

SEC: Symmetric Encrypted Computation via Fast Look-ups

Debadrita Talapatra Nimish Mishra Arnab Bag
IIT Kharagpur, India IIT Kharagpur, India IIT Kharagpur, India

Sikhar Patranabis Debdeep Mukhopadhyay
IBM Research, India IIT Kharagpur, India

July 26, 2024

Abstract

Encrypted computation allows a client to securely outsource the storage and processing of sensitive private data to an untrusted third party cloud server. Fully homomorphic encryption (FHE) allows computing arbitrary functions over encrypted data, but incurs huge overheads and does not practically scale to large databases. Whereas, slightly weaker yet efficient constructions- Searchable Symmetric Encryption (SSE)-support lookup-based evaluations of a *restricted* class of Boolean circuits over symmetrically encrypted data. In this paper, we investigate the use of SSE to *efficiently* perform *arbitrary* Boolean circuit evaluations over symmetrically encrypted data via look-ups.

To this end, in this work, we propose Symmetric Encrypted Computation (SEC): the first practically efficient and provably secure lookup-based construction, analogous to traditional FHE, that supports evaluation of arbitrary Boolean circuits over symmetrically encrypted data. SEC relies on purely symmetric-key cryptoprimitives and achieves flexible performance versus leakage trade-offs. SEC extends and generalizes the functional capabilities of SSE, while inheriting its data privacy guarantees and desirable performance benefits. We provide a concrete construction of SEC and analyze its security with respect to a rigorously defined and thoroughly analyzed leakage profile. We also present a prototype implementation of SEC and experimentally validate its practical efficiency. Our experiments show that SEC outperforms state-of-the-art FHE schemes (such as Torus FHE) substantially, with around $1000\times$ speed-up in basic Boolean gate evaluations. We further showcase the scalability of SEC for functions with multi-bit inputs via experiments performing encrypted evaluation of the entire AES-128 circuit, as well as three max-pooling layers of AlexNet architecture. For both sets of experiments, SEC outperforms state-of-the-art and accelerated FHE implementations by $1000\times$ in terms of processing time, while incurring $250\times$ lower storage.

Contents

1	Introduction	3
1.1	Our Contributions	4
1.2	Technical Overview	6
1.3	Related Work	12
2	Preliminaries and Background	12
2.1	Conjunctive SSE: Syntax and Security Model	13
2.2	Adaptive Security of CSSE	14
3	Symmetric Encrypted Computation	14
3.1	Syntax of SEC	15
3.2	SEC Construction	16
3.3	Proof of Correctness of SEC	19
3.4	Practical Instantiation of SEC	20
3.5	Complexity Analysis of SEC	21
4	Security and Leakage Profile Analysis of SEC	22
4.1	Leakage Profile of SEC_{OXT}	23
4.2	Analysis of Potential Leakages in SEC_{OXT}	23
4.3	Proof of Theorem 1	24
4.4	\mathcal{L}_{SEC} and Reusability of Look-up Table	27
4.5	Statistical Analysis of Leakage Due to Reusability	31
5	Experimental Results	32
6	Discussion	35

1 Introduction

Outsourced Storage. The substantial upswing in data production in today’s digitally-driven world has motivated the concept of outsourcing data to third-party cloud servers for storage. However, such *outsourced storage* solutions are often plagued by security incidents that lead to disclosure of client data [22, 46, 48]. Without specific privacy mechanisms, third-party cloud servers gain access to sensitive user data, thus leading to serious privacy concerns. This establishes the requirement of adopting secure and scalable privacy mechanisms for protecting sensitive outsourced data from unauthorized access. A straightforward solution to this problem is to encrypt this data before offloading to the third-party server, thereby ensuring data privacy and preventing disclosures. However, this leads to the challenge of securely computing queries (or more generally, executing functions/programs) directly on the encrypted data without decrypting it first.

“Secure Computation” on Outsourced Data. The question of privacy-preserving computation on encrypted, outsourced data has been studied extensively in the cryptographic literature. There exist elegant solutions such as Fully Homomorphic Encryption (FHE), Functional Encryption (FE), and Multi-party computation (MPC), all of which vary in terms of adversarial structure, communication models, and security guarantees. Classic FHE [26, 12, 32] works in the single client, single (adversarial) server setting and supports evaluating *any* (poly-sized) circuit directly over encrypted data, but the adversary does not learn anything about the data without the knowledge of the secret decryption key. Traditional FE [10, 39] operates in a similar setting, but offers the capability of more fine-grained query evaluation on encrypted data, while only leaking the output of the computation to an adversarial server. While significant improvements have been made to FHE schemes and implementations in recent years [11, 31, 30, 29, 28, 27, 3, 2, 19, 24, 36, 9, 1, 18, 20], FHE and FE solutions remain computationally expensive and do not scale efficiently to large datasets in practice. MPC operates in a different setting where the client “shares” its data across *multiple* servers, with the guarantee that these servers learn nothing about the client’s data apart from the output of the computation so long as the adversary does not corrupt more than a threshold number of parties. Certain MPC protocols are based on garbled circuits (GCs) [5, 34, 37], which allow hiding a circuit/program as long it is evaluated on only a single input. There exist practically efficient implementations of MPC and GCs [42, 6, 45, 44, 35], particularly in the setting where the adversary corrupts a minority of the parties (the “honest majority” setting).

Encrypted Computation via Table Lookups. In this paper, we focus on applications that adhere to the single server, single client setting of outsourced computation (unlike MPC). In addition, we consider applications where the program/circuit being evaluated on the encrypted data does not need to be private (unlike GCs). This is in line with most traditional FHE applications, where the focus is to maintain the privacy of the inputs to as well as the outputs from the function, against a (semi-honest¹/corrupt²) server. In this setting, we explore the possibility of replacing FHE-style evaluation of Boolean circuits

¹We follow the traditional semi-honest server model considered in FHE/SSE literature, where the server does not maliciously deviate from the protocol. Although standard techniques like zero-knowledge proofs could enforce semi-honest server behavior, this is currently outside the scope of this work.

²Verifiable FHE [25] considers malicious servers, but all known constructions are inefficient in practice.

with an alternative, *table look-up* based approach of evaluating Boolean gates over encrypted binary inputs. Concretely, we ask the following question:

Can we support arbitrary Boolean circuit evaluation over encrypted data via efficient table look-ups?

We answer this question in the affirmative, in the case of symmetrically encrypted outsourced data. We leverage existing approaches for evaluating restricted class of Boolean circuits over symmetrically encrypted structured databases (called Searchable Symmetric Encryption or SSE [23, 17, 14]) and show, for the first time, how to elevate/exploit such look-up based approaches to evaluate *arbitrary* Boolean circuits over encrypted Boolean inputs (with no additional structure whatsoever). We ensure reusability of the same look-up table for multiple gate evaluations without compromising the privacy of data being computed upon. Concretely, even though traditional SSE constructions incur certain non-trivial leakages (captured by a well defined leakage profile), our construction ensures such leakages do not reveal any information about the (encrypted) input/output bits to the adversarial server. Therefore, by allowing for practically viable fine-grained trade-offs between leakage and efficiency, we design a novel framework for symmetric encrypted computation that significantly outperforms its (symmetric-key) FHE-based counterparts in terms of computational efficiency and storage requirements. We validate the same via practical experiments.

1.1 Our Contributions

We introduce *Symmetric Encrypted Computation* (SEC) – a novel framework for practically efficient evaluation of arbitrary Boolean functions over symmetrically encrypted data. The technical centerpiece of SEC is a mapping of Boolean gate computations over encrypted Boolean inputs to look-ups over encrypted tables. We then show how to realize such encrypted lookup computations by leveraging existing practically efficient SSE schemes that support searching for conjunctive predicates over encrypted structured databases (several such schemes exist in the SSE literature, e.g., [14, 40]). Since SEC replaces certain algebraic operations that are inherent to any FHE scheme by table look-ups, it avoids many computationally expensive operations that limit the scalability of existing FHE solutions (most notably, bootstrapping). To the best of our knowledge, SEC is the *first* provably secure framework capable of arbitrary function evaluation over encrypted data using fast and efficient look-ups.

Overview of SEC. SEC relies on “encoding” Boolean functions as encrypted lookup tables for the universal Boolean gate set $\{\text{XOR}, \text{AND}, \text{OR}\}$. Then for any arbitrary Boolean circuit, SEC replaces explicit circuit computations with encrypted look-up operations by deploying a fast and encrypted search operations over the encrypted tables. For this, we leverage, in a black-box way, existing SSE techniques for fast conjunctive look-ups over such encrypted lookup tables for the primitive operations (which we model as “encrypted search indices” as in standard SSE scheme). As it turns out, this enables extremely fast circuit evaluation, surpassing the performance of the most efficient symmetric-key FHE schemes by several orders of magnitude.

We present two concrete instantiations of the above SEC framework based on two practically efficient conjunctive SSE constructions: Oblivious Cross Tags (OXT) [15] and CONJ-FILTER [40]. The former is chosen since it is the first practically efficient conjunctive SSE scheme to be proposed, while latter is chosen as it supports particularly fast searches based upon purely symmetric-key cryptoprimitives. We call the resulting schemes SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$, respectively.

Supporting Arbitrary Boolean Functions. A pivotal feature of SEC is its ability to evaluate function compositions in a single round of communication between the client and server (for one function evaluation). This allows decomposing any arbitrary function into an expression comprising composition of trivial functionality (i.e. $\{\text{XOR}, \text{AND}, \text{OR}\}$), which are then *solved* by issuing search queries to the underlying SSE scheme. Consequently, any arbitrary function can be easily evaluated using SEC with optimal computation overhead proportional to the circuit size. This is achieved by incurring a small amount of additional storage server-side; while the communication cost for the client (*amortized* over several circuit evaluations) is proportional to the function description length. Empirically, SEC’s storage requirement for realistic circuits is much less than that of the bootstrapping key in FHE schemes. Collectively all these features render SEC practically ideal for encrypted outsourced computation frameworks.

Security Analysis. SEC relies on an efficient, adaptively secure conjunctive SSE scheme for efficient encrypted lookup-based function evaluation. Thus, SEC inherits the adaptive security properties of the underlying SSE construction. However, the inherent design of SEC averts the direct extrapolation of several non-trivial leakages of the underlying SSE scheme. This in turn enables SEC to guarantee privacy of the input bits given to a function as well as the output of the function evaluation. *We emphasize, that the same look-up tables can be used for evaluation of multiple gates, while ensuring the adversary can infer no correlation between two gates in the same circuit, or two isomorphic gates in different circuits with same/different encrypted inputs.* We present a detailed security and leakage profile analysis of SEC_{OXT} following the security properties of the underlying conjunctive SSE scheme OXT³. Also we elucidate the improvements in leakage due to our improvised construction of the lookup tables in Section 4.

Performance And Scalability. We demonstrate the efficacy and scalability of SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$ by evaluating basic Boolean gates and cascaded gates as function composition. Section 5 gives a detailed analysis of our experimental evaluations. SEC decomposes any arbitrary (> 2 input-bits) circuit into universal binary gates XOR, AND, OR (similar to state-of-the-art FHE schemes) and reuses this storage across multiple computations of the same gate, thus preventing exponential storage-blowup. We showcase scalability of SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$ for functions with multi-bit inputs by using it for encrypted evaluation of the entire AES-128 circuit and three max-pooling layers of AlexNet architecture⁴. The experimental results involving these intricate circuit evaluations using SEC is demonstrated in Section 5. In Table 1 below, we show a practical use case of the AES SBox⁵,

³Analysis of SEC’s instantiations wrt. different conjunctive SSE constructions shall follow a similar security analysis based on real/ideal world simulation paradigm, as shown for SEC_{OXT} .

⁴KSH17 Imagenet classification with deep convolutional neural networks.

⁵For fair evaluation, we use *unparallelized* version of the SBox [38], which involves 5 XORs per bit in the output, thereby totaling 40 XORs for the entire byte.

wherein SEC outperforms various TFHE [20] backends by several orders of magnitude.

Table 1: Time taken (in seconds) and storage overhead (in MB) for evaluation of one byte AES SBox by SEC_{OXT} and SEC_{CONJFILTER} against different TFHE backends. Both FMA and AVX requires specialized execution units, and thus may not be executable on all hardware.

Scheme	Time taken (in seconds)	Storage (in MB)
TFHE-Nayuki AVX	14.85	24
TFHE-Nayuki Portable	22.862	24
TFHE-Spqlios AVX	4.21	24
TFHE-Spqlios FMA	2.57	24
SEC _{CONJFILTER} (This work)	0.1013	0.449
SEC _{OXT} (This work)	0.96	0.098

1.2 Technical Overview

Efficient “search” over Encrypted Lookup Tables. The first step towards constructing SEC is to create efficient mechanisms for encrypted table lookups, for which we rely upon Searchable Symmetric Encryption (SSE) schemes [43, 23, 17, 14]. The class of SSE constructions offers a *weaker* yet *efficient* set of capabilities than FHE: SSE schemes allow fast and efficient searches over symmetrically encrypted data. Although traditionally, SSE schemes have been used exclusively for searches, we however show that SSE data structures are amenable to design modifications such that *efficient searches can be used to compute on encrypted data*. To elaborate, any function computation can be modeled into an equivalent Boolean circuit composed of Boolean variables and universal basic logic gate set (XOR, AND, OR). Hence, ① creating encrypted lookup tables for these basic gates and ② using efficient *search* capabilities of SSE constructions over these tables, is equivalent to *computing* these logic gates over encrypted inputs.

We exemplify the end-to-end evaluation of XOR (AND and OR follow suit) using SEC. As aforementioned, the first step is to create a ① representative *encrypted lookup table* corresponding to XOR(x, y), as in Table 2 (where x and y are plaintext input bits to XOR). Using generic SSE notation, every plaintext bit (x or y) is *mapped* to alphanumeric strings (called *keywords* hereafter). Concretely, for $x = 1$, the corresponding keyword used by the underlying SSE scheme to *search* over Table 2 is \mathbf{w}_1 . Likewise, $x = 0$ is mapped to the keyword $\bar{\mathbf{w}}_1$. A similar mapping follows from input bit y to keywords \mathbf{w}_2 or $\bar{\mathbf{w}}_2$ depending on whether y is 1 or 0 respectively. The output of the evaluation XOR(x, y) is encrypted and stored as *documents* by the underlying SSE scheme, indexed by identifiers $\text{doc_id} (D_i) : i \in \{0, 1, 2, 3\}$. The second step is to use ② the efficient *search* capabilities of the underlying SSE scheme over a conjunctive predicate of *input* keywords, to retrieve an *output* encrypted document. This output document, upon decryption on client-side, gives a plaintext bit equal to the actual plaintext evaluation of XOR(x, y). We note that an additional keyword **KW 1** (hereafter mentioned as “special” term) is added to *all* documents in the encrypted lookup table. However, it does not play an explicit role in computation (since it is not mapped to input bits x or y). This allows reusing the same lookup table for computing the same/different circuits more than once without revealing any correlation for the adversary to infer between the computations. We elaborate on this subsequently.

Table 2: Contents of documents related to the functional evaluation of 2-bit XOR and mapping of document identifiers to their corresponding keywords. Here, Enc_k refers to any generic symmetric encryption scheme with secret key k .

