

SQL on Structurally-Encrypted Databases

Seny Kamara *
Brown University

Tarik Moataz†
Telecom Bretagne &
Colorado State University

Abstract

We show how to encrypt a relational database in such a way that it can efficiently support a large class of SQL queries. Our construction is based solely on structured encryption and does not make use of any property-preserving encryption (PPE) schemes such as deterministic and order-preserving encryption. As such, our approach leaks considerably less than PPE-based solutions which have recently been shown to reveal a lot of information in certain settings (Naveed et al., *CCS '15*). Our construction achieves asymptotically optimal query complexity under very natural conditions on the database and queries.

1 Introduction

The problem of encrypted search has received a lot of attention from Industry, Academia and Government due to its potential applications to cloud computing and database security. Most of the progress in this area, however, has been in the setting of keyword search on encrypted documents. While this has many applications in practice (e.g., email, NoSQL databases, desktop search engines, cloud document storage), much of the data produced and consumed in practice is stored and processed in *relational databases*. A relational database is, roughly speaking, a set of tables with rows representing entities/items and columns representing their attributes. The relational database model was proposed by Codd [14] and most relational DBs are queried using the structured query language (SQL) which is a special-purpose declarative language introduced by Chamberlain and Boyce [10].

The problem of encrypted relational DBs is the “holy-grail” of database security. As far as we know, it was first explicitly considered by Hacigümüs, Iyer, Li and Mehrotra [20] who described a quantization-based approach which leaks the range within which an item falls. In [29], Popa, Redfield, Zeldovich and Balakrishnan describe a system called CryptDB that can support a non-trivial subset of SQL on encrypted relational DBs without quantization. CryptDB achieves this in part by making use of property-preserving encryption (PPE) schemes like deterministic and order-preserving (OPE) encryption, which reveal equality and order, respectively. The high-level technique is to replace the plaintext operations needed to execute a SQL query (e.g., equality tests

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

†tarik.moataz@colostate.edu. Work done in part at Microsoft Research.

and comparisons) by the same operations on PPE-encrypted ciphertexts. This approach was later adopted by other systems including Cipherbase [2] and SEED [19]. While this leads to systems that are efficient and legacy-friendly, it was recently shown by Naveed, Kamara and Wright [26], that such PPE-based EDB systems can leak a lot of information when used in certain settings like electronic medical records (EMRs). In light of this result, the major open problem in encrypted search and, more generally, in database security is whether it is possible to efficiently execute SQL queries on encrypted DBs with less leakage than the PPE-based solutions.

Our contributions. In this work, we address this problem and propose the first solution for SQL on encrypted DBs that does not make use of either PPE or general-purpose primitives like fully-homomorphic encryption (FHE) or oblivious RAM (ORAM). As such, our scheme leaks considerably less than any of the previously-known practical approaches and is much more practical than any leakage-free solution. Our approach is efficient and handles a large sub-class of SQL queries and an even larger class if we allow for a small amount of post-processing at the client.

More precisely, our construction handles the class of *conjunctive queries*¹ [11] which corresponds to SQL queries of the form

$$\text{Select attributes From tables Where } (\mathbf{a}_1 = X_1 \wedge \dots \wedge \mathbf{a}_\ell = X_\ell),$$

where \mathbf{a}_1 through \mathbf{a}_n are attributes in the DB schema and X_1 through X_n are either attributes or constants. While the class of conjunctive queries is smaller than the class supported by the PPE-based solutions, it is one of the most well-studied and useful classes of queries. Furthermore, as mentioned above, if one allows for a small amount of post-processing at the client, we show how to extend the expressiveness of our solution to a much wider sub-class.

With respect to efficiency, we show that the query complexity of our scheme is asymptotically optimal in time and space when $t, s_1, \dots, s_t, h \ll m$, where t denotes the number of tables in the query, s_i denotes the number of columns in the i th table, h denotes the number of attributes in the Select term of the query and m is the number of rows in the tables (assuming, for ease of exposition, that all tables have m rows). Note that this is a very natural condition which is satisfied in practice by most SQL databases and queries.²

1.1 Our Techniques

The PPE-based approach to EDBs essentially replaces the plaintext execution of a SQL query with an encrypted execution of the query by executing the server’s low-level operations (i.e., comparisons and equality tests) directly on the encrypted cells. This can be done thanks to the properties of PPE which guarantee that operations on plaintexts can be done on ciphertexts as well. This “plug-and-play” approach makes the design of EDBs relatively straightforward since the only requirement is to replace plaintext cells with PPE-encrypted cells. Given the complexity of relational DBs and of SQL queries it is not a-priori clear how to solve this problem without PPE or without resorting to general-purpose solutions like fully-homomorphic encryption (FHE) or oblivious RAM (ORAM).

¹We stress that conjunctive queries in the context of relational databases (and as used throughout this work) is completely unrelated to conjunctive *keyword* queries as studied in the searchable encryption literature (e.g., in [8]). In particular, our scheme does not make use of any searchable encryption schemes for conjunctive keyword queries.

²Our approach incurs only a small asymptotic overhead even when these parameters are unrealistically large and on the order of m

Conceptual approach. Our first step towards a solution is in isolating some of the conceptual difficulties of the problem. Relational DBs are relatively simple from a data structure perspective since they just consist of a set of two-dimensional arrays. The high-level challenge stems from SQL and, in particular, from its complexity (it can express first-order logic) and the fact that it is *declarative*. To overcome this we restrict ourselves to a simpler but widely applicable and well-studied subset of SQL queries (see above) and we take a more procedural view of SQL. More precisely, we work with the *relational algebra* formulation of SQL which is more amenable to cryptographic techniques. The relational algebra was introduced by Codd [14] as a way to formalize queries on relational databases. Roughly speaking, it consists of all the queries that can be expressed from a set of basic operations. It was later shown by Chandra and Merlin [11] that three of these operations (selection, projection and cross product) capture a large class of useful queries called *conjunctive queries* that have particularly nice theoretical properties. Since their introduction, conjunctive queries have been studied extensively in the database literature.

The subset of the relational algebra expressed by the selection, projection and Cartesian product operators is also called the *SPC algebra*. By working in the SPC algebra, we not only get a procedural representation of SQL queries, but we also reduce the problem to handling just three basic operations. Conceptually, this is reminiscent of the benefits one gets by working with circuits in secure multi-party computation and FHE. Another important advantage of working in the SPC algebra is that it admits a *normal form*; that is, every SPC query can be written in a standard form. By working with this normal form, we get another benefit of general-purpose solutions which are that we can design and analyze a *single* construction that handles *all* SPC queries. Note, however, that like circuit representations the SPC normal form is not always guaranteed to be the most efficient.

The SPC algebra. As mentioned, the SPC algebra consists of all queries that can be expressed by a combination of the select, project and cross product operators which, at a high-level, work as follows. The select operator σ_Ψ takes as input a table T and outputs the rows of T that satisfy the predicate Ψ . The project operator π_{a_1, \dots, a_h} takes as input a table T and outputs the columns of T indexed by a_1, \dots, a_h . Finally, the cross product operator $T_1 \times T_2$ takes two tables as input and outputs a third table consisting of rows in the Cartesian product of T_1 and T_2 when viewed as sets of rows. An SPC query in normal form over a database $DB = (T_1, \dots, T_n)$ has the form,

$$\pi_{a_1, \dots, a_h} \left([a_1] \times \dots \times [a_f] \times \sigma_\Psi(T_{i_1} \times \dots \times T_{i_t}) \right),$$

where Ψ is of the form $a_1 = X_1 \wedge \dots \wedge a_\ell = X_\ell$ where a_1, \dots, a_ℓ are attributes in the schema of DB and X_1, \dots, X_ℓ are either attributes or constants. So, concretely, our problem reduces to the problem of encrypting a relational database $DB = (T_1, \dots, T_n)$ in such a way that it can support SPC queries in normal form.

Structured encryption & constructive queries. The concrete problem stated above is a structured encryption problem: namely, how can we encrypt a data structure (i.e., DB) such that it can support a specific type of query (i.e., SPC normal form queries). Structured encryption (STE) is a generalization of searchable symmetric encryption (SSE) introduced by Chase and Kamara [12] which allows a client to encrypt a data structure so that it can be privately queried. As discussed above, relational DBs are relatively simple structures (i.e., two-dimensional arrays)

but the queries they support are complex. This is in opposition to most STE schemes which handle relatively complex structures such as dictionaries [12, 7], multi-maps (i.e., inverted indexes) [15, 23, 8], trees [22], graphs [12, 25] and web graphs [12], but support relatively simple queries. The main difficulty in the case of relational DBs and, in particular, in handling SPC queries is that queries are *constructive* in the sense that they produce new data structures from the original base structure. Intuitively, handling constructive queries (without interaction) is particularly challenging because the intermediate and final structures that have to be created by the server to answer the query are *dependent* on the query and, therefore, cannot be constructed by the client in the setup/pre-processing phase.

An important observation about relational DBs that underlies our approach, however, is that while SPC queries are constructive, they are not arbitrarily so. In other words, the tables needed to answer an SPC query are not completely arbitrary but are structured in a way that can be predicted at setup. What is query-dependent is the *content* of these tables but, crucially, all of that content is already stored in the original database. So the challenge then is to provide the server with the means to construct the appropriate intermediate and final tables and to design encrypted structures that will allow it to efficiently find the (encrypted) content it needs to populate those tables.

Cross products & the location map. By taking a closer look at the SPC normal form, one can see that the first intermediate table needed to answer a query is the cross product $T' = T_{i_1} \times \dots \times T_{i_t}$. Ignoring the cross products with $[a_1], \dots, [a_f]$ for ease of exposition, the remaining intermediate tables as well as the final table are “sub-tables” of T' that result from selecting a subset of rows (according to Ψ) and keeping a subset of columns (according to $\mathbf{a}_1, \dots, \mathbf{a}_h$). As we show in Section 3, it turns out that from only t and the dimensions of T_{i_1} through T_{i_t} , we can derive not only the dimensions of T' but also a map Λ that maps any cell in DB to its location in T' . We refer to this map as the *location map* and provide it to the server so that it can place the necessary encrypted content in T' . To help the server find the necessary content in the first place, we create a set of encrypted structures which, roughly speaking, can be viewed as different representations of the database. For example, one of the encrypted structures we build provides a row-wise representation of the database whereas another one provides a column-wise representation. By using these various representations and by “combining” them in an appropriate manner, we can generate tokens for the server to recover the (encrypted) database cells that result from applying the select and projection operators to T' . Once the server can recover those encrypted cells it can place them in T' using the location map and prune T' to produce the final table.

