

Boolean Searchable Symmetric Encryption with Worst-Case Sub-Linear Complexity

Seny Kamara* Tarik Moataz**
Brown University Brown University

Abstract. Recent work on searchable symmetric encryption (SSE) has focused on increasing its expressiveness. A notable example is the OXT construction (Cash et al., *CRYPTO '13*) which is the first SSE scheme to support conjunctive keyword queries with sub-linear search complexity. While OXT efficiently supports disjunctive and boolean queries that can be expressed in searchable normal form, it can only handle *arbitrary* disjunctive and boolean queries in linear time. This motivates the problem of designing expressive SSE schemes with *worst-case* sub-linear search; that is, schemes that remain highly efficient for *any* keyword query.

In this work, we address this problem and propose non-interactive highly efficient SSE schemes that handle *arbitrary* disjunctive and boolean queries with worst-case sub-linear search and optimal communication complexity. Our main construction, called IEX, makes black-box use of an underlying single keyword SSE scheme which we can instantiate in various ways. Our first instantiation, IEX-2Lev, makes use of the recent 2Lev construction (Cash et al., *NDSS '14*) and is optimized for search at the expense of storage overhead. Our second instantiation, IEX-ZMF, relies on a new single keyword SSE scheme we introduce called ZMF and is optimized for storage overhead at the expense of efficiency (while still achieving asymptotically sub-linear search). Our ZMF construction is the first adaptively-secure highly compact SSE scheme and may be of independent interest. At a very high level, it can be viewed as an encrypted version of a new Bloom filter variant we refer to as a Matryoshka filter. In addition, we show how to extend IEX to be dynamic and forward-secure. To evaluate the practicality of our schemes, we designed and implemented a new encrypted search framework called *Clusion*. Our experimental results demonstrate the practicality of IEX and of its instantiations with respect to either search (for IEX-2Lev) and storage overhead (for IEX-ZMF).

*seny@brown.edu. Work done in part at Microsoft Research.

**tarik_moataz@brown.edu. Work done in part at IMT Atlantique and Colorado State.

1 Introduction

A structured encryption (STE) scheme encrypts a data structure in such a way that it can be privately queried. An STE scheme is secure if it reveals nothing about the structure and query beyond a well-specified and “reasonable” leakage profile [15,13]. STE schemes come in two forms: response-revealing and response-hiding. The former reveals the query response in plaintext whereas the latter does not. An important special case of STE is searchable symmetric encryption (SSE) which encrypts search structures such as inverted indexes [15,13,24,23,11,10] or search trees [19,23]. Another example is graph encryption which encrypts various kinds of graphs [13,27]. STE has received a lot of attention from Academia and Industry due to: (1) its potential applications to cloud storage and database security; and (2) the fact that, among a host of different encrypted search solutions (e.g., property-preserving encryption, fully-homomorphic encryption, oblivious RAM, functional encryption) it seems to provide the best tradeoffs between security and efficiency.

In recent years, much of the work on STE has focused on supporting more complex structures and queries. A notable example in the setting of SSE is the work of Cash et al. which proposed the first SSE scheme to support conjunctive queries in sub-linear time [11]. Their scheme, OXT, is also shown to support disjunctive and even boolean queries. Faber et al. later showed how to extend OXT to achieve even more complex queries including range, substring, wildcard and phrase queries. Another example is the BlindSeer project from Pappas et al. [30] and Fisch et al. [17] which present a solution that supports boolean and range queries as well as stemming in sub-linear time.

Naive solutions. Any boolean query $\phi(w_1, \dots, w_q)$, where w_1, \dots, w_q are keywords and ϕ is a boolean formula, can be handled using a single-keyword SSE scheme in a naive way. In the case of response-revealing schemes it suffices to search for each keyword and have the server take the intersection and unions of the result sets appropriately. The issue with this approach, of course, is that the server learns more information than necessary: namely, it learns the result sets $\text{DB}(w_1), \dots, \text{DB}(w_q)$ whereas it should only learn the set $\text{DB}(\phi(w_1, \dots, w_q))$. For response-hiding schemes, one can search for each keyword and compute the intersections and unions at the client. The problem with this approach is that the parties communicate more information than necessary: namely, the server sends elements within the intersections of the result sets multiple times. With this in mind, any boolean SSE solution should improve on one of the naive approaches depending on whether it is response-hiding or response-revealing.

Worst-case sub-linear search complexity. While OXT achieves sub-linear search complexity for conjunctive queries, its extension to disjunctive and arbitrary boolean queries does not. More precisely, OXT remains sub-linear only for queries in *searchable normal form* (SNF) which have the form $w_1 \wedge \phi(w_2, \dots, w_q)$, where w_1 through w_q are keywords and ϕ is an arbitrary boolean formula. For non-SNF queries, OXT requires linear time in the number of documents. This motivates the following natural question: *can we design SSE schemes that support arbitrary disjunctive and arbitrary boolean queries with sub-linear search*

complexity? In other words, can we design solutions for these queries that are efficient even in the worst-case?

1.1 Our Contributions and Techniques

In this work, we address this problem and propose efficient disjunctive and boolean SSE schemes with worst-case sub-linear search complexity and optimal communication overhead. Our schemes are non-interactive and, as far as we know, the first to achieve optimal communication complexity. To do this we make several contributions which we summarize below

Worst-case disjunctive search. Our first solution, which we call IEX, is a worst-case sub-linear disjunctive SSE scheme. While it leaks more than the naive response-hiding solution, we stress that it achieves *optimal* communication complexity which, for response-hiding schemes, is the main tradeoff we seek. In addition, it leaks *less* than OXT (when used for disjunctive queries) while achieving worst-case efficiency.

The underlying idea behind IEX’s design is best expressed in set-theoretic terms where we view the result of a disjunctive query $w_1 \vee \dots \vee w_q$ as the union of the results of each individual term. More precisely, if we denote by $\text{DB}(w)$ the set of document identifiers that contain the query w , then $\text{DB}(w_1 \vee \dots \vee w_q) = \text{DB}(w_1) \cup \dots \cup \text{DB}(w_q)$. Using the naive response-hiding approach, one could use a single-keyword response-hiding scheme to query each keyword and compute the union at the client but, as discussed above, this would incur poor communication complexity. Our approach is different and, intuitively speaking, makes use of the *inclusion-exclusion* principle as follows. Consider a three-term query $w_1 \vee w_2 \vee w_3$. Instead of searching for $\text{DB}(w_1)$, $\text{DB}(w_2)$, $\text{DB}(w_3)$ and computing the union, we compute $\text{DB}(w_1)$ and remove from it

$$\text{DB}(w_1) \cap \text{DB}(w_2) \quad \text{and} \quad \text{DB}(w_1) \cap \text{DB}(w_3).$$

We then compute $\text{DB}(w_2)$ and remove from it $\text{DB}(w_2) \cap \text{DB}(w_3)$. Finally, we take the union of the remaining sets and add $\text{DB}(w_3)$. It follows by the inclusion-exclusion principle that this results in exactly $\text{DB}(w_1) \cup \text{DB}(w_2) \cup \text{DB}(w_3)$. If we could somehow support the intersection and removal operations at the server, then we could achieve optimal communication complexity. Note that this high-level approach is “purely disjunctive” in the sense that it does not rely on transforming the query into another form as done in OXT. The avoidance of SNF in particular is what enables us to achieve worst-case efficiency.

We stress that the intuition provided thus far is only a very high-level conceptual explanation of our approach and cannot be translated directly to work on encrypted data. The challenge is that no SSE scheme we are aware of directly supports the kind of set operations needed to implement this idea. Therefore, a major part of our contribution is in designing and analyzing such a scheme.

Boolean search. While IEX is naturally disjunctive, we show that it also supports boolean queries. Similarly to the disjunctive case, we explain our high-level approach in set-theoretic terms. First, recall that any boolean query can be written in conjunctive normal form (CNF) so it has the form $\Delta_1 \wedge \dots \wedge \Delta_\ell$,

where each $\Delta_i = w_{i,1} \vee \dots \vee w_{i,q}$ is a disjunction. Given a response-hiding *disjunctive*-search scheme like IEX, a naive approach for CNF queries is to execute disjunctive searches for each disjunction $\Delta_1, \dots, \Delta_\ell$ and have the client perform the intersection of the results. This approach is problematic, however, because it requires more communication than necessary. To avoid this we take the following alternative approach. We note that the result $\text{DB}(\Delta_1 \wedge \dots \wedge \Delta_\ell)$ is a subset of $\text{DB}(\Delta_1)$ and that it can be computed by progressively keeping only the identifiers in $\text{DB}(\Delta_1)$ that are also included in $\text{DB}(\Delta_2)$ through $\text{DB}(\Delta_\ell)$. Again, we stress that this description is only a high-level conceptual explanation of our approach and requires more work to instantiate over encrypted data.

The IEX structure. As mentioned above, a major challenge in this work is the design of an encrypted structure that supports the set-theoretic operations needed to implement the strategies discussed above. To achieve this, IEX makes use of a more complex structure than the traditional encrypted inverted index. In particular, IEX combines several instantiations of two kinds of structures: dictionaries and multi-maps. A dictionary (i.e., a key-value store) maps labels to values whereas a multi-map (i.e., an inverted index) maps labels to tuples of values. More precisely, the IEX design consists of an encrypted *global* multi-map that maps every keyword w to its document identifiers $\text{DB}(w)$ and an encrypted dictionary that maps every keyword to a *local* multi-map for w . The local multi-map of a keyword w maps all the keywords v that co-occur with w to the identifiers of the documents that contain both v and w . At a high-level, with the encrypted global multi-map we can recover $\text{DB}(w_1)$. With the encrypted dictionary, we can recover the encrypted local multi-map for keywords w_2 through w_ℓ . And, finally, by querying the (encrypted) local multi-map of a keyword w with a keyword v , we can recover the identifiers of the documents that contain both w and v . With these basic operations, we can then execute a full disjunctive query as discussed above.

Instantiations. IEX is an abstract construction that makes black-box use of encrypted multi-maps and dictionaries which, in turn, can be instantiated with several concrete constructions, e.g., [15,13,23,10].¹ While its asymptotic complexity is not affected by how the building blocks are instantiated, its concrete efficiency is so we consider this choice carefully—especially how the local multi-maps are instantiated. We consider two instantiations. The first, IEX-2Lev, uses the 2Lev construction of Cash et al. [10] to encrypt the multi-maps (local and global). This particular instantiation is very efficient with respect to search time but produces large encrypted structures (e.g., 9.8GB for datasets of 34M keyword/id pairs).

To address this we propose a second instantiation called IEX-ZMF which trades off efficiency for compactness. In fact, we show that IEX-ZMF is an order of magnitude more compact than IEX-2Lev (e.g., producing 0.9GB EDBs for datasets with 34M keyword/id pairs). This compactness is achieved by encrypting

¹ Other constructions such as [24,11,28,32] could also be used but these are either dynamic or conjunctive which is not needed for the IEX.

IEX’s local multi-maps with a new construction called ZMF which may be of independent interest and that we detail below.²

The ZMF scheme. ZMF is a multi-map encryption scheme that is inspired by and has similarities to the classic Z-IDX construction of Goh [19]. Its core design as well as its security are very different, however. While Z-IDX produces a collection of non-adaptively-secure *fixed*-size encrypted Bloom filters, ZMF produces a collection of *adaptively*-secure *variable*-sized encrypted Bloom filters. In addition, the hash functions used for each filter can all be derived from a fixed set of hash functions (even though the filters store a different number of elements). This last property is non-standard but is crucial for our approach to be practical as it allows us to generate constant-size tokens that can be used with every filter in the collection. We refer to such collections of Bloom filters as *matryoshka filters* and, as far as we know, they have not been considered in the past. As we detail in Section 7, encrypting matryoshka filters with adaptive security is quite challenging. For this, we rely on the random oracle model and on a non-standard use of online ciphers [4] which are streaming block ciphers in the sense that every ciphertext block depends only on the previous plaintext blocks. Note that like Z-IDX, ZMF has linear search time but we use it in our IEX construction only to encrypt the *local* multi-maps which guarantees that IEX-ZMF is still sub-linear.