KW 1	x	KW 2	y	KW 3	XOR(x, y)	doc_id	Doc. content
\mathbf{w}_d	0	$\bar{\mathbf{w}}_1$	0	$\bar{\mathbf{w}}_2$	0	D_0	$Enc_k(0)$
\mathbf{w}_d	0	$\bar{\mathbf{w}}_1$	1	\mathbf{w}_2	1	D_1	$Enc_k(1)$
\mathbf{w}_d	1	\mathbf{w}_1	0	$\bar{\mathbf{w}}_2$	1	D_2	$Enc_k(1)$
\mathbf{w}_d	1	\mathbf{w}_1	1	\mathbf{w}_2	0	D_3	$Enc_k(0)$

In order to map a conjunctive predicate of input keywords to an encrypted output document (i.e. step ② aforementioned), Table 2 needs to be converted into a SSE specific data structure (as shown in Table 3), which maps a given keyword to the doc_ids which is related to the keyword. The underlying SSE scheme then encrypts this mapping before offloading the same to the server, thereby constituting the encrypted lookup table for XOR. Moreover, in contrast to Table 2, we have added n “special” terms (for a circuit with n gates); this allows reusability of Table 3 over evaluations of *multiple* circuits of size n^6 . Also note that each *special* term is mapped to four documents, while the keywords participating in the computation ($\bar{\mathbf{w}}_1, \bar{\mathbf{w}}_2, \mathbf{w}_1, \mathbf{w}_2$) are mapped to two. Hence, we add *dummy* documents ($\{D'_0, \dots, D'_7\}$; each comprising of unique random alphanumeric values) to ensure that the frequency of all keywords remains same in the final encrypted database. This keeps the adversarial server’s view during searches over Table 3 uniform wrt. all keywords. Finally, we note that at no point in this entire process have we performed an explicit computation of XOR gate (unlike as done in FHE). The entire computation is completed by *searching* over encrypted lookup tables (search index) which is extremely fast and efficient.

Table 3: An inverted search index representation of the database (search index) for XOR function.

Keywords	doc_id
$\bar{\mathbf{w}}_1$	D_0, D_1, D'_0, D'_1
\mathbf{w}_1	D_2, D_3, D'_2, D'_3
$\bar{\mathbf{w}}_2$	D_0, D_2, D'_4, D'_5
\mathbf{w}_2	D_1, D_3, D'_6, D'_7
\mathbf{w}_d^1	D_0, D_1, D_2, D_3
\mathbf{w}_d^2	D_0, D_1, D_2, D_3
\vdots	\vdots
\mathbf{w}_d^n	D_0, D_1, D_2, D_3

Without loss of generality, assume $XOR(1, 1)$ is to be evaluated using SEC. Considering the offloaded encrypted database in Table 3, SEC computes a conjunctive query by choosing *any*⁷ “special” term (say \mathbf{w}_d^1), along with the correct keywords for plaintext input bits $x = 1$ and $y = 1$. Concretely, SEC translates $XOR(1, 1)$ to a conjunctive query $\mathbf{q} =$

⁶The lookup tables need only setup once (by the client), after which it is reused. This amortizes the client’s communication overhead to function description length (over multiple circuit evaluations).

⁷The choice of the “special” term is dependent on the underlying SSE scheme: generally in OXT or CONJFILTER the “special” term is selected according to the least frequent keyword/conjunct in a given

$(\mathbf{w}_d^1 \wedge \mathbf{w}_1 \wedge \mathbf{w}_2)$. We now note the sequence of operations that are consequently executed (assuming a black-box conjunctive SSE search), and draw parallels with conjunctive SSE terminology:

1. **Processing “special” Term:** Use the first keyword in \mathbf{q} (i.e. \mathbf{w}_d^1), generate an address to index into the encrypted lookup table, and retrieve an encrypted list of `doc.id` mapped to \mathbf{w}_d^1 . For consistency hereafter, we denote $\mathbf{sval}_{\mathbf{w}, D_j}$ as the retrieved encrypted output from this phase for some arbitrary keyword \mathbf{w} and document D_j . In our example, this phase shall return the encrypted list $[\mathbf{sval}_{\mathbf{w}_d^1, D_0}, \mathbf{sval}_{\mathbf{w}_d^1, D_1}, \mathbf{sval}_{\mathbf{w}_d^1, D_2}, \mathbf{sval}_{\mathbf{w}_d^1, D_3}]$.
2. **Processing “cross” Term.** In this phase, some auxiliary information⁸ dependent on the tuples $(\mathbf{w}_d^1, \mathbf{w}_1)$ and $(\mathbf{w}_d^1, \mathbf{w}_2)$ is used in conjunction with retrieved \mathbf{sval} to create a *search token*. For consistency, we denote $\mathbf{Token}_{\mathbf{w}_d^1, \mathbf{w}_i, D_j}$ to denote a search token generated upon combination of $\mathbf{sval}_{\mathbf{w}_d^1, D_j}$ and auxiliary information dependent on tuple $(\mathbf{w}_d^1, \mathbf{w}_i)$. Upon being indexed into SSE specific data structures, $\mathbf{Token}_{\mathbf{w}_d^1, \mathbf{w}_i, D_j}$ returns a binary decision on whether keyword \mathbf{w}_i is present in document D_j .
3. **Query Result.** Any document identifier which is included in the output of step ①, and for which *all* search tokens generated in step ② give a positive binary result, is included in the *result* of the query. Note that, for our example, only document D_3 contains the keywords \mathbf{w}_d^1 , \mathbf{w}_1 , and \mathbf{w}_2 . Hence, the output of the query $\mathbf{q} = (\mathbf{w}_d^1 \wedge \mathbf{w}_1 \wedge \mathbf{w}_2)$, for this example, is $\mathbf{sval}_{\mathbf{w}_d^1, D_3}$. This result is returned back to the client, which upon decryption obtains D_3 . It follows trivially that since $D_3 = \mathit{Enc}_k(0)$, query \mathbf{q} and associated *search* over the encrypted lookup table has effectively *computed* $\mathbf{XOR}(1, 1)$.

Functional Correctness. Concretely, we *design* the encrypted lookup table (search index) in a way such that for arbitrary encrypted input bits x and y , *exactly* one encrypted document is returned, which upon decryption reveals a single bit $b = \mathbf{XOR}(x, y)$, (thereby guaranteeing correctness).

“special” Terms and Reusability of Lookup Table for Multiple Evaluations. In the previous sequence of operations, the server’s view (informally) consists of *memory access patterns* wrt. $\mathbf{sval}_{\mathbf{w}_d, D_j}$ and $\mathbf{Token}_{\mathbf{w}_d, \mathbf{w}_i, D_j}$; note that the “special” term is shared commonly between the two. Hence, for multiple evaluations of the same gate, should the client choose a new “special” term, it obfuscates the *access patterns* of the ① index into the encrypted look-up table and ② retrieval of the respective encrypted documents. Thus even if the client computes the same function, say $\mathbf{XOR}(x, y)$ twice on the same encrypted input, the server (adversary) cannot infer any correlation between the two computations. This ensures the reusability of the same encrypted look-up table for multiple circuit evaluations without any significant information leakage to the semi-honest server. To achieve this, the client maintains a $O(1)$ -state that is used to *permute* the allotment of the n “special” terms to the

conjunctive query for efficiency purposes. In our instantiation of $\mathit{SEC}_{\mathbf{OXT}}$ and $\mathit{SEC}_{\mathbf{CONJFILTER}}$, it is set to the first keyword/conjunct (since the frequency of all keywords is the same).

⁸Abstracted in the underlying black-box SSE scheme used.

Table 4: Mapping between result of the inner function, to keywords used in querying the outer function. For the sake of clarity, we demonstrate this mapping in plaintext. In actuality, this table shall be encrypted before being offloaded.

Inner $f(\cdot)$ result	Outer $fn(\cdot)$ input order	Keywords for outer $f(\cdot)$	Mapped doc.id for outer $f(\cdot)$
D_0	First input (corr. bit $x = 0$)	$\bar{\mathbf{w}}_1$	D_0, D_1
D_3	First input (corr. bit $x = 0$)	$\bar{\mathbf{w}}_3$	D_0, D_1
D_1	First input (corr. bit $x = 1$)	\mathbf{w}_1	D_2, D_3
D_2	First input (corr. bit $x = 1$)	\mathbf{w}_3	D_2, D_3
D_0	Second input (corr. bit $y = 0$)	$\bar{\mathbf{w}}_2$	D_0, D_2
D_3	Second input (corr. bit $y = 0$)	$\bar{\mathbf{w}}_4$	D_0, D_2
D_1	Second input (corr. bit $y = 1$)	\mathbf{w}_2	D_1, D_3
D_2	Second input (corr. bit $y = 1$)	\mathbf{w}_4	D_1, D_3

different gates in a circuit. Reusability is then derived from such client-controlled “permutation” of assigning unique “special” term to isomorphic gates across multiple runs, ensuring unique access-pattern across multiple gate evaluations. Concretely, the upper bound on reusability is then the number of possible derangements, given by $n! - \sum_{i=0}^n \frac{(-1)^{i+1}}{i!} (O(n!))$ asymptotically). “special” terms are permuted using a Pseudo-Random Permutation primitive (which requires client-side $O(1)$ state). Note that since the same encrypted lookup table is *reused*, the *amortized* communication cost across multiple circuit evaluations is proportional to the function description length. Finally, the client does not require prior knowledge of circuit composition to set the “special” terms. More details follow in Section 4.

Composable Function Evaluation. Note that so far, we have assumed that a SEC query is constructed at the client-side. Extending this version of SEC to an n -depth circuit is non-trivial because of the need for repeated involvement of the client. Concretely, for SEC to evaluate compositions of form $f_k(f_i(\cdot, \cdot), f_j(\cdot, \cdot))$ (where $f_i, f_j, f_k \in \{\text{XOR}, \text{AND}, \text{OR}\}$), first evaluate $f_i(\cdot, \cdot), f_j(\cdot, \cdot)$, send the result to the client, who in turn constructs the query for the *outer* function, and then execute the outer function query. However, this incurs extra rounds of communication that scale linearly with the circuit size.

To circumvent this issue, we design SEC to perform query construction for the outer function on the (semi-honest) server itself. That is, instead of the need to decrypt the result (i.e. $\text{sval}_{\mathbf{w}_d, D_j}$) of the inner function to recover D_j , we use mechanisms to directly map $\text{sval}_{\mathbf{w}_d, D_j}$ to construct relevant query used by the outer function. To do so, we extend Table 2 to allow compositions of XOR (Table 4). The first major change is an increase in the number of keywords to 8: there are four possible outcomes for the inner XOR evaluation (denoted by plaintext document identifiers D_0, D_1, D_2, D_3), and each outcome can either behave as the first input (corresponding to bit x) or the second input (corresponding to bit y) for the outer function. For example, should the inner function output D_0 (corresponding to evaluation of $\text{XOR}(0, 0)$; Table 2), then for the outer function, either $x = 0$ or $y = 0$ depending on specific wiring topology of the these gates. Therefore, according to Table 4, either keyword $\bar{\mathbf{w}}_1$ or $\bar{\mathbf{w}}_2$ shall be involved in the outer function’s query. This mapping hence allows query construction of the outer function evaluation on the server side itself. We note that compositions of AND/OR, as well as intermixing of XOR/AND/OR follow suit.

Illustrative Example 1: Functional Composition. For this example, consider a client who wishes to evaluate the circuit of Figure 1 using SEC without involving mul-




Figure 1: Evaluation of isomorphic circuits multiple times using the same lookup table

multiple communication rounds with the client. SEC uses the *same* lookup table as in Table 4, that is encrypted and offloaded to the server *once* for evaluating all the gates at every depth of the circuit. For ensuring a single communication round between the client and server, the query construction for outer XOR needs to happen server-side. For now, we assume $\text{ConstructQuery}(\mathbf{w}_d^k, \text{sval}_{\mathbf{w}_d^k, D_j}, b)$ to abstract the following mapping: Given that $\text{sval}_{\mathbf{w}_d^k, D_j}$ is output by the inner function evaluation, ConstructQuery returns the search tokens of the form $\text{Token}_{\mathbf{w}_d^k, \mathbf{w}_i, D_j}$ depending on the mapping in Table 4. Bit b denotes whether $\text{sval}_{\mathbf{w}_d^k, D_j}$ is first or second input to the outer function. We defer details of the exact operation of ConstructQuery to Section 3.2. We also note that since *three* evaluations of XOR occur, we use a *distinct* “special” term for each. We now explain how the evaluation proceeds.

To compute, SEC constructs a query $\mathbf{q}_1 = (\mathbf{w}_d^1 \wedge \mathbf{w}_1 \wedge \bar{\mathbf{w}}_2)$ for $f_{\text{XOR}}(1, 0)$, resulting in $\text{sval}_{\mathbf{w}_d^1, D_1}$. Likewise, SEC constructs a query $\mathbf{q}_2 = (\mathbf{w}_d^2 \wedge \bar{\mathbf{w}}_1 \wedge \bar{\mathbf{w}}_2)$ for $f_{\text{XOR}}(0, 0)$, resulting in $\text{sval}_{\mathbf{w}_d^2, D_0}$. Note that the result of $f_{\text{XOR}}(1, 0)$ drives the first input of outer XOR. Thus, SEC invokes $\text{ConstructQuery}(\mathbf{w}_d^3, \text{sval}_{\mathbf{w}_d^1, D_1}, 1)$ to obtain $\text{Token}_{\mathbf{w}_d^3, \mathbf{w}_1, D_2}$ and $\text{Token}_{\mathbf{w}_d^3, \mathbf{w}_1, D_3}$ (corresponding to row entry 3 in Table 4). Likewise, since $f_{\text{XOR}}(0, 0)$ drives the second input of outer XOR, SEC invokes $\text{ConstructQuery}(\mathbf{w}_d^3, \text{sval}_{\mathbf{w}_d^2, D_0}, 0)$ to obtain $\text{Token}_{\mathbf{w}_d^3, \bar{\mathbf{w}}_2, D_0}$ and $\text{Token}_{\mathbf{w}_d^3, \bar{\mathbf{w}}_2, D_2}$ (corresponding to row 5 in Table 4). Overall, for the outer XOR, the constructed query is $\mathbf{q}_3 = (\mathbf{w}_d^3 \wedge \mathbf{w}_1 \wedge \bar{\mathbf{w}}_2)$. It is straightforward to see that the result of the functional composition, by design of Table 4 shall be $\text{sval}_{\mathbf{w}_d^3, D_2}$. The server sends this *final* result to the client. Upon decryption, the client obtains the output plaintext bit 1, which is the correct evaluation of the example. From the adversarial server’s perspective, since the “special” term changes across all gate evaluations, the corresponding access patterns are different for all gates, preventing any correlation across evaluations.

