math>:  <math>D \leftarrow (id : D)</math>  For <math>w \in W_q</math>:  <math>k_w \leftarrow \mathcal{F}.Run(k_f, w)</math>  <math>c \leftarrow \Gamma.Get(W, w)</math>  For <math>k \in [0 \dots c]</math>:  <math>l \leftarrow \mathcal{F}.Run(k_w, k)</math>  <math>L \leftarrow l : L; n \leftarrow n + 1</math>  <math>(\Delta, L_p) \leftarrow \text{Sort}(L)</math>  <math>(st_A, m) \leftarrow \mathcal{A}_3(st_A, (L_p, q^*))</math>  For <math>i \in [0 \dots n]</math>:  <math>(l, \star) \leftarrow \mathcal{O}.Dec(k_e, m[i])</math>  If <math>L_p[i] \neq l</math>: abort  <math>r \leftarrow \text{Resolve}(\phi, D, nDocs)</math>  <math>r^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{l r })</math>  <math>(st_A) \leftarrow \mathcal{A}_4(st_A, r^*)</math>  Return <math>(r, st_A)</math></p>
---	--

Fig. 16: Game 4.

- Update: Each label is now generated by running  $F$  on the unique insertion counter. Observe that this is only distinguishable from the alternative if it is possible to run  $F$  on duplicate  $(id, c)$  pairs. From the construction of Update on  $G_4^A$ , this is not the case. The structure  $W$  is also correctly updated whenever a new  $w$  is inserted.
- Search: We no longer execute the  $F$  for each unique word to determine  $|L|$ , but this can be computed by multiplying the fixed size of the output of  $F$  and size of counter  $c$  by the number of unique words in structure  $W$ :  $N * f_s$ . Furthermore, where  $G_4^A$  computed the labels for all  $(w, c)$  in  $W$  for which  $w \in W_q$ , the same can be achieved by computing all  $c'$  in  $A$  for which  $w \in W_q$ . This ensures consistency in the document identifiers retrieved in Search.

<p><b>Game <math>G_5^A(1^\lambda)</math>:</b>  <math>(k_c, t) \leftarrow \mathcal{S}_{IEE}(1^\lambda)</math>  <math>k_f \leftarrow \mathcal{F}.Gen(1^\lambda)</math>  <math>k_e \leftarrow \mathcal{E}.Gen(1^\lambda)</math>  <math>l \leftarrow \Gamma.ulnit()</math>  <math>l' \leftarrow \Gamma.New()</math>  <math>st \leftarrow (k_c, k_e, k_f, 0, 0, l', [])</math>  <math>st_A \leftarrow (t, l)</math>  Return <math>\mathcal{A}_1^{Update, Search}(1^\lambda, st_A)</math></p> <p><b>Oracle Update(st, id, w):</b>  <math>(k_c, k_e, k_f, nDocs, c', l', W) \leftarrow st</math>  <math>c^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{U_s})</math>  <b>If <math>w \notin W</math>: <math>W \leftarrow w : W</math></b>  <math>l' \leftarrow \Gamma.Put(l', c', (op, w, id, c'))</math>  <math>c' \leftarrow c' + 1</math>  <math>l \leftarrow \mathcal{F}.Run(k_f, c')</math>  <b>If <math>id &gt; nDocs</math>: <math>nDocs \leftarrow nDocs + 1</math></b>  <math>id^* \leftarrow \mathcal{E}.Enc(k_e, (l, \{0\}^{id_s + op_s}))</math>  <math>(l, l_t) \leftarrow \Gamma.uPut(l, l, id^*)</math>  <math>st \leftarrow (k_c, k_e, k_f, nDocs, W)</math>  <math>st_A \leftarrow \mathcal{A}_2(st_A, (c^*, l_t))</math>  Return <math>st_A</math></p>	<p><b>Oracle Search(st, q):</b>  <math>(k_c, k_e, k_f, nDocs, c', l', W) \leftarrow st</math>  <math>L \leftarrow []; n \leftarrow 0; D \leftarrow []</math>  <math>N \leftarrow 0</math>  <math>(W_q, \phi) \leftarrow PBQ(q)</math>  <b>For <math>w \in W_q</math>: <b>If <math>w \in W</math>: <math>N \leftarrow N + 1</math></b></b>  <math>qi_{size} \leftarrow  \phi  + N * f_s</math>  <math>q^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{qi_{size}})</math>  <b>For <math>(op, w, id, c') \in A \wedge w \in W_q</math>:</b>  <math>D \leftarrow (id : D)</math>  <math>l \leftarrow \mathcal{F}.Run(k_f, c')</math>  <math>L \leftarrow (l : L)</math>  <math>n \leftarrow n + 1</math>  <math>(\Delta, L_p) \leftarrow \text{Sort}(L)</math>  <math>(st_A, m) \leftarrow \mathcal{A}_3(st_A, (L_p, q^*))</math>  <b>For <math>i \in [0..n]</math>:</b>  <math>(l, \star) \leftarrow \mathcal{E}.Dec(k_e, m[i])</math>  <b>If <math>L_p[i] \neq l</math>: abort</b>  <math>r \leftarrow \text{Resolve}(\phi, D, nDocs)</math>  <math>r^* \leftarrow \mathcal{S}_{IEE}(k_c, \{0\}^{ r })</math>  <math>(st_A) \leftarrow \mathcal{A}_4(st_A, r^*)</math>  Return <math>(r, st_A)</math></p>
---	--