Dynamism and forward-security. We extend IEX to be dynamic resulting in a new scheme DIEX. An important security property for dynamic SSE schemes is *forward security* which guarantees that updates to an encrypted structure cannot be correlated with previous queries. Forward security was introduced by Stefanov, Papamanthou and Shi [32] and recent work of Zhang, Katz and Papamanthou [34] has shown that it mitigates certain injection attacks on SSE schemes. One advantage of our DIEX construction is that it naturally inherits the forward-security of its underlying encrypted multi-maps and dictionaries. That is, if the underlying structures are forward-secure then so is DIEX.

Reduced leakage. As we mentioned above, IEX leaks more than the naive response-hiding solution *while achieving optimal communication complexity*. We stress, however, that it leaks less than the naive response-revealing solution and than OXT. As an example, consider that if OXT is used to search for two conjunctions $\mathbf{w} = w_1 \wedge w_2$ and $\mathbf{w}' = w_3 \wedge w_2$ which share a common term, the server can recover the results for $\mathbf{w}'' = w_1 \wedge w_2 \wedge w_3$. In the case of disjunctions, OXT’s leakage is equivalent to the naive response-revealing solution.

Experiments. To evaluate the efficiency of IEX and its instantiations we designed and built a new encrypted search framework called *Clusion* [22]. It is written in Java and leverages the Apache Lucene search library [1]. It also includes a Hadoop-based distributed parser and indexer we implemented to handle massive datasets. Our experiments show that IEX—specifically our IEX-2Lev instantiation—is very efficient and even achieves faster search times than those

² Multi-map encryption schemes are equivalent to SSE schemes so ZMF is an *adaptively*-secure compact SSE scheme with linear-time search.

reported for a C++ implementation of OXT [11] on a comparable system. For example, for conjunctive, disjunctive and boolean queries with selectivity on the order of thousands, IEX-2Lev takes 12, 14.8 and 23.7ms, respectively. For the same conjunctive query, OXT is reported to take 200ms on a comparable system. Clearly, a C/C++ implementation of IEX would perform even better.

We also implemented IEX-ZMF to evaluate its efficiency and compactness. In our experiments, it produced EDBs of size 198MB and 0.9GB from datasets with 1.5M and 34M keyword/id pairs, respectively. This is highly compact in comparison to IEX-2Lev which produced 1.6GB, 9.8GB EDBs for 1.5M and 34M keyword/id pairs, respectively. We also evaluated the efficiency of IEX-ZMF and, as expected, its performance for setup, search and token size are worse than IEX-2Lev. For example, for a dataset with 34M keyword/id pairs, EDB setup takes 7.58 hours to process compared to 31 minutes for IEX-2Lev.

On a boolean query of the form $(w \vee x) \wedge (y \vee z)$, where the disjunctions had selectivity 2K and 10K, respectively, IEX-ZMF took 1610ms whereas IEX-2Lev took only 23.7ms. As expected due to its high degree of compactness, IEX-ZMF is slower than IEX-2Lev (this is the exact tradeoff we seek).

2 Related Work

SSE was first considered by Song, Wagner and Perrig [31]. Curtmola, Garay, Kamara and Ostrovsky [15] introduced the notion of adaptive-security for SSE and presented the first constructions that achieved optimal search time with a space-efficient index. STE was introduced by Chase and Kamara [13] who proposed constructions for two-dimensional arrays, graphs and web graphs.

In [19], Goh introduced the Z-IDX construction which has linear search complexity and produces highly compact indexes due to its use of Bloom filters. Here, we extract a general transformation implicitly used in the Z-IDX construction and use it in part to construct our ZMF scheme. Kamara, Papamanthou and Roeder gave the first optimal-time dynamic SSE scheme [24]. Cash et al. [11] proposed OXT; the first optimal-time conjunctive keyword search scheme. Faber et al. [16] extend OXT to handle range, substring, wildcard and phrase queries. Pappas et al. [30] and Fisch et al. [17] present solutions based on garbled circuits and Bloom filters that can support boolean formulas, ranges and stemming. In [30], the authors show how to build the first worst-case sub-linear time boolean encrypted search solution. Like Goh’s Z-IDX construction and our ZMF scheme, the solution makes use of Bloom filters. In addition, it is the first adaptively-secure construction based on Bloom filters. For a disjunctive query \mathbf{w} , the scheme has search complexity $O(\log(n) \cdot C \cdot \text{DB}(\mathbf{w}))$, where n is the number of documents and C is the cost of a 2-party secure function evaluation of a function that takes as input a Bloom filter of size $O(\#W)$ (i.e., the number of unique keywords in DB) and a q -term disjunctive query. We note that unlike IEX and OXT, it does not achieve optimal communication complexity. Also, while its search is sub-linear it involves multiple rounds of interactions.

Ishai, Kushilevitz, Lu and Ostrovsky propose a two-server SSE scheme that hides the access pattern and supports various complex queries including ranges,

stemming and substring [21]. Cash et al. [10] design several I/O-efficient SSE schemes including the 2Lev construction which we use in one of our IEX instantiations. Kurosawa and Ohtaki [26] designed the first UC secure SSE scheme. Kurosawa [25] designed a linear-time construction that handles arbitrary boolean queries while not disclosing the structure of the boolean query itself. Forward Secrecy was first considered by Stefanov, Papamanthou and Shi [32]. In [9], Bost introduced an efficient forward secure construction. In [12], Cash and Tessaro give lower bounds on the locality of SSE by showing tradeoffs between locality, space overhead and read efficiency. Recently, Asharov, Naor, Segev and Shahaf gave SSE constructions with optimal locality, optimal space overhead and nearly-optimal read efficiency [3]. Encrypted search can also be achieved with other primitives like property-preserving encryption [5,6], functional encryption [7,8,29], oblivious RAM [20], full-homomorphic encryption [18] and multi-party computation [33].

Online ciphers were introduced by Bellare, Boldyreva, Knudsen and Nam-prepre [4], where they propose several schemes including the HCB1 construction which we make use of in our ZMF implementation. More efficient constructions were later proposed by Andreeva et al. [2].

3 Preliminaries

Notation. The set of all binary strings of length n is denoted as $\{0, 1\}^n$, and the set of all finite binary strings as $\{0, 1\}^*$. $[n]$ is the set of integers $\{1, \dots, n\}$, and $2^{[n]}$ is the corresponding power set. We write $x \leftarrow \chi$ to represent an element x being sampled from a distribution χ , and $x \stackrel{\$}{\leftarrow} X$ to represent an element x being sampled uniformly at random from a set X . The output x of an algorithm \mathcal{A} is denoted by $x \leftarrow \mathcal{A}$. Given a sequence \mathbf{v} of n elements, we refer to its i th element as v_i or $\mathbf{v}[i]$. If S is a set then $\#S$ refers to its cardinality. If s is a string then $|s|$ refers to its bit length and s_i to its i th bit. s^n denotes the string s padded with $n - |s|$ 0's and $s|_n$ represents the first n bits of s . Given strings s and r , we refer to their concatenation as either $\langle s, r \rangle$ or $s||r$. For an n -bit string s and for all nonnegative d , we denote by $s^{\parallel d}$ the string $\langle s_1^d, \dots, s_n^d \rangle$. In this work, padding takes precedence over truncation; that is, $s^{\parallel d}_p = (s^{\parallel d})|_p$.

Data types. An *abstract data type* is a collection of objects together with a set of operations defined on those objects. Examples include sets, dictionaries (also known as key-value stores or associative arrays) and graphs. The operations associated with an abstract data type fall into one of two categories: query operations, which return information about the objects; and update operations, which modify the objects. If the abstract data type supports only query operations it is *static*, otherwise it is *dynamic*.

Data structures. A *data structure* for a given data type is a representation in some computational model³ of an object of the given type. Typically, the representation is optimized to support the type's query operation as efficiently as possible. For data types that support multiple queries, the representation is often optimized to efficiently support as many queries as possible. As a concrete

³ In this work, the underlying model will always be the word RAM.

example, the dictionary *type* can be represented using various data structures depending on which queries one wants to support efficiently. Hash tables support **Get** and **Put** in expected $O(1)$ time whereas balanced binary search trees support both operations in worst-case $\log(n)$ time. For ease of understanding and to match colloquial usage, we will sometimes blur the distinction between data types and structures. So, for example, when referring to a *dictionary structure* or a *multi-map structure* what we are referring to is an unspecified instantiation of the dictionary or multi-map data type.

Basic structures. We make use of several basic data types including arrays, dictionaries and multi-maps which we recall here. An array A of capacity n stores n items at locations 1 through n and supports read and write operations. We write $v = A[i]$ to denote reading the item at location i and $A[i] = v$ the operation of storing an item at location i . A dictionary DX of capacity n is a collection of n label/value pairs $\{(\ell_i, v_i)\}_{i \leq n}$ and supports **Get** and **Put** operations. We write $v_i = DX[\ell_i]$ to denote getting the value associated with label ℓ_i and $DX[\ell_i] = v_i$ to denote the operation of associating the value v_i in DX with label ℓ_i . A multi-map MM with capacity n is a collection of n label/tuple pairs $\{(\ell_i, V_i)\}_{i \leq n}$ that supports **Get** and **Put** operations. Similarly to dictionaries, we write $V_i = MM[\ell_i]$ to denote getting the tuple associated with label ℓ_i and $MM[\ell_i] = V_i$ to denote operation of associating the tuple V_i to label ℓ_i . We sometimes write $MM^{-1}[v]$ to refer to the set of labels in MM associated with tuples that include the value v . Multi-maps are the abstract data type instantiated by an inverted index. In the encrypted search literature multi-maps are sometimes referred to as indexes, databases or tuple-sets (T-sets) [11,10].

Document collections. A document collection is a set of documents $\mathbf{D} = (D_1, \dots, D_n)$, each document consisting of a set of keywords from some universe W . We assume the universe of keywords is totally ordered (e.g., using lexicographic order) and denote by $W[i]$ the i th keyword in W . We assume every document has an identifier that is independent of its contents and denote it $\text{id}(D_i)$. We assume the existence of an efficient indexing algorithm that takes as input a data collection \mathbf{D} and outputs a multi-map that maps every keyword w in W to the identifiers of the documents that contain w . In previous work, this multi-map is referred to as an inverted index or as a database. For consistency, we refer to any multi-map derived in this way from a document collection as a database and denote it DB . Given a keyword w , we denote by $\text{co}_{DB}(w) \subseteq W$ the set of keywords in W that co-occur with w ; that is, the keywords that are contained in documents that contain w . When DB is clear from the context we omit DB and write only $\text{co}(w)$.

3.1 Cryptographic Primitives

Basic cryptographic primitives. A private-key encryption scheme is a set of three polynomial-time algorithms $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ such that **Gen** is a probabilistic algorithm that takes a security parameter k and returns a secret key K ; **Enc** is a probabilistic algorithm takes a key K and a message m and returns a ciphertext c ; **Dec** is a deterministic algorithm that takes a key K and a ciphertext

c and returns m if K was the key under which c was produced. Informally, a private-key encryption scheme is secure against chosen-plaintext attacks (CPA) if the ciphertexts it outputs do not reveal any partial information about the plaintext even to an adversary that can adaptively query an encryption oracle. We say a scheme is random-ciphertext-secure against chosen-plaintext attacks (RCPA) if the ciphertexts it outputs are computationally indistinguishable from random even to an adversary that can adaptively query an encryption oracle.⁴ In addition to encryption schemes, we also make use of pseudo-random functions (PRF) and permutations (PRP), which are polynomial-time computable functions that cannot be distinguished from random functions by any probabilistic polynomial-time adversary.