Illustrative Example 2: Reusability and Leakage. We consider a slightly more complex evaluation in this example, where we assume that the client wishes to evaluate the circuit twice. Assume that for Evaluation 1, SEC constructs a query similar to the previous example. For Evaluation 2, the client assigns a different permutation of “special” terms and SEC constructs corresponding query as: $\mathbf{q}_1 = (\mathbf{w}_d^3 \wedge \mathbf{w}_1 \wedge \bar{\mathbf{w}}_2)$ for $f_{\text{XOR}}(1, 0)$, resulting in $\text{sval}_{\mathbf{w}_d^3, D_1}$. Likewise, a query \mathbf{q}_2 is constructed as $(\mathbf{w}_d^1 \wedge \bar{\mathbf{w}}_1 \wedge \bar{\mathbf{w}}_2)$ for $f_{\text{XOR}}(0, 0)$, resulting in $\text{sval}_{\mathbf{w}_d^1, D_0}$. As the result of $f_{\text{XOR}}(1, 0)$ drives the first input of outer XOR, SEC invokes $\text{ConstructQuery}(\mathbf{w}_d^2, \text{sval}_{\mathbf{w}_d^3, D_1}, 1)$ to obtain $\text{Token}_{\mathbf{w}_d^2, \mathbf{w}_1, D_2}$ and $\text{Token}_{\mathbf{w}_d^2, \mathbf{w}_1, D_3}$ (cor-

responding to row entry 3 in Table 4). Likewise, since $f_{\text{XOR}}(0,0)$ drives the second input of outer XOR, SEC invokes $\text{ConstructQuery}(\mathbf{w}_d^2, \text{sval}_{\mathbf{w}_d^1, D_0}, 0)$ to obtain $\text{Token}_{\mathbf{w}_d^2, \bar{\mathbf{w}}_2, D_0}$ and $\text{Token}_{\mathbf{w}_d^2, \bar{\mathbf{w}}_2, D_2}$ (corresponding to row 5 in Table 4). Overall, for the outer XOR, the constructed query is $\mathbf{q}_3 = (\mathbf{w}_d^2 \wedge \mathbf{w}_1 \wedge \bar{\mathbf{w}}_2)$, and the final result is $\text{sval}_{\mathbf{w}_d^2, D_2}$. Since the “special” terms do not contribute in the actual computation, even in different permutations of “special” terms, the result of the final evaluation will be the same (this ensures functional correctness).

From the server’s perspective (in Figure 1), for gate 1, the server observes access patterns related to $\{\text{sval}_{\mathbf{w}_d^1, D_1}\}$ in Evaluation 1 and $\{\text{sval}_{\mathbf{w}_d^3, D_1}\}$ in Evaluation 2. Similarly, for gate 2, access patterns observed are $\{\text{sval}_{\mathbf{w}_d^2, D_0}\}$ in Evaluation 1 and $\{\text{sval}_{\mathbf{w}_d^1, D_0}\}$ in Evaluation 2. Likewise, for gate 3, the server observes access patterns for $\{\text{Token}_{\mathbf{w}_d^3, \mathbf{w}_1, D_2}, \text{Token}_{\mathbf{w}_d^3, \mathbf{w}_1, D_3}, \text{Token}_{\mathbf{w}_d^3, \bar{\mathbf{w}}_2, D_0}, \text{Token}_{\mathbf{w}_d^3, \bar{\mathbf{w}}_2, D_2}, \text{sval}_{\mathbf{w}_d^3, D_2}\}$ in Evaluation 1 and $\{\text{Token}_{\mathbf{w}_d^2, \mathbf{w}_1, D_2}, \text{Token}_{\mathbf{w}_d^2, \mathbf{w}_1, D_3}, \text{Token}_{\mathbf{w}_d^2, \bar{\mathbf{w}}_2, D_0}, \text{Token}_{\mathbf{w}_d^2, \bar{\mathbf{w}}_2, D_2}, \text{sval}_{\mathbf{w}_d^2, D_2}\}$ in Evaluation 2. Briefly, through client-controlled permutation, assignment of “special” terms to *non-isomorphic* gates⁹ prevents the server from correlating between two (similar/different) circuit evaluations, thereby making the reusability of the same lookup table for multiple evaluations secure, in terms of data privacy.

Note. It is important to note that, for ease of exposition we exemplified the working mechanism of SEC for computing a circuit with XOR gates only. An exactly similar execution methodology can be used for evaluating any arbitrary circuit with any combination of gates, $f \in \{\text{XOR}, \text{AND}, \text{OR}\}$, since both sval and Token are function agnostic.

Communication Complexity and Security in Reusability. We stress that the same look-up tables can be used for evaluation of multiple circuits. We do this through client-controlled *permutation* of the pool of “special” terms available in SEC’s encrypted look-up tables. Like in the aforementioned example, the client assigns a permutation of “special” terms $\mathbf{w}_d^1, \mathbf{w}_d^2, \mathbf{w}_d^3$ to the three gates respectively for Evaluation 1. For Evaluation 2, however, the client chooses a *new* permutation: $\mathbf{w}_d^3, \mathbf{w}_d^1, \mathbf{w}_d^2$. This ensures *reusability* of SEC’s look-up tables across different circuits, while ensuring ① no two gates in the *same* circuit share a “special” term, and ② two isomorphic gates in *different* circuits also do not share corresponding “special” terms, thereby preventing the server from inferring any correlation between two computations. Hence, the client’s *amortized* communication overhead is proportional to the function description length (i.e. for a circuit evaluation with n gates). Post the *one-time* setup, the client uses derangements of $\{\mathbf{w}_d^1, \mathbf{w}_d^2, \dots, \mathbf{w}_d^n\}$ to evaluate subsequent circuit (assuming there are n gates in the circuit). Asymptotically, the number of derangements (because of the need to avoid isomorphic gate assignments) is $\left\{n! - \sum_{i=0}^n \frac{(-1)^{i+1}}{i!}\right\} = O(n!)$. We elaborately discuss and analyze the leakage profile of SEC in Section 4.

⁹If in Evaluation 2, the client had assigned \mathbf{w}_d^2 to gate 2, then this assignment is isomorphic to that in Evaluation 1. Consequently, the server’s leakage profile for gate 2 is $\{\text{sval}_{\mathbf{w}_d^2, D_0}\}$ for both evaluations, hence leaking that both evaluations have the *same type of gate*. Note that, however, neither the gate description (i.e. whether it is AND/OR/XOR) nor the exact plaintext bit in $\{\text{sval}_{\mathbf{w}_d^2, D_0}\}$ is leaked, since this information is encrypted, thereby still protecting data privacy.

1.3 Related Work

Garbled Circuits. Garbled circuits [5, 34, 37] typically provide privacy for the input data and entire circuit that is being computed. However, such constructions have the disadvantage of not being reusable. As such, there have been attempts to construct reusable Garbled Circuits [34], but the underlying primitive used is Functional Encryption (using FHE as a black box) which is computationally intensive. SEC differs from GCs in the sense that it guarantees input/output privacy, but not circuit privacy.

Fully Homomorphic Encryption (FHE). In practice, computing over encrypted out-sourced databases has used sophisticated and highly structured but typically expensive primitives, such as Fully Homomorphic Encryption (FHE) [26, 30, 29, 28, 32, 3, 2, 19] which views generic computations as either arithmetic or Boolean circuits, and provides an all-or-nothing flavor of security. FHE has recently gained traction in the cryptographic literature for privacy-preserving computation with rich functionalities along with the *ideal* notion of privacy. However, as each primitive function evaluation is realized by explicit circuit evaluation followed by an expensive *bootstrapping* operation. FHE has prohibitively high computation costs and storage overheads. In this work, we find a solution to bypassing explicit circuit evaluation and realizing any generic function by efficient encrypted look-up operations.

Searchable Symmetric Encryption (SSE). SSE schemes [43, 23, 17, 14] provision users with *search* capabilities over symmetrically encrypted data. There exist today efficient SSE schemes that support conjunctive (and more general Boolean) queries [14, 13, 41]. An *ideal* SSE construction using Oblivious RAM (ORAM) promises oblivious memory access patterns (and hence *no* leakage) but is hard to actualize in hardware, is closed-source, and has not been tested against scaled databases [16]. Modern SSE schemes thus trade-off security for efficiency. These schemes allow the server to learn “some” information during query execution, detailed by their *leakage* profile. While SSE schemes are extremely fast and highly scalable with arbitrarily large real-world datasets, their restricted functionality (to only *search*) renders them inapt for practical deployment in an encrypted computation framework. In this work, we leverage the efficient look-up capabilities of SSE while extending the limited functionality of existing SSE schemes to supporting encrypted computations of arbitrary Boolean functions.

2 Preliminaries and Background

We present preliminary concepts and background in this section. Table 5 lists basic notations used in this paper. Any other notation used is defined in-place within the context of the main text.

Table 5: Summary of notations

Notations	Meaning
λ	security parameter
id/doc_id	document identifier
\mathbf{w}	a keyword
\mathcal{W}	dictionary of keywords $\mathcal{W} = \{\mathbf{w}_1, \dots, \mathbf{w}_N\}$
\mathbf{DB}	database $(\text{id}_j, \mathbf{w}_i)_{i=1}^N _{j=1}^{ \mathbf{DB}(\mathbf{w}) } \in \mathbf{DB}$
$\mathbf{DB}(\mathbf{w})$	all documents containing \mathbf{w}
n	max. number of keywords per conjunctive query.
$x \stackrel{\$}{\leftarrow} \chi$	uniformly sampling x from χ
$x = \mathcal{A}$	x is output of a deterministic algorithm
$x \leftarrow \mathcal{A}'$	x is output of a randomized algorithm

2.1 Conjunctive SSE: Syntax and Security Model

A *Conjunctive Searchable Symmetric Encryption* scheme (CSSE) provisions the client with *conjunctive* search capability (i.e. search Boolean queries of the form $\mathbf{w}_1 \wedge \mathbf{w}_2 \wedge \dots \wedge \mathbf{w}_n$) over an encrypted database. A CSSE scheme can be formally defined as an ensemble of four polynomial-time algorithms $\{\text{KEYGEN}, \text{ENCRYPT}, \text{GENTOKEN}, \text{SEARCH}\}$ ¹⁰ such that:

- $\text{KEYGEN}(\lambda)$ is a probabilistic algorithm that takes the security parameter λ as input. The output of this algorithm is the client’s secret key sk .
- $\text{ENCRYPT}(\text{sk}, \mathbf{DB})$ is a probabilistic algorithm that takes as input the client secret key sk and plain database \mathbf{DB} . The output is encrypted database \mathbf{EDB} .
- $\text{GENTOKEN}(\text{sk}, q = \mathbf{w}_1 \wedge \dots \wedge \mathbf{w}_n)$ is a deterministic algorithm executed by the client that takes as input secret key sk and a conjunctive query $q = \mathbf{w}_1 \wedge \dots \wedge \mathbf{w}_n$. It generates *search tokens* (st_q) corresponding to the conjunctive query q as the output.
- $\text{SEARCH}(\mathbf{EDB}, \text{st}_q)$ is a deterministic algorithm executed by the *server* that takes as input \mathbf{EDB} and the search token st_q corresponding to a conjunctive query q . It returns the encrypted document identifiers $\mathbf{DB}(\text{st}_q)$ corresponding to the conjunction $q = \mathbf{w}_1 \wedge \dots \wedge \mathbf{w}_n$ as output.

Correctness. A CSSE scheme is said to be correct if for an \mathbf{EDB} generated from a \mathbf{DB} using ENCRYPT , for a search token st_q generated by GENTOKEN from any conjunctive Boolean query q formed over the keywords \mathbf{w}_i in \mathcal{W} , the SEARCH routine returns a set of ids as result which is the same as $\mathbf{DB}(q)$ with high probability.

Security. The security of a CSSE scheme is parameterized by a *leakage function* \mathcal{L} , which encapsulates the information that can be learnt (potentially by an adversary) from the encrypted database and query transcripts. Formally, the security notion says that the server’s view during an adaptive attack (where the server selects the database and queries) can be simulated given only the output of \mathcal{L} .

Let $\text{CSSE} = \{\text{KEYGEN}, \text{ENCRYPT}, \text{GENTOKEN}, \text{SEARCH}\}$ be a CSSE scheme, and let

¹⁰Our syntax for a conjunctive SSE is different from the syntax of traditional conjunctive SSE scheme $\{\text{SETUP}, \text{GENTOKEN}, \text{SEARCH}\}$, but the underlying functionality is exactly similar

\mathcal{L} be a stateful algorithm. For algorithms \mathcal{A} (denoting the adversary) and SIM (denoting a simulator), we define the experiments (algorithms) $\mathbf{Real}_A^{\text{CSSE}}(\lambda)$ and $\mathbf{Ideal}_{A,\text{SIM}}^{\text{CSSE}}(\lambda)$, as in Algorithm 1 and Algorithm 2, respectively (see Section 2.2). We say that CSSE is \mathcal{L} -semantically-secure against adaptive attacks if for all adversaries \mathcal{A} there exists an algorithm SIM such that

$$|\Pr[\mathbf{Real}_A^{\text{CSSE}}(\lambda) = 1] - \Pr[\mathbf{Ideal}_{A,\text{SIM}}^{\text{CSSE}}(\lambda) = 1]| \leq \text{negl}(\lambda).$$

In these experiments, the leakage function for CSSE is expressed as

$$\mathcal{L}_{\text{CSSE}} = (\mathcal{L}_{\text{CSSE}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{CSSE}}^{\text{GENTOKEN}}, \mathcal{L}_{\text{CSSE}}^{\text{SEARCH}}),$$

where $\mathcal{L}_{\text{CSSE}}^{\text{ENCRYPT}}$ encapsulates the leakage to an adversarial server during the ENCRYPT phase, $\mathcal{L}_{\text{CSSE}}^{\text{GENTOKEN}}$ encapsulates the leakage to an adversarial server during the GENTOKEN phase, and $\mathcal{L}_{\text{CSSE}}^{\text{SEARCH}}$ encapsulates the leakage to an adversarial server during each execution of the SEARCH protocol.

2.2 Adaptive Security of CSSE

We present the **Real** and **Ideal** experiments for the security analysis of a conjunctive SSE scheme CSSE in this Section. In these experiments, the leakage function for CSSE is expressed as

$$\mathcal{L}_{\text{CSSE}} = (\mathcal{L}_{\text{CSSE}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{CSSE}}^{\text{GENTOKEN}}, \mathcal{L}_{\text{CSSE}}^{\text{SEARCH}}),$$

where $\mathcal{L}_{\text{CSSE}}^{\text{ENCRYPT}}$ encapsulates the leakage to an adversarial server during the ENCRYPT phase, $\mathcal{L}_{\text{CSSE}}^{\text{GENTOKEN}}$ encapsulates the leakage to an adversarial server during the GENTOKEN phase, and $\mathcal{L}_{\text{CSSE}}^{\text{SEARCH}}$ encapsulates the leakage to an adversarial server during each execution of the SEARCH protocol.

Algorithm 1 Experiment $\mathbf{Real}_A^{\text{CSSE}}(\lambda)$

```

1: function  $\mathbf{Real}_A^{\text{CSSE}}(\lambda)$ 
2:    $N \leftarrow \mathcal{A}(\lambda)$ 
3:    $(\text{sk}, \text{s}_0, \mathbf{EDB}_0) \leftarrow \text{CSSE.ENCRYPT}(\lambda, N)$ 
4:   for  $k \leftarrow 1$  to  $Q$  do
5:     Let  $q_k \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
6:     Let  $(\text{s}_k, \mathbf{EDB}_k, \mathbf{DB}(q_k)) \leftarrow$ 
       CSSE.SEARCH( $\text{sk}, \text{s}_{k-1}, q_k; \mathbf{EDB}_{k-1}$ )
7:     Let  $\tau_k$  denote the view of the adversary after
       the  $k^{\text{th}}$  query
8:    $b \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_Q, \tau_1, \dots, \tau_Q)$ 
9:   return  $b$ 

```

3 Symmetric Encrypted Computation

We introduce the concept of Symmetric Encrypted Computation (SEC) as an efficient framework for fast outsourced privacy-preserved Boolean circuit evaluation in the symmetric-key setting. SEC supports arbitrary circuit depth via encrypted look-up and a single round

Algorithm 2 Experiment $\text{Ideal}_{\mathcal{A}, \text{SIM}}^{\text{CSSE}}(\lambda, Q, \mathcal{L})$

```

1: function  $\text{Ideal}_{\mathcal{A}, \text{SIM}}^{\text{CSSE}}(\lambda, Q, \mathcal{L})$ 
2:   Parse the leakage function  $\mathcal{L}$  as:
      $\mathcal{L} = (\mathcal{L}_{\text{CSSE}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{CSSE}}^{\text{SEARCH}})$ .
3:    $(s_{\text{SIM}}, \mathbf{EDB}_0) \leftarrow \text{SIM}_{\text{SETUP}}(\mathcal{L}^{\text{ENCRYPT}}(\lambda, N))$ 
4:   for  $k \leftarrow 1$  to  $Q$  do
5:     Let  $q_k \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_{k-1}, \tau_1, \dots, \tau_{k-1})$ 
6:     Let  $(s_{\text{SIM}}, \mathbf{EDB}_k, \tau_k) \leftarrow \text{SIM}_{\text{SEARCH}}$ 
        $(s_{\text{SIM}}, \mathcal{L}_{\text{CSSE}}^{\text{SEARCH}}(q_k); \mathbf{EDB}_{k-1})$ 
7:     Let  $\tau_k$  denote the view of the adversary after
       the  $k^{\text{th}}$  query
8:    $b \leftarrow \mathcal{A}(\lambda, \mathbf{EDB}_Q, \tau_1, \dots, \tau_Q)$ 
9:   return  $b$ 

```

of communication with the remote server. We start by presenting the high-level syntax of SEC in this section before delving into the elaborate technical details of the construction subsequently.