Note that, as described, our solution is not practical as the server needs to generate and operate on a m^t -by- $\sum_{j=1}^t s_{i_j}$ table, where s_{i_j} denotes the number of columns in T_{i_j} . Fortunately, this is not necessary (we describe this naive approach in Section 5 only for ease of exposition) and queries can be processed in optimal time and space using a more complex algorithm which we describe in Section 5.1.

A note on our techniques. We stress that our approach to handling the SPC algebra is completely novel and has no resemblance to how these queries are usually handled on plaintext databases. In other words, our approach does not simply replicate standard data structures and algorithms from databases. In fact, as far as we know, currently the only way to re-use standard data structures and algorithms on encrypted data is to use PPE as done, for example, in [29, 2, 19].

2 Related Work

Searchable & structured encryption. Encrypted search was first considered explicitly by Song, Wagner and Perrig in [30] which introduced the notion of searchable symmetric encryption (SSE). Goh provided the first security definition for SSE and a solution based on Bloom filters with linear search complexity. Chang and Mitzenmacher proposed an alternative security definition and construction, also with linear search complexity. Curtmola, Garay, Kamara and Ostrovsky introduced and formulated the notion of adaptive semantic security for SSE [15] together with optimal-time and optimal-space constructions. Chase and Kamara introduced the notion of structured encryption which generalizes SSE to arbitrary data structures [12]. Kurosawa and Ohtaki proposed a universally composable variant of adaptive semantic security [24]. Kamara, Papamanthou and Roeder [23] gave the first optimal-time dynamic SSE scheme. Cash, Jutla, Jarecki, Krawczyk, Rosu and Steiner [8] proposed the first optimal-time scheme that handles conjunctive keyword search and Faber, Krawczyk, Jarecki, Nguyen, Rosu and Steiner show how to extend it to rich queries (e.g., range, substring and wildcard queries) [16]. Naveed, Prabakaran and Gunther [27] propose an optimal-time dynamic SSE scheme based on blind storage. Cash, Jaeger, Jarecki, Jutla, Krawczyk, Rosu and Steiner [7] show how to construct optimal-time SSE schemes with low I/O complexity and Cash and Tessaro [9] gave lower bounds on the locality of adaptively-secure SSE schemes. SSE has also been considered in the multi-user setting [15, 21]. Pappas, Krell, Vo, Kolesnikov, Malkin, Choi, George, Keromytis, Bellovin [28] present another approach for multi-user encrypted search based on garbled circuits and Bloom filters that can support Boolean formulas, ranges and stemming. Other approaches for encrypted search include oblivious RAMs (ORAM) [18], secure multi-party computation [4], functional encryption [6] and fully-homomorphic encryption [17] as well as solutions based on deterministic encryption [3] and order-preserving encryption (OPE) [5].

Encrypted relational databases. As far as we know the first encrypted relational DB solution was proposed by Hacigümüs, Iyer, Li and Mehrotra [20] and was based on quantization. Roughly speaking, the attribute space of each column is partitioned into bins and each element in the column is replaced with its bin number. Popa, Redfield, Zeldovich and Balakrishnan proposed CryptDB [29]. CryptDB was the first non-quantization-based solution and can handle a large subset of SQL. Instead of quantization, CryptDB relies on PPE like deterministic [3] and OPE [1, 5]. The CryptDB design influenced the Cipherbase system from Arasu *et al.* [2] and the SEED system from Grofig *et al.* [19]. In [26], Naveed, Kamara and Wright study the security of these PPE-based solutions in the context of medical data.

3 Preliminaries

Notation. The set of all binary strings of length n is denoted as $\{0, 1\}^n$, and the set of all finite binary strings as $\{0, 1\}^*$. $[n]$ is the set of integers $\{1, \dots, n\}$. We write $x \leftarrow \chi$ to represent an element x being sampled from a distribution χ , and $x \xleftarrow{\$} X$ to represent an element x being sampled uniformly at random from a set X . The output x of an algorithm \mathcal{A} is denoted by $x \leftarrow \mathcal{A}$. Given a sequence \mathbf{v} of n elements, we refer to its i th element as v_i or $\mathbf{v}[i]$. If S is a set then $\#S$ refers to its cardinality. If s is a string then $|s|$ refers to its bit length.

Basic cryptographic primitives. A private-key encryption scheme is a set of three polynomial-time algorithms $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ such that Gen is a probabilistic algorithm that takes a security parameter k and returns a secret key K ; Enc is a probabilistic algorithm that takes a key K and a message m and returns a ciphertext c ; Dec is a deterministic algorithm that takes a key K and a ciphertext c and returns m if K was the key under which c was produced. Informally, a private-key encryption scheme is secure against chosen-plaintext attacks (CPA) if the ciphertexts it outputs do not reveal any partial information about the plaintext even to an adversary that can adaptively query an encryption oracle.

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

Data structures. A *data structure* for a given data type is a representation in some computational model³ of an object of the given type. Typically, the representation is optimized to support the type’s query operation as efficiently as possible. For data types that support multiple queries, the representation is often optimized to efficiently support as many queries as possible. As a concrete example, the dictionary *type* can be represented using various data structures depending on which queries one wants to support efficiently. Hash tables support Get and Put in expected $O(1)$ time whereas balanced binary search trees support both operations in worst-case $\log(n)$ time. For ease of understanding and to match colloquial usage, we sometimes blur the distinction between data types and structures. So, for example, when referring to a *dictionary structure* or a *multi-map structure* what we are referring to is an unspecified instantiation of the dictionary or multi-map data type.

We make use of several basic data types including arrays, dictionaries and multi-maps which we recall here. An array \mathbf{A} of capacity n stores n items at locations 1 through n and supports read and write operations. We write $v := \mathbf{A}[i]$ to denote reading the item at location i and $\mathbf{A}[i] := v$ the operation of storing an item at location i . A dictionary \mathbf{DX} of capacity n is a collection of n label/value pairs $\{(\ell_i, v_i)\}_{i \leq n}$ and supports get and put operations. We write $v_i := \mathbf{DX}[\ell_i]$ to denote getting the value associated with label ℓ_i and $\mathbf{DX}[\ell_i] := v_i$ to denote the operation of associating the value v_i in \mathbf{DX} with label ℓ_i . A multi-map \mathbf{MM} with capacity n is a collection of n label/tuple pairs $\{(\ell_i, \mathbf{t}_i)\}_{i \leq n}$ that supports get and put operations. Similarly to dictionaries, we write $\mathbf{t}_i := \mathbf{MM}[\ell_i]$ to denote getting the tuple associated with label ℓ_i and $\mathbf{MM}[\ell_i] := \mathbf{t}_i$ to denote operation of associating the tuple \mathbf{t}_i to label ℓ_i . Multi-maps are the abstract data type instantiated by an inverted index. In the encrypted search literature multi-maps are sometimes referred to as indexes, databases or tuple-sets (T-sets). We refer to the set of all possible queries a data structure supports as its *query space* and to the set of its possible responses as its *response space*. For some data structure \mathbf{DS} we sometimes write $\mathbf{DS} : \mathbf{Q} \rightarrow \mathbf{R}$ to mean that \mathbf{DS} has query and response spaces \mathbf{Q} and \mathbf{R} , respectively.

Relational databases. A relational database $\mathbf{DB} = (\mathbf{T}_1, \dots, \mathbf{T}_n)$ is a set of *tables* where each table \mathbf{T}_i is a two-dimensional array with rows corresponding to an entity (e.g., a customer or an employee) and columns corresponding to attributes (e.g., age, height, salary). For any given

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

attribute, we refer to the set of all possible values that it can take as its *space* (e.g., integers, booleans, strings). We define the *schema* of a table T to be its set of attributes and denote it $\mathbb{S}(T)$. The schema of a database $DB = (T_1, \dots, T_n)$ is then the set $\mathbb{S}(DB) = \bigcup_i \mathbb{S}(T_i)$. We assume the attributes in $\mathbb{S}(DB)$ are unique and represented as positive integers. We denote a table T 's number of rows as $\|T\|_r$ and its number of columns as $\|T\|_c$.

We sometimes view tables as a tuple of rows and write $\mathbf{r} \in T$ and sometimes as a tuple of columns and write $\mathbf{c} \in T^\top$. Similarly, we write $\mathbf{r} \in DB$ and $\mathbf{c} \in DB^\top$ for $\mathbf{r} \in \bigcup_i T_i$ and $\mathbf{c} \in \bigcup_i T_i^\top$, respectively. For a row $\mathbf{r} \in T_i$, its table identifier $\text{tbl}(\mathbf{r})$ is i and its row rank $\text{rrk}(\mathbf{r})$ is its position in T_i when viewed as a tuple of rows. Similarly, for a column $\mathbf{c} \in T_i^\top$, its table identifier $\text{tbl}(\mathbf{c})$ is i and its column rank $\text{crk}(\mathbf{c})$ is its position in T_i when viewed as a tuple of columns. For any row $\mathbf{r} \in DB$ and column $\mathbf{c} \in DB^\top$, we refer to the pairs $(\text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}))$ and $(\text{tbl}(\mathbf{c}), \text{crk}(\mathbf{c}))$, respectively, as their *coordinates* in DB . We write $\mathbf{r}[i]$ and $\mathbf{c}[i]$ to refer to the i th element of a row \mathbf{r} and column \mathbf{c} . The coordinate of the j th cell in row $\mathbf{r} \in T_i$ is the triple $(i, \text{rrk}(\mathbf{r}), j)$. Given a column $\mathbf{c} \in DB^\top$, we denote its corresponding attribute by $\text{att}(\mathbf{c})$. For any pair of attributes $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{S}(DB)$ such that $\text{dom}(\mathbf{a}_1) = \text{dom}(\mathbf{a}_2)$, $DB_{\mathbf{a}_1=\mathbf{a}_2}$ denotes the set of row pairs $\{(\mathbf{r}_1, \mathbf{r}_2) \in DB^2 : \mathbf{r}_1[\mathbf{a}_1] = \mathbf{r}_2[\mathbf{a}_2]\}$. For any attribute $\mathbf{a} \in \mathbb{S}(DB)$ and constant $a \in \bigcup_{\mathbf{a} \in \mathbb{S}(DB)} \text{dom}(\mathbf{a})$, $DB_{\mathbf{a}=a}$ is the set of rows $\{\mathbf{r} \in DB : \mathbf{r}[\mathbf{a}] = a\}$.