Online ciphers. An online cipher (OC) is a block cipher that can encrypt data streams. In particular, with an OC the encryption of the i th block in a stream depends only on the 1st through i th message blocks. OCs were introduced by Bellare, Boldyreva, Knudsen and Namprempre [4]. More formally, we say that a cipher $\text{OC} : \{0, 1\}^k \times \{0, 1\}^{n \times B} \rightarrow \{0, 1\}^{n \times B}$, where $B > 1$ is the block length, is B -online if there exists a function $X : \{0, 1\}^k \times \{0, 1\}^{n \times B} \rightarrow \{0, 1\}^B$ such that for any $\mathbf{m} \in \{0, 1\}^{n \times B}$,

$$\text{OC}_K(\mathbf{m}) = \text{OC}_K^1(\mathbf{m}) \parallel \dots \parallel \text{OC}_K^n(\mathbf{m}),$$

where $\text{OC}_K^i(\mathbf{m}) = X(K, m_1, \dots, m_i)$ for all $i \in [n]$ and where m_i is the i th block of \mathbf{m} . OCs cannot be pseudo-random permutations (see [4] for a simple distinguisher) but can satisfy the weaker requirement of being computationally indistinguishable from a random *online* permutation. An online permutation is simply a permutation on a domain $\{0, 1\}^{n \times B}$ whose i th block depends only on the first i blocks of its input. We denote by $\text{OPerm}_{n,B}$ the set of all online permutations over $\{0, 1\}^{n \times B}$. Security for an online cipher $\text{OC} : \{0, 1\}^k \times \{0, 1\}^{n \times B} \rightarrow \{0, 1\}^{n \times B}$ then holds if for all PPT adversaries \mathcal{A} ,

$$\left| \Pr \left[\mathcal{A}^{\text{OC}_K(\cdot)} = 1 : K \xleftarrow{\$} \{0, 1\}^k \right] - \Pr \left[\mathcal{A}^{f(\cdot)} = 1 : f \xleftarrow{\$} \text{OPerm}_{n,B} \right] \right| \leq \text{negl}(k).$$

4 Definitions

Structured encryption schemes encrypt data structures in such a way that they can be privately queried. There are several natural forms of structured encryption. The original definition of [13] considered schemes that encrypt both a structure and a set of associated data items (e.g., documents, emails, user profiles etc.). In [14], the authors also describe *structure-only* schemes which only encrypt structures. Another distinction can be made between *interactive* and *non-interactive* schemes. Interactive schemes produce encrypted structures that are queried through an interactive two-party protocol, whereas non-interactive schemes produce structures that can be queried by sending a single message,

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

i.e, the token. One can also distinguish between *response-hiding* and *response-revealing* schemes: the former reveal the response to queries whereas the latter do not.

STE schemes are used as follows. During a setup phase, the client constructs an encrypted data structure EDS under a key K . The client then sends EDS to the server. During the query phase, the client constructs and sends a token tk generated from its query q and the key K . The server then uses the token tk to query EDS. If the scheme is response-revealing, it recovers a response r . On the other hand, if the scheme is response-hiding it recovers a message that it returns to the client who in turn decrypts it with a resolving algorithm.

Definition 1 (Structured encryption). *A single-round response-hiding structured encryption scheme $\Sigma_{\mathcal{T}} = (\text{Setup}, \text{Token}, \text{Query}, \text{Resolve})$ for data type \mathcal{T} consists of four polynomial-time algorithms that work as follows:*

- $(K, \text{EDS}) \leftarrow \text{Setup}(1^k, \text{DS})$: *is a probabilistic algorithm that takes as input a security parameter 1^k and a structure DS of type \mathcal{T} and outputs a secret key K and an encrypted structure EDS.*
- $\text{tk} \leftarrow \text{Token}(K, q)$: *is a (possibly) probabilistic algorithm that takes as input a secret key K and a query q and returns a token tk .*
- $c \leftarrow \text{Query}(\text{EDS}, \text{tk})$: *is a (possibly) probabilistic algorithm that takes as input an encrypted structure EDS and a token tk and outputs a message c .*
- $r \leftarrow \text{Resolve}(K, c)$: *is a deterministic algorithm that takes as input a secret key K and a message c and outputs a response r .*

We say that a structured encryption scheme Σ is correct if for all $k \in \mathbb{N}$, for all $\text{poly}(k)$ -size structures DS of type \mathcal{T} , for all (K, EDS) output by $\text{Setup}(1^k, \text{DS})$ and all sequences of $m = \text{poly}(k)$ queries q_1, \dots, q_m , for all tokens tk_i output by $\text{Token}(K, q_i)$, for all messages c output by $\text{Query}(\text{EDS}, \text{tk}_i)$, $\text{Resolve}(K, c)$ returns the correct response with all but negligible probability. The syntax of a response-revealing STE scheme can be recovered by omitting the Resolve algorithm and having Query output the response.

Security. The standard notion of security for STE guarantees that an encrypted structure reveals no information about its underlying structure beyond the setup leakage \mathcal{L}_S , and that the query algorithm reveals no information about the structure and the queries beyond the query leakage \mathcal{L}_Q . If this holds for non-adaptively chosen operations then this is referred to as non-adaptive security. If, on the other hand, the operations are chosen adaptively, this leads to the stronger notion of adaptive security [15]. This notion of security was first formalized by Curtmola *et al.* in the context of searchable encryption [15] and later generalized to structured encryption in [13].

Definition 2 (Adaptive security [15,13]). *Let $\Sigma_{\mathcal{T}} = (\text{Setup}, \text{Token}, \text{Query})$ be a structured encryption scheme for type \mathcal{T} and consider the following probabilistic experiments where \mathcal{A} is a stateful adversary, \mathcal{S} is a stateful simulator, \mathcal{L}_S and \mathcal{L}_Q are leakage profiles and $z \in \{0, 1\}^*$:*

Real $_{\Sigma, \mathcal{A}}(k)$: given z the adversary \mathcal{A} outputs a structure DS of type \mathcal{T} and receives EDS from the challenger, where $(K, \text{EDS}) \leftarrow \text{Setup}(1^k, \text{DS})$. The adversary then adaptively chooses a polynomial number of queries q_1, \dots, q_m . For all $i \in [m]$, the adversary receives $\text{tk}_i \leftarrow \text{Token}(K, q_i)$. Finally, \mathcal{A} outputs a bit b that is output by the experiment.

Ideal $_{\Sigma, \mathcal{A}, \mathcal{S}}(k)$: given z the adversary \mathcal{A} generates a structure DS of type \mathcal{T} which it sends to the challenger. Given z and leakage $\mathcal{L}_S(\text{DS})$ from the challenger, the simulator \mathcal{S} returns an encrypted data structure EDS to \mathcal{A} . The adversary then adaptively chooses a polynomial number of operations q_1, \dots, q_m . For all $i \in [m]$, the simulator receives query leakage $\mathcal{L}_Q(\text{DS}, q_i)$ and returns a token tk_i to \mathcal{A} . Finally, \mathcal{A} outputs a bit b that is output by the experiment.

We say that Σ is adaptively $(\mathcal{L}_S, \mathcal{L}_Q)$ -secure if for all PPT adversaries \mathcal{A} , there exists a PPT simulator \mathcal{S} such that for all $z \in \{0, 1\}^*$,

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

5 IEX: A Worst-Case Sub-Linear Disjunctive SSE Scheme

Our main construction, IEX, makes black-box use of a dictionary encryption scheme $\Sigma_{\text{DX}} = (\text{Setup}, \text{Token}, \text{Get})$, a multi-map encryption scheme $\Sigma_{\text{MM}} = (\text{Setup}, \text{Token}, \text{Get})$, a pseudo-random function F , and of a private-key encryption scheme $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$. The details of the scheme are provided in Fig. 1. At a high-level, it works as follows.

Setup. The **Setup** algorithm takes as input a security parameter k and an index DB . It makes use of two data structures: a dictionary DX and a *global* multi-map MM_g . MM_g maps every keyword in $w \in \mathcal{W}$ to an encryption of the identifiers in $\text{DB}(w)$. We refer to these encryptions as *tags* and they are computed by evaluating SKE.Enc using as coins the evaluation of F on keyword w and the identifier. The global multi-map MM_g is then encrypted using Σ_{MM} , resulting in EMM_g .

For each keyword $w \in \mathcal{W}$, the algorithm creates a *local* multi-map MM_w , that maps the keywords $v \in \text{co}(w)$ to tags of identifiers in $\text{DB}(v) \cap \text{DB}(w)$. Intuitively, the purpose of the local multi-map MM_w is to quickly find out which documents contain both w and v , for any $v \neq w$. The local multi-maps MM_w are then encrypted with Σ_{MM} . This results in encrypted multi-maps EMM_w which are then stored in the dictionary DX such that $\text{DX}[w] = \text{EMM}_w$. In other words, it stores label/value pairs (w, MM_w) in DX . Finally, DX is encrypted with Σ_{DX} , resulting in an encrypted dictionary EDX . The output of **Setup** includes the encrypted structures $(\text{EDX}, \text{EMM}_g)$ as well as their keys.

There are several optimizations possible for **Setup** that we omit in our formal description for ease of exposition. The first is that the encrypted local multi-maps can be stored “by reference” in the encrypted dictionary EDX instead of “by value”. More precisely, instead of storing the actual encrypted local multi-maps EMM_w in EDX one can just store a *pointer* to them. Another optimization is that, depending on how Σ_{MM} is designed, the keys for the local encrypted multi-maps could all be generated from a single key using a PRF (with a counter). This

would reduce the size of K . This optimization can be easily applied to most known encrypted multi-map schemes including the ones from [15,13,23,11,10].

Token. The Token algorithm takes as input a key and a vector of keywords $\mathbf{w} = (w_1, \dots, w_q)$. For all $i \in [q - 1]$ it creates a “sub-token” $\mathbf{tk}_i = (\mathbf{dtk}_i, \mathbf{gtk}_i, \mathbf{ltk}_{i+1}, \dots, \mathbf{ltk}_q)$ composed of a dictionary token \mathbf{dtk}_i , a global token \mathbf{gtk}_i for w_i and, for all keywords w_{i+1} through w_q in the disjunction, a local token \mathbf{ltk}_j for w_j , with $i + 1 \leq j \leq q$. Intuitively, the global token will allow the server to query the encrypted global multi-map \mathbf{EMM}_g to recover tags of the ids in $\text{DB}(w_i)$. The dictionary token for w_i will then allow the server to query the encrypted dictionary \mathbf{EDX} to recover w_i ’s local multi-map \mathbf{EMM}_i . Finally, the local tokens will allow the server to query w_i ’s encrypted local multi-map \mathbf{EMM}_i to recover the tags of the ids of the documents that contain both w_i and w_{i+1} , w_i and w_{i+2} , etc. As we will see next, this information will be enough for the server to find the relevant documents. For the last keyword w_q in the disjunction, the algorithm only needs to create a global token.

Search. The Search algorithm takes as input $\mathbf{EDB} = (\mathbf{EDX}, \mathbf{EMM}_g)$ and a token $\mathbf{tk} = (\mathbf{tk}_1, \dots, \mathbf{tk}_{q-1}, \mathbf{gtk}_q)$. For each sub-token $\mathbf{tk}_i = (\mathbf{dtk}_i, \mathbf{gtk}_i, \mathbf{ltk}_{i+1}, \dots, \mathbf{ltk}_q)$, the server does the following. It first uses \mathbf{gtk}_i to query the global multi-map \mathbf{EMM}_g and recover a set of identifier tags T_i for $\text{DB}(w_i)$. It then uses \mathbf{dtk}_i to query the encrypted dictionary \mathbf{EDX} to recover the local multi-map \mathbf{EMM}_i for w_i and uses \mathbf{ltk}_{i+1} to query \mathbf{EMM}_i to recover the tags T' for identifiers of the documents that contain both w_i and w_{i+1} ; that is, the tags for the set $I' = \text{DB}(w_i) \cap \text{DB}(w_{i+1})$. The server then removes T'_i from T_i . It then repeats this process for all local tokens \mathbf{ltk}_{i+2} to \mathbf{ltk}_q . Once it finishes processing all local tokens in \mathbf{tk}_i , it holds the set of tags for the set

$$\text{DB}(w_i) \setminus \bigcup_{j=i}^{q-1} \left(\text{DB}(w_i) \cap \text{DB}(w_{j+1}) \right). \quad (1)$$

Once it finishes processing all the sub-tokens, the server holds tags T_1 through T_{q-1} . For \mathbf{gtk}_q , the server just queries the global multi-map to recover T_q . Finally, it outputs the set

$$T = \bigcup_{i=1}^q T_i. \quad (2)$$

5.1 Correctness and Efficiency

We now analyze the correctness and efficiency of our construction. The correctness of IEX follows from Eqs. (1) and (2) and from the inclusion-exclusion principle. Given a disjunctive query $\mathbf{w} = (w_1, \dots, w_q)$, by Eq. (2), $\text{IEX.Search}(\mathbf{EDB}, \text{Token}(K, \mathbf{w}))$ will output

$$T = \bigcup_{i=1}^q T_i$$

Let F be a pseudo-random function, $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a private-key encryption scheme, $\Sigma_{\text{DX}} = (\text{Setup}, \text{Token}, \text{Get})$ be a dictionary encryption scheme and $\Sigma_{\text{MM}} = (\text{Setup}, \text{Token}, \text{Get})$ be a multi-map encryption scheme. Consider the disjunctive SSE scheme $\text{IEX} = (\text{Setup}, \text{Token}, \text{Search})$ defined as follows:

- $\text{Setup}(1^k, \text{DB})$:
 1. sample $K_1, K_2 \xleftarrow{\$} \{0, 1\}^k$;
 2. initialize a dictionary DX and a multi-map MM_g ;
 3. for all $w \in \text{W}$,
 - (a) for all $\text{id} \in \text{DB}(w)$, let $\text{tag}_{\text{id}} := \text{Enc}_{K_1}(\text{id}; F_{K_2}(\text{id}||w))$;
 - (b) set $\text{MM}_g[w] := (\text{tag}_{\text{id}})_{\text{id} \in \text{DB}(w)}$;
 - (c) initialize a multi-map MM_w of size $\#\text{co}(w)$;
 - (d) for all $v \in \text{co}(w)$,
 - i. for all $\text{id} \in \text{DB}(v) \cap \text{DB}(w)$, let $\text{tag}_{\text{id}} := \text{Enc}_{K_1}(\text{id}; F_{K_2}(\text{id}||w))$;
 - ii. set $\text{MM}_w[v] := (\text{tag}_{\text{id}})_{\text{id} \in \text{DB}(v) \cap \text{DB}(w)}$;
 - (e) compute $(K_w, \text{EMM}_w) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_w)$;
 - (f) set $\text{DX}[w] := \text{EMM}_w$;
 4. compute $(K_g, \text{EMM}_g) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_g)$;
 5. compute $(K_d, \text{EDX}) \leftarrow \Sigma_{\text{DX}}.\text{Setup}(1^k, \text{DX})$;
 6. set $K = (K_g, K_d, \{K_w\}_{w \in \text{W}})$ and $\text{EDB} = (\text{EMM}_g, \text{EDX})$;
 7. output (K, EDB) .
- $\text{Token}(K, \mathbf{w})$:
 1. parse \mathbf{w} as (w_1, \dots, w_q) ;
 2. for all $i \in [q - 1]$,
 - (a) compute $\text{gtk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}(K_g, w_i)$;
 - (b) compute $\text{dtk}_i \leftarrow \Sigma_{\text{DX}}.\text{Token}(K_d, w_i)$;
 - (c) for all $i + 1 \leq j \leq \#\mathbf{w}$, compute $\text{ltk}_j \leftarrow \Sigma_{\text{MM}}.\text{Token}(K_{w_i}, w_j)$;
 - (d) set $\mathbf{tk}_i = (\text{dtk}_i, \text{gtk}_i, \text{ltk}_{i+1}, \dots, \text{ltk}_{\#\mathbf{w}})$;
 3. compute $\text{gtk}_q \leftarrow \Sigma_{\text{MM}}.\text{Token}(K_g, w_q)$;
 4. output $\mathbf{tk} = (\mathbf{tk}_1, \dots, \mathbf{tk}_{q-1}, \text{gtk}_q)$.
- $\text{Search}(\text{EDB}, \mathbf{tk})$:
 1. parse EDB as $(\text{EMM}_g, \text{EDX})$;
 2. parse \mathbf{tk} as $(\mathbf{tk}_1, \dots, \mathbf{tk}_{q-1}, \text{gtk}_q)$;
 3. for all $i \in [q - 1]$,
 - (a) parse \mathbf{tk}_i as $(\text{dtk}_i, \text{gtk}_i, \text{ltk}_{i+1}, \dots, \text{ltk}_q)$;
 - (b) compute $T_i \leftarrow \Sigma_{\text{MM}}.\text{Get}(\text{EMM}_g, \text{gtk}_i)$;
 - (c) compute $\text{EMM}_i \leftarrow \Sigma_{\text{DX}}.\text{Get}(\text{EDX}, \text{dtk}_i)$;
 - (d) for all $i + 1 \leq j \leq q$,
 - i. compute $T'_j \leftarrow \Sigma_{\text{MM}}.\text{Get}(\text{EMM}_i, \text{ltk}_j)$;
 - ii. set $T_j = T_j \setminus T'_j$;
 4. compute $T_q \leftarrow \Sigma_{\text{MM}}.\text{Get}(\text{EMM}_g, \text{gtk}_q)$;
 5. output $\bigcup_{i \in [q]} T_i$;

Fig. 1. Our disjunctive SSE scheme IEX.

$$\begin{aligned}
&= \left(\bigcup_{i=1}^{q-1} T_i \right) \cup T_q \\
&= \left(\bigcup_{i=1}^{q-1} \left(\text{DB}(w_i) \setminus \bigcup_{j=i}^{q-1} (\text{DB}(w_i) \cap \text{DB}(w_{j+1})) \right) \right) \cup \text{DB}(w_q) \\
&= \left(\bigcup_{i=1}^{q-2} \left(\text{DB}(w_i) \setminus \bigcup_{j=i}^{q-1} (\text{DB}(w_i) \cap \text{DB}(w_{j+1})) \right) \right) \\
&\quad \underbrace{\bigcup (\text{DB}(w_{q-1}) \setminus (\text{DB}(w_{q-1}) \cap \text{DB}(w_q))) \cup \text{DB}(w_q)}_U
\end{aligned} \tag{3}$$

where the first and third equalities hold by Eqs. (2) and (1), respectively. Note, however, that U equals $\text{DB}(w_{q-1}) \cup \text{DB}(w_q)$:

$$\begin{aligned}
U &= \text{DB}(w_{q-1}) \cap \left(\overline{\text{DB}(w_{q-1}) \cup \text{DB}(w_q)} \right) \cup \text{DB}(w_q) \\
&= \left(\text{DB}(w_{q-1}) \cap \overline{\text{DB}(w_{q-1})} \right) \cup \left(\text{DB}(w_{q-1}) \cap \overline{\text{DB}(w_q)} \right) \cup \text{DB}(w_q) \\
&= \left(\text{DB}(w_{q-1}) \cap \overline{\text{DB}(w_q)} \right) \cup \text{DB}(w_q) \\
&= \left(\text{DB}(w_{q-1}) \cup \text{DB}(w_q) \right) \cap \left(\overline{\text{DB}(w_q)} \cup \text{DB}(w_q) \right) \\
&= \text{DB}(w_{q-1}) \cup \text{DB}(w_q)
\end{aligned}$$

Repeating the same argument for $q-2$, $q-3$ and so on and plugging into Eq. (3), we get that $T = \bigcup_{i=1}^q \text{DB}(w_i)$.

Efficiency. The search complexity of IEX is $O(q^2 \cdot M)$, where $M = \max_{i \in [q]} \#\text{DB}(w_i)$ and q is the number of terms in the disjunction. Tokens are of size $O(q)$. We also note that unlike BXT and OXT [11], IEX tokens are *selectivity-independent* in the sense that they do not depend on the size of the result. The IEX storage complexity is,

$$O\left(\text{strg}\left(\sum_w \#\text{DB}(w)\right) + \sum_w \text{strg}\left(\sum_{v \in \text{co}(w)} \#\text{DB}(v) \cap \text{DB}(w)\right)\right),$$

where strg is the storage complexity of the underlying encrypted multi-map encryption scheme Σ_{MM} .

A storage optimization. As we can see, the storage complexity of IEX can be large, especially if the underlying encrypted multi-maps are. This is indeed the case when they are instantiated with standard sub-linear constructions. We observe, however, that we can tradeoff storage complexity (and setup time) for the communication complexity of search as follows. When constructing a local multi-map EMM_w for a keyword w , we normally insert tags for the identifiers in

$\text{DB}(w) \cap \text{DB}(v)$ for all $v \in \text{co}(w)$. This is not necessary for correctness, however, so we can omit some of the co-occurring keywords from w 's local multi-map. The tradeoff is that this will increase the communication complexity of IEX's search operation and, in particular, make it non-optimal.

To do this, we suggest using the following approach to decide whether to add a keyword $v \in \text{co}(w)$ or not. Let $p < 1$ be a *filtering* parameter and let

$$T_{w,v} \stackrel{\text{def}}{=} \frac{\#\text{DB}(v) \cap \text{DB}(w)}{\max(\#\text{DB}(w), \#\text{DB}(v))}.$$

If $T_{w,v} > p$, then add v to EMM_w otherwise do not. With this filtering in place, the storage complexity of IEX is now

$$O\left(\text{strg}\left(\sum_w \#\text{DB}(w)\right) + \sum_w \text{strg}\left(\sum_{\substack{v \in \text{co}(w) \\ T_{w,v} > p}} \#\text{DB}(v) \cap \text{DB}(w)\right)\right).$$

In our experiments we set $p = 0.2$.

Remark. We note that when all the terms of the disjunctive query have selectivity $O(n)$, IEX has linear search complexity. This is, however, the best one can do. On the other hand, the communication complexity of IEX remains optimal independently of the selectivity of the terms. This similarly applies to OXT but not to BlindSeer since it induces a logarithmic (multiplicative) overhead.

5.2 Security

The setup leakage of IEX consists of the setup leakage of its underlying building blocks. In particular, this includes the setup leakage of the encrypted global multi-map and of the encrypted dictionary. Assuming the use of standard optimal-time multi-map and dictionary encryption schemes [15,13,23,10], this reveals the size of the database DB as well as the total size of the local multi-maps stored in the dictionary. The query leakage of IEX for a query \mathbf{w} includes, for each keyword $w_i \in W$, the query leakage of the encrypted dictionary and of the encrypted global multi-map. It also includes the query leakage of every queried local multi-map as well as their setup leakage. Again if instantiated with standard constructions, this will consist of the search and access patterns which, respectively, capture whether or not the same query has been searched for and (in our case) the tags. Finally, the query leakage also includes the number of documents containing $\text{DB}(w_i) \cap \text{DB}(w_{j+1})$, for all $j \geq i$ and $i \in [q-1]$.

We now give a precise description of IEX's leakage profile and show that it is adaptively-secure with respect to it. Its setup leakage is

$$\mathcal{L}_S^{\text{ieX}}(\text{DB}) = \left(\mathcal{L}_S^{\text{dX}}(\text{DX}), \mathcal{L}_S^{\text{mm}}(\text{MM}_g) \right),$$

where $\mathcal{L}_S^{\text{dX}}(\text{DX})$ and $\mathcal{L}_S^{\text{mm}}(\text{MM}_g)$ are the setup leakages of the underlying dictionary and multi-map encryption schemes, respectively. Its query leakage is

$$\mathcal{L}_Q^{\text{ieX}}(\text{DB}, \mathbf{w}) = \left(\left(\mathcal{L}_Q^{\text{dX}}(\text{DX}, w_i), \mathcal{L}_S^{\text{mm}}(\text{MM}_i) \right) \right),$$

$$\mathcal{L}_Q^{\text{mm}}(\text{MM}_g, w_i), \dots, \mathcal{L}_Q^{\text{mm}}(\text{MM}_i, w_q), \text{TagPat}_i(\text{DB}, \mathbf{w}) \Big)_{i \in [q-1]},$$

$$\mathcal{L}_Q^{\text{mm}}(\text{MM}_g, w_q), \text{TagPat}_q(\text{DB}, \mathbf{w}) \Big),$$

where, for all $i \in [q]$,

$$\text{TagPat}_i(\text{DB}, \mathbf{w}) = \left(\left(f_i(\text{id}) \right)_{\text{id} \in \text{DB}(w_i) \cap \text{DB}(w_{i+1})}, \dots, \left(f_i(\text{id}) \right)_{\text{id} \in \text{DB}(w_i) \cap \text{DB}(w_q)} \right),$$

and f_i is a random function from $\{0, 1\}^{|\text{id}| + \log \#W}$ to $\{0, 1\}^k$.

Theorem 1. *If Σ_{DX} is adaptively $(\mathcal{L}_S^{\text{dx}}, \mathcal{L}_Q^{\text{dx}})$ -secure, Σ_{MM} is adaptively $(\mathcal{L}_S^{\text{mm}}, \mathcal{L}_Q^{\text{mm}})$ -secure, SKE is RCPA-secure and F is pseudo-random, then IEX is $(\mathcal{L}_S^{\text{ieX}}, \mathcal{L}_Q^{\text{ieX}})$ -secure.*

The proof of Theorem 1 is deferred to the full version of the paper.