3.1 Syntax of SEC

We briefly explain a general syntax of our proposed primitive here. It uses a static conjunctive SSE construction CSSE as a black-box. We assume a single (honest) client and a (semi-honest) server in SEC. SEC is abstracted as a tuple of four polynomial-time algorithms $\{\text{KEYGEN}, \text{ENCRYPT}, \text{EVALUATE}, \text{DECRYPT}\}$, as defined below:

- $\text{KEYGEN}(\lambda)$: A probabilistic algorithm executed by the client that takes as input the security parameter λ . It outputs a client secret key sk and a public parameter pp .
- $\text{ENCRYPT}(\text{sk}, x_1, \dots, x_p)$: A probabilistic algorithm executed by the client. It takes the client's secret key sk and a p -bit input $\{x_1, \dots, x_p\}$. The output is a ciphertext $c = \{c_1, \dots, c_p\}$.
- $\text{EVALUATE}(f_{\text{desc}}, c_1, \dots, c_p, \text{pp})$: A deterministic algorithm executed by the server that takes as input a description of a circuit f_{desc} that is to be evaluated, the encrypted input bits $c = \{c_1, \dots, c_p\}$ and the public parameter pp . The server returns the encrypted evaluation of the circuit eval to the client.
- $\text{DECRYPT}(\text{sk}, \text{eval})$: A deterministic algorithm executed by the client with its secret-key sk and an encrypted evaluation eval as input, which outputs the decrypted result.

Correctness. SEC is said to be functionally correct if for security parameter λ , and for the following sequence of operations:

$$\begin{aligned}
\text{sk, pp} &\leftarrow \text{SEC.KEYGEN}(\lambda) \\
\{c_1, \dots, c_p\} &\leftarrow \text{SEC.ENCRYPT}(\text{sk}, x_1, \dots, x_p) \\
\text{eval} &= \text{SEC.EVALUATE}(f_{\text{desc}}, c_1, \dots, c_p, \text{pp}) \\
\text{result} &= \text{SEC.DECRYPT}(\text{sk}, \text{eval}),
\end{aligned}$$

The following holds with certainty:

$$\Pr[\text{result} = f_{\text{desc}}(x_1, \dots, x_p)] = 1,$$

Security. SEC guarantees privacy of inputs to a function and output of the evaluation by encrypting them using an IND-CPA secure symmetric-key encryption. Formally, SEC is said to be adaptively secure with respect to a leakage function $\mathcal{L}_{\text{SEC}} = \{\mathcal{L}_{\text{SEC}}^{\text{KEYGEN}}, \mathcal{L}_{\text{SEC}}^{\text{ENCRYPT}}\}$ if for any PPT adversary \mathcal{A} that encrypts a p -bit input x_1, \dots, x_p , there exists a PPT simulator $\text{SIM} = \{\text{SIM}_{\text{KEYGEN}}, \text{SIM}_{\text{ENCRYPT}}\}$ such that the following holds:

$$|\Pr[\mathbf{Real}_{\mathcal{A}}^{\text{SEC}}(\lambda) = 1] - \Pr[\mathbf{Ideal}_{\mathcal{A}, \text{SIM}}^{\text{SEC}}(\lambda, \mathcal{L}) = 1]| \leq \text{negl}(\lambda),$$

where the “real” experiment $\mathbf{Real}_{\mathcal{A}}^{\text{SEC}}$ and the “ideal” experiment $\mathbf{Ideal}_{\mathcal{A}}^{\text{SEC}}$ are as described in Algorithm 3 and Algorithm 4, $\mathcal{L}_{\text{SEC}}^{\text{KEYGEN}}$ captures the leakage from SEC.KEYGEN, and $\mathcal{L}_{\text{SEC}}^{\text{ENCRYPT}}$ captures the leakage from SEC.ENCRYPT.

Algorithm 3 Experiment $\mathbf{Real}_{\mathcal{A}}^{\text{SEC}}(\lambda)$

```

1: function  $\mathbf{Real}_{\mathcal{A}}^{\text{SEC}}(\lambda)$ 
2:    $(\text{sk}, \text{pp}) \leftarrow \text{KEYGEN}(\lambda)$ 
3:   for  $k \leftarrow 1$  to  $y$   $\triangleright y = \text{poly}(\lambda)$ 
4:     do
5:       Let  $x_1, \dots, x_p \leftarrow \mathcal{A}(\lambda, \text{pp}, \tau_1, \dots, \tau_{k-1})$ 
6:       Let  $c_1, \dots, c_p \leftarrow \text{SEC.ENCRYPT}(\text{sk}, x_1, \dots, x_p)$ 
7:       Let  $f(c_1, \dots, c_p) \leftarrow \text{SEC.EVALUATE}(f_{\text{desc}},$ 
8:          $\text{pp}, c_1, \dots, c_p)$ 
9:       ( $\tau_k$  denote  $\mathcal{A}$ 's view after the  $k^{\text{th}}$  evaluation)
10:   $b \leftarrow \mathcal{A}(\lambda, \text{pp}, f_{\text{desc}}, \tau_1, \dots, \tau_y)$ 
11:  return  $b$ 

```

3.2 SEC Construction

The fundamental goal of SEC is to compute arbitrary depth Boolean circuits using encrypted lookup tables while ensuring data-privacy, in a single round of communication between the client and server. The three universal Boolean function set supported by SEC are $f = \{f_{\text{XOR}}, f_{\text{AND}}, f_{\text{OR}}\}$. SEC uses black-box conjunctive SSE constructions to *compute* by performing searches over encrypted lookup tables on encrypted inputs. We refer the reader to Section 1.2 for an overview of SEC, and proceed with the construction here.

Algorithm 4 Experiment $\text{Ideal}_{\mathcal{A}, \text{SIM}}^{\text{SEC}}(\lambda, \mathcal{L})$

```

1: function  $\text{Ideal}_{\mathcal{A}, \text{SIM}}^{\text{SEC}}(\lambda, \mathcal{L})$ 
2:   Parse the leakage function  $\mathcal{L}$  as:
      $\mathcal{L} = (\mathcal{L}_{\text{SEC}}^{\text{KEYGEN}}, \mathcal{L}_{\text{SEC}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{SEC}}^{\text{EVALUATE}})$ .
3:    $(\mathbf{s}_{\text{SIM}}, \mathbf{pp}) \leftarrow \text{SIM}_{\text{KEYGEN}}(\mathcal{L}_{\text{SEC}}^{\text{KEYGEN}}(\lambda))$ 
4:   for  $k \leftarrow 1$  to  $y$   $\triangleright y = \text{poly}(\lambda)$ 
     do
5:     Let  $x_1, \dots, x_p \leftarrow \mathcal{A}(\lambda, \mathbf{pp}, \tau_1, \dots, \tau_{k-1})$ 
6:     Let  $(\mathbf{s}_{\text{SIM}}, c_1, \dots, c_p) \leftarrow \text{SIM}_{\text{ENCRYPT}}(\mathbf{s}_{\text{SIM}},$ 
        $\mathcal{L}_{\text{SEC}}^{\text{ENCRYPT}}(x_1, \dots, x_p); \mathbf{pp})$ 
7:     Let  $f(c_1, \dots, c_p) \leftarrow \text{SIM}_{\text{EVALUATE}}(\mathbf{s}_{\text{SIM}},$ 
        $\mathcal{L}_{\text{SEC}}^{\text{EVALUATE}}(f_{\text{desc}}, \mathbf{pp}, c_1, \dots, c_p))$ 
      $(\tau_k \text{ denote } \mathcal{A}'\text{s view after the } k^{\text{th}} \text{ evaluation})$ 
8:    $b \leftarrow \mathcal{A}(\lambda, \mathbf{pp}, f_{\text{desc}}, \tau_1, \dots, \tau_q)$ 
9:   return  $b$ 

```

SEC.KeyGen. Algorithm 5 summarizes SEC.KEYGEN. It is executed by the client and is responsible for creating the client secret key sk and a public parameter \mathbf{pp} . \mathbf{pp} constitutes ① encrypted lookup tables (i.e. encrypted search indexes) for $\{f_{\text{XOR}}, f_{\text{AND}}, f_{\text{OR}}\}$ and ② encrypted search tokens st (i.e. the token set). \mathbf{pp} is then offloaded to the server.

Concretely, first, a helper function GENDB generates \mathbf{DB} , which essentially comprises the plaintext mappings between keywords and encrypted documents (Table 2 and Table 4). \mathbf{DB} is then encrypted into SSE specific data structures using CSSE.ENCRYPT to generate the encrypted search index or \mathbf{EDB} ¹¹. Following the discussion established in Section 1.2, \mathbf{EDB} contains data structures relevant ① to process “special” terms (i.e. set of $\text{sval}_{\mathbf{w}, D_j}$ against each keyword \mathbf{w}), as well as ② to process “cross” terms (i.e. data structures to return a binary decision on whether $(\mathbf{w}_d^k \wedge \mathbf{w}_i)$ is present in D_j , upon being queried by search tokens of form $\text{Token}_{\mathbf{w}_d^k, \mathbf{w}_i, D_j}$). Thereafter, the specific search tokens st are generated. We recall from SEC’s general overview from Section 1.2 that search tokens of form $\text{Token}_{\mathbf{w}_d^k, \mathbf{w}_i, D_j}$ are used to index into SSE specific data structures, and a binary decision is returned as to whether *both* \mathbf{w}_d^k and \mathbf{w}_i are present in D_j . We also recall that security considerations require \mathbf{w}_d^k to be *refreshed* for every lookup. The size of TokenSet is henceforth $\mathcal{O}(n)$, assuming a set of n distinct “special” terms $\{\mathbf{w}_d^1, \mathbf{w}_d^2, \dots, \mathbf{w}_d^n\}$ for a circuit with n gates¹².

Moreover, recall from Section 1.2 the need of ConstructQuery that maps the output of one SEC evaluation to obtain search tokens for the *next* query (outer function evaluation). SEC hence additionally generates a lookup table that returns a binary decision on whether $\text{sval}_{\mathbf{w}_d^k, D_j}$ and $\text{Token}_{\mathbf{w}_d^{k'}, \mathbf{w}_i, D_j}$ is a valid mapping based on query parameters. ConstructQuery is thus realised as an efficient membership test (and implemented using Bloom Filters), thereby ensuring space $\mathcal{O}(n^2)$ bits and $\mathcal{O}(1)$ lookup. Concretely, *any* “special” term \mathbf{w}_d^k (used in $\text{sval}_{\mathbf{w}_d^k, D_j}$) can be used to obtain search tokens belonging to any other “special” term $\mathbf{w}_d^{k'}$ (used in $\text{Token}_{\mathbf{w}_d^{k'}, \mathbf{w}_i, D_j}$).

¹¹Refer to Section 2.1 for conjunctive SSE syntax and security model.

¹²The analysis subsumes $\mathcal{O}(1)$ number of keywords for input combinations (i.e. \mathbf{w}_1 and \mathbf{w}_2), as well as $\mathcal{O}(1)$ number of documents ($D_j : j \in \{0, 1, 2, 3\}$)

We emphasize that by trading off storage, SEC achieves the capability of randomizing the order of “special” terms $\{\mathbf{w}_d^1, \mathbf{w}_d^2, \dots, \mathbf{w}_d^n\}$ across different executions of circuits, thereby achieving reusability. Finally, all the encrypted look-up tables are then offloaded to the server. Note that KEYGEN is a one-time routine, executed only once at the beginning.

Algorithm 5 SEC.KEYGEN

Input: Security parameter λ

Output: Client’s secret key \mathbf{sk} and a public parameter \mathbf{pp}

```

1: function SEC.KEYGEN( $\lambda$ )
2:   Samples a uniformly random key  $\mathbf{sk}$  for an IND-CPA Symmetric-key Encryption scheme:
    $\Pi_{\text{sym}} = (\text{KEYGEN}, \text{ENCRYPT}_{\mathbf{sk}}, \text{DECRYPT}_{\mathbf{sk}})$ 
3:    $\mathbf{DB} \leftarrow \text{GENDB}(\lambda)$ 
4:    $\mathbf{EDB} \leftarrow \text{CSSE.ENCRYPT}(\mathbf{sk}, \mathbf{DB})$ 
5:   for all conjunctive query  $q$  of keywords in  $\mathbf{DB}$  do
6:      $\{\mathbf{st}_q\} = \text{CSSE.GENTOKEN}(\mathbf{sk}, q)$ 
7:      $\text{TokenSet} = \text{TokenSet} \cup \{\mathbf{st}_q\}$ 
8:   return  $\mathbf{sk}, \mathbf{pp} = \{\mathbf{EDB} \cup \text{TokenSet}\}$ 

```

SEC.Encrypt. This routine is executed on the client’s end and is responsible for encrypting and offloading the actual data. As shown in Algorithm 6, the client uses its secret key \mathbf{sk} to encrypt plaintext bits x_1, \dots, x_p using an IND-CPA secure symmetric-key encryption scheme $(\Pi_{\text{sym}} = (\text{KEYGEN}, \text{ENCRYPT}_{\mathbf{sk}}, \text{DECRYPT}_{\mathbf{sk}}))$, and returns a ciphertext c_1, \dots, c_p .