SQL. In practice, relational databases are queried using the special-purpose language SQL, introduced by Chamberlain and Boyce [10]. SQL is a declarative language and can be used to modify and query a relational DB. In this work, we only focus on its query operations. Informally, SQL queries typically have the form

Select *attributes* From *tables* Where *condition*,

where *attributes* is a set of attributes/columns, *tables* is a set of tables and *condition* is a predicate over the rows of *tables* and can itself contain a nested SQL query. More complex queries can be obtained using Group-by, Order-by and aggregate operators (i.e., max, min, average etc.) but the simple form above already captures a large subset of SQL. The most common class of queries on relational DBs are *conjunctive queries* [11] which have the above form with the restriction that *condition* is a conjunction of equalities over attributes and constants. In particular, this means there are no nested queries in *condition*. More precisely, conjunctive queries have the form

Select *attributes* From *tables* Where $(\mathbf{a}_1 = X_1 \wedge \dots \wedge \mathbf{a}_\ell = X_\ell)$,

where \mathbf{a}_i is an attribute in $\mathbb{S}(DB)$ and X_i can be either an attribute or a constant.

The SPC algebra. It was shown by Chandra and Merlin [11] that conjunctive queries could be expressed as a subset of Codd's relational algebra which is an imperative query language based on a set of basic operators. In particular, they showed that three operators *select*, *project* and *cross product* were enough. The *select* operator σ_Ψ is parameterized with a predicate Ψ and takes as input a table T and outputs a new table T' that includes the rows of T that satisfy the predicate Ψ . The *projection* operator $\pi_{\mathbf{a}_1, \dots, \mathbf{a}_h}$ is parameterized by a set of attributes $\mathbf{a}_1, \dots, \mathbf{a}_h$ and takes as input a table T and outputs a table T' that consists of the columns of T indexed by \mathbf{a}_1 through \mathbf{a}_h . The *cross product* operator \times takes as input two tables T_1 and T_2 and outputs a new table $T' = T_1 \times T_2$ such that each row of T' is an element of the Cartesian product between the set of rows of T_1 and the set of rows of T_2 . The query language that results from any combination of select, project and cross product is referred to as the *SPC algebra*. We formalize this in Definition 3.1 below.

Definition 3.1 (SPC algebra). Let $DB = (T_1, \dots, T_n)$ be a relational database. The SPC algebra consists of any query that results from the combination of the following operators:

- $T' \leftarrow \sigma_\Psi(T)$: the select operator is parameterized with a predicate Ψ of form $a_1 = X_1 \wedge \dots \wedge a_\ell = X_\ell$, where $a_i \in \mathbb{S}(DB)$ and X_i is either a constant a in the domain of a_i (type-1) or an attribute $x_j \in \mathbb{S}(DB)$ (type-2). It takes as input a table $T \in DB$ and outputs a table $T' = \{\mathbf{r} \in T : \Psi(\mathbf{r}) = 1\}$, where terms of the form $a_i = x_j$ are satisfied if $\mathbf{r}(a_i) = \mathbf{r}(x_j)$ and terms of the form $a_i = a$ are satisfied if $\mathbf{r}(a_i) = a$.
- $T' \leftarrow \pi_{a_1, \dots, a_h}(T)$: the project operator is parameterized by a set of attributes $a_1, \dots, a_h \in \mathbb{S}(DB)$. It takes as input a table $T \in DB$ and outputs a table $T' = \{\langle \mathbf{r}(a_1), \dots, \mathbf{r}(a_h) \rangle : \mathbf{r} \in T\}$.
- $T_3 \leftarrow T_1 \times T_2$: the cross product operator takes as input two tables T_1 and T_2 and outputs a table $T_3 = \{\langle \mathbf{r}, \mathbf{v} \rangle : \mathbf{r} \in T_1 \text{ and } \mathbf{v} \in T_2\}$, where $\langle \mathbf{r}, \mathbf{v} \rangle$ is the concatenation of rows \mathbf{r} and \mathbf{v} .

Intuitively, the connection between conjunctive SQL queries and the SPC algebra can be seen as follows: Select corresponds to the projection operator, From to the cross product and Where to the (SPC) select operator.

SPC normal form. Any query in the SPC algebra can be reduced to a *normal form* using a certain set of well-known identities. The normal form of an SPC query over a relational database $DB = (T_1, \dots, T_n)$ has the form:

$$\pi_{a_1, \dots, a_h} \left([a_1] \times \dots \times [a_f] \times \sigma_\Psi(T_{i_1} \times \dots \times T_{i_\ell}) \right),$$

where $a_1, \dots, a_f \in \bigcup_{a \in \mathbb{S}(DB)} \text{dom}(a)$ and $[a_j]$ is the 1×1 table that holds a_j . Here, the attributes a_1, \dots, a_h in the projection are either in $\mathbb{S}(DB)$ or refer to the columns generated by $[a_1]$ through $[a_h]$. In the latter case, we say that they are *virtual attributes* and are in $\mathbb{S}(VDB)$, where VDB is the *virtual database* defined as $VDB = ([a_1], \dots, [a_f])$.

We note that converting SQL queries of the above form to SPC queries is a well-studied problem with highly-optimized solutions. In particular, the SPC queries that result from such a translation are “compact” in the sense that the number of projects, selects and cross products in the resulting SPC query is the same as the number of attributes, tables and conditions, respectively, in the original SQL query (for an overview of SQL-to-SPC translation we refer the reader to [31]). In addition, the transformation of SPC queries into their normal form can also be done efficiently based on a standard set of identities and also results in “compact” normal form queries in the sense above.

The location map. An important notion for our purposes is what we refer to as the *location map* which, intuitively, maps the cell coordinates in DB to their location in the cross product $T_{i_1} \times \dots \times T_{i_t}$. Here, we assume the tables T_{i_1}, \dots, T_{i_t} all have the same number of rows m . More precisely, let M be a $m^t \times \sum_i s_i$ two-dimensional array. The location map $\Lambda_{t,m,s}$ is parameterized by the number of tables in the cross product $t \geq 1$, the size of the tables $m \geq 1$ and a vector $\mathbf{s} = (s_1, \dots, s_t)$ such that $s_i = \|T_{i_i}\|_c$. It takes as input the coordinates of any cell in DB and outputs

a set of locations $(\alpha, \beta) \in [m^t] \times [\sum_i s_i]$ in \mathbf{M} . Concretely, the location map is defined as

$$\Lambda_{t,m,s}(z, i, j) = \left\{ \left((i-1) \cdot m^{t-z} + h \cdot m^{t-z+1}, j + \sum_{\ell=1}^{z-1} \|\mathbb{T}_\ell\|_c \right), \dots, \right. \\ \left. \left(i \cdot m^{t-z} + h \cdot m^{t-z+1}, j + \sum_{\ell=1}^{z-1} \|\mathbb{T}_\ell\|_c \right) \right\}_{h \in \{1, \dots, m^{z-1}\}},$$

where $z \in \{i_1, \dots, i_t\}$, $i \in [m]$ and $j \in \|\mathbb{T}_z\|_c$.

Note that given a position $(\alpha, \beta) \in [m^t] \times [\sum_i s_i]$, and a specific order of tables $\mathbb{T}_{i_1}, \dots, \mathbb{T}_{i_z}$, the original coordinates of the cell in the tables can be computed *efficiently* such that

$$(z, i, j) = \Lambda_{t,m,s}^{-1}(\alpha, \beta),$$

where $i = \alpha \bmod m^{t-z}$ and $j = \beta \bmod \sum_{\ell=1}^z \|\mathbb{T}_\ell\|_c$.

4 Definitions

In this Section, we define the syntax and security of STE schemes. A STE scheme encrypts data structures in such a way that they can be privately queried. There are several natural forms of structured encryption. The original definition of [12] considered schemes that encrypt both a structure and a set of associated data items (e.g., documents, emails, user profiles etc.). In [13], the authors also describe *structure-only* schemes which only encrypt structures. Another distinction can be made between *interactive* and *non-interactive* schemes. Interactive schemes produce encrypted structures that are queried through an interactive two-party protocol, whereas non-interactive schemes produce structures that can be queried by sending a single message, i.e, the token. One can also distinguish between *response-hiding* and *response-revealing* schemes: the latter reveal the query response to the server whereas the former do not.

In this work, we focus on non-interactive structure-only schemes. Our main construction, SPX, is response-hiding but makes use of response-revealing schemes as building blocks. As such, we define both forms below. At a high-level, non-interactive STE works as follows. During a setup phase, the client constructs an encrypted structure EDS under a key K from a plaintext structure DS. The client then sends EDS to the server. During the query phase, the client constructs and sends a token tk generated from its query q and secret key K . The server then uses the token tk to query EDS and recover either a response r or an encryption ct of r depending on whether the scheme is response-revealing or response-hiding.

Definition 4.1 (Response-revealing structured encryption [12]). *A response-revealing structured encryption scheme $\Sigma = (\text{Setup}, \text{Token}, \text{Query})$ consists of three polynomial-time algorithms that work as follows:*

- $(K, \text{EDS}) \leftarrow \text{Setup}(1^k, \text{DS})$: is a probabilistic algorithm that takes as input a security parameter 1^k and a structure DS and outputs a secret key K and an encrypted structure EDS.
- $\text{tk} \leftarrow \text{Token}(K, q)$: is a (possibly) probabilistic algorithm that takes as input a secret key K and a query q and returns a token tk.

- $\{\perp, r\} \leftarrow \text{Query}(\text{EDS}, \text{tk})$: is a deterministic algorithm that takes as input an encrypted structure EDS and a token tk and outputs either \perp or a response.