6 Boolean Queries with IEX

While IEX is naturally disjunctive, it can also support boolean queries. The boolean variant is similar to IEX in that it uses the same encrypted structures (i.e., it has the same `Setup` algorithm) but different `Token` and `Search` algorithms. We refer to the boolean variant of IEX as BIEX. We now provide an overview of how BIEX works.

Overview of BIEX. Recall that any query can be written in conjunctive normal form (CNF) so it has the form $\Delta_1 \wedge \dots \wedge \Delta_\ell$, where each $\Delta_i = w_{i,1} \vee \dots \vee w_{i,q}$ is a disjunction. Note that the result $\text{DB}(\Delta_1 \wedge \dots \wedge \Delta_\ell)$ is the intersection of $\text{DB}(\Delta_1)$ through $\text{DB}(\Delta_\ell)$. But this intersection does not have to be computed “directly” by executing a naive intersection operation. A better alternative (from a leakage point of view) is to compute the intersection by starting with $\text{DB}(w_1)$, keeping only the subset of identifiers of $\text{DB}(w_1)$ that are also in $\text{DB}(w_2)$, then keeping only the subset of identifiers that are also in $\text{DB}(w_3)$ and so on. This alternative approach requires only information about $\text{DB}(w_1)$ and the progressive subsets. Moreover, it uses operations that are already supported by the IEX structures.

How we do this exactly, is best explained through a concrete example. Suppose we have a CNF query with $\Delta_1 = w_1 \vee w_2$ and $\Delta_2 = w_3$. The first step would be to perform a disjunctive query for Δ_1 , resulting in tags for the identifiers in $\text{DB}(\Delta_1)$. In the second step, we want to filter out and keep the tags of identifiers in $\text{DB}(\Delta_1) \cap \text{DB}(w_3)$. To find these tags, it suffices to query the local multi-maps of w_1 and w_2 on w_3 . In the first case, EMM_{w_1} will return tags for $\text{DB}(w_1) \cap \text{DB}(w_3)$ and in the second case EMM_{w_2} will return tags for $\text{DB}(w_2) \cap \text{DB}(w_3)$. Finally, we take the union of both of these intersections and perform a final intersection with $\text{DB}(\Delta_1)$. The final result equals $\text{DB}(\Delta_1) \cap \text{DB}(w_3)$. Fig. 2 describes this process in more detail and for arbitrary boolean queries.

Correctness. To show correctness we need to show that, given a boolean query in CNF form $\Delta_1 \wedge \dots \wedge \Delta_\ell$ such that $\Delta_i = w_{i,1} \vee \dots \vee w_{i,q}$ (for simplicity we assume the disjunctions all have q terms), `BIEX.Search` outputs

$$\bigcap_{i \in [\ell]} \bigcup_{j \in [q]} \text{DB}(w_{i,j}). \quad (4)$$

Looking at the description of `BIEX.Search` in Fig. 2, one can see that every time Step 4(d)i is invoked it outputs

$$\bigcup_{j \in [q]} \text{DB}(w_{i,t}) \cap \text{DB}(w_{1,j}),$$

for all $t \in [q]$ and $i \in [\ell]$. Note that this stems from the fact that $\Sigma_{\text{MM}}.\text{Get}(\text{EMM}_j, \text{ltk}_{t,i,j})$ outputs $\text{DB}(w_{i,t}) \cap \text{DB}(w_{1,j})$ for every $j \in [q]$.

Also, based on the correctness of `IEX` we know that the search for the first disjunction will output $\bigcup_{j \in [q]} \text{DB}(w_{1,j})$ (with no redundant identifiers). So we have the final result of the query

$$I_\ell = \bigcup_{j \in [q]} \text{DB}(w_{1,j}) \cap \underbrace{\left(\bigcup_{j,t \in [q]} (\text{DB}(w_{2,j}) \cap \text{DB}(w_{1,t})) \right) \cap \dots \cap \left(\bigcup_{j,t \in [q]} (\text{DB}(w_{\ell,j}) \cap \text{DB}(w_{1,t})) \right)}_{\ell-1 \text{ terms}} \quad (5)$$

On the other hand, note that for all $i \in [\ell]$ we have by Morgan's laws that

$$\begin{aligned} \bigcup_{j \in [q]} \text{DB}(w_{1,j}) \cap \bigcup_{j \in [q]} \text{DB}(w_{i,j}) &= \bigcup_{j \in [q]} \text{DB}(w_{1,j}) \cap \bigcup_{j \in [q]} \text{DB}(w_{1,j}) \cap \bigcup_{j \in [q]} \text{DB}(w_{i,j}) \\ &= \bigcup_{j \in [q]} \text{DB}(w_{1,j}) \cap \left(\bigcup_{j,t \in [q]} (\text{DB}(w_{i,j}) \cap \text{DB}(w_{1,t})) \right) \end{aligned}$$

That is, we can recursively apply the above result on Eq. (5) for all $l \in [\ell]$ to obtain Eq. (4).

Efficiency. The storage complexity of `BIEX` is the same as `IEX`. Its search complexity is

$$O\left(q^2 \cdot \left(\max_{w \in \Delta_1} \#\text{DB}(w) + \ell \cdot \#\text{DB}(\Delta_1) \right)\right).$$

The term $q^2 \cdot \max_{w \in \Delta_1} \#\text{DB}(w)$ is the time to search for the first disjunction and the second term $q^2 \cdot \ell \cdot \#\text{DB}(\Delta_1)$ is the total number of local multi-map queries.

We can clearly see from the search complexity of `BIEX` that we can achieve better efficiency if the selectivity of the first disjunction is as small as possible. In practice, therefore, the first disjunction should be the one with the smallest selectivity; similarly to how the first keyword is chosen in `OXT`. Note that if the first disjunction in the CNF form of the boolean query matches the entire database then the search complexity of `BIEX` will be linear while the optimal complexity might be sub-linear (the communication complexity of `BIEX` will

remain optimal, however). It is not obvious to us how to improve this without pre-computing every possible query as it seems almost inherent to the query itself. With this in mind, it follows that BIEX has a sub-linear worst-case search complexity when the first disjunction's selectivity is sub-linear.

The communication complexity of BIEX is optimal since the final set I_ℓ does not contain any redundant identifiers. Finally, note that it is non-interactive and token size is independent of the query's selectivity.

Security. The setup leakage of BIEX is the same as IEX's. Its query leakage includes the query leakage of IEX on the first disjunction and the query leakage of the encrypted local multi-maps when queried on all the terms of disjunctions $\Delta_2, \dots, \Delta_\ell$. Finally, it also includes the number of documents that match the terms of the first disjunction and the terms of remaining disjunctions.

We now give a precise description of the leakage profile of BIEX and show that it is adaptively-secure with respect to it. The setup leakage is

$$\mathcal{L}_S^{\text{iexb}}(\text{DB}) = \mathcal{L}_S^{\text{ie x}}(\text{DB}),$$

where $\mathcal{L}_S^{\text{ie x}}(\text{DB})$ is the setup leakages of IEX. Given a CNF query $\Delta_1 \wedge \dots \wedge \Delta_\ell$, the query leakage is

$$\mathcal{L}_Q^{\text{iexb}}\left(\text{DB}, \bigwedge_{i=1}^{\ell} \Delta_i\right) = \left(\mathcal{L}_Q^{\text{ie x}}(\text{DB}, \Delta_1), \left(\mathcal{L}_Q^{\text{mm}}(\text{MM}_i, w_{l,1}), \dots, \mathcal{L}_Q^{\text{mm}}(\text{MM}_i, w_{l,q}) \right), \right. \\ \left. \text{TagPat}_{i,l}\left(\text{DB}, \bigwedge_{i=1}^{\ell} \Delta_i\right)_{\substack{i \in [q] \\ l \in [2, \dots, \ell]}} \right).$$

where,

$$\text{TagPat}_{i,l}\left(\text{DB}, \bigwedge_{i=1}^{\ell} \Delta_i\right) = \left(\left(f_i(\text{id}) \right)_{\text{DB}(w_{1,i}) \cap \text{DB}(w_{1,1})}, \dots, \left(f_i(\text{id}) \right)_{\text{DB}(w_{1,i}) \cap \text{DB}(w_{l,q})} \right)$$

and f_i is a random function from $\{0, 1\}^{n+\log \#W}$ to $\{0, 1\}^k$.

Theorem 2. *If Σ_{DX} is adaptively $(\mathcal{L}_S^{\text{dx}}, \mathcal{L}_Q^{\text{dx}})$ -semantically secure and Σ_{MM} is adaptively $(\mathcal{L}_S^{\text{mm}}, \mathcal{L}_Q^{\text{mm}})$ -secure, then BIEX is adaptively $(\mathcal{L}_S^{\text{iexb}}, \mathcal{L}_Q^{\text{iexb}})$ -secure.*

The proof of Theorem 2 is similar (at a high-level) to the proof of Theorem 1.

7 ZMF: A Compact and Adaptively-Secure SSE Scheme

The main limitation of IEX is its storage complexity of

$$O\left(\text{strg}_g(\text{MM}_g) + \sum_w \text{strg}_\ell(\text{MM}_w)\right),$$

Let $\text{IEX} = (\text{Setup}, \text{Token}, \text{Search})$ be the IEX scheme described in Figure 1 and let $\Sigma_{\text{DX}} = (\text{Setup}, \text{Token}, \text{Get})$ and $\Sigma_{\text{MM}} = (\text{Setup}, \text{Token}, \text{Get})$ be its underlying dictionary and multi-map encryption schemes, respectively. Consider the boolean SSE encryption scheme $\text{BIEX} = (\text{Setup}, \text{Token}, \text{Search})$ defined as follows:

- $\text{Setup}(1^k, \text{DB})$: output $(K, \text{EDB}) \leftarrow \text{IEX.Setup}(1^k, \text{DB})$.
- $\text{Token}(K, \mathbf{w})$:
 1. parse K as $(K_g, K_d, \{K_w\}_{w \in W})$;
 2. parse \mathbf{w} as $(\Delta_1 \wedge \cdots \wedge \Delta_\ell)$ where for all $i \in [\ell]$, $\Delta_i = (w_{i,1} \vee \cdots \vee w_{i,d})$;
 3. compute $\mathbf{tk}_1 \leftarrow \text{IEX.Token}_K(\Delta_1)$;
 4. for all $2 \leq i \leq \ell$ and all $j \in [q]$,
 - (a) for all $1 \leq s \leq q$, compute $\text{ltk}_{s,i,j} \leftarrow \Sigma_{\text{MM}}.\text{Token}(K_{w_{1,s}}, w_{i,j})$;
 - (b) set $\mathbf{tk}_{i,j} = (\text{ltk}_{1,i,j}, \dots, \text{ltk}_{q,i,j})$;
 - (c) set $\mathbf{tk}_i = (\mathbf{tk}_{i,1}, \dots, \mathbf{tk}_{i,q})$;
 5. output $\mathbf{tk} = (\mathbf{tk}_1, \dots, \mathbf{tk}_\ell)$.
- $\text{Search}(\text{EDB}, \mathbf{tk})$:
 1. parse EDB as $(\text{EMM}_g, \text{EDX})$;
 2. parse \mathbf{tk} as $(\mathbf{tk}_1, \dots, \mathbf{tk}_\ell)$;
 3. compute $I_1 \leftarrow \text{IEX.Search}(\text{EDB}, \mathbf{tk}_1)$;
 4. for all $2 \leq i \leq \ell$,
 - instantiate an empty set I_i ;
 - parse $\mathbf{tk}_i = (\mathbf{tk}_{i,1}, \dots, \mathbf{tk}_{i,q})$;
 - for $j \in [q]$,
 - (a) get dtk_j from \mathbf{tk}_1 ;
 - (b) compute $\text{EMM}_j \leftarrow \Sigma_{\text{DX}}.\text{Get}(\text{EDX}, \text{dtk}_j)$;
 - (c) parse $\mathbf{tk}_{i,j} = (\text{ltk}_{1,i,j}, \dots, \text{ltk}_{q,i,j})$;
 - (d) for $s \in [q]$,
 - i. compute $I' \leftarrow \Sigma_{\text{MM}}.\text{Get}(\text{EMM}_j, \text{ltk}_{s,i,j})$;
 - ii. compute $I_i = I_i \cup (I_{i-1} \cap I')$;
 5. output I_ℓ ;

Fig. 2. The scheme BIEX.

where strg_g and strg_ℓ are the storage complexity of the global and local EMMs, respectively. If the latter are instantiated with standard sub-linear-time constructions such as [15,24,11,10], we have

$$O\left(\sum_w \#\text{DB}(w) + \sum_w \sum_{v \in \text{co}(w)} \#\text{DB}(v) \cap \text{DB}(w)\right), \quad (6)$$

which does not compare favorably to standard single-keyword search solutions which require only $O(\sum_w \#\text{DB}(w))$, or to the OXT construction of [11] which requires

$$O\left(\sum_w \#\text{DB}(w) + \log\left(\frac{1}{\varepsilon}\right) \cdot \sum_w \#\text{DB}(w)\right)$$

when XSet is instantiated with a Bloom filter with a false positive rate of ε . In particular, note that the second term in the asymptotic expression above hides a constant of 1, which makes OXT reasonably compact.