Algorithm 6 SEC.ENCRYPT

Input: $\mathbf{sk}, x_1, \dots, x_p$

Output: c_1, \dots, c_p

```

1: function SEC.ENCRYPT( $\mathbf{sk}, x_1, \dots, x_p$ )
2:   Encrypt input bits  $\{x_1, \dots, x_p\}$  using an IND-CPA secure symmetric-key encryption scheme
    $Enc$  with client secret-key  $\mathbf{sk}$ 
3:    $\{c_1, \dots, c_p\} \leftarrow Enc_{\mathbf{sk}}(x_1, \dots, x_p)$ 
4:   return  $\{c_1, \dots, c_p\}$ 

```

SEC.Evaluate. SEC.EVALUATE is executed by the server and is mainly responsible for computing a circuit (whose description is provided by the client) on encrypted inputs. This is summarized in Algorithm 7. We assume the client wishes to compute an arbitrary Boolean circuit:

$f_{\text{cir_depth}}(f_{\text{cir_depth}-1}(x_{\text{cir_depth}-1}, y_{\text{cir_depth}-1}), \dots, f_1(x_1, y_1))$, where $f_j \in \{f_{\text{XOR}}, f_{\text{AND}}, f_{\text{OR}}\}$ and (x_j, y_j) are inputs to the j -th function f_j (where $1 \leq j \leq \text{cir_width}$; cir_width being the maximal width of the circuit) at depth i (for $1 \leq i \leq \text{cir_depth}$; cir_depth being the depth of the entire circuit). The server runs the SEC.EVALUATE algorithm that takes as input a circuit description f_{desc} , encrypted inputs c_1, \dots, c_p , and the public parameter \mathbf{pp} (already offloaded to the server at the end of SEC.KEYGEN). It parses \mathbf{pp} as $\{\mathbf{EDB}, \text{TokenSet}\}$ and then uses the encrypted inputs c_1, \dots, c_p to retrieve the corresponding *search tokens* (\mathbf{st}_{q_f}) for all j functions at depth 1 of the circuit. Next, for each of the j -th function at depth i ($1 \leq i \leq \text{cir_depth}$), it invokes the CSSE.SEARCH algorithm using a specific search token and the \mathbf{EDB} as input. The output obtained (i.e. $\text{eval}_{j,i}$) is in turn used to retrieve the search tokens for the functions at depth $(i + 1)$ of the circuit by calling $\text{ConstructQuery}(\text{eval}_{j,i})$. It is to be noted that $\text{eval}_{j,i}$ encapsulates $(\mathbf{w}_d^k, \text{sval}_{\mathbf{w}_d^i, D_j}, b)$ and

Algorithm 7 SEC.EVALUATE

Input: $f_{\text{desc}}, c_1, \dots, c_p, \text{pp}$ **Output:** eval

```
1: function SEC.EVALUATE( $f_{\text{desc}}, c_1, \dots, c_p, \text{pp}$ )
2:   Parse  $\text{pp} = \{\mathbf{EDB}, \text{TokenSet}\}$ 
3:   Retrieve search tokens  $\text{st}_{q_f}$  from TokenSet using  $c_1, \dots, c_p \triangleright \text{ConstructQuery}(\text{TokenSet}, c_1, \dots, c_p)$ 
4:   for  $i = 1$  to  $\text{cir\_depth} - 1$  do  $\triangleright$  cir\_depth is depth of the circuit  $f_{\text{desc}}$ 
5:      $\text{eval}_{j,i} = \text{CSSE.SEARCH}(\mathbf{EDB}, \text{st}_{q_f}) \triangleright$  for a circuit of width  $j$  at depth  $i$ , run  $j$  instances of CSSE.SEARCH in parallel
6:     Retrieve search tokens  $\text{st}_{q_f} \leftarrow \text{ConstructQuery}(\text{eval}_{j,i})$ 
7:      $\text{eval} = \text{CSSE.SEARCH}(\mathbf{EDB}, \text{st}_{q_f})$ 
8:   return eval to the client at the end of the protocol
```

retrieves the search token st_{q_f} (which is of the form $\text{Token}_{\mathbf{w}_d^k, \mathbf{w}_i, D_j}$). Bit b is used to determine whether to use the retrieved search token as the first or second input to the outer function. Again CSSE.SEARCH is called with the new search token query and \mathbf{EDB} as the input. This process continues till the last level of the circuit. Note, that each call to CSSE.SEARCH can be made parallelly for all j functions at a certain depth of the circuit, since all functions at depth/level i are pairwise independent wrt. required inputs. The final encrypted evaluation eval returned by the last level of the circuit contains the encrypted bit corresponding to the actual output of the entire circuit evaluation. eval is returned to the client as the final encrypted output of the circuit f_{desc} .

Algorithm 8 SEC.DECRYPT

Input: sk, eval**Output:** result

```
1: function SEC.DECRYPT(sk, eval)
2:   result =  $\text{Dec}_k(\text{eval})$ 
3:   return result
```

SEC.Decrypt. The client runs the SEC.DECRYPT algorithm to decrypt the encrypted evaluation eval. The decrypted value result is equal to the evaluation of the circuit f_{desc} on (x_1, \dots, x_p) .

3.3 Proof of Correctness of SEC

The proof of correctness for SEC follows from the correctness of CSSE. The correctness of CSSE ensures that a conjunctive query $q = \mathbf{w}_1 \wedge \dots \wedge \mathbf{w}_n$ over an encrypted database satisfies the following relations (we refer to Section 2.1 for generic conjunctive SSE syntax):

$$\begin{aligned} \text{sk} &\leftarrow \text{CSSE.KEYGEN}(\lambda) \\ \mathbf{EDB} &\leftarrow \text{CSSE.ENCRYPT}(\text{sk}, \mathbf{DB}) \\ \text{st}_q &= \text{CSSE.GENTOKEN}(\text{sk}, q) \\ \mathbf{DB}(\mathbf{w}_1) \cap \dots \cap \mathbf{DB}(\mathbf{w}_n) &= \text{CSSE.SEARCH}(\mathbf{EDB}, \text{st}_q) \end{aligned}$$

Proof. By deploying CSSE as a black-box, SEC generates the encrypted search index

EDB specific to the function set $f = \{f_{\text{XOR}}, f_{\text{AND}}, f_{\text{OR}}\}$ supported by the scheme. The encrypted search index **EDB** consists of keywords corresponding to encrypted input bits and documents that encapsulate the encrypted output of the function evaluation. It also invokes the CSSE.GENTOKEN algorithm that generates search tokens corresponding to the keywords (that maps to encrypted input bits of a function). The search tokens are stored in a **TokenSet** and then $\text{pp} = \{\text{EDB} \cup \text{TokenSet}\}$ is offloaded to the server. For evaluating a function $f(x, y)$, the function description f_{desc} (which is essentially a single binary function in this case) and encrypted inputs $\{c_1, c_2\}$ (corresponding to plaintext bits x and y respectively) are sent to the server. The server retrieves the search tokens from **TokenSet** and invokes CSSE.SEARCH protocol with the search token query st_{qf} and **EDB** as input. It retrieves a document from **EDB** that consists of the encrypted output of the computation, (denoted as *eval*) and sends it to the client. The client decrypts *eval* locally using the SEC.DECRYPT function and obtains the output bit *result* which is equal to $f(x, y)$. Correctness is hence a combination of correctness of CSSE, along with SEC specific mappings (Table 3 and Table 4) that ensure *exactly* one document being returned by CSSE.SEARCH, which is the *correct* result of respective gate evaluation.

Following the functionally correct evaluation of a single binary gate $f(\cdot, \cdot)$ over encrypted inputs, the correct encrypted evaluation *eval* returned by SEC.EVALUATE for any arbitrary circuit f_{desc} over encrypted inputs $\{c_1, \dots, c_p\}$ can be asserted transitively. This is because f_{desc} is essentially viewable as a collection of several binary gates $f(\cdot, \cdot)$ which are functionally *correct* as established above, and thereby ensures that the output *eval* returned by SEC.EVALUATE on f_{desc} on decryption is equal to $f_{\text{desc}}(x_1, \dots, x_p)$ for any circuit f_{desc} and unencrypted input set $\{x_1, \dots, x_p\}$. Thereby, we conclude that *for a functionally correct and exact conjunctive SSE scheme CSSE, a set of functions $f = \{f_{\text{XOR}}, f_{\text{AND}}, f_{\text{OR}}\}$, and encrypted values $\{c_1, \dots, c_p\}$ of input $\{x_1, \dots, x_p\}$, SEC is functionally correct, since for the following sequence of operations:*

$$\begin{aligned}
 \text{sk}, \text{pp} &\leftarrow \text{SEC.KEYGEN}(\lambda) \\
 \{c_1, \dots, c_p\} &\leftarrow \text{SEC.ENCRYPT}(\text{sk}, x_1, \dots, x_p) \\
 \text{eval} &= \text{SEC.EVALUATE}(f_{\text{desc}}, c_1, \dots, c_p, \text{pp}) \\
 \text{result} &= \text{SEC.DECRYPT}(\text{sk}, \text{eval})
 \end{aligned}$$

result is the decrypted output of the evaluation of a circuit as specified by f_{desc} on encrypted inputs c_1, \dots, c_p , such that $\text{result} = f_{\text{desc}}(x_1, \dots, x_p)$ (where, x_1, \dots, x_p is the unencrypted input), i.e.,

$$\Pr[\text{result} = f_{\text{desc}}(x_1, \dots, x_p)] = 1,$$

3.4 Practical Instantiation of SEC

Our generic privacy-preserving computation framework SEC can be practically implemented by deploying any conjunctive SSE scheme as a black-box. In this work, we provide concrete

constructions of SEC using two conjunctive SSE schemes:

SEC_{OXT}. We analyze a concrete instantiation of SEC based on the OXT protocol [14], abbreviated SEC_{OXT}. Our analysis fundamentally covers the complexity of SEC_{OXT} in terms of performance and storage overhead, and a formal security analysis based on a well-defined leakage profile. We provide a detailed complexity analysis of SEC_{OXT} based on our experimental results in Section 5. The leakage profile and security proof of SEC_{OXT} is elaborated in Section 4.3.

SEC_{CONJFILTER}. We also demonstrate the scalability and complexity overhead of a second instantiation of SEC by deploying a purely symmetric-key based (plausibly quantum-safe) conjunctive SSE, CONJFILTER [40]. We provide performance and storage overhead analysis of SEC_{CONJFILTER} in Section 5.

3.5 Complexity Analysis of SEC

Storage Overhead. The storage requirement of SEC depends upon the number of functions it supports along with the number of search tokens (input combinations) for every function. The final **DB** is collectively composed of three sub-databases (search indices) $\mathbf{DB} = \{\mathbf{DB}_{\text{AND}}, \mathbf{DB}_{\text{OR}}, \mathbf{DB}_{\text{XOR}}\}$ encrypted and offloaded to the server. As discussed in Section 1.2 and Section 3.2, each function-specific sub-database contains exactly *eight* keywords¹³, which are mapped to exactly *two* documents (out of four possible documents). The bulk of storage overhead comes from the set of n distinct “special” terms $\{\mathbf{w}_d^1, \mathbf{w}_d^2, \dots, \mathbf{w}_d^n\}$ (for a circuit with n binary gates), for which SSE specific data structures occupy $\mathcal{O}(n)$ space. Some constant number of *dummy* documents are also added to maintain a uniform frequency of each keyword in the final encrypted database. Concretely, there are exactly eight keywords for one binary function, four documents (including dummy documents) per keyword, and n “special” terms. Total storage for three gates can be calculated as -

$$\text{Total Storage} = \{[(8 + n) \times 4] \times 3\} \cdot b \text{ bytes} = \mathcal{O}(n)$$

where b is a constant.

Computation and Communication Overhead. The evaluation time of SEC for computing arbitrary depth Boolean circuit over encrypted data scales linearly with the search time complexity of the underlying CSSE scheme times some constant which depends upon the depth of the circuit. The crux of SEC is to bypass the explicit circuit evaluation as done in state-of-the-art encrypted computation schemes like FHE and leverage the extremely efficient encrypted search capability of a CSSE scheme to evaluate functions on encrypted inputs. Detail Analysis is given in Section 5.

By constructing the mapping for all pairwise combinations in $\{\mathbf{w}_d^1, \mathbf{w}_d^2, \dots, \mathbf{w}_d^n\}$, multiple circuits can be executed securely without the need to refresh SEC’s data structures. This is achieved by ensuring ① no two gates in the *same* circuit share a “special” term, and ② two isomorphic gates in *different* circuits also do not share corresponding “special” terms. Hence, the client’s *amortized* communication overhead is proportional to the function description

¹³We consider this as $\mathcal{O}(1)$ overhead in our analysis.

length (we assume a circuit evaluation with n gates). Post the *one-time* setup, the client uses derangements of $\{\mathbf{w}_d^1, \mathbf{w}_d^2, \dots, \mathbf{w}_d^n\}$ to evaluate subsequent circuits. Asymptotically, the number of derangements (because of the need to avoid isomorphic gate assignments) is $\left\{n! - \sum_{i=0}^n \frac{(-1)^{i+1}}{i!}\right\} = O(n!)$.

4 Security and Leakage Profile Analysis of SEC

We analyze the security of SEC in this section. We follow a semi-honest adversarial setting for our security analysis where the remote server is assumed to be honest-but-curious. That implies the untrusted server follows the algorithmic specification exactly, but can also observe and record additional information for analysis. Using a simulation-based approach, we establish formally that a probabilistic polynomial-time (PPT) simulator can simulate the view of the adversarial server in an indistinguishable manner given only the leakage profile of SEC. In this framework, a PPT adversary is required to distinguish between the real world (where the adversary interacts with a *real* execution of SEC) and the *ideal* world (where the adversary interacts with a simulator that only has access to the leakage profile for SEC, described below). We assert this framework to be provably secure when no PPT adversary can distinguish between the two scenarios with a significant advantage over random guessing.

SEC inherits security properties and leakage profile from the underlying CSSE construction. We note that CSSE is an adaptively secure conjunctive SSE scheme that is secure against a semi-honest adversary \mathcal{A} . The leakage of CSSE is characterized by the leakage function $\mathcal{L}_{\text{CSSE}}$ which is an ensemble of the leakage functions for ENCRYPT, GENTOKEN and SEARCH individually, expressed in the following way.

$$\mathcal{L}_{\text{CSSE}} = \{\mathcal{L}_{\text{CSSE}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{CSSE}}^{\text{GENTOKEN}}, \mathcal{L}_{\text{CSSE}}^{\text{SEARCH}}\},$$

Given the above CSSE leakage functions, security of SEC can be analyzed using SEC leakage function \mathcal{L}_{SEC} in the same adaptive semi-honest adversarial model. \mathcal{L}_{SEC} is composed of two separate leakage functions for KEYGEN and ENCRYPT (as expressed below), that capture the leakage from SEC.KEYGEN and SEC, ENCRYPT execution respectively.

$$\mathcal{L}_{\text{SEC}} = \{\mathcal{L}_{\text{SEC}}^{\text{KEYGEN}}, \mathcal{L}_{\text{SEC}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{SEC}}^{\text{EVALUATE}}\},$$

Security of SEC_{OXT} . We provide the security analysis of SEC_{OXT} , which follows from the security notions of generic SEC and underlying OXT protocol.¹⁴

Theorem 1 *Given that OXT is an adaptively secure SSE scheme with respect to the leakage function $\mathcal{L}_{\text{OXT}} = \{\mathcal{L}_{\text{OXT}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{OXT}}^{\text{GENTOKEN}}, \mathcal{L}_{\text{OXT}}^{\text{SEARCH}}\}$ against a polynomially-bounded adaptive adversary, SEC_{OXT} is also an adaptively secure encrypted computation framework with respect to the leakage function $\mathcal{L}_{\text{SEC}_{\text{OXT}}} = \{\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN}}, \mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{EVALUATE}}\}$, where the SEC_{OXT}*

¹⁴Analysis for $\text{SEC}_{\text{CONJFILTER}}$ follows suit, since the leakage profile of SEC is agnostic of the CSSE used.

instantiation encrypts an input x_1, \dots, x_p using an IND-CPA secure symmetric-key encryption scheme to obtain corresponding encrypted bits c_1, \dots, c_p over which a (publicly known) function f is evaluated, where $f \in \{f_{\text{XOR}}, f_{\text{AND}}, f_{\text{OR}}\}$.

Proof 1 We give an extensive security analysis of SEC_{OXT} through formal proof of Theorem 1 in Section 4.3.

4.1 Leakage Profile of SEC_{OXT}

We formally explain the leakage profile for the specific instantiation of SEC based on the OXT scheme, namely SEC_{OXT} .