Fig. 17: Game 5.

In game  $G_6^A$ , we replace all calls to  $F$  by calls to a randomly generated function  $g$ . Since  $\mathcal{A}$  does not have access to  $k_f$ , we upper bound the distance between these two games, by constructing an adversary  $\mathcal{D}$  against the prf-security of  $F$  such that

$$|\Pr[G_5^A(1^\lambda) \Rightarrow T] - \Pr[G_6^A(1^\lambda) \Rightarrow T]| = \text{Adv}_{F, \mathcal{D}}^{\text{prf}}(\lambda)$$

Adversary  $\mathcal{D}$  simulates the environment of  $G_6^A$  as follows. Whenever a label has to be produced, the result is retrieved by calling the Oracle of  $\text{prf}_{F, \mathcal{D}}$ .  $\mathcal{D}$  presents the result of  $G_6^A$  as the guessing bit of  $\text{prf}_{F, \mathcal{D}}$ . Given that the difference between the two games is exactly that of presenting either the result of  $F(c)$  or  $g(c)$ , the advantage of  $\mathcal{A}$  distinguishing between  $G_6^A$  and  $G_5^A$  is exactly that of breaking the prf security of  $F$ .

Finally,  $G_7^A$  matches the ideal world of Figure 8, extended with the behaviour of the leakage function of Figure 10 and the simulator detailed in Figure 11. This final game is achieved by reorganizing the code of  $G_6^A$ , and thus

$$\Pr[G_6^A(1^\lambda) \Rightarrow T] = \Pr[G_7^A(1^\lambda) \Rightarrow T]$$

Let

$$\text{Adv}_{II, S, \mathcal{A}}^{\text{Att}}(\lambda) = |\Pr[\text{Real}_{II, \mathcal{A}}(1^\lambda) \Rightarrow T] - \Pr[\text{Ideal}_{S, \mathcal{A}}(1^\lambda) \Rightarrow T]|$$

To conclude, we have that

$$\begin{aligned} \text{Adv}_{II, S, \mathcal{A}}^{\text{Att}} &= \sum_{i=0}^7 |\Pr[G_i^A(1^\lambda) \Rightarrow T] - \Pr[G_{i+1}^A(1^\lambda) \Rightarrow T]| \\ &\leq \frac{\text{Adv}_{\mathcal{E}, \mathcal{B}}^{\text{uf}}(\lambda)}{s * i} + \text{Adv}_{\text{AttKE}, \mathcal{A}}^{\text{UT}}(\lambda) + \\ &\quad \frac{\text{Adv}_{\mathcal{E}, \mathcal{C}}^{\text{IND-CCA}}(\lambda)}{u} + \text{Adv}_{F, \mathcal{D}}^{\text{prf}}(\lambda) \\ &\leq \mu(\lambda) \end{aligned}$$

and Theorem 1 follows.