Our approach. The main storage inefficiency in IEX comes from the local EMMs which contribute the second term in Eq. (6). Ideally, we could improve things if we could use more compact local EMMs. Unfortunately, all known sub-linear constructions require $O(\sum_w \#\text{DB}(w))$ storage. We observe, however, that for local EMMs sub-linear search is not necessary since in practice the number of label/tuple pairs they store is small in comparison to the total number of documents n . So, for our purposes, a linear-time construction would work as long as it was compact. In [19], Goh proposed a very compact construction called Z-IDX based on Bloom filters. Specifically, it needs only

$$O\left(\log\left(\frac{1}{\varepsilon}\right) \cdot \sum_{v \in \mathbf{V}} \#\text{MM}^{-1}[v]\right)$$

bits of storage, where \mathbf{V} is the value space of the multi-map and ε is the false positive rate. If we could encrypt the local EMMs of IEX with Z-IDX, the former's storage would be

$$O\left(\sum_w \#\text{DB}(w) + \log\left(\frac{1}{\varepsilon}\right) \cdot \sum_w \#\text{co}(w)\right),$$

which is much more competitive with OXT (note that the second term here also has a constant of 1). Unfortunately, this approach does not work because Z-IDX is not adaptively secure. Nevertheless, we show how to construct a highly compact scheme that is. In the following, we first recall how Z-IDX works.

Goh's Z-IDX scheme. Like any SSE scheme, Z-IDX can be viewed as a STE scheme and, in particular, as a multi-map encryption scheme. Conceptually, we observe that Z-IDX can be abstracted into two parts: (1) a compiler that transforms an underlying set encryption scheme into a multi-map encryption scheme; and (2) a concrete set encryption scheme based on Bloom filters and PRFs. We refer to the former as the Z-IDX transformation and describe it in

Let $\Sigma_{\text{SET}} = (\text{Gen}, \text{Enc}, \text{Token}, \text{Test})$ be a multi-structure set encryption scheme and consider the multi-map encryption scheme $\Sigma_{\text{MM}} = (\text{Setup}, \text{Token}, \text{Get})$ defined as follows:

- **Setup**($1^k, \text{MM}$):
 1. compute $K \leftarrow \text{Gen}(1^k)$;
 2. let \mathbf{V} be the range of MM ;
 3. for all $v \in \mathbf{V}$,
 - (a) let $S_v = \text{MM}^{-1}(v)$;
 - (b) compute $\text{ESET}_v \leftarrow \Sigma_{\text{SET}}.\text{Enc}(K, S_v)$;
 4. output $\text{EMM} = (\text{ESET}_v)_{v \in \mathbf{V}}$.
- **Token**(K, ℓ): output $\text{tk} \leftarrow \Sigma_{\text{SET}}.\text{Token}(K, \ell)$
- **Get**(EMM, tk):
 1. let $I = \emptyset$;
 2. for all $v \in \mathbf{V}$,
 - (a) if $\Sigma_{\text{SET}}.\text{Test}(\text{ESET}_v, \text{tk})$ outputs 1, set $I = I \cup \{v\}$;
 3. output I .

Fig. 3. The Z-IDX transformation.

detail in Fig. 3. Given a set encryption scheme Σ_{SET} , it produces a multi-map encryption scheme Σ_{MM} that works as follows. The $\Sigma_{\text{MM}}.\text{Setup}$ algorithm takes as input a multi-map MM that maps labels to tuples of values from \mathbf{V} . It creates $\#\mathbf{V}$ sets $(S_v)_{v \in \mathbf{V}}$ such that S_v holds the labels in MM that map to v . It then encrypts each set S_v with Σ_{SET} resulting in an encrypted set ESET_v . The encrypted multi-map EMM is simply the collection of encrypted sets $(\text{ESET}_v)_{v \in \mathbf{V}}$. A Σ_{MM} token for a label ℓ is a Σ_{SET} token for ℓ and $\Sigma_{\text{MM}}.\text{Get}$ uses the token to test each set in $\text{EMM} = (\text{ESET}_v)_{v \in \mathbf{V}}$ and outputs v if the test succeeds.

Note that for Σ_{MM} to work, Σ_{SET} must satisfy a stronger STE form than what is described in Definition 1. In particular, it must be what we call *multi-structure* in the sense that the tokens produced with a key K can be used to query all the structures encrypted under K . We provide formal syntax and security definitions of multi-structure STE schemes in the full version of the paper. The main difference between standard and multi-structure STE schemes are that in the latter the **Setup** algorithm is replaced with a key generation algorithm $\text{Gen}(1^k)$ that takes as input a security parameter and outputs a secret key K ; and an encryption algorithm $\text{Enc}(K, \text{DS})$ that takes as input a secret key K and a data structure DS and outputs an encrypted structure EDS .

Adaptive security. From our abstract perspective, the reason Z-IDX is not adaptively-secure is because the Z-IDX transformation is (implicitly) applied to a set encryption scheme that is not adaptively-secure. We show in Theorem 3 below, however, that if the transformation is applied to an adaptively-secure set encryption scheme then the result is adaptively-secure as well. More precisely, we show that if the set encryption scheme is adaptively $(\mathcal{L}_S^{\text{set}}, \mathcal{L}_Q^{\text{set}})$ -secure then the Z-IDX transformation yields a multi-map encryption scheme with the following

leakage profile:

$$\mathcal{L}_S^{\text{mm}}(\text{MM}) = \left((\mathcal{L}_S^{\text{set}}(\text{MM}^{-1}[v]))_{v \in \mathbf{V}}, \#\mathbf{V} \right),$$

and

$$\mathcal{L}_Q^{\text{mm}}(\text{MM}, q) = \mathcal{L}_Q^{\text{set}} \left((\text{MM}^{-1}[v])_{v \in \mathbf{V}}, q \right).$$

Theorem 3. *If Σ_{SET} is adaptively $(\mathcal{L}_S^{\text{set}}, \mathcal{L}_Q^{\text{set}})$ -secure then the scheme Σ_{MM} that results from applying the Z-IDX transformation to it is adaptively $(\mathcal{L}_S^{\text{mm}}, \mathcal{L}_Q^{\text{mm}})$ -secure.*

Due to space constraints, the proof of Theorem 3 will appear in the full version of the paper.

7.1 An Adaptively-Secure and Multi-Structure Set Encryption Scheme

In this Section, we construct an adaptively-secure, highly-compact and multi-structure set encryption scheme. Then, by applying the Z-IDX transformation to it we get an adaptively-secure and highly-compact multi-map encryption scheme which we then use in IEX.

Adaptive security. The main difficulty in designing adaptively-secure encrypted structures is supporting *equivocation* during simulation. Roughly speaking, the issue is that during the $\mathbf{Ideal}(k)$ experiment the simulator first needs to simulate an encrypted structure for the adversary and later needs to be able to simulate tokens that work correctly with the simulated structure produced in the first step. The challenge in supporting equivocation is that at the time the encrypted structure is simulated, the simulator has no information about the adversary’s queries so it is not clear how to simulate the structure in a way that will work correctly at query time. So to handle equivocation, the construction needs to be carefully designed and, typically, needs expensive cryptographic primitives. Fortunately, as first shown by Chase and Kamara [13], in the setting of symmetric STE, equivocation can be achieved very efficiently based only on XOR and PRF operations.

Our base scheme. One possible way to design an encrypted Bloom filter is as follows. Let \mathbf{U} be a universe of elements. Given a set $S \subseteq \mathbf{U}$, insert the value $F_K(a)$, for all $a \in S$, in a standard Bloom filter, where F is a pseudo-random function. The token for an element $a \in \mathbf{U}$ is $\text{tk} = F_K(a)$ and the Bloom filter can be queried by doing a standard Bloom filter test on tk .

The main problems with this construction are that: (1) it reveals information about the size of S ; and (2) it is not adaptively-secure. To achieve adaptive security, we can encrypt the Bloom filter by XORing each of its bits with a pad generated from another pseudo-random function G . This encryption step both hides the size of S and allows for equivocation. Now the token tk for an element $a \in \mathbf{U}$ includes $F_{K_1}(a)$ and the pads for locations $H_1(F_{K_1}(a))$ through $H_\lambda(F_{K_1}(a))$, where (H_1, \dots, H_λ) are the hash functions used for the Bloom filter.

For this to work, however, the pads have to be designed carefully. More precisely, correctness requires that the pads only depend on the locations that they mask otherwise two (or more) elements a_1 and a_2 that collide under one of the hash functions will produce different masks for the same location. To get such *location-dependent* pads we compute them as $G_{K_2}(\ell)$, where ℓ is the ℓ th bit of the filter. Now, a token for element a is set to

$$\text{tk} = \left(F_{K_1}(a), G_{K_2} \left(H_1(F_{K_1}(a)) \right), \dots, G_{K_2} \left(H_\lambda(F_{K_1}(a)) \right) \right).$$

The base construction described so far is *compact* and *adaptively-secure* but not multi-structure.

Reusability. Recall that a multi-structure STE scheme can produce *multiple* encrypted structures $(\text{EDS}_1, \dots, \text{EDS}_n)$ under a single key K in such a way that a *single* (constant-size) token tk can be used to query all the structures generated under key K . So to make our base scheme multi-structure, the pads have to be filter-dependent in addition to being location-dependent so that different pads are used for different filters *even if they mask the same location*. We do this by setting the pads to be the output of a random oracle applied to the pair $(G_{K_2}(\ell), \text{id})$ where id is the identifier of the filter. The purpose of the random oracle here is twofold. First, it enables the extraction of n (random) pads from pairs $(G_{K_2}(\ell), \text{id}_1)$ through $(G_{K_2}(\ell), \text{id}_n)$ without relying on n secret keys. This, in turn, means the tokens can be of size independent of n . Second, it allows the simulator to equivocate on the pads while, again, keeping the tokens independent of n .

While the base scheme is now compact, adaptively-secure and multi-structure, it produces very large tokens. The problem is that if two sets S_1 and S_2 have different sizes, then the parameters of their Bloom filters (i.e., the array sizes, number of hash functions and hash function ranges) have to be different. The consequence is that in our encrypted set scheme, we will need different sets of hash functions for *each* filter which, in turn, means the tokens will have to include multiple pads for *every* filter.

Matryoshka filters. We solve this problem as follows. Instead of encrypting a set of standard Bloom filters as in our base construction, we encrypt a new filter-based structure we refer to as *matryoshka filters* (MF).⁵ MFs are essentially a set of *nested* Bloom filters of varying sizes whose hash functions are all derived from a fixed set of hash functions. More precisely, consider a sequence of sets $S_1, \dots, S_n \subseteq \mathbf{U}$ not necessarily of the same size. We assume for simplicity that the sets have size a multiple of 2. For some false negative rate $2^{-\lambda}$, choose λ independent and ideal random hash functions (H_1, \dots, H_λ) from \mathbf{U} to $[(\lambda/\ln 2) \cdot \max_i \#S_i]$. We refer to these functions as the *maximal* hash functions and to their associated filter as the *maximal* filter. For every set S_i , construct a Bloom filter of size $\lceil [(\lambda/\ln 2) \cdot \#S_i] \rceil$ with hash functions (H_1^i, \dots, H_h^i) where, for all $j \in [\lambda]$, $H_j^i(a) = H_j(a)_{\parallel p_i}$ with

⁵ The term matryoshka here refers to Russian nested dolls which are called matryoshka dolls.

$p_i = \lceil \log((\lambda/\ln 2) \cdot \#S_i) \rceil$. We refer to these hash functions as the *derived* functions and to their associated filters as the *derived* filters. Note that if the maximal hash functions are ideal random functions then so are the derived functions so the standard Bloom filter analysis holds.