The significance of each component of the leakage function in SEC_{OXT} is comparable to that of OXT. We define each leakage component as follows.

- $N = \sum_{i=1}^d |\mathcal{W}_i|$ - the total number of appearances of keywords in documents. The parameter N signifies an upper bound which is equivalent to the total size of **EDB**. Leaking such a bound is unavoidable and is considered a trivial leakage in the literature of SSE.
- **SP** - size pattern of the queries i.e., the number of documents matching the **stern** in each query. Formally, $\text{SP} \in [d]^n$ and $\text{SP}[i] = |\mathbf{DB}(\text{stern}[i])|$. It leaks the number of documents satisfying the **stern** in a query. In SEC this is always constant (equal to 2), since each keyword (encrypted input bit) maps to exactly four documents (encrypted output of function evaluated on the particular encrypted input) in **EDB**.
- **RP** - result pattern of the queries or the indices of documents matching the entire conjunction. Formally, **RP** is vector of size n with $\text{RP}[i] = \mathbf{DB}(\text{stern}[i]) \cap \mathbf{DB}(\text{xstern}[i])$ for each $i = 1, \dots, n$ where **xstern** refers to all the other keywords in the conjunctive query other than the **stern**. It is the final output of the search query and is not considered a real leakage in the context of SSE. This is always a single document in SEC.

4.2 Analysis of Potential Leakages in SEC_{OXT}

The database (search index) generation process and algorithmic design of SEC_{OXT} , renders certain non-trivial information leakage insignificant or redundant, which is otherwise considered crucial by the underlying OXT scheme and could lead to potential correlation inference by the server between two encrypted function computation. Due to the uniform keyword frequency and selection of random “special” term for each query by the client, SEC_{OXT} restricts certain non-trivial leakages like *size-pattern*, *result-pattern*, *equality-pattern*, *conditional intersection pattern* leakages that analyze the frequency pattern of the “special” terms and “cross” terms in OXT over multiple conjunctive queries. This makes OXT vulnerable to certain state-of-the-art leakage-abuse attacks [47, 8].

- *Size Pattern Leakage (SP)*. It leaks the number of documents satisfying the “special” term in a query. In SEC_{OXT} since every keyword (input bit) maps to exactly four documents

(encapsulating encrypted output bits), this leakage reveals no significant information.

- *Result Pattern Leakage (RP)*. It is the final output of the search query i.e. indices of documents matching the entire conjunction. By the design of SEC_{OXT} the result of a conjunctive query (binary function evaluation) is always a single document (single encrypted output bit) and hence this does not reveal any significant information to the server.
- *Equality pattern (EP)*. It indicates which queries have the equal “special” terms. This occurs due to the optimization technique devised in OXT in order to ensure sub-linear search complexity by filtering out the least frequent term (*sterm*) during the search. In SEC_{OXT} , since the frequency of all keywords is the same, the client chooses a different “special” term from a pool of dummy keywords ($\mathbf{w}_d^k: k \in \{1, \dots, n\}$) for different queries (gate evaluation). Hence, the adversary will not be able to infer any correlation for multiple gate evaluation over multiple (repeated/non-repeated) inputs.
- *Conditional Intersection Pattern Leakage (IP)*. It is a subtle leakage in OXT that occurs when two distinct queries have a common “cross” term but a different “special” term and there exists a document that satisfies both the “special” terms. In such a scenario the set of document indices matching both “special” terms is leaked (if no document matching both “special” terms exists then nothing is leaked). In SEC_{OXT} this leakage will not reveal any significant information about the underlying input bits to a function or the output of a function evaluation, since all the input/output bits are encrypted using an IND-CPA secure symmetric-key encryption scheme. The server cannot even correlate between the matched documents of the “special” terms because for every query a unique “special” term is chosen by the client.

All these leakages except N (total size of **EDB**) are essentially encapsulated by $\mathcal{L}_{\text{OXT}}^{\text{SEARCH}}$, and hence are leaked by the SEC.EVALUATE algorithm (encapsulated by $\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{EVALUATE}}$). As explained above none of these leakages have a significant impact on the data privacy guarantees of SEC.

4.3 Proof of Theorem 1

The security analysis of SEC_{OXT} (proof of Theorem 1) stems from the provable security guarantee of OXT. We first outline the leakage sources of OXT, with respect to which OXT is simulation secure. Subsequently, we show that adaptive semantic security of OXT implies adaptive security guarantees of SEC_{OXT} .

We resort to the same simulation-based security analysis approach for SEC_{OXT} as of OXT. We show that SEC_{OXT} is secure against an adaptive semi-honest adversary \mathcal{A} , which has access to leakages from $\mathcal{L}_{\text{SEC}_{\text{OXT}}}$. We build a simulator $\text{SIM} = \{\text{SIM}_{\text{KEYGEN}}, \text{SIM}_{\text{ENCRYPT}}, \text{SIM}_{\text{EVALUATE}}\}$ for SEC_{OXT} where the simulator emulates SEC_{OXT} execution just from the knowledge of public information and leakage $\mathcal{L}_{\text{SEC}_{\text{OXT}}}$.

Leakage Cover. We briefly describe why each of the individual leakage components ($\mathcal{L}_{\text{SEC}_{\text{OXT}}} = \{\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN}}, \mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{ENCRYPT}}, \mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{EVALUATE}}\}$) are necessary for a simulator to produce correct results. To simulate SEC_{OXT} correctly each of the leakage components are critically

analyzed and their significance is justified. N or the total number of appearances of keywords in the database gives the size of the **EDB**, which is encapsulated by the public parameter \mathbf{pp} along with the size of search token set $|\text{TokenSet}|$.

Simulating SEC_{OXT} KeyGen and Encrypt. In OXT the **EDB** comprises of two data structures $\mathbf{EDB} = \{\text{TSet}, \text{XSet}\}$. The main crux of our adaptive security proof is that the simulator for SEC_{OXT} initializes the XSet and TokenSet to consist entirely of uniformly random elements from a *discrete log* hard group initially (while relying on the DDH assumption for indistinguishability of the real and simulated XSet and TokenSet entries). Additionally, the simulator for SEC_{OXT} can directly invoke the simulator for the adaptively secure TSet to simulate the TSet entries at ENCRYPT . Overall $\text{SIM} = \{\text{SIM}_{\text{KEYGEN}}, \text{SIM}_{\text{ENCRYPT}}, \text{SIM}_{\text{EVALUATE}}\}$ takes as input the leakage components as defined by $\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN}}$. N and $|\text{TokenSet}|$ is essentially learned from the public parameter \mathbf{pp} , hence, $\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN}} = \perp$, $\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{ENCRYPT}}$ that reveals the length n of the ciphertext c_1, \dots, c_n , and $\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{EVALUATE}}$ reveals $\{\text{SP}, \text{RP}, \text{EP}, \text{CIP}\}$. Using the leakages and the public parameter the SIM then produces the ciphertext c_1, \dots, c_n which is indistinguishable from the encrypted output returned by the original SEC.ENCRYPT algorithm on an input x_1, \dots, x_n .

Simulating SEC_{OXT} .KeyGen. We observe that, SEC_{OXT} . KEYGEN comprises of the GENDB , OXT.ENCRYPT (CSSE.ENCRYPT), OXT.GENTOKEN (CSSE.GENTOKEN) (this phase is encapsulated in the OXT.SEARCH protocol in [14], and is entirely executed by the client). Without loss of generality, we extract the search token generation phase (OXT.GENTOKEN) and store all possible search tokens in TokenSet during SEC_{OXT} . KEYGEN . A number of values generated by pseudo-random functions (PRF) and group operations, are inserted into TSet using TSet.SETUP [14] and XSet respectively during OXT.ENCRYPT . Note that, GENDB routine creates the plain look-up table for the supported primitive operations, and it is executed on the client side. Hence, the adversarial server learns no information from the GENDB execution itself and thus the leakage from GENDB can be expressed as null.

$$\mathcal{L}^{\text{GENDB}} = \perp,$$

Thus, the simulator $\text{SIM}_{\text{KEYGEN}}$ can exactly simulate GENDB execution straightforwardly.

Subsequently, the OXT.ENCRYPT is invoked with the plain \mathbf{DB} generated by GENDB . Since the OXT.ENCRYPT algorithm is executed in a black-box way, the leakage from SEC_{OXT} . KEYGEN is same as the OXT.ENCRYPT executed over \mathbf{DB} . Also, OXT.GENTOKEN phase is invoked on all possible queries q of keywords (alphanumeric translations of encrypted input bits) in \mathbf{DB} . This algorithm is also used a black-box and is entirely executed by the client. Thus, the leakage can be expressed as below.

$$\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN}} = \{\mathcal{L}_{\text{OXT}}^{\text{ENCRYPT}}(\mathbf{DB}), \mathcal{L}_{\text{OXT}}^{\text{GENTOKEN}}(q)\},$$

Finally, the TSet SETUP execution does not leak additional information apart from already known public information (the size of the database $|\mathcal{W}| = N$ is known). The leakage for this part can be expressed as below.

$$\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN, TSet}} = N = \perp \text{ as this is a public information}$$

$\text{SIM}_{\text{KEYGEN}}$ can run the **TSet** simulator (as discussed in the original paper [15]). Combined all, the simulator for **KEYGEN** ($\text{SIM}_{\text{KEYGEN}}$) simulates $\text{SEC}_{\text{OXT}}.\text{KEYGEN}$ with access to following the leakage.

$$\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN}} = \{\mathcal{L}^{\text{GENDB}}, \mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN}}, \mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{KEYGEN, TSet}}\},$$

Simulating $\text{SEC}_{\text{OXT}}.\text{ENCRYPT}$. For simulating $\text{SEC}_{\text{OXT}}.\text{ENCRYPT}$ the simulator $\text{SIM}_{\text{ENCRYPT}}$ observes $\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{ENCRYPT}}$ which is equal to the length of the ciphertext c_1, \dots, c_n .

$$\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{ENCRYPT}} = |c_1, \dots, c_n|,$$

The $\text{SEC}_{\text{OXT}}.\text{ENCRYPT}$ algorithm invokes an IND-CPA secure symmetric-key encryption scheme which is used to encrypt an input bit x_1, \dots, x_n . The $\text{SIM}_{\text{ENCRYPT}}$ produces a ciphertext c_1, \dots, c_n corresponding to the input only with the information from $\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{ENCRYPT}}$ and its state (s_{SIM}). The ciphertext thus produced by $\text{SIM}_{\text{ENCRYPT}}$ is indistinguishable from the ciphertext produced by $\text{SEC}_{\text{OXT}}.\text{ENCRYPT}$ in the real scheme.

Simulating $\text{SEC}_{\text{OXT}}.\text{EVALUATE}$. The **EVALUATE** function takes as input encrypted bits c_1, \dots, c_n and retrieves *search tokens* from the **TokenSet** using the encrypted input bits (as input to the **ConstructQuery** subroutine). It then invokes the **CSSE.SEARCH** function using the search tokens and gets an encrypted bit as output. This process is repeated for all gates at every depth of the circuit being evaluated. The **ConstructQuery** subroutine leaks essentially no information to the server. This is because the input bits are encrypted using an IND-CPA symmetric-key encryption algorithm and the search tokens to be searched are dependent on the “special” term for that particular function/query (which can be selected at random by the client and the number of possible permutation of “special terms” for a circuit with n gates is upper bounded by $O(n!)$) and also the function to be evaluated. Therefore, the search pattern in the **TokenSet** for any repeated input bits cannot be correlated by the server. This proves that -

$$\mathcal{L}^{\text{ConstructQuery}} = \perp,$$

The leakage $\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{EVALUATE}}$ is therefore exactly similar to $\mathcal{L}_{\text{OXT}}^{\text{SEARCH}}$. Hence, we can write the following.

$$\mathcal{L}_{\text{SEC}_{\text{OXT}}}^{\text{EVALUATE}} = \{\mathcal{L}_{\text{OXT}}^{\text{SEARCH}}(\text{EDB}, st_q)\},$$

By the simulation security guarantee of OXT, SEC_{OXT} is secure against these leakages. We show that the leakages from **OXT.SEARCH** phase do not have any detrimental effect on the encrypted function evaluations in **SEC**.

Let a client evaluate a circuit of the form $f_3(f_2(\mathbf{x}_1, \mathbf{x}_2), f_1(\mathbf{x}_3, \mathbf{x}_4))$. The client sends the encrypted input bits to **SEC.EVALUATE**. A two-input function evaluation is translated to a three-keyword conjunctive query, where the first keyword is a “special” term (chosen randomly by the client), the second keyword corresponds to the first input bit, and the third keyword to the second input bit. The client maintains a state of record of all “special terms” used. For a circuit with n gates, the client chooses n “special” terms, which can be permuted and selected in $n!$ different ways (upper bounded by $O(n!)$). The client sends a set of “special” terms to be used at each gate in a circuit. Since the “special terms” are

used to fetch the records from memory, permuting “special terms” for consecutive (same) gate evaluation ensures no repetition of the same “special” term hence, the same memory location is not accessed twice.

CASE-I: If the same circuit is evaluated twice, the “special” terms are permuted (which is upper bounded by $O(n!)$). Hence the server cannot distinguish between two identical gate evaluations.

CASE-II: The search tokens generated corresponding to encrypted input bits are a function of the “special” term and the function being evaluated. Thus, for each gate, the search tokens depend on “special” terms ($n!$ possible combinations). This ensures that an adversary cannot distinguish between two isomorphic gate evaluation.

This proves that the leakages incurred by the underlying CSSE.SEARCH algorithm does not compromise the security of SEC.EVALUATE. The output of SEC.EVALUATE is indistinguishable from random by the security guarantees of an IND-CPA secure symmetric-key encryption scheme. Since, OXT is proven simulation secure it follows from the simulation security guarantee that \mathcal{A} no additional advantage over the real experiment. This implies the **Real** experiment of SEC (Algorithm 3) is indistinguishable from the **Ideal** experiment (Algorithm 4), and proves Theorem 1.

4.4 \mathcal{L}_{SEC} and Reusability of Look-up Table

The design rationale of SEC guarantees privacy of the input bits to a function and the output bit returned after function evaluation. SEC leverages the encrypted search capability of an efficient CSSE scheme to perform encrypted computation. It might seem that the inherent leakages of a typical CSSE should therefore directly affect the leakage profile of SEC. However, pre-processing and configuring the encrypted search index (as done in SEC) prevents direct extrapolation of the leakages of the underlying CSSE in SEC. We discuss this in detail below.

Details on $\mathcal{L}_{\text{CSSE}}$. The most generic notion of SSE with optimal guarantees on security is achievable through Oblivious RAM [33], which allows evaluation of search queries without leaking *anything* to the server¹⁵. However, such *ideal* security guarantees come at immense computational/communication overheads. Hence, most modern SSE constructions trade-off security for efficiency by allowing calculated, acceptable leakages. Some usual leakages:

- **sval Access Pattern:** Two queries having the *same* “special” term can be correlated by the server by the same set of $\text{sval}_{\mathbf{w}, D_j}$ returned. We emphasize that the server learns not the plaintext alphanumeric value of \mathbf{w} , but rather the fact that two queries share the same “special” term.
- **“Cross” term Access Pattern:** Two queries having a common document matched to their “special” term and the same cross term tuples (for example, $(\mathbf{w}_d^k, \mathbf{w}_i)$) are leaked since the same $\text{Token}_{\mathbf{w}_d^k, \mathbf{w}_i, D_j}$ is generated for both queries. As before, the

¹⁵As with state-of-the-art FHE and SSE constructions, we assume semi-honest server. That is, the server acts as a passive adversary which does not deviate from the protocol.

server can not learn the underlying plaintext alphanumeric value of \mathbf{w}_i ; it can simply correlate same cross-terms across two queries.