<p><b>Game <math>G_6^A(1^\lambda)</math>:</b></p> <p><math>(k_c, t) \leftarrow \mathcal{S}_{\text{IEE}}(1^\lambda)</math>  <math>g \leftarrow \text{Func}(D, R)</math>  <math>k_e \leftarrow \Theta.\text{Gen}(1^\lambda)</math>  <math>l \leftarrow \Gamma.\text{ulnit}()</math>  <math>l' \leftarrow \Gamma.\text{New}()</math>  <math>\text{st} \leftarrow (k_c, k_e, g, 0, 0, l', [])</math>  <math>\text{st}_A \leftarrow (t, l)</math>  Return <math>\mathcal{A}_1^{\text{Update, Search}}(1^\lambda, \text{st}_A)</math></p> <p><b>Oracle Update(st, id, w):</b></p> <p><math>(k_c, k_e, g, \text{nDocs}, c', l', W) \leftarrow \text{st}</math>  <math>c^* \leftarrow \mathcal{S}_{\text{IEE}}(k_c, \{0\}^{U_s})</math>  If <math>w \notin W</math>: <math>W \leftarrow w : W</math>  <math>l' \leftarrow \Gamma.\text{Put}(l', c', (\text{op}, w, \text{id}, c'))</math>  <math>c' \leftarrow c' + 1</math>  <math>l \leftarrow g(c')</math>  If <math>\text{id} &gt; \text{nDocs}</math>: <math>\text{nDocs} \leftarrow \text{nDocs} + 1</math>  <math>\text{id}^* \leftarrow \Theta.\text{Enc}(k_e, (l, \{0\}^{\text{id}_s + \text{ops}}))</math>  <math>(l, l_t) \leftarrow \Gamma.\text{uPut}(l, l, \text{id}^*)</math>  <math>\text{st} \leftarrow (k_c, k_e, k_f, \text{nDocs}, W)</math>  <math>\text{st}_A \leftarrow \mathcal{A}_2(\text{st}_A, (c^*, l_t))</math>  Return <math>\text{st}_A</math></p>	<p><b>Oracle Search(st, q):</b></p> <p><math>(k_c, k_e, g, \text{nDocs}, c', l', W) \leftarrow \text{st}</math>  <math>L \leftarrow []; n \leftarrow 0; D \leftarrow []</math>  <math>N \leftarrow 0</math>  <math>(W_q, \phi) \leftarrow \text{PBQ}(q)</math>  For <math>w \in W_q</math>: If <math>w \in W</math>: <math>N \leftarrow N + 1</math>  <math>\text{qi}_{\text{size}} \leftarrow  \phi  + N * f_s</math>  <math>q^* \leftarrow \mathcal{S}_{\text{IEE}}(k_c, \{0\}^{\text{qi}_{\text{size}}})</math>  For <math>(\text{op}, w, \text{id}, c') \in A \wedge w \in W_q</math>:  <math>D \leftarrow (\text{id} : D)</math>  <math>l \leftarrow g(c')</math>  <math>L \leftarrow (l : L)</math>  <math>n \leftarrow n + 1</math>  <math>(\Delta, L_p) \leftarrow \text{Sort}(L)</math>  <math>(\text{st}_A, m) \leftarrow \mathcal{A}_3(\text{st}_A, (L_p, q^*))</math>  For <math>i \in [0..n]</math>:  <math>(l, \star) \leftarrow \Theta.\text{Dec}(k_e, m[i])</math>  If <math>L_p[i] \neq l</math>: abort  <math>r \leftarrow \text{Resolve}(\phi, D, \text{nDocs})</math>  <math>r^* \leftarrow \mathcal{S}_{\text{IEE}}(k_c, \{0\}^{l^*})</math>  <math>(\text{st}_A) \leftarrow \mathcal{A}_4(\text{st}_A, r^*)</math>  Return <math>(r, \text{st}_A)</math></p>
---	--

Fig. 18: Game 6.