We say that a response-revealing structured encryption scheme Σ is correct if for all $k \in \mathbb{N}$, for all $\text{poly}(k)$ -size structures $\text{DS} : \mathbf{Q} \rightarrow \mathbf{R}$, for all (K, EDS) output by $\text{Setup}(1^k, \text{DS})$ and all sequences of $m = \text{poly}(k)$ queries q_1, \dots, q_m , for all tokens tk_i output by $\text{Token}(K, q_i)$, $\text{Query}(\text{EDS}, \text{tk}_i)$ returns $\text{DS}(q_i)$ with all but negligible probability.

Definition 4.2 (Response-hiding structured encryption [12]). A response-hiding structured encryption scheme $\Sigma = (\text{Setup}, \text{Token}, \text{Query}, \text{Dec})$ consists of four polynomial-time algorithms such that Setup and Token are as in Definition 4.1 and Query and Dec are defined as follows:

- $\{\perp, \text{ct}\} \leftarrow \text{Query}(\text{EDS}, \text{tk})$: is a deterministic algorithm that takes as input an encrypted structured EDS and a token tk and outputs either \perp or a ciphertext ct.
- $r \leftarrow \text{Dec}(K, \text{ct})$: is a deterministic algorithm that takes as input a secret key K and a ciphertext ct and outputs a response r .

We say that a response-hiding structured encryption scheme Σ is correct if for all $k \in \mathbb{N}$, for all $\text{poly}(k)$ -size structures $\text{DS} : \mathbf{Q} \rightarrow \mathbf{R}$, for all (K, EDS) output by $\text{Setup}(1^k, \text{DS})$ and all sequences of $m = \text{poly}(k)$ queries q_1, \dots, q_m , for all tokens tk_i output by $\text{Token}(K, q_i)$, $\text{Dec}_K\left(\text{Query}\left(\text{EDS}, \text{tk}_i\right)\right)$ returns $\text{DS}(q_i)$ with all but negligible probability.

Security. The standard notion of security for structured encryption guarantees that an encrypted structure reveals no information about its underlying structure beyond the setup leakage \mathcal{L}_S and that the query algorithm reveals no information about the structure and the queries beyond the query leakage \mathcal{L}_Q . If this holds for non-adaptively chosen operations then this is referred to as non-adaptive semantic security. If, on the other hand, the operations are chosen adaptively, this leads to the stronger notion of adaptive semantic security. This notion of security was introduced by Curtmola *et al.* in the context of SSE [15] and later generalized to structured encryption in [12].

Definition 4.3 (Adaptive semantic security [15, 12]). Let $\Sigma = (\text{Setup}, \text{Token}, \text{Query})$ be a response-revealing structured encryption scheme and consider the following probabilistic experiments where \mathcal{A} is a stateful adversary, \mathcal{S} is a stateful simulator, \mathcal{L}_S and \mathcal{L}_Q are leakage profiles and $z \in \{0, 1\}^*$:

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

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

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

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

The security definition for *response-hiding* schemes can be derived from Definition 4.3 by giving the simulator $(\perp, \mathcal{L}_Q(\text{DS}, q_i))$ instead of $(\text{DS}(q_i), \mathcal{L}_Q(\text{DS}, q_i))$.

5 Our Construction

We describe SPX, our structured encryption scheme for SPC queries on relational DBs. More precisely, the scheme encrypts relational database structures DB which we define as tuples of two-dimensional arrays $(\mathsf{T}_1, \dots, \mathsf{T}_n)$ which support SPC queries in normal form. The scheme makes black-box use of a response-revealing multi-map encryption scheme $\Sigma_{\text{MM}} = (\text{Setup}, \text{Token}, \text{Get})$, of a response-revealing dictionary encryption scheme $\Sigma_{\text{DX}} = (\text{Setup}, \text{Token}, \text{Get})$ and of a symmetric-key encryption scheme $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$. Note that encrypted multi-maps and dictionaries can be instantiated using a variety of schemes [15, 12, 23, 8, 7, 27].

Our construction is detailed in Figs. 1, 2 and 3 and we provide high-level overviews of its setup, token generation and query algorithms. We stress that, here, we only describe a *naive* query algorithm that is not efficient but that is relatively easy to understand. In Section 5.1 below, we describe our optimized query algorithm which is optimal under natural conditions.

Setup. The Setup algorithm takes as input a relational database $\text{DB} = (\mathsf{T}_1, \dots, \mathsf{T}_n)$ and starts by creating three multi-maps MM_R , MM_C and MM_V , each of which stores a different representation of the database. MM_R stores its row-wise representation, by which we mean that it maps row coordinates (i.e., the row table and rank) to encrypted rows. Analogously, MM_C stores a column-wise representation of DB and maps column coordinates (i.e., the column table and rank) to encrypted columns. Setup then encrypts MM_R and MM_C with Σ_{MM} , resulting in EMM_R and EMM_C . It then uses MM_V to store a value-wise representation of the database in the sense that the latter maps the value of a cell to *tokens* for the rows in DB that store that same value (where the tokens are for EMM_R). Again, it uses Σ_{MM} to encrypt MM_V , resulting in EMM_V . Now, recall the SPC normal form:

$$\pi_{\mathbf{a}_1, \dots, \mathbf{a}_h} \left([a_1] \times \dots \times [a_f] \times \sigma_{\Psi}(\mathsf{T}_{i_1} \times \dots \times \mathsf{T}_{i_t}) \right).$$

At a high-level, EMM_C will enable projection operations and EMM_V combined with EMM_R will enable type-1 select operations (i.e., of the form $\mathbf{a}_i = a_i$). To support type-2 select operations, however, we need additional structures.

For this, Setup will create, for all columns $\mathbf{c} \in \text{DB}^\top$, a multi-map $\text{MM}_{\mathbf{c}}$ that maps pairs of the form

$$\left\langle \langle \text{tbl}(\mathbf{c}), \text{crk}(\mathbf{c}) \rangle, \langle \text{tbl}(\mathbf{c}'), \text{crk}(\mathbf{c}') \rangle \right\rangle$$

to tokens for the rows $(\mathbf{r}_1, \mathbf{r}_2) \in \text{DB}_{\text{att}(\mathbf{c})=\text{att}(\mathbf{c}')}$, where $\mathbf{c}' \in \text{DB}^\top$ is a column with the same domain as \mathbf{c} . It then encrypts each of these multi-maps with Σ_{MM} , resulting in a set $\{\text{EMM}_{\mathbf{c}}\}_{\mathbf{c} \in \text{DB}^\top}$, and creates a dictionary DX that maps the attributes of each column $\text{att}(\mathbf{c})$, for all $\mathbf{c} \in \text{DB}^\top$, to $\text{EMM}_{\mathbf{c}}$. Finally, it encrypts DX with Σ_{DX} , resulting in an encrypted dictionary EDX.

Setup outputs the key $K = (K_R, K_C, K_V, K_D, \{K_{\mathbf{c}}\}_{\mathbf{c} \in \text{DB}^\top}, K_1)$, where K_R, K_C, K_V, K_D and $\{K_{\mathbf{c}}\}_{\mathbf{c} \in \text{DB}^\top}$ are the keys for $\text{EMM}_R, \text{EMM}_C, \text{EMM}_V, \text{EDX}$ and $\{\text{EMM}_{\mathbf{c}}\}_{\mathbf{c} \in \text{DB}^\top}$, respectively and K_1 is a key for SKE. The encrypted database is

$$\text{EDB} = (\text{EMM}_R, \text{EMM}_C, \text{EMM}_V, \text{EDX}).$$

Token. The Token algorithm takes as input a secret key K and an SPC query q in normal form. The algorithm first generates tokens for the projections. Recall that in the SPC normal form, projections can occur either on virtual or real attributes. So, for all $i \in [h]$, if $\mathbf{a}_i \in \mathbb{S}(\text{VDB})$, Token sets $\text{ytk}_i = (1, \mathbf{a}_i)$; otherwise, if $\mathbf{a}_i \in \mathbb{S}(\text{DB})$, it creates a projection token $\text{ptk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_C} \left(\left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right)$ and sets $\text{ytk}_i = (0, \text{ptk}_i)$.

For every constant a_1 through a_f it computes $e_{a_1} \leftarrow \text{Enc}_{K_1}(a_1)$ through $e_{a_f} \leftarrow \text{Enc}_{K_1}(a_f)$. Then, for each term $\mathbf{a}_i = X_i$ of Ψ it does the following. If the term is a type-1 query of the form $\mathbf{a}_i = a_i$, it computes an EMM_V token vtk_i for the pair that consists of the constant a_i and the coordinates of \mathbf{a}_i and sets $\text{stk}_i = \text{vtk}_i$. Throughout, we refer to s-tokens of this form as type-1 tokens. On the other hand, if the term is a type-2 query of the form $\mathbf{a}_i = \mathbf{x}_i$, then it computes two tokens dtk_i and jtk_i . The first, dtk_i is an EDX token for the coordinates of \mathbf{a}_i and the second, jtk_i , is an $\text{EMM}_{\mathbf{a}_i}$ token for the pair that consists of the coordinates of \mathbf{a}_i and \mathbf{x}_i . It then sets $\text{stk}_i = (\text{dtk}_i, \text{jtk}_i)$. We refer to s-tokens of this form as type-2 tokens. Token then creates a vector \mathbf{s} that holds the (column) size of T_{i_1}, \dots, T_{i_t} and outputs the token

$$\text{tk} = \left(t, \mathbf{s}, (e_{a_i})_{i \in [f]}, (\text{ytk}_i)_{i \in [h]}, (\text{stk}_i)_{i \in [\ell]} \right).$$

We first present an *inefficient* but algorithmic clear description of Query in 3. We then provide a realistic *efficient* instantiation of Query in 4 and 5.

Naive Query. We begin by describing naive version of the Query algorithm. While this version is not efficient, it is easier to understand and will serve as an introduction to the more efficient but more complex variant we describe in below. The naive version uses the basic location map defined in Section 3, whereas the efficient variant below, uses the extended version.

The Query algorithm takes as input the encrypted database EDB and a token tk . Given the arity of the cross product t , the size of the columns m and the column-sizes \mathbf{s} of the tables in the product, it instantiates an empty $m^t \times \sum_i s_i$ matrix \mathbf{M} .