- **Query Result Pattern:** Two queries having the same result (i.e. $\text{sval}_{\mathbf{w}, D_j}$) can be correlated by the server.

Unlinking “Computation” from Generic SSE Leakages. To the best of our knowledge, state-of-the-art SSE constructions *tolerate* such correlations made by the semi-honest server. However, when we use the search capabilities of SSE to *compute*, leaking these correlations essentially allows a server to learn: ① when inputs of two different gates evaluations are same, and ② when the output of two different gates is same. We stress that while the *exact* plaintext input/output bit can not be leaked (thereby not violating data privacy guaranteed by the IND-CPA symmetric-key encryption scheme used to encrypt the data); still, such correlations between different computations are undesirable non-trivial leakages.

As such, we focus not on plugging these leakages, but rather on *unlinking* the computation from such leakages. In other words, we allow the server to learn these leakages; but embed no critical information in such leakages. To do so, we first observe the aforementioned leakages: *the “special” term is present in all leakage functions*, be it whether the leakage occurs through $\text{sval}_{\mathbf{w}, D_j}$ or through $\text{Token}_{\mathbf{w}_d^k, \mathbf{w}_i, D_j}$. Hence, in our construction, we *do not embed any computation-related information in the “special” term*. From Table 2 it is observed that the actual bits participating in computing XOR are independent of the choice of \mathbf{w}_d^k . This **design choice allows SEC to change the “special” term across multiple queries**. The server still learns the aforementioned leakages; however, no useful correlations as to the underlying computation are revealed.

Example. Re-consider the same problem of computing XOR(1, 1) through SEC, but now the computations happen twice. As such, two queries $\mathbf{q}_1 = (\mathbf{w}_d^1 \wedge \mathbf{w}_1 \wedge \mathbf{w}_2)$ and $\mathbf{q}_2 = (\mathbf{w}_d^2 \wedge \mathbf{w}_1 \wedge \mathbf{w}_2)$ are issued by a black-box conjunctive SSE scheme. We enumerate the leakages visible across these queries:

- **sval Access Pattern Leakage:** \mathbf{q}_1 leaks $[\text{sval}_{\mathbf{w}_d^1, D_j} : j \in \{0, 1, 2, 3\}]$. Likewise, \mathbf{q}_2 leaks $[\text{sval}_{\mathbf{w}_d^2, D_j} : j \in \{0, 1, 2, 3\}]$
- **“Cross” term Access Pattern Leakage:** \mathbf{q}_1 leaks accesses made by $\text{Token}_{\mathbf{w}_{d_1}, \mathbf{w}_1, D_j}$ and by $\text{Token}_{\mathbf{w}_d^1, \mathbf{w}_2, D_j}$. Likewise, \mathbf{q}_2 leaks accesses make by $\text{Token}_{\mathbf{w}_d^2, \mathbf{w}_1, D_j}$ and by $\text{Token}_{\mathbf{w}_{d_1}, \mathbf{w}_2, D_j}$. In all cases, $j \in \{0, 1, 2, 3\}$.
- **Query Result Pattern Leakage:** \mathbf{q}_1 leaks $\text{sval}_{\mathbf{w}_d^1, D_3}$, while \mathbf{q}_2 leaks $\text{sval}_{\mathbf{w}_d^2, D_3}$ to the server.

Hence, even though the same gate with same inputs is being evaluated, by unlinking the “special” term from the actual computation, the view of the server is *different* in both cases. \mathcal{L}_{SEC} still contains these leakages (as with $\mathcal{L}_{\text{CSSE}}$), but the server can still not correlate multiple evaluations of the same gate.

Concrete Realization of \mathcal{L}_{OXT} . We now discuss how \mathcal{L}_{SEC} extends to a concrete instantiation using OXT. We first recall that SEC ensures data privacy by supporting function evaluation over symmetrically encrypted (IND-CPA secure symmetric-key encryption)

data (input bits) and returning the encrypted output after evaluation. In addition, it also prevents the server (adversary) to correlate or infer any statistical information about the input/output bits by observing the leakage profile of SEC. To extend \mathcal{L}_{SEC} to a concrete instantiation using OXT, we begin with the question:

For SEC_{OXT} , what aspects of \mathcal{L}_{OXT} allow correlation of two different queries pertaining to the same inputs?

The answer is by observing the *access pattern* leakage. In our instantiation of SEC_{OXT} , every input to a function is mapped to a keyword, and a function evaluation’s encrypted output is encapsulated in a document. Also there exist some “*dummy*” keywords (equal to the number of gates in a circuit) and some “*dummy*” documents which are meticulously added to segregate any correlation between multiple gate evaluation from the leakages incurred by the underlying SSE scheme (OXT). Each keyword-document pair for a particular function is stored in a memory location pointed to by an address that is generated by a pseudo-random function. An adversary can correlate the (encrypted) input bits being evaluated over multiple queries/function evaluations, by observing the pattern of the memory locations that are being accessed for different queries/function evaluations. For the same input combination and same function, the server will access the same memory location to fetch the document that stores the encrypted output of the function evaluation. A straightforward mitigation of access-pattern leakage in a SSE scheme is non-trivial and to the best of our knowledge, has not been explored in the literature. We provide a potential solution in SEC_{OXT} such that an adversary cannot infer the access pattern over multiple searches, and hence any statistical analysis of the encrypted input bits.

Reusability of Lookup Table. The unique selection of “special” terms by a client for each gate evaluation guarantees resistance of SEC_{OXT} to any statistical analyses of input bits by restricting the adversary’s advantage of reverse-engineering inputs to negligible. This design choice also ensures that the same function invocation across different input sets has a non-identical access pattern. Concretely, for an adversary to correlate the encrypted input bits in two same function evaluations of the form $f_i(x, y)$ (where, x, y are all encrypted), it requires correlating the “special” term used for each query (evaluation). Let in this case the client chooses two different “special” terms for each gate evaluation, the corresponding conjunctive query translates to - $q_{f_i} = \mathbf{w}_d^1 \wedge \mathbf{w}_x \wedge \mathbf{w}_y$ for first evaluation and $q_{f'_i} = \mathbf{w}_d^2 \wedge \mathbf{w}_x \wedge \mathbf{w}_y$ for second evaluation (where, $\mathbf{w}_x, \mathbf{w}_y$ are keywords that map to the corresponding input bits). By the design of OXT the access pattern leakage is dependent on the “special” term because the memory location of the documents corresponding to the “special” term is accessed and only those documents are fetched during a conjunctive search. In the example above the probability of an “special” term being repeated is upper bounded by $O(n!)$ (n is the number of gates in a circuit) because the client chooses unique “special” terms for every query (function evaluation). This is unlike the original OXT scheme where the “special” term is determined based on the least frequent keyword in the query and hence for the example above it would be the same for both functions. Since the frequency of all keywords in SEC_{OXT} is equal, we can leverage the unique selection of a “special” term thereby preventing an adversary from potentially guessing the encrypted input bits to a function by observing the memory access pattern. We re-iterate that both queries q_{f_i} and $q_{f'_i}$ to the encrypted lookup table evaluate the *same* function $f_i(x, y)$, hence the output of both evaluations will be equal.

However, because of *refreshing* the “special” term across two queries, the adversarial view of the server is unique for both queries. Thus, *using the same look-up table that is encrypted and offloaded to the server once during CSSE.SETUP phase, SEC can evaluate any arbitrary Boolean circuit multiple times.* This is guaranteed both in a single circuit across multiple gate evaluation as well as across multiple circuit evaluation. Hence, even though the generic SSE leakages still occur, nothing significant to the underlying computation is compromised.

Theorem 2 (Reusability of Lookup Table) *Given that SEC_{OXT} encrypts an input x_1, \dots, x_p using an IND-CPA secure symmetric-key encryption scheme to obtain corresponding encrypted bits c_1, \dots, c_p that is used as an input to a (publicly known) function f , and all information leaked from the underlying CSSE.SEARCH (OXT.SEARCH) phase is encapsulated by $\mathcal{L}_{\text{OXT}}^{\text{SEARCH}}$, SEC_{OXT} ensures reusability of the same look-up table for multiple (similar/different) gate evaluations without leaking any extra information than that encapsulated by $\mathcal{L}_{\text{SEC}_{\text{OXT}}}$, while guaranteeing input and output data privacy from semantic security guarantees of an IND-CPA secure encryption scheme.*

Proof 2 *We prove Theorem 2 via a sequence of games between a challenger and an adversary, where the first game (G_0) is identical to the real experiment $\mathbf{Real}_A^{\text{SEC}}$ and the final game (Simulator) is identical to the simulation experiment $\mathbf{Ideal}_{\text{SIM},A}^{\text{SEC}}$. We establish formally that the view of the adversary \mathcal{A} in each pair of consecutive experiments is computationally indistinguishable. For ease of exposition, we consider the client computes a binary function f on encrypted input bits $\{c_1, c_2\}$.*

Game G_0 . This game is identical to $\mathbf{Real}_A^{\text{SEC}}$, where the challenger generates transcripts for the encrypted input bits $\{c_1, c_2\}$ by invoking SEC.ENCRIPT and transcripts for the output $f(c_1, c_2)$ by invoking SEC.EVALUATE .

$$\Pr[G_0 = 1] \leq \Pr[\mathbf{Real}_A^{\text{SEC}}(\lambda) = 1] - \text{negl}(\lambda),$$

Game G_1 . This game is identical to G_0 except for the fact that the challenger changes the encrypted input of function f to $\{c_3, c_4\}$. The evaluation of the function is done by invoking SEC.EVALUATE in the same way as done in G_0 , i.e. lookup is performed on the same encrypted lookup table. The adversary cannot distinguish between G_0 and G_1 due to the use of different “special” terms for both queries.

$$\begin{aligned} f(c_1, c_2) &\xrightarrow{\text{translated}} \text{OXT.SEARCH}(\mathbf{EDB}, \{\mathbf{w}_d^1 \wedge \mathbf{w}_1 \wedge \mathbf{w}_2\}) \\ f(c_3, c_4) &\xrightarrow{\text{translated}} \text{OXT.SEARCH}(\mathbf{EDB}, \{\mathbf{w}_d^2 \wedge \mathbf{w}_3 \wedge \mathbf{w}_4\}) \end{aligned}$$

where, $\{\mathbf{w}_d^1, \mathbf{w}_d^2\}$ are randomly chosen “special” terms, and $\{\mathbf{w}_1, \dots, \mathbf{w}_4\}$ are translated keywords from input bits $\{c_1, \dots, c_4\}$. We say that by IND-CPA security guarantees the output of both G_0 and G_1 are indistinguishable from random.

$$|\Pr[G_1 = 1] - \Pr[G_0 = 1]| \leq \text{negl}(\lambda),$$

Game G_2 . This game is identical to G_1 except for the fact that the function being evaluated is changed to f' by the challenger. Since the search tokens generated depend on the “special” term and the function being evaluated, and since for every function evaluation the “special” term is chosen randomly by the client, for every function (same/different) evaluation the search tokens generated are different. The adversary can therefore not distinguish between two isomorphic gate evaluations.

$$|\Pr[G_2 = 1] - \Pr[G_1 = 1]| \leq \text{negl}(\lambda),$$

Simulator. The simulator SIM generates similar transcripts for SEC.EVALUATE using the leakages from $\mathcal{L}_{\text{SEC}_{\text{OXT}}}$ and this experiment is similar to $\mathbf{Ideal}_{\text{SIM},\mathcal{A}}^{\text{SEC}}$. How the transcripts are generated from each leakage component is discussed in detail above. We state here, that the SIM generates the transcripts from the corresponding leakages correctly, and the output of $\mathbf{Ideal}_{\text{SIM},\mathcal{A}}^{\text{SEC}}$ is indistinguishable from G_2 .

4.5 Statistical Analysis of Leakage Due to Reusability

We demonstrated a proof of computational indistinguishability (from an adversarial perspective) that establishes the secure reusability of SEC’s lookup tables in Section 4.3. In this section, we provide statistical analysis of the leakage from SEC’s lookup tables to validate the same. For this, we closely follow the non-interference security notion well established in several side-channel analysis paradigms [4, 21].

Abstractly, the *non-interference property* ensures no sensitive information flow to the output of a system, given the system’s inputs. In context of SEC, non-interference between inputs to SEC.EVALUATE and $\mathcal{L}_{\text{SEC}}^{\text{EVALUATE}}$ (i.e. the observable leakage) translates directly to the server’s inability to infer (with statistical significance) anything about the inputs purely from $\mathcal{L}_{\text{SEC}}^{\text{EVALUATE}}$. Concretely, non-interference can be defined as [4, 21]:

Definition 1 (Non-Interference) For a probabilistic program P , consider the set of secret (“high”) inputs as H , the set of public (“low”) inputs as L , and the output as \mathcal{L} . Then, P is said to be non-interfering if and only if the mutual information $\mathcal{I}(\mathcal{L}; H \mid L) = 0$.

In other words, this information-theoretic definition captures the mutual information (or the mutual dependence) of \mathcal{L} and H (or the critical, secret input to P), given knowledge of non-secret L . P is considered to be non-interfering if variations in H do not (statistically) affect \mathcal{L} (given knowledge of L). This is captured by the mutual information (conditioned on L) being 0. However, estimating conditional mutual information in an information-theoretic setting is a difficult problem in general. Thus, a slightly “relaxed” definition for non-interference can be used instead [21]:

Definition 2 (“Relaxed” Non-Interference) For a probabilistic program P , consider the set of secret (“high”) inputs as H , the set of public (“low”) inputs as L , and the output as \mathcal{L} . Then, P is said to be non-interfering if the marginal distribution of \mathcal{L} is independent of the distribution of H .

Concretely, from the point of view of reusability in SEC (as in Section 4.3), the random choice of “special” terms by the client leads to a computationally indistinguishable view of the server for repeated evaluations. From the perspective of non-interference definitions established, given a query executed in SEC.EVALUATE, the “high” inputs \mathbf{H} correspond to the actual tokens that map inputs to the function being evaluated, while the “low” input is the “special” term (that does not participate in actual functional evaluation). Evidently, \mathcal{L} is then essentially the leakage observed from SEC’s evaluate phase (i.e. $\mathcal{L}_{\text{SEC}}^{\text{EVALUATE}}$).

To establish statistical independence between \mathbf{H} and \mathcal{L} , we rely on Welch’s t-test (that is naturally applied when two populations have unequal variances). In our experiments, we allow the server to learn memory access patterns wrt. the aforementioned leakage profile of SEC.EVALUATE. The null hypothesis (for a two-tailed test) then tests whether the population means for two evaluations of SEC (while reusing SEC’s data structures) are indistinguishable. We initialize the “special” terms during SEC.KEYGEN as usual, and the client controls the permutations of the same over the execution of 1 million queries (i.e. a circuit consisting of 1 million gates). Our α value (i.e. the probability of incorrectly rejecting the null hypothesis when it is instead true) is 1%. Empirically, we observe a t-statistic of -1.7321 and a p-value of 0.0832. We hence conclude that there is not sufficient evidence to reject the null hypothesis. In other words, statistically, the server’s view (given SEC.EVALUATE’s leakage profile) of reusable executions in SEC is statistically indistinguishable from the execution of a randomly sampled circuit of the same size (for client-controlled permutation of “special” terms). This statistically establishes non-interference of \mathbf{H} in the leakage $\mathcal{L}_{\text{SEC}}^{\text{EVALUATE}}$.