<p><b>Game <math>G_7^A(1^\lambda)</math>:</b></p> <p><math>\text{st}_C \leftarrow ([], \Gamma.\text{New}(), 0, 0)</math>  <math>g \leftarrow \text{Func}(D, R)</math>  <math>(k_c, t) \leftarrow \mathcal{S}_{\text{IEE}}(1^\lambda)</math>  <math>k_e \leftarrow \Theta.\text{Gen}(1^\lambda)</math>  <math>l \leftarrow \Gamma.\text{ulnit}()</math>  <math>\text{st}_S \leftarrow (k_c, k_e, g, 0)</math>  <math>\text{st}_A \leftarrow (t, l)</math>  Return <math>\mathcal{A}_1^{\text{Update, Search}}(1^\lambda, \text{st}_A)</math></p> <p><b>Oracle Update(st, id, w):</b></p> <p><math>(W, l', c', \text{nDocs}) \leftarrow \text{st}_C</math>  <math>c^* \leftarrow \mathcal{S}_{\text{IEE}}(k_c, \{0\}^{U_s})</math>  If <math>w \notin W</math>: <math>W \leftarrow w : W</math>  <math>l' \leftarrow \Gamma.\text{Put}(l', c', (\text{op}, w, \text{id}, c'))</math>  If <math>\text{id} &gt; \text{nDocs}</math>: <math>\text{nDocs} \leftarrow \text{nDocs} + 1</math>  <math>\text{st}_C \leftarrow (W, l', c' + 1, \text{nDocs})</math>  <math>(k_c, k_e, g, c) \leftarrow \text{st}_S</math>  <math>l \leftarrow g(c)</math>  <math>\text{id}^* \leftarrow \Theta.\text{Enc}(k_e, (l, \{0\}^{\text{id}_s + \text{ops}}))</math>  <math>(l, l_t) \leftarrow \Gamma.\text{uPut}(l, l, \text{id}^*)</math>  <math>\text{st}_S \leftarrow (k_c, k_e, g, c + 1)</math>  <math>\text{st}_A \leftarrow \mathcal{A}_2(\text{st}_A, (c^*, l_t))</math>  Return <math>\text{st}_A</math></p>	<p><b>Oracle Search(st, q):</b></p> <p><math>(W, l', c', \text{nDocs}) \leftarrow \text{st}_C</math>  <math>C \leftarrow []; D \leftarrow []; n \leftarrow 0</math>  <math>(W_q, \phi) \leftarrow \text{PBQ}(q)</math>  For <math>w \in W_q</math>: If <math>w \in W</math>: <math>N \leftarrow N + 1</math>  <math>\text{qi}_{\text{size}} \leftarrow  \phi  + N * f_s</math>  <math>\text{qi} \leftarrow \{0\}^{\text{qi}_{\text{size}}}</math>  <math>q^* \leftarrow \mathcal{S}_{\text{IEE}}(k_c, \text{qi})</math>  For <math>(\text{op}, w, \text{id}, c') \in A \wedge w \in W_q</math>:  <math>C \leftarrow (c' : C); D \leftarrow (\text{id} : D)</math>  <math>(\cdot, C_p) \leftarrow \text{Sort}(C)</math>  <math>r \leftarrow \text{Resolve}(\phi, D, \text{nDocs})</math>  <math>\text{qo}_{\text{size}} \leftarrow  r </math>  <math>\text{st}_C \leftarrow (W, l', c', \text{nDocs})</math>  <math>(k_c, k_e, g, c) \leftarrow \text{st}_S</math>  <math>L \leftarrow []; n' \leftarrow 0</math>  For <math>k \in C_p</math>:  <math>l \leftarrow g(k)</math>  <math>L \leftarrow l : L; n' \leftarrow n' + 1</math>  <math>(\text{st}_A, m) \leftarrow \mathcal{A}_3(\text{st}_A, (L, q^*))</math>  <math>\text{qo} \leftarrow \{0\}^{\text{qo}_{\text{size}}}</math>  <math>r^* \leftarrow \mathcal{S}_{\text{IEE}}(k_c, \text{qo})</math>  For <math>i \in [0..n']</math>:  <math>(l, \star) \leftarrow \Theta.\text{Dec}(k_e, m[i])</math>  If <math>L[i] \neq l</math>: abort  <math>\text{st}_S \leftarrow (k_c, k_e, g, c)</math>  <math>(\text{st}_A) \leftarrow \mathcal{A}_4(\text{st}_A, r^*)</math>  Return <math>(r, \text{st}_A)</math></p>
---	---

Fig. 19: Extended ideal world.