It then processes the s-tokens stk_1 through stk_ℓ . If stk_i is type-1, it uses it to query EMM_V . This results in a tuple \mathbf{t} of row tokens which it then uses to query the EMM_R structure. More precisely, for all $\text{rtk} \in \mathbf{t}$, it computes

$$\left\langle \text{ct}_1, \dots, \text{ct}_d, \text{tb}, \text{rrk} \right\rangle \leftarrow \Sigma_{\text{MM}}.\text{Get}(\text{EMM}_R, \text{rtk}),$$

where $\langle \text{ct}_1, \dots, \text{ct}_d \rangle$ is an encrypted row and (tb, rrk) are its coordinates. Note that these will be encryptions of the rows in DB that hold the constant a_i at column \mathbf{a}_i , where \mathbf{a}_i and a_i are the attribute/constant pair in the type-1 select query that underlies vtk_i . For each ciphertext ct_i of every encrypted row above, it then computes $(\alpha, \beta) \leftarrow \Lambda_{t,m,\mathbf{s}}(\text{tb}, \text{rrk}, i)$ and sets $\mathbf{M}[\alpha, \beta] := \text{ct}_i$. On the other hand, if stk_i is type-2, the algorithm uses dtk_i to query EDX, resulting in an encrypted multi-map EMM_c which it then queries with jtk_i . This last operation results in a tuple \mathbf{t} of *pairs* of row tokens for EMM_R . For each pair of tokens, it queries EMM_R , resulting in a pair of encrypted rows with their coordinates. It uses the location map on each pair of coordinates to recover locations (α, β) and (α', β') . If $\alpha = \alpha'$, it further uses the location map to place each ciphertext in the encrypted rows in \mathbf{M} .

After processing the s-tokens, it processes the y-tokens ($\text{ytk}_1, \dots, \text{ytk}_h$) as follows. It starts by initializing a 1-by-1 matrix VA . If ytk_i has the form $(1, \mathbf{a})$, for some attribute \mathbf{a} , it sets $\text{VA} = \text{VA} \times \mathbf{e}_a$. On the other hand, if $\text{ytk}_i = (0, \text{ptk}_i)$, it computes

$$\langle \text{ct}_1, \dots, \text{ct}_m, \text{tb}, \text{crk} \rangle \leftarrow \Sigma_{\text{MM}}.\text{Get}(\text{EMM}_C, \text{ptk}_i),$$

and for all $j \in [m]$ it sets $\text{M}[\alpha_j, \beta_j] := \text{ct}_j$, where $(\alpha_j, \beta_j) \leftarrow \Lambda_{t,m,s}(\text{tb}, \text{crk}, j)$. It then sets $\text{ET} := \text{VA} \times \text{M}$ and outputs ET after eliminating all columns that contain at least one empty cell.

Decryption. The Dec algorithm takes as input a secret key K and the encrypted table ET returned by the server and simply decrypts each cell of ET . Note that we omit the formal description of Dec from Figs. 1, 2, 3, 4 and 5.

5.1 An Optimal Query Algorithm

The naive Query algorithm above makes use of a m^t -by- $\sum_{i=1}^t s_i$ array which can be very large. Fortunately, we do not need to process queries in this manner and we now describe a more efficient approach that builds the result table directly *without needing to generate the cross product of the tables*. Our optimized Query algorithm is relatively complex compared to the naive algorithm but considerably more efficient. An important component in this variant of the query algorithm is an additional map we refer to as the *row location map*.

The row location map. The row location map takes as input a tuple of b table-identifier/row-rank pairs $((\text{tb}_1, \text{rrk}_1), \dots, (\text{tb}_b, \text{rrk}_b))$, where $b \leq t$, and outputs a set of locations for these rows in the cross product $\text{M} = T_{i_1} \times \dots \times T_{i_t}$. Here, we assume the input to the row location map has the form $((i_1, \text{rrk}_1), \dots, (i_b, \text{rrk}_b))$; that is, the rows belong to tables T_{i_1} to T_{i_b} (all our inputs to the row location map will have this form). It outputs a row location $\alpha \in [m^t]$ in M . More precisely, the row location map is defined as

$$\Gamma_{t,m,b}((z_1, l_1), \dots, (z_b, l_b)) = \left\{ (\gamma - 1) \cdot m^{t-b} + 1, \dots, \gamma \cdot m^{t-b} \right\},$$

where $\gamma = l_1 \cdot m^{b-1} + l_2 \cdot m^{b-2} + \dots + l_b$, $z_i (l_1, \dots, l_b) \in [m]^b$. Note that if $b = t$, $\Gamma_{t,m,b}$ outputs a single location γ in M .

Overview. Our optimized query algorithm is described in detail in Figs. (4) and (5). At a high-level, there are four steps: (1) a recovery step where we process all the s-tokens and recover sets of rows for each token; (2) an intersection step, where we take per-table intersections of these sets (i.e., we take intersections of all sets that hold rows belonging to the same table); (3) an augmented cross-product step where we combine the remaining rows across sets; and (4) a placement step, where we use the location map to place the merged rows in the results table.

We first process the s-tokens. If a token is type-1, we recover a set of encrypted rows together with their coordinates (i.e., their table identifiers and row ranks). If, on the other hand, it is type-2 then we recover *pairs* of coordinates. We denote by R_i the i th *result set*; that is, the set of coordinates (or pair of coordinates) recovered from processing stk_i . If stk_i is type-1 we say that R_i is type-1 and if stk_i is type-2 then we say that R_i is type-2. Note that if stk_i is type-1 then all the coordinates we recover from processing it have the same table identifier. If stk_i is type-2, then the

Let $\Sigma_{DX} = (\text{Setup}, \text{Token}, \text{Get})$ be a response-revealing dictionary encryption scheme, $\Sigma_{MM} = (\text{Setup}, \text{Token}, \text{Get})$ be a response-revealing multi-map encryption scheme and $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a symmetric-key encryption scheme. Consider the DB encryption scheme $\text{SPX} = (\text{Setup}, \text{Token}, \text{Query}, \text{Dec})$ defined as follows ^a:

- $\text{Setup}(1^k, \text{DB})$:

1. initialize a dictionary DX ;
2. initialize multi-maps MM_R, MM_C and MM_V ;
3. initialize multi-maps $(\text{MM}_a)_{a \in \mathbb{S}(\text{DB})}$;
4. for all $\mathbf{r} \in \text{DB}$ set

$$\text{MM}_R \left[\left\langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \right\rangle \right] := \left(\text{Enc}_{K_1}(r_1), \dots, \text{Enc}_{K_1}(r_{\#\mathbf{r}}), \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \right);$$

5. compute $(K_R, \text{EMM}_R) \leftarrow \Sigma_{MM}.\text{Setup}(1^k, \text{MM}_R)$;
6. for all $\mathbf{c} \in \text{DB}^\top$, set

$$\text{MM}_C \left[\left\langle \text{tbl}(\mathbf{c}), \text{crk}(\mathbf{c}) \right\rangle \right] := \left(\text{Enc}_{K_1}(c_1), \dots, \text{Enc}_{K_1}(c_{\#\mathbf{c}}), \text{tbl}(\mathbf{c}), \text{crk}(\mathbf{c}) \right);$$

7. compute $(K_C, \text{EMM}_C) \leftarrow \Sigma_{MM}.\text{Setup}(1^k, \text{MM}_C)$;
8. for all $\mathbf{c} \in \text{DB}^\top$,

(a) for all $v \in \mathbf{c}$ and $\mathbf{r} \in \text{DB}_{\mathbf{c}=v}$,

i. compute $\text{rtk}_{\mathbf{r}} \leftarrow \Sigma_{MM}.\text{Token}_{K_R} \left(\left\langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \right\rangle \right)$,

(b) set

$$\text{MM}_V \left[\left\langle v, \left\langle \text{tbl}(\text{att}(\mathbf{c})), \text{crk}(\text{att}(\mathbf{c})) \right\rangle \right\rangle \right] := \left(\text{rtk}_{\mathbf{r}} \right)_{\mathbf{r} \in \text{DB}_{\mathbf{c}=v}};$$

9. compute $(K_V, \text{EMM}_V) \leftarrow \Sigma_{MM}.\text{Setup}(1^k, \text{MM}_V)$;
10. for all $\mathbf{c} \in \text{DB}^\top$,

(a) for all $\mathbf{c}' \in \text{DB}^\top$ such that $\text{dom}(\text{att}(\mathbf{c}')) = \text{dom}(\text{att}(\mathbf{c}))$,

i. initialize an empty tuple \mathbf{t} ;

ii. for all $i, j \in [m]$ such that $\mathbf{c}[i] = \mathbf{c}'[j]$,

A. compute $\text{rtk}_i \leftarrow \Sigma_{MM}.\text{Token}_{K_R} \left(\left\langle \text{tbl}(\mathbf{c}), i \right\rangle \right)$;

B. compute $\text{rtk}_j \leftarrow \Sigma_{MM}.\text{Token}_{K_R} \left(\left\langle \text{tbl}(\mathbf{c}'), j \right\rangle \right)$;

C. add $(\text{rtk}_i, \text{rtk}_j)$ to \mathbf{t} ;

iii. set

$$\text{MM}_{\mathbf{c}} \left[\left\langle \left\langle \text{tbl}(\mathbf{c}), \text{crk}(\mathbf{c}) \right\rangle, \left\langle \text{tbl}(\mathbf{c}'), \text{crk}(\mathbf{c}') \right\rangle \right\rangle \right] := \mathbf{t};$$

(b) compute $(K_{\mathbf{c}}, \text{EMM}_{\mathbf{c}}) \leftarrow \Sigma_{MM}.\text{Setup}(1^k, \text{MM}_{\mathbf{c}})$;

(c) set $\text{DX}[\left\langle \text{tbl}(\mathbf{c}), \text{crk}(\mathbf{c}) \right\rangle] = \text{EMM}_{\mathbf{c}}$;

11. compute $(K_D, \text{EDX}) \leftarrow \Sigma_{DX}.\text{Setup}(1^k, \text{DX})$;
12. output $K = (K_R, K_C, K_V, K_D, \{K_{\mathbf{c}}\}_{\mathbf{c} \in \text{DB}^\top})$ and $\text{EDB} = (\text{EMM}_R, \text{EMM}_C, \text{EMM}_V, \text{EDX})$;

^aNote that we omit the description of Dec since it simply decrypts every cell of ET .