5 Experimental Results

In this section, we report on a prototype implementation of SEC_{OXT} and compare it with a prototype implementation of the TFHE library [20], which implements an efficient and fast gate-by-gate bootstrapping [19].

Implementation Details. Our prototype implementations are developed in C++ and we use Redis as the database backend. More specifically, we realize all PRF operations using AES-256 in counter mode, BLAKE3 hash function for computing all hash operations, and all group operations over the elliptic curve Curve25519 [7].

Platform. For our experiments, we used a *single* node with 64-bit Intel Xeon Silver 4214R v4 3.27GHz processors, running Ubuntu 20.04.4 LTS, with 128GB RAM and 1TB SSD hard disk.

Evaluation Of Storage Overhead. As discussed in Section 3.5, the storage required for SEC scales with the number of keyword-document pairs and the number of search tokens for a particular function. In our implementations, the server storage required for SEC_{OXT} to store TSet¹⁶ and XSet¹⁷ and the TokenSet is around 43 KB while that for SEC_{CONJFILTER} is 26 KB. We thus note that SEC is highly optimized and scalable with significantly fewer storage

¹⁶OXT specific data structure to store inverted index for the “special” term.

¹⁷OXT specific data structure to check presence of “cross” terms in respective document identifiers.

requirements than state-of-the-art FHE schemes. Table 6 offers a quantifiable comparison.

Table 6: Storage Overhead comparison (in MB) of SEC with existing FHE schemes in literature. Storage overhead of FHE scheme typically indicates the bootstrapping key size whereas for SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$ it implies the size of the encrypted search index stored at the cloud server.

Scheme	Storage Overhead (in MB)
Gentry et. al[31]	3700
Gentry et. al[27]	2300
Halevi et. al.[36]	1600
Ducas et. al[24]	1000
Chillotti et. al.[19]	24
$\text{SEC}_{\text{CONJFILTER}}$ (This work)	0.449
SEC_{OXT} (This work)	0.098

Evaluation of Computation time. The evaluation time of SEC_{OXT} for computing arbitrary depth Boolean circuit over encrypted data scales linearly with the search time complexity of OXT times some constant which depends upon the depth of the circuit. The time required to retrieve the documents corresponding to a conjunctive query scales with the least frequent keyword in the query in OXT. Since the database (search index) in SEC is extremely small, the time taken by OXT.SEARCH is significantly less. Hence, the average time required by SEC_{OXT} is around 10 milliseconds for one binary function evaluation which is remarkably fast. On a similar note $\text{SEC}_{\text{CONJFILTER}}$ scales with the search complexity of CONJFILTER, which is dependent on the least frequent conjunct in the query. Notably $\text{SEC}_{\text{CONJFILTER}}$ exhibits even faster performance with an average evaluation time of 40 microseconds for one binary function evaluation. Our experimental results validate that SEC is highly efficient and extremely fast while evaluating arbitrary Boolean functions over encrypted data. Figure 3 compares the execution time of SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$ with different TFHE backends for varying depth of circuits.




Figure 2: Time taken (in seconds) for 1000 invocations of SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$ against different TFHE backends.

Comparison With FHE. We compare SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$ with different variations of TFHE in Figure 2. One variation is Nayuki portable (non-AVX) and AVX builds, which implement very efficient versions of Fast Fourier Transform. Another back-end family is spqlios AVX and spqlios FMA back-ends, which are efficient assembly implementations of ring operations. It is observed from Figure 2 that SEC_{OXT} is $10^3\times$ and six to seven




Figure 3: Time taken (in seconds) for different circuit depths of SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$ against different TFHE backends.

Table 7: Time taken (in minutes) and Storage overhead (in MB) for evaluation of AES-128 circuit and Maxpool function (AlexNet) by SEC_{OXT} and $\text{SEC}_{\text{CONJFILTER}}$ against different TFHE backends.

Scheme	Time (in minutes)		Storage (in MB)
	AES-128	Maxpool function	
TFHE-Nayuki Portable	336.03	2920.52	24
TFHE-Nayuki AVX	179.02	1441.18	24
TFHE-Spqlios AVX	57.37	523.62	24
TFHE-Spqlios FMA	41.87	349.30	24
$\text{SEC}_{\text{CONJFILTER}}$ (This work)	1.02	6.18	0.449
SEC_{OXT} (This work)	6.57	48.17	0.098

times faster; $\text{SEC}_{\text{CONJFILTER}}$ is $10^6 \times$ and $10^3 \times$ faster than the portable TFHE backend and the fastest (non-portable) TFHE backend Spqlios AVX, respectively. Figure 3 compares the increase in execution time with an increase in the depth of the circuit. Both instantiations of SEC outperforms the fastest TFHE backend using Spqlios AVX optimization for function evaluation of arbitrary depth.

We showcase SEC’s scalability for functions with multi-bit inputs by using it for encrypted evaluation of (i) the entire AES-128 circuit (with XOR/AND/NOT-gate count of 25124/6800/1692) and (ii) three max-pooling layers of AlexNet architecture¹⁸ (a circuit with OR-gate count of 289060). This requires no extra storage (since we still only require storage for three extra gates), and the performance figures (as well as a comparison with Torus-FHE) are described in Table 7. For both circuits, a (non-parallelized) implementation of SEC_{OXT} outperforms a (non-parallelized) implementation of Torus-FHE by six to seven orders of magnitude, while a (non-parallelized) implementation of $\text{SEC}_{\text{CONJFILTER}}$ shows an improvement of $10^3 \times$ in computation time (we expect the relative comparison to remaining unchanged with parallelization and additional hardware/software-level optimizations). SEC_{OXT} requires around $250 \times$ less storage while $\text{SEC}_{\text{CONJFILTER}}$ requires $50 \times$ times less storage, which are remarkably less. These results clearly showcase the efficiency and scalability of SEC for circuits with multi-bit inputs.

¹⁸KSH17 Imagenet classification with deep convolutional neural networks

6 Discussion

We conclude with a brief discussion comparing SEC with traditional FHE. The core technical difference between SEC and FHE is as follows: SEC models each Boolean gate as a truth table, and leverages encrypted look-ups for evaluating this truth table on an encrypted input, while FHE models each Boolean gate as an algebraic operation over some appropriate algebraically structured mathematical object (e.g., polynomial rings [27, 11, 32] or the Torus [19, 20]), and exploits the algebraic structure underlying each encrypted input to evaluate the gate. This offers an efficiency vs functionality tradeoff. As demonstrated empirically, SEC outperforms traditional FHE schemes significantly, both in terms of computation time and storage requirements, when operating over symmetrically encrypted data. On the other hand, the algebraic structure underlying FHE allows it to operate over publicly encrypted data, and we leave it as an interesting open question to extend the lookup-based approach underlying SEC to computing over publicly encrypted data.

We note, however, that in many practical applications (e.g., querying over outsourced encrypted databases), it suffices to support evaluation of arbitrary Boolean circuits over symmetrically encrypted data, since the data owner is also the primary entity querying the (encrypted) data after outsourcing it to an untrusted server for storage and processing. Indeed, this setting motivates the entire literature on SSE [43, 23, 17, 14], albeit for restricted classes of functions. To the best of our knowledge, our work is the first to establish the possibility of supporting arbitrary Boolean circuit evaluation efficiently over encrypted data using purely symmetric-key encryption techniques on top of lookup-based gate evaluation. Indeed, as demonstrated by our theoretical analysis and practical evaluation, the usage of purely symmetric-key primitives is what enables the highly desirable efficiency and compactness guarantees of SEC, allowing it to scale over extremely large symmetrically encrypted datasets while outperforming FHE.

References

- [1] Al Badawi, A., Bates, J., Bergamaschi, F., Cousins, D.B., Erabelli, S., Genise, N., Halevi, S., Hunt, H., Kim, A., Lee, Y., et al.: Openfhe: Open-source fully homomorphic encryption library. In: WAHC (2022)
- [2] Alperin-Sheriff, J., Peikert, C.: Practical bootstrapping in quasilinear time. In: Advances in Cryptology–CRYPTO 2013: 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2013. Proceedings, Part I (2013)
- [3] Alperin-Sheriff, J., Peikert, C.: Faster bootstrapping with polynomial error. In: Advances in Cryptology–CRYPTO 2014: 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17–21, 2014, Proceedings, Part I 34 (2014)
- [4] Barthe, G., Belaïd, S., Dupressoir, F., Fouque, P.A., Grégoire, B., Strub, P.Y.: Verified proofs of higher-order masking. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 457–485. Springer (2015)

- [5] Bellare, M., Hoang, V.T., Rogaway, P.: Foundations of garbled circuits. In: the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA (2012)
- [6] Ben-Efraim, A., Lindell, Y., Omri, E.: Efficient scalable constant-round MPC via garbled circuits. In: Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Proceedings, Part II (2017)
- [7] Bernstein, D.J.: Curve25519: New diffie-hellman speed records. In: Public Key Cryptography - PKC. Lecture Notes in Computer Science (2006)
- [8] Blackstone, L., Kamara, S., Moataz, T.: Revisiting leakage abuse attacks. In: 27th Annual Network and Distributed System Security Symposium, NDSS 2020, San Diego, California, USA (2020)
- [9] Boemer, F., Kim, S., Seifu, G., DM de Souza, F., Gopal, V.: Intel hexl: accelerating homomorphic encryption with intel avx512-ifma52. In: WAHC (2021)
- [10] Boneh, D., Sahai, A., Waters, B.: Functional encryption: Definitions and challenges. In: Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings 8 (2011)
- [11] Brakerski, Z., Gentry, C., Halevi, S.: Packed ciphertexts in lwe-based homomorphic encryption. In: Public-Key Cryptography–PKC 2013 (2013)
- [12] Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011 (2011)
- [13] Cash, D., Jaeger, J., Jarecki, S., Jutla, C.S., Krawczyk, H., Rosu, M., Steiner, M.: Dynamic searchable encryption in very-large databases: Data structures and implementation. In: NDSS 2014 (2014)
- [14] Cash, D., Jarecki, S., Jutla, C., Krawczyk, H., Roşu, M.C., Steiner, M.: Highly-scalable searchable symmetric encryption with support for boolean queries. In: CRYPTO (2013)
- [15] Cash, D., Jarecki, S., Jutla, C.S., Krawczyk, H., Rosu, M., Steiner, M.: Highly-scalable searchable symmetric encryption with support for boolean queries. In: CRYPTO 2013 (2013)
- [16] Chang, Z., Xie, D., Li, F.: Oblivious ram: A dissection and experimental evaluation. Proceedings of the VLDB Endowment (2016)
- [17] Chase, M., Kamara, S.: Structured encryption and controlled disclosure. In: ASIACRYPT 2010 (2010)
- [18] Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: ASIACRYPT 2017 (2017)
- [19] Chillotti, I., Gama, N., Georgieva, M., Izabachene, M.: Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In: ASIACRYPT 2016 (2016)

- [20] Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Tfhe: Fast fully homomorphic encryption library (2019)
- [21] Clark, D., Hunt, S., Malacaria, P.: Quantified interference: Information theory and information flow. In: Workshop on Issues in the Theory of Security (WITS'04) (2004)
- [22] Clearinghouse., P.R.: Chronology of data breaches. <https://privacyrights.org/data-breaches> (2024)
- [23] Curtmola, R., Garay, J., Kamara, S., Ostrovsky, R.: Searchable symmetric encryption: improved definitions and efficient constructions. In: ACM CCS (2006)
- [24] Ducas, L., Micciancio, D.: FHEW: bootstrapping homomorphic encryption in less than a second. In: EUROCRYPT 2015 (2015)
- [25] El-Yahyaoui, A., Kettani, M.D.E.E.: A verifiable fully homomorphic encryption scheme for cloud computing security. CoRR (2018)
- [26] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Proceedings of the forty-first annual ACM symposium on Theory of computing (2009)
- [27] Gentry, C., Halevi, S.: Implementing gentry's fully-homomorphic encryption scheme. In: EUROCRYPT 2011 (2011)
- [28] Gentry, C., Halevi, S., Peikert, C., Smart, N.P.: Ring switching in bgv-style homomorphic encryption. In: SCN (2012)
- [29] Gentry, C., Halevi, S., Smart, N.P.: Better bootstrapping in fully homomorphic encryption. In: PKC 2012 (2012)
- [30] Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. In: Eurocrypt (2012)
- [31] Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the aes circuit. In: CRYPTO 2012 (2012)
- [32] Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In: Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I (2013)
- [33] Goldreich, O., Ostrovsky, R.: Software protection and simulation on oblivious rams. Journal of the ACM (JACM) (1996)
- [34] Goldwasser, S., Kalai, Y.T., Popa, R.A., Vaikuntanathan, V., Zeldovich, N.: Reusable garbled circuits and succinct functional encryption. In: Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA. ACM (2013)
- [35] Goyal, V., Li, H., Ostrovsky, R., Polychroniadou, A., Song, Y.: ATLAS: efficient and scalable MPC in the honest majority setting. In: Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, Proceedings, Part II (2021)

- [36] Halevi, S., Shoup, V.: Design and implementation of helib: a homomorphic encryption library. Cryptology ePrint Archive (2020)
- [37] Kamara, S., Wei, L.: Garbled circuits via structured encryption. In: Financial Cryptography and Data Security - FC 2013 Workshops, USEC and WAHC 2013, Okinawa, Japan (2013)
- [38] Mentens, N., Batina, L., Preneel, B., Verbauwhede, I.: A systematic evaluation of compact hardware implementations for the rijndael s-box. In: Topics in Cryptology–CT-RSA 2005 (2005)
- [39] Okamoto, T., Takashima, K.: Fully secure functional encryption with general relations from the decisional linear assumption. In: Annual cryptology conference (2010)
- [40] Patel, S., Persiano, G., Seo, J.Y., Yeo, K.: Efficient boolean search over encrypted data with reduced leakage. In: ASIACRYPT 2021 (2021)
- [41] Patranabis, S., Mukhopadhyay, D.: Forward and backward private conjunctive searchable symmetric encryption. In: NDSS 2021 (2021)
- [42] Smart, N.P.: Practical and efficient fhe-based MPC. In: Quaglia, E.A. (ed.) Cryptography and Coding - 19th IMA International Conference, IMACC 2023, London, UK, Proceedings (2023)
- [43] Song, D.X., Wagner, D., Perrig, A.: Practical techniques for searches on encrypted data. In: Proceeding 2000 IEEE symposium on security and privacy. S&P 2000 (2000)
- [44] Yang, K., Wang, X., Zhang, J.: More efficient MPC from improved triple generation and authenticated garbling. In: CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, 2020 (2020)
- [45] Yang, Q., Peng, G., Gasti, P., Balagani, K.S., Li, Y., Zhou, G.: MEG: memory and energy efficient garbled circuit evaluation on smartphones. IEEE Trans. Inf. Forensics Secur. (2019)
- [46] Yuan, B., Jia, Y., Xing, L., Zhao, D., Wang, X., Zou, D., Jin, H., Zhang, Y.: Shattered chain of trust: Understanding security risks in cross-cloud iot access delegation. In: USENIX Security Symposium (2020)
- [47] Zhang, Y., Katz, J., Papamanthou, C.: Queries are belong to us: The power of file-injection attacks on searchable encryption (2016)
- [48] Zhou, W., Jia, Y., Yao, Y., Zhu, L., Guan, L., Mao, Y., Liu, P., Zhang, Y.: Discovering and understanding the security hazards in the interactions between iot devices, mobile apps, and clouds on smart home platforms. In: 28th USENIX Security Symposium (2019)