Encrypting matryoshka filters. As mentioned above, our final solution consists of adapting our base scheme to encrypt matryoshka filters instead of standard Bloom filters. In other words, we XOR each bit of each matryoshka filter with location- and filter-dependent pads. The main difference with the base scheme is that here the pads also need to be nested; that is, given a pad for the maximal filter we need to be able to construct the pads for the derived filters. To support this, we make use of the properties of online ciphers; namely, that given an n -bit string s and a B -online cipher OC , the following equality holds:

$$\text{OC}_K\left(s\Big|_{p \times B}^{\parallel B}\right) = \text{OC}_K\left(s\Big|_{p \times B}^{\parallel B}\right), \quad (7)$$

where $p < n$. This can be derived as follows. From the correctness property of online ciphers, we have

$$\begin{aligned} \text{OC}_K\left(s\Big|_{p \times B}^{\parallel B}\right) &= \text{OC}_K^1\left(s\Big|_{p \times B}^{\parallel B}\right) \parallel \cdots \parallel \text{OC}_K^p\left(s\Big|_{p \times B}^{\parallel B}\right) \\ &= X\left(K, s\Big|_B^{\parallel B}\right) \parallel \cdots \parallel X\left(K, s\Big|_{p \times B}^{\parallel B}\right), \end{aligned}$$

and

$$\begin{aligned} \text{OC}_K\left(s\Big|_B^{\parallel B}\right) &= \text{OC}_K^1\left(s\Big|_B^{\parallel B}\right) \parallel \cdots \parallel \text{OC}_K^n\left(s\Big|_B^{\parallel B}\right) \\ &= X\left(K, s\Big|_B^{\parallel B}\right) \parallel \cdots \parallel X\left(K, s\Big|_{n \times B}^{\parallel B}\right), \end{aligned}$$

for some function X . It follows then that

$$\begin{aligned} \text{OC}_K\left(s\Big|_{p \times B}^{\parallel B}\right) &= X\left(K, s\Big|_B^{\parallel B}\right) \parallel \cdots \parallel X\left(K, s\Big|_{p \times B}^{\parallel B}\right) \\ &= \text{OC}_K\left(s\Big|_{p \times B}^{\parallel B}\right). \end{aligned}$$

Now, to encrypt the ℓ th bit of a matryoshka filter, we use a pad constructed as

$$\text{R}\left(\text{OC}_K\left(\ell\Big|_{p \times B}^{\parallel B}\right), \text{id}(S)\right),$$

where R is a random oracle and $\text{id}(S)$ is the identifier of the filter. Note that the pad is both filter- and location-dependent. In addition, if the server is provided the value $\text{OC}_K(\ell^{\parallel B})$ for the maximal filter, it follows by Eq. (7) that it can derive the above pad as

$$\text{R}\left(\text{OC}_K(\ell^{\parallel B})\Big|_{p \times B}, \text{id}(S)\right).$$

The detailed description of our set encryption scheme is given in Fig. 4. In the Theorem below we show that it is adaptively-secure with the following leakage profile:

$$\mathcal{L}_S^{\text{est}}(S) = \#S \quad \text{and} \quad \mathcal{L}_Q^{\text{set}}(S_1, \dots, S_n, q) = (b_1, \dots, b_n, \text{SP}(q)),$$

where SP is the *search pattern*; that is, if and when two queries are the same. More formally, if t queries have been made, SP(q) outputs a t -bit string with a 1 at location i if q is equal to the i th query.

Theorem 4. *If OC is secure, then the multi-structure set encryption scheme described in Fig. 4 is adaptively $(\mathcal{L}_S^{\text{set}}, \mathcal{L}_Q^{\text{set}})$ -secure in the random oracle model.*

The proof of Theorem 4 is in the full version of the paper.

7.2 The ZMF Multi-Map Encryption Scheme

By applying the Z-IDX transformation to our multi-structure set encryption scheme from Fig. 4, we get a new adaptively-secure multi-map encryption scheme we call ZMF. We state its security formally in the following Corollary of Theorems 3 and 4. Its leakage profile is,

$$\mathcal{L}_S^{\text{zmf}}(\text{MM}) = \left(\left(\# \text{MM}^{-1}[v] \right)_{v \in \mathbf{V}}, \# \mathbf{V} \right) \quad \text{and} \quad \mathcal{L}_Q^{\text{zmf}}(\text{MM}, q) = (b_1, \dots, b_{\# \mathbf{V}}, \text{SP}(q))$$

where b_i is 1 if $q \in \text{MM}^{-1}[v_i]$ and 0 otherwise, and v_i is the i th value in \mathbf{V} .

Corollary 1. *The ZMF multi-map encryption scheme which results from applying the Z-IDX transformation to the set encryption scheme of Fig. 4 is $(\mathcal{L}_S^{\text{zmf}}, \mathcal{L}_Q^{\text{zmf}})$ -adaptively secure.*

8 DIEX: A Dynamic SSE Scheme

We describe our dynamic SSE construction DIEX. As far as we know, it is the first adaptively-secure dynamic SSE scheme that is forward-secure and supports Boolean search queries in sub-linear time. In particular, it supports the addition, deletion and editing of files. In the full version of the paper, we recall the syntax and security definitions for dynamic STE.

Overview. As a starting point, we describe a dynamic version of IEX that is *not* forward-secure. For this, we make two changes to our static construction. First, we replace the encrypted dictionary EDX and the global encrypted multi-map EMM_g with a dynamic encrypted dictionary EDX^+ and a dynamic global encrypted multi-map EMM_g^+ . The encrypted local multi-maps remain static. Second, we require that these new structures be *response-hiding*. We provide a high level description of our construction which is described in detail in Fig. 5.

The DIEX.Setup algorithm is the same as the IEX.Setup with the exception that it uses a dynamic encrypted dictionary and a dynamic encrypted multi-map and outputs state information st . The DIEX.Token^{sr} algorithm is similar

Let F be a pseudo-random function family, $\mathcal{H} : \{0, 1\}^* \rightarrow [\sigma]$ be a family of hash functions modeled as random oracles where σ is a public upper bound, and $\mathbf{R} : \{0, 1\}^* \rightarrow \{0, 1\}$ be a random oracle. Let $\mathbf{OC} : \{0, 1\}^{g(k)} \times \{0, 1\}^{\gamma \times B} \rightarrow \{0, 1\}^{\gamma \times B}$ a B -online cipher with $\gamma = \log \sigma$ blocks. Let $\varepsilon \in [0, 1]$ be a false positive rate that is hardcoded in each algorithm. Set $\lambda = \log(1/\varepsilon)$ and set $\mathbf{H}_1, \dots, \mathbf{H}_\lambda \leftarrow \mathcal{H}$. Consider the set encryption scheme $\Sigma = (\text{Gen}, \text{Enc}, \text{Token}, \text{Test})$ defined as follows:

- **Gen**(1^k):
 1. sample $K_1 \xleftarrow{\$} \{0, 1\}^k$ and $K_2 \xleftarrow{\$} \{0, 1\}^{g(k)}$;
 2. output $K = (K_1, K_2)$;
- **Enc**(K, S):
 1. let \mathbf{A} be a binary array of size $m = \lceil \lambda \cdot \#S / \ln 2 \rceil$ initialized to all 0's;
 2. for all items $a \in S$ and all $i \in [\lambda]$,
 - (a) compute $T = F_{K_1}(a)$;
 - (b) compute $\ell = \mathbf{H}_i(T)$;
 - (c) compute $s = \mathbf{OC}_{K_2}(\ell_{|\log m \times B}^{\parallel B})$;
 - (d) set $\mathbf{A}[\ell_{|\log m}] = 1 \oplus \mathbf{R}(s, \text{id}(S))$;
 3. for all $i \in [m]$ such that $\mathbf{A}[i] = 0$,
 - (a) compute $s = \mathbf{OC}_{K_2}(i_{|\log m \times B}^{\parallel B})$;
 - (b) set $\mathbf{A}[i] = 0 \oplus \mathbf{R}(s, \text{id}(S))$;
 4. set $\mathbf{ESET} = \mathbf{A}$;
 5. output \mathbf{ESET} .
- **Token**(K, a):
 1. compute $T = F_{K_1}(a)$;
 2. for all $i \in [\lambda]$,
 - (a) compute $\ell = \mathbf{H}_i(T)$;
 - (b) compute $s_i = \mathbf{OC}_{K_2}(\ell^{\parallel B})$;
 3. output $\text{tk} = (T, s_1, \dots, s_\lambda)$.
- **Test**(\mathbf{ESET}, tk):
 1. parse tk as $(T, s_1, \dots, s_\lambda)$;
 2. parse \mathbf{ESET} as \mathbf{A} ;
 3. set $m = |\mathbf{A}|$;
 4. for all $i \in [\lambda]$,
 - (a) compute $b_i = \mathbf{A}[\mathbf{H}_i(T)_{|\log m}] \oplus \mathbf{R}\left((s_i)_{|\log m \times B}, \text{id}(\mathbf{ESET})\right)$;
 5. if, for all $i \in [\lambda]$, $b_i = 1$ output 1, otherwise output 0.

Fig. 4. An adaptively-secure multi-structure set encryption scheme.

to IEX.Token with the exception that it is stateful. Here, the state is just used to generate tokens for the underlying dynamic dictionary and global multi-map encrypted structures. The DIEX.Token^{up} algorithm works as follows. It takes as inputs the key K , the state st and an update $u = (\text{op}, \text{id}, W_{\text{id}})$ that consists of an operation $\text{op} \in \{\text{edit}^+, \text{edit}^-\}$, the document identifier id being edited and a set of keywords W_{id} to add or delete based on op . We have the following cases:

- if $u = (\text{edit}^+, \text{id}, W_{\text{id}})$, the client will update the global multi-map EMM_g with pairs $(w, \text{tag}_{\text{id}})$ for all $w \in W_{\text{id}}$. Here, $\text{tag}_{\text{id}} := \text{Enc}_{K_1}(\text{id}; F_{K_2}(\text{id}||w))$ as in IEX. This is done by generating update tokens $(\text{utk}_g^w)_{w \in W_{\text{id}}}$ for EMM_g using $\Sigma_{\text{MM}}.\text{Token}^{\text{up}}$. For all $w \in W_{\text{id}}$, the client generates a new local multi-map MM_w that maps all $v \in W_{\text{id}} \setminus \{w\}$ to tag_{id} . It encrypts all these local multi-maps $(\text{MM}_w)_{w \in W_{\text{id}}}$ with $\Sigma_{\text{DX}}.\text{Setup}$, resulting in $(\text{EMM}_w)_{w \in W_{\text{id}}}$ and creates update tokens $(\text{utk}_d^w)_{w \in W_{\text{id}}}$ for EDX. The algorithm outputs an update token

$$\text{utk} = \left(\text{op}, (\text{utk}_d^w)_{w \in W_{\text{id}}}, (\text{utk}_g^w)_{w \in W_{\text{id}}} \right),$$

and $st = (st_d, st_g)$, where the former is the state maintained by Σ_{DX} and the latter is the state maintained by Σ_{MM} .

- if $u = (\text{edit}^-, \text{id}, W_{\text{id}})$, the client only updates EMM_g . Specifically, it removes all pairs $(w, \text{tag}_{\text{id}})$ for $w \in W_{\text{id}}$. This can be done by computing tags as above and generating update tokens $(\text{utk}_g^w)_{w \in W_{\text{id}}}$ using $\Sigma_{\text{MM}}.\text{Token}^{\text{up}}$. The algorithm outputs the update token

$$\text{utk} = \left(\text{op}, (\text{utk}_g^w)_{w \in W_{\text{id}}} \right),$$

and $st = st_g$ where st_g is the state maintained by Σ_{MM} .

The Update algorithm takes as input EDB and an update token utk and outputs EDB'. If $\text{op} = \text{edit}^+$, it uses the sub-tokens in utk to update EMM_g and EDX. If $\text{op} = \text{edit}^-$, it only updates EMM_g . The Search algorithm is the same as IEX.Search. Recall that we do not update the local multi-maps already in EDX. This is not necessary to for correctness because, during search, the server will take the intersection of the tags returned from the global multi-map EMM_g and from the appropriate local multi-maps. However, because EMM_g is properly updated, the intersection operation will filter out the old/stale tags from the local multi-map.