• $\text{Token}(K, q)$:

1. parse q as $\pi_{\mathbf{a}_1, \dots, \mathbf{a}_h} \left([a_1] \times \dots \times [a_f] \times \sigma_{\Psi}(\mathbb{T}_{i_1} \times \dots \times \mathbb{T}_{i_t}) \right)$;

2. for all $i \in [h]$,

(a) if $\mathbf{a}_i \in \mathbb{S}(\text{DB})$, compute

$$\text{ptk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_C} \left(\left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right);$$

and set $\text{ytk}_i := (0, \text{ptk}_i)$;

(b) if $\mathbf{a}_i \in \mathbb{S}(\text{VDB})$, set $\text{ytk}_i := (1, \mathbf{a}_i)$;

3. parse Ψ as $\mathbf{a}_1 = X_1 \wedge \dots \wedge \mathbf{a}_\ell = X_\ell$;

4. for all $i \in [\ell]$,

(a) if $\mathbf{a}_i = X_i$ is type-1, compute

$$\text{vtk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_V} \left(\left\langle X_i, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right\rangle \right),$$

and set $\text{stk}_i = \text{vtk}_i$;

(b) if $\mathbf{a}_i = X_i$ is type-2, compute

$$\text{jtk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_{a_i}} \left(\left\langle \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle, \left\langle \text{tbl}(X_i), \text{crk}(X_i) \right\rangle \right\rangle \right)$$

and

$$\text{dtk}_i \leftarrow \Sigma_{\text{DX}}.\text{Token}_{K_D} \left(\left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right)$$

and set $\text{stk}_i = (\text{dtk}_i, \text{jtk}_i)$

5. for all $i \in [f]$, compute $\mathbf{e}_i \leftarrow \text{Enc}_{K_1}(a_i)$;

6. let $\mathbf{s} = (\|\mathbb{T}_{i_1}\|_c, \dots, \|\mathbb{T}_{i_t}\|_c)$;

7. output $\text{tk} = (t, \mathbf{s}, (\mathbf{e}_i)_{i \in [f]}, (\text{ytk})_{i \in [h]}, (\text{stk}_i)_{i \in [\ell]})$;

Figure 2: SPX: a relational DB encryption scheme (Part 2).

- Query(EDB, tk):
 1. parse EDB as $(EMM_R, EMM_C, EMM_V, EDX)$;
 2. parse $tk = (t, \mathbf{s}, (e_i)_{i \in [f]}, (ytk)_{i \in [h]}, (stk_i)_{i \in [\ell]})$;
 3. initialize an m^t -by- $\sum_{i=1}^t s_i$ matrix M and a 1-by-1 matrix VA ;
 4. for all $i \in [\ell]$,
 - (a) if stk_i is type-1, parse it as $stk_i = vtk_i$ and,
 - i. compute $\mathbf{t} \leftarrow \Sigma_{MM}.Get(vtk_i, EMM_V)$;
 - ii. for all $rtk \in \mathbf{t}$,
 - A. $\langle ct_1, \dots, ct_d, tb, rrk \rangle \leftarrow \Sigma_{MM}.Get(EMM_R, rtk)$;
 - B. set $M[\alpha, \beta] := ct_j$ for all $j \in [d]$ and all $(\alpha, \beta) \in \Lambda_{t,m,s}(tb, rrk, j)$;
 - (b) if stk_i is type-2, parse it as $stk_i = (dtk_i, jtk_i)$ and,
 - i. compute $EMM_c \leftarrow \Sigma_{DX}.Get(EDX, dtk_i)$;
 - ii. compute $\mathbf{t} \leftarrow \Sigma_{MM}.Get(EMM_c, jtk_i)$;
 - iii. for all $(rtk, rtk') \in \mathbf{t}$,
 - A. compute $\langle ct_1, \dots, ct_d, tb, rrk \rangle \leftarrow \Sigma_{MM}.Get(EMM_R, rtk)$;
 - B. compute $\langle ct'_1, \dots, ct'_d, tb', rrk' \rangle \leftarrow \Sigma_{MM}.Get(EMM_R, rtk')$;
 - C. for all $(\alpha, \beta) \in \Lambda_{t,m,s}(tb, rrk, 1)$ and $(\alpha', \beta') \in \Lambda_{t,m,s}(tb', rrk', 1)$ s.t. $\alpha = \alpha'$,
 - set $M[\alpha, \beta] := ct_j$ for all $j \in [d]$ and all $(\alpha, \beta) \in \Lambda_{t,m,s}(tb, rrk, j)$;
 - set $M[\alpha', \beta'] := ct'_z$ for all $z \in [d']$ and all $(\alpha', \beta') \in \Lambda_{t,m,s}(tb', rrk', z)$;
 5. for all $i \in [h]$,
 - (a) if $ytk_i = (1, \mathbf{a})$, compute $VA := VA \times [e_a]$;
 - (b) if $ytk_i = (0, \mathbf{ptk}_i)$,
 - i. compute $\langle \mathbf{c}, (tb, crk) \rangle \leftarrow \Sigma_{MM}.Get(EMM_C, \mathbf{ptk}_i)$;
 - ii. for all $\gamma \in [m]$ and all $(\alpha, \beta) \in \Lambda_{t,m,s}(tb, \gamma, crk)$, set $M[\alpha, \beta] := \mathbf{c}[\gamma]$;
 6. compute $ET := VA \times M$;
 7. delete all columns with a least one empty row in ET ;
 8. output ET

Figure 3: SPX: a relational DB encryption scheme (Part 3).

first coordinates of the all the coordinate pairs in R_i have the same table identifier and, similarly, the second coordinates of all the coordinate pairs have the same table identifiers. As such, we can assign a unique table identifier to each type-1 set and two table identifiers to each type-2 set. To make this explicit, we write type-1 sets as R_{tb} and type-2 sets as R_{tb_1, tb_2} . In the following, when referring to the table identifier of a type-2 set, we mean the table identifier of its first elements.

Our next step is to take per-table intersections. More precisely, let $T \subseteq [n]$ be the set of table identifiers of the recovered sets. Now, for all identifiers $tb \in T$, we generate the set I_{tb} as the intersection between all the result sets with table identifier tb (this could include type-1 and type-2 sets). Note that we could be taking intersections of type-1 sets, which hold coordinates, and type-2 sets, which hold pairs of coordinates. This is, of course, undefined so in such a case the intersection is defined to be the set of coordinate pairs in the type-2 set whose first coordinate is in the type-1 set. This results in one set I_{tb} per table, which holds either coordinates (if it results from the intersection of type-1 sets) or pairs of coordinates (if it results from the intersection of two type-2 sets or the intersection of one type-1 set and one type-2 set).

The next step is to take a cross product between the sets $\{I_{tb}\}_{tb \in T}$. Note that when doing this, there could be instances where we need to create a new tuple from two *pairs* of coordinates. In such a case, we only create the new tuple if the first elements of the pairs are equal. If they are not, we discard the two pairs. We refer to this process as an *augmented cross-product* and denote its result as \mathcal{C} , which is a set of tuples of coordinates. Now, for every coordinate tuple in \mathcal{C} , we do the following. We recover the encrypted rows associated with the coordinates in the tuple. We then query the location map on the coordinates in the tuple to get a location i in M . We then store this information in a multi-map TR_3 ; that is, we set TR_3 to map the location i to the encrypted rows associated with the coordinates in the tuple. For correctness, we add dummy elements to the encrypted rows. This is necessary in case the project operation is on an attribute that is not included in the select operation.

The final step consists of retrieving the desired columns based on the y-tokens. More precisely, we use the y-token to query EMM_C which returns an encrypted column together with its table identifier tb and column rank crk . Then, for every row i in the column, we lookup $TR_3[i]$ to recover an encrypted row. If the element at location crk in the encrypted row is not a dummy, we add it to our result table. On the other hand, if it is a dummy, we use i and the column rank crk to generate a location in M on which we apply the inverse of the location map. This returns a cell coordinate which we use to fetch an encrypted cell from the results of the EMM_C query. Finally, we replace the dummy element with this cell.

5.2 Efficiency

We now turn to analyzing the search and storage efficiency of our construction with the optimized Query algorithm.

Search complexity. Given an SPC query

$$\pi_{a_1, \dots, a_h} \left([a_1] \times \dots \times [a_f] \times \sigma_{\Psi}(T_{i_1} \times \dots \times T_{i_\ell}) \right),$$

the size of the result is linear in

$$Q = h \cdot (m^{h'} \cdot \prod_{i=1}^{\ell} \#R_i),$$

- Query(EDB, tk):
 1. parse EDB as $(EMM_R, EMM_C, EMM_V, EDX)$;
 2. parse $tk = (t, s, (e_i)_{i \in [f]}, (ytk)_{i \in [h]}, (stk_i)_{i \in [\ell]})$;
 3. initialize six multi-maps $TR_1, TR_2, TR_3, CT_1, CT_2$, and **Result**;
 4. initialize two non-duplicate sets TB_1 , and TB_2 ;
 5. for all $i \in [\ell]$,
 - (a) initialize an empty set R_i ;
 - (b) if stk_i is type-1, parse it as $stk_i = vtk_i$ and,
 - i. compute $\mathbf{t} \leftarrow \Sigma_{MM}.Get(EMM_V, vtk_i)$;
 - ii. for all $rrk \in \mathbf{t}$,
 - A. $\langle ct_1, \dots, ct_d, tb, rrk \rangle \leftarrow \Sigma_{MM}.Get(EMM_R, rrk)$;
 - B. set $R_i = R_i \cup \{rrk\}$;
 - C. set $CT_1[(tb, rrk)] := \langle ct_1, \dots, ct_d \rangle$;
 - D. add tb to TB_1 , and set $TR_1[tb] := i$;
 - (c) if stk_i is type-2, parse it as $stk_i = (dtk_i, jtk_i)$ and,
 - i. compute $EMM_c \leftarrow \Sigma_{DX}.Get(EDX, dtk_i)$;
 - ii. compute $\mathbf{t} \leftarrow \Sigma_{MM}.Get(EMM_c, jtk_i)$;
 - iii. for all $(rrk, rrk') \in \mathbf{t}$,
 - A. compute $\langle ct_1, \dots, ct_d, tb, rrk \rangle \leftarrow \Sigma_{MM}.Get(EMM_R, rrk)$;
 - B. compute $\langle ct'_1, \dots, ct'_d, tb', rrk' \rangle \leftarrow \Sigma_{MM}.Get(EMM_R, rrk')$;
 - C. set $R_i = R_i \cup \{(rrk, rrk')\}$;
 - D. set $CT_2[(tb, rrk)] := \langle ct_1, \dots, ct_d \rangle$;
 - E. set $CT_2[(tb', rrk')] := \langle ct'_1, \dots, ct'_d \rangle$;
 - F. add (tb, tb') to TB_2 , and set $TR_2[(tb, tb')] := i$;
 6. for all $tb \in TB_1$, set $R_{tb} = \bigcap_{i \in \mathbf{g}} R_i$, where $\mathbf{g} := TR_1[tb]$;
 7. for all $(tb_1, tb_2) \in TB_2$, set $R_{(tb_1, tb_2)} = \bigcap_{i \in \mathbf{g}} R_i$, where $\mathbf{g} := TR_2[(tb_1, tb_2)]$;
 8. for all (tb_1, tb_2) in TB_2 ,
 - (a) for all tb in TB_1 , if $tb = tb_i$, for $i = 1$ or 2 ,
 - set $S_{(tb_1, tb_2)} = \{(rrk_1, rrk_2) \mid (rrk_1, rrk_2) \in R_{(tb_1, tb_2)} \wedge rrk_1 \in R_{tb}\}$, and delete R_{tb} ;
 - (b) for all (tb'_1, tb'_2) in TB_2 such that $tb_1 = tb'_1$ and $tb_2 \neq tb'_2$,
 - set $P = \{rrk_1 \mid (rrk_1, rrk_2) \in R_{(tb_1, tb_2)}\}$;
 - set $S_{(tb'_1, tb'_2)} = \{(rrk_1, rrk_2) \mid (rrk_1, rrk_2) \in R_{(tb'_1, tb'_2)} \wedge rrk_1 \in P\}$;

Figure 4: An optimized Query algorithm (Part 1).

• Continuation of Query(EDB, tk):

9. initialize counter $\mu = 0$;
10. for all $(tb_1, tb_2) \in \text{TB}_2$,
 - (a) initialize two sets \mathcal{T}_μ and \mathcal{C}_μ ;
 - (b) set $\mathcal{C}_\mu = S_{(tb_1, tb_2)}$;
 - (c) add tb_1 , and tb_2 to \mathcal{T}_μ ;
 - (d) for all $(tb'_1, tb'_2) \in \text{TB}_2$,
 - i. if $(tb'_1, tb'_2) \neq (tb_1, tb_2)$ and $tb'_i = tb_i$ for $i = 1$ or 2 ,
 - A. set $P = \{\text{rrk}_i \mid (\text{rrk}_1, \text{rrk}_2) \in R_{(tb'_1, tb'_2)}\}$;
 - B. set $\mathcal{C}_i = \{(\text{rrk}_1, \dots, \text{rrk}_e, \text{rrk}) \mid (\text{rrk}_1, \dots, \text{rrk}_e) \in \mathcal{C}_i \wedge (\exists j \in [e] \text{ s.t. } \text{rrk} = \text{rrk}_j \wedge \text{rrk} \in P)\}$;
 - C. remove (tb'_1, tb'_2) from TB_2 and set $\mu = \mu + 1$;
 - D. add tb'_i to \mathcal{T}_μ ;
11. set $w = \kappa + h_1 + \dots + h_\mu$, where $h_i = \#\mathcal{C}_i$;
12. for all $\text{rrk}_1, \dots, \text{rrk}_\kappa \in R_{tb_1}, \dots, R_{tb_\kappa}$ and for all $(\text{rrk}_1^1, \dots, \text{rrk}_{h_1}^1), \dots, (\text{rrk}_1^\mu, \dots, \text{rrk}_{h_\mu}^\mu) \in \mathcal{C}_1, \dots, \mathcal{C}_\mu$
 - (a) for all i in

$$\Gamma_{t,m,w} \left((tb_1, \text{rrk}_1), \dots, (tb_\kappa, \text{rrk}_\kappa), (\mathcal{T}_1[1], \text{rrk}_1^1), \dots, (\mathcal{T}_1[h_1], \text{rrk}_{h_1}^1), \dots, (\mathcal{T}_\mu[1], \text{rrk}_1^\mu) \dots (\mathcal{T}_\mu[h_\mu], \text{rrk}_{h_\mu}^\mu), 1 \right),$$

set,

$$\text{TR}_3[i] := \left\langle \text{CT}_1[(tb_1, \text{rrk}_1)], \dots, \text{CT}_1[(tb_h, \text{rrk}_h)], \right. \\ \text{CT}_2[(\mathcal{T}_1[1], \text{rrk}_1^1)], \dots, \text{CT}_2[(\mathcal{T}_1[h_1], \text{rrk}_{h_1}^1)], \\ \dots, \\ \left. \text{CT}_2[(\mathcal{T}_\mu[1], \text{rrk}_1^\mu)], \dots, \text{CT}_2[(\mathcal{T}_\mu[h_\mu], \text{rrk}_{h_\mu}^\mu)], \underbrace{\text{dummy}, \dots, \text{dummy}}_{\sum_{i=w+1}^t s_i \text{ time}} \right\rangle$$

13. for all $i \in [h]$,
 - (a) if $\text{ytk}_i = (0, \text{ptk}_i)$,
 - i. compute $\langle \mathbf{c}, \langle \text{tb}, \text{crk} \rangle \rangle \leftarrow \Sigma_{\text{MM}}.\text{Get}(\text{EMM}_C, \text{ptk}_i)$;
 - ii. set Θ be the range of TR_3 ;
 - iii. set $\text{crk} = \text{crk} + \sum_{i=1}^{\text{tb}} s_i$;
 - iv. if $\text{crk} \leq \sum_{i=1}^w s_i$, set

$$\text{Result}[\text{crk}] := \langle \text{TR}_3[\Theta_1]_{\text{crk}}, \dots, \text{TR}_3[\Theta_{|\Theta|}]_{\text{crk}} \rangle;$$

- v. if $\text{crk} > \sum_{i=1}^w s_i$, set

$$\text{Result}[\text{crk}] := \langle \mathbf{c}[\Lambda_{t,m,s}^{-1}(\Theta_1, \text{crk})_2], \dots, \mathbf{c}[\Lambda_{t,m,s}^{-1}(\Theta_{|\Theta|}, \text{crk})_2] \rangle;$$

- (b) if $\text{ytk}_i = (1, \mathbf{a})$, set $\text{Result}[\mathbf{a}] := \underbrace{\langle \mathbf{e}_a, \dots, \mathbf{e}_a \rangle}_{|\Theta|}$;

14. output ET such that every row equals an entry of a label in **Result**;

Figure 5: An optimized Query algorithm (Part 2).

where R_i is the set of rows that match the i th term in the select (i.e., the i th result set), and $0 \leq h' \leq h$ is the number of attributes in the project operator π_{a_1, \dots, a_h} that are *not* included in the select operator σ_Ψ . Note that here we assume the worst-case where these projection attributes (i.e., the ones not involved in σ_Ψ) all belong to distinct tables. In \mathcal{Q} , the term $\prod_{i=1}^\ell \#R_i$ is the number of rows in the tables $T_{i_1}, \dots, T_{i_\ell}$ (not in their cross product) that satisfy σ_Ψ . The term $m^{h'}$ is the number of rows in $T_{i_1}, \dots, T_{i_\ell}$ that include one of the attributes in π_{a_1, \dots, a_h} that are not involved in σ_Ψ (assuming all tables have m rows). Finally, h is the number of columns/attributes in the projection operation.

\mathcal{Q} represents a worst-case lower bound on both the running time⁴ and the space complexity of a plaintext SPC query. Note that in practice the number h' of attributes in π_{a_1, \dots, a_h} that are not involved in the select operator σ_Ψ is usually very small. The best and most common case is when $h' = 0$ which leads to $\mathcal{Q} = h \cdot \prod_{i=1}^\ell \#R_i$.

We now show that the optimized Query algorithm presented in Section 5.1, requires $O(\mathcal{Q})$ time and space when $t, s_1, \dots, s_n, h, \ell \ll m$ which is common in practice. Consider a token

$$\text{tk} = \left(t, \mathbf{s}, (e_{a_i})_{i \in [f]}, (\text{ytk}_i)_{i \in [h]}, (\text{stk}_i)_{i \in [\ell]} \right),$$

and an encrypted database $\text{EDB} = (\text{EMM}_R, \text{EMM}_C, \text{EMM}_V, \text{EDX})$. For ease of exposition, all tables have the same number of rows m .

When processing the s-tokens we use four multi-maps $\text{TR}_1, \text{TR}_2, \text{CT}_1, \text{CT}_2$, and two sets TB_1 and TB_2 that store the table identifiers in the result sets. Recall that TR_1, TR_2 both have size $O(\ell)$ since there are at most ℓ s-tokens in tk . TB_1 and TB_2 have respectively size $O(t), O(\max(t, \ell))$ since there are at most t tables involved in the query, and ℓ possible type-2 pairs. The multi-maps CT_1 and CT_2 store the encrypted rows that are recovered by processing the s-tokens. That is, they are of size $O(\sum_{i=1}^\ell \#R_i \cdot s_i)$, where s_i is the number of cells in the i th table. Note that the size of CT_1 and CT_2 dominates the other data structures. Also, note that the running time is $O(\sum_{i=1}^\ell \#R_i \cdot s_i)$.

Computing intersections is performed in $O(\tau^2 \cdot \max_{i \in [\ell]} \#R_i)$ steps, where $\tau = \max(t, \ell)$. The intersection goes over all possible tables identifiers in TB_1 and TB_2 . Note that the worst case *running time* occurs when the terms of the select operation are all type-2 and have the same attributes, e.g., $\mathbf{a}_1 = a_1 \wedge \dots \wedge a_1 = a_1$. On the other hand, the worst case *space complexity* occurs when there are no common attributes between terms in the select operation and this requires $O(\sum_{i=1}^\ell \#R_i \cdot s_i)$ space.