Forward security. We note that DIEX is forward secure if its underlying structures are. Specifically, if Σ_{MM} and Σ_{DX} are forward secure then so is DIEX. This is easy to see from the fact the DIEX tokens only consist of Σ_{DX} and Σ_{MM} tokens so if the former can be simulated from the security parameter, then the latter can. Due to space constraints, the definition of forward security is deferred to the full version of the paper. As a possible instantiation of a forward secure multi-map and dictionary encryption scheme, one can use the Sophos scheme of Bost [9].

Let $\Sigma_{\text{DX}}^+ = (\text{Setup}, \text{Token}^{\text{sr}}, \text{Get}, \text{Token}^{\text{up}}, \text{Update})$ and $\Sigma_{\text{MM}}^+ = (\text{Setup}, \text{Token}^{\text{sr}}, \text{Get}, \text{Token}^{\text{up}}, \text{Update})$ be dynamic dictionary and multi-map encryption schemes, respectively. Let $\text{IEX}^+ = (\text{Setup}^+, \text{Token}^+, \text{Search}^+)$ be the IEX scheme described in Fig. 1 with Σ_{MM} and Σ_{DX} replaced with Σ_{MM}^+ and Σ_{DX}^+ , respectively, and let $\Sigma_{\text{MM}} = (\text{Setup}, \text{Token}, \text{Query})$ be the static multi-map encryption scheme used to encrypt the local multi-maps. Consider the dynamic disjunctive SSE scheme $\text{DIEX} = (\text{Setup}, \text{Token}^{\text{sr}}, \text{Search}, \text{Token}^{\text{up}}, \text{Update})$ defined as follows:

- $\text{Setup}(1^k, \text{DB})$: output $(K, st, \text{EDB}) \leftarrow \text{IEX}^+. \text{Setup}(1^k, \text{DB})$;
- $\text{Token}^{\text{sr}}(K, \mathbf{w})$: output $\text{tk} \leftarrow \text{IEX}^+. \text{Token}(K, st, \mathbf{w})$;
- $\text{Token}^{\text{up}}(K, st, u)$
 1. parse u as $(\text{op}, \text{id}, \text{W}_{\text{id}})$ and st as (st_g, st_d)
 2. if $\text{op} = \text{edit}^+$,
 - (a) for all $w \in \text{W}_{\text{id}}$,
 - i. let $\text{tag}_{\text{id}} := \text{Enc}_{K_1}(\text{id}; F_{K_2}(\text{id}||w))$;
 - ii. compute $(\text{utk}_g^w, st_g) \leftarrow \Sigma_{\text{MM}}^+. \text{Token}^{\text{up}}(K, st_g, (\text{op}, w, \text{tag}_{\text{id}}))$;
 - iii. initialize a multi-map MM_w of size $\#\text{W}_{\text{id}}$;
 - iv. for all $v \in \text{W}_{\text{id}} \setminus \{w\}$, set $\text{MM}_w[v] = \text{tag}_{\text{id}}$;
 - v. compute $(K_w, \text{EMM}_w) \leftarrow \Sigma_{\text{MM}}. \text{Setup}(1^k, \text{MM}_w)$;
 - vi. compute $(\text{utk}_d^w, st_d) \leftarrow \Sigma_{\text{MM}}^+. \text{Token}^{\text{up}}(K, st_d, (\text{op}, w, \text{EMM}_w))$;
 - (b) output $\text{utk} = (\text{op}, (\text{utk}_d^w)_{w \in \text{W}_{\text{id}}}, (\text{utk}_g^w)_{w \in \text{W}_{\text{id}}})$;
 3. if $\text{op} = \text{edit}^-$,
 - (a) for all $w \in \text{W}_{\text{id}}$
 - i. let $\text{tag}_{\text{id}} := \text{Enc}_{K_1}(\text{id}; F_{K_2}(\text{id}||w))$;
 - ii. compute $(\text{utk}_g^w, st_g) \leftarrow \Sigma_{\text{MM}}^+. \text{Token}^{\text{up}}(K, st_g, (\text{op}, \text{W}_{\text{id}}, \text{tag}_{\text{id}}))$;
 - (b) output $\text{utk} = (\text{op}, (\text{utk}_g^w)_{w \in \text{W}_{\text{id}}})$ and the updated state $st = (st_g, st_d)$;
- $\text{Update}(\text{EDB}, \text{utk})$
 1. parse utk as $(\text{op}, (\text{utk}_i)_{i \in [\#\text{tk}]})$ and $\text{EDB} = (\text{EDX}, \text{EMM}_g)$;
 2. if $\text{op} = \text{edit}^-$, then for all $i \in [\#\text{utk}]$ compute $\text{EMM}_g \leftarrow \Sigma_{\text{MM}}^+. \text{Update}(\text{EMM}_g, \text{utk}_i, \text{op})$;
 3. if $\text{op} = \text{edit}^+$, then for all $i \in [\#\text{utk}/2]$, compute $\text{EMM}_g \leftarrow \Sigma_{\text{MM}}^+. \text{Update}(\text{EMM}_g, \text{utk}_i, \text{op})$ and $\text{EDX} \leftarrow \Sigma_{\text{DX}}^+. \text{Update}(\text{EDX}, \text{utk}_{i+\#\text{utk}/2+1}, \text{op})$;
 4. output $\text{EDB} = (\text{EDX}, \text{EMM}_g)$;
- $\text{Search}(\text{EDB}, \text{tk})$: output $\bigcup_{i \in [q]} T_i \leftarrow \text{IEX}^+. \text{Search}(\text{EDB}, \text{tk})$.

Fig. 5. The scheme DIEX.

Efficiency. The efficiency of DIEX depends on the underlying multi-map and dictionary encryption schemes. Using optimal constructions, the search complexity of DIEX is the same as IEX; that is, $O(q^2 \cdot M)$, where $M = \max_{i \in [q]} \#\text{DB}(w_i)$ and q is the number of terms in the disjunction.

Security. We show that DIEX is adaptively secure with respect to the following well-defined leakage profile. The setup and query leakages are the same as IEX so we only describe the update leakage. For an update $u = (\text{edit}^+, \text{id}, W_{\text{id}})$,

$$\mathcal{L}_U(\text{DB}, u) = \left(\mathcal{L}_U^{\text{mm}}(\text{MM}_g, (\text{op}, w, \text{id})), \mathcal{L}_U^{\text{dx}}(\text{DX}, (\text{op}, w, \text{id})), \mathcal{L}_S^{\text{mm}}(\text{MM}_w) \right)_{w \in W_{\text{id}}}.$$

$$\text{If } u = (\text{edit}^-, \text{id}, W_{\text{id}}): \mathcal{L}_U^{\text{diex}}(\text{DB}, u) = \left(\mathcal{L}_U^{\text{mm}}(\text{MM}_g, (\text{op}, w, \text{id})) \right)_{w \in W_{\text{id}}}.$$

Theorem 5. *If Σ_{DX} is adaptively $(\mathcal{L}_S^{\text{dx}}, \mathcal{L}_Q^{\text{dx}}, \mathcal{L}_U^{\text{dx}})$ -semantically secure and Σ_{MM} is adaptively $(\mathcal{L}_S^{\text{mm}}, \mathcal{L}_Q^{\text{mm}}, \mathcal{L}_U^{\text{mm}})$ -secure, then DIEX is adaptively $(\mathcal{L}_S^{\text{diex}}, \mathcal{L}_Q^{\text{diex}}, \mathcal{L}_U^{\text{diex}})$ -secure.*

We defer the proof to the final version of this work.

9 Empirical Evaluation

Due to space limitations, the empirical evaluation of our constructions is deferred to the full version. Clusion [22], our new open source encrypted search framework, is publicly available. We evaluate IEX-2Lev and IEX-ZMF which are instantiations of IEX with 2Lev and ZMF, respectively. We also evaluate our Boolean scheme BIEX. Our experiments report setup time, search time, storage and token size for all our constructions.

References

1. Apache lucene. <http://lucene.apache.org>.
2. E. Andreeva, A. Bogdanov, A. Luykx, B. Mennink, E. Tischhauser, and K. Yasuda. Parallelizable and authenticated online ciphers. In *ASIACRYPT*, 2013.
3. G. Asharov, M. Naor, G. Segev, and I. Shahaf. Searchable symmetric encryption: Optimal locality in linear space via two-dimensional balanced allocations. In *STOC*, 2016.
4. M. Bellare, A. Boldyreva, L. R. Knudsen, and C. Namprempre. On-line ciphers and the hash-cbc constructions. *IACR Cryptology ePrint Archive*, 2007:197, 2007.
5. M. Bellare, A. Boldyreva, and A. O’Neill. Deterministic and efficiently searchable encryption. In *CRYPTO*, 2007.
6. A. Boldyreva, N. Chenette, Y. Lee, and A. O’neill. Order-preserving symmetric encryption. In *EUROCRYPT*, 2009.
7. D. Boneh, G. di Crescenzo, R. Ostrovsky, and G. Persiano. Public key encryption with keyword search. In *EUROCRYPT*, 2004.
8. D. Boneh, A. Sahai, and B. Waters. Functional encryption: Definitions and challenges. In *TCC*, 2011.
9. R. Bost. Sophos - forward secure searchable encryption. In *ACM CCS*, 2016.

10. 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*, 2014.
11. D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *CRYPTO*, 2013.
12. D. Cash and S. Tessaro. The locality of searchable symmetric encryption. In *EUROCRYPT*, 2014.
13. M. Chase and S. Kamara. Structured encryption and controlled disclosure. In *ASIACRYPT*, 2010.
14. M. Chase and S. Kamara. Structured encryption and controlled disclosure. Technical Report 2011/010.pdf, IACR Cryptology ePrint Archive, 2010.
15. R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky. Searchable symmetric encryption: Improved definitions and efficient constructions. In *ACM CCS*, 2006.
16. S. Faber, S. Jarecki, H. Krawczyk, Q. Nguyen, M. Rosu, and M. Steiner. Rich queries on encrypted data: Beyond exact matches. In *ESORICS*, 2015.
17. B. A. Fisch, B. Vo, F. Krell, A. Kumarasubramanian, V. Kolesnikov, T. Malkin, and S. M. Bellovin. Malicious-client security in blind seer: A scalable private dbms. In *IEEE S&P*, 2015.
18. C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, 2009.
19. E.-J. Goh. Secure indexes. Technical Report 2003/216, IACR ePrint Cryptography Archive, 2003. See <http://eprint.iacr.org/2003/216>.
20. O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious RAMs. *Journal of the ACM*, 43(3):431–473, 1996.
21. Y. Ishai, E. Kushilevitz, S. Lu, and R. Ostrovsky. Private large-scale databases with distributed searchable symmetric encryption. In *CT-RSA*, 2016.
22. S. Kamara and T. Moataz. Clusion. <https://github.com/orochi89/Clusion>.
23. S. Kamara and C. Papamanthou. Parallel and dynamic searchable symmetric encryption. In *FC*, 2013.
24. S. Kamara, C. Papamanthou, and T. Roeder. Dynamic searchable symmetric encryption. In *ACM CCS*, 2012.
25. K. Kurosawa. Garbled searchable symmetric encryption. In *FC*, 2014.
26. K. Kurosawa and Y. Ohtaki. Uc-secure searchable symmetric encryption. In *FC*, 2012.
27. X. Meng, S. Kamara, K. Nissim, and G. Kollios. GRECS: graph encryption for approximate shortest distance queries. In *ACM CCS*, 2015.
28. M. Naveed, M. Prabhakaran, and C. Gunter. Dynamic searchable encryption via blind storage. In *IEEE S&P*, 2014.
29. A. O’Neill. Definitional issues in functional encryption, 2010. Cryptology ePrint Archive, Report 2010/556.
30. V. Pappas, F. Krell, B. Vo, V. Kolesnikov, T. Malkin, S.-G. Choi, W. George, A. Keromytis, and S. Bellovin. Blind seer: A scalable private dbms. In *IEEE S&P*, 2014.
31. D. Song, D. Wagner, and A. Perrig. Practical techniques for searching on encrypted data. In *IEEE S&P*, 2000.
32. E. Stefanov, C. Papamanthou, and E. Shi. Practical dynamic searchable encryption with small leakage. In *NDSS*, 2014.
33. A. Yao. Protocols for secure computations. In *FOCS*, 1982.
34. Y. Zhang, J. Katz, and C. Papamanthou. All your queries are belong to us: The power of file-injection attacks on searchable encryption. In *USENIX*, 2016.