When computing the augmented cross product, the worst case with respect to space occurs when the result sets have no elements in common so we need to keep all pairs. The cost in space of this cross product will be

$$O\left(m^{h'} \cdot (\prod_{i=1}^\ell \#R_i) \cdot \sum_{i=1}^t s_i\right) = O\left(\frac{\mathcal{Q}}{h} \cdot \sum_{i=1}^t s_i\right).$$

On the other hand, the worst-case with respect to time occurs when all terms in the select operation share an attribute in common. In such a case the running time is,

$$O\left(m^{h'} \cdot (\prod_{i=1}^\ell \#R_i) \cdot \sum_{i=1}^t s_i + \tau^2 \cdot \max_{i \in [\ell]} \#R_i\right) = O\left(\frac{\mathcal{Q}}{h} \cdot \sum_{i=1}^t s_i\right).$$

⁴This is assuming, of course, one is not using pre-computation.

The generation of the final encrypted table ET requires us to move items from TR₃ to Result and then into ET. The space and time complexity of this operation is dominated by the size of TR₃ which is $O(h^{-1} \cdot \mathcal{Q} \cdot \sum_{i=1}^t s_i)$.

In summary, Query requires $O(h^{-1} \cdot \mathcal{Q} \cdot \sum_{i=1}^t s_i)$ space and time. If we view $h^{-1} \cdot \sum_{i=1}^t s_i$ as a constant since, in practice $t, s_1, \dots, s_t, h \ll m$, both the running time and the space complexity are $O(\mathcal{Q})$.

Storage complexity. For a database $DB = (T_1, \dots, T_n)$, SPX produces four encrypted multi-maps EMM_R , EMM_C , EMM_V and EDX. Let $s = \sum_i \|T_i\|_c$. For ease of exposition, we again assume each table has m rows. Finally, note that standard multi-map encryption schemes [15, 23, 8, 7] produce encrypted structures with storage overhead that is linear in sum of the tuple sizes. Using such a scheme as the underlying multi-map encryption scheme, we have that EMM_R and EMM_C will be $O(\sum_{\mathbf{r} \in DB} \#\mathbf{r})$ and $O(\sum_{\mathbf{c} \in DB^\top} \#\mathbf{c})$, respectively, since the former maps the coordinates of each row in DB to their (encrypted) row and the latter maps the coordinates of very column to their (encrypted) columns. Since EMM_V maps each cell in DB to tokens for the rows that contain the same value, it requires $O(\sum_{\mathbf{c} \in DB^\top} \sum_{v \in \mathbf{c}} \#DB_{\mathbf{c}=v})$ storage. EDX maps the coordinates of each column $\mathbf{c} \in DB^\top$ to an encrypted multi-map $EMM_{\mathbf{c}}$ which in turn maps each pair of form $(\mathbf{c}, \mathbf{c}')$ such that $\text{dom}(\mathbf{c}) = \text{dom}(\mathbf{c}')$ to a tuple of tokens for rows in $DB_{\mathbf{c}=\mathbf{c}'}$. As such, EDX will have size

$$O\left(\sum_{\mathbf{c} \in DB^\top} \sum_{\mathbf{c}': \text{dom}(\mathbf{c}') = \text{dom}(\mathbf{c})} \#DB_{\mathbf{c}=\mathbf{c}'}\right).$$

Note that the expression will vary greatly depending on the number of columns in DB with the same domain. In the worst case, all columns will have a common domain and the expression will be a sum of $s(s-1)$ terms of the form $\#DB_{\mathbf{c}=\mathbf{c}'}$. In the best case, none of the columns will share a domain and EDX will be empty. In practice, however, we expect there to be some relatively small number of columns with common domains.

6 Security and Leakage

We now show that our construction is adaptively-semantically secure with respect to a well-specified leakage profile. Part of the subtlety in our security analysis is that some of the leakage is “black-box” in the sense that it comes from the underlying schemes and part of it “non-black-box” in the sense that it comes directly from the SPX construction itself. Throughout our discussion of SPX’s leakage, we consider both its black-box leakage (i.e., when the underlying schemes are left abstract) and its concrete leakage (i.e., when the underlying schemes are instantiated). To instantiate the underlying schemes, we consider any of a standard set of SSE constructions from the literature [15, 23, 12, 8, 7, 27] which all have the same leakage profile, i.e., the search pattern which reveals if and when a query is repeated.⁵ In particular, these SSE schemes can be used to instantiate both Σ_{MM} and Σ_{DX} since the former is a generalization of the latter.

⁵Since the schemes are response-revealing, we do not consider the “access pattern leakage” of SSE schemes (which is defined as the search response) to be leakage in our context.

Setup leakage. The setup leakage of SPX captures what an adversary can learn before performing any query operation. The setup leakage of SPX is

$$\mathcal{L}_S^{\text{SPX}}(\text{DB}) = \left(\mathcal{L}_S^{\text{dx}}(\text{DX}), \mathcal{L}_S^{\text{mm}}(\text{MM}_R), \mathcal{L}_S^{\text{mm}}(\text{MM}_C), \mathcal{L}_S^{\text{mm}}(\text{MM}_V) \right),$$

where $\mathcal{L}_S^{\text{dx}}$ and $\mathcal{L}_S^{\text{mm}}$ are the setup leakages of Σ_{DX} and Σ_{MM} , respectively. If the latter are instantiated with standard encrypted multi-map constructions, the setup leakage of SPX will consist of the number of rows and columns in DB and the size of the dictionary. Note that standard encrypted dictionary constructions leak only the maximum size of the values they store so the size of the $\text{EMM}_{\mathbf{c}}$'s will be hidden (up to the maximum size).

Query leakage. The query leakage is more complex and is defined as follows,

$$\mathcal{L}_Q^{\text{SPX}}(\text{DB}, q) = \left(\text{XPP}(\text{DB}, q), \text{PrP}(\text{DB}, q), \text{SelP}(\text{DB}, q) \right),$$

where each individual pattern is described next.

Cross product. The first leakage pattern is the *cross product* pattern which is defined as

$$\text{XPP}(\text{DB}, q) = \left\{ t, \|\mathbf{T}_{i_1}\|_c, \dots, \|\mathbf{T}_{i_\ell}\|_c, (|a_i|)_{i \in [f]} \right\},$$

and includes the number and size of the tables needed in the query as well as the size of the virtual attributes.

Projection. The second leakage pattern is the *projection pattern* which is defined as

$$\text{PrP}(\text{DB}, q) = \left(\mathcal{P}(\mathbf{a}_1), \dots, \mathcal{P}(\mathbf{a}_h) \right),$$

where

$$\mathcal{P}(\mathbf{a}_i) = \begin{cases} \left(\text{real}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_C, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right), (|c_j|)_{j \in [\#\mathbf{c}_i]}, \text{AccP}(\mathbf{c}_i) \right) & \text{if } \mathbf{a}_i \in \mathbb{S}(\text{DB}); \\ \left(\text{virtual}, \mathbf{a}_i \right) & \text{if } \mathbf{a}_i \in \mathbb{S}(\text{VDB}), \end{cases}$$

where $\mathbf{c}_i \in \text{DB}^\top$ is the column with attribute \mathbf{a}_i and $\text{AccP}(\mathbf{c}_i)$ indicates if and when the column \mathbf{c}_i has been accessed before. PrP captures the leakage produced when the server queries MM_C and for every attribute \mathbf{a}_i reveals whether the attribute is real or virtual. If it is real, it also reveals the size of the items in the projected column and if and when this column has been accessed before. Notice that it also reveals the Σ_{MM} query leakage on the coordinates of the projected attribute. If the latter is instantiated with any of the standard multi-map encryption schemes then this leakage will reveal whether the attribute \mathbf{a}_i has appeared in a previous query. If the attribute is virtual, it just reveals the attribute.

Selection. The third leakage pattern is the *selection pattern* which is defined as

$$\text{SelP}(\text{DB}, q) = \left(\mathcal{Z}(\mathbf{a}_1, X_1), \dots, \mathcal{Z}(\mathbf{a}_\ell, X_\ell) \right).$$

If $\mathbf{a}_i = X_i$ is type-1, then $\mathcal{Z}(\mathbf{a}_i, X_i)$ is defined as

$$\mathcal{Z}(\mathbf{a}_i, X_i) = \left(\text{type-1}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_V, \left\langle X_i, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right\rangle \right), \left. \left\{ \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_R, \left\langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \right\rangle \right), \text{AccP}(\mathbf{r}) \right\}_{\mathbf{r} \in \text{DB}_{\mathbf{a}_i = X_i}} \right).$$

where $\text{AccP}(\mathbf{r})$ indicates whether the row \mathbf{r} has been accessed before. $\mathcal{Z}(\mathbf{a}_i, X_i)$ captures the leakage produced when the server queries MM_V and uses the resulting row tokens to then query MM_R . It reveals whether the selection term is of type-1 and the Σ_{MM} query leakage on the constant X_i and the coordinates of the attribute \mathbf{a}_i . In addition, it also leaks the Σ_{MM} query leakage on the coordinates of the rows in $\text{DB}_{\mathbf{a}_i = X_i}$ as well as if and when they have been accessed before. If the encrypted multi-maps are instantiated with standard constructions, $\mathcal{Z}(\mathbf{a}_i, X_i)$ amounts to: if and when the pair (X_i, \mathbf{a}_i) has been queried before and if and when any of the rows in $\text{DB}_{\mathbf{a}_i = X_i}$ have been accessed in the past.

If, on the other hand, $\mathbf{a}_i = X_i$ is type-2, then $\mathcal{Z}(\mathbf{a}_i, X_i)$ is defined as

$$\mathcal{Z}(\mathbf{a}_i, X_i) = \left(\text{type-2}, \mathcal{L}_Q^{\text{dx}} \left(\text{DX}, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right), \mathcal{L}_S^{\text{mm}}(\text{EMM}_{\mathbf{a}_i}), \text{AccP}(\text{EMM}_{\mathbf{a}_i}), \right. \\ \left. \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_{\mathbf{a}_i}, \left\langle \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle, \left\langle \text{tbl}(X_i), \text{crk}(X_i) \right\rangle \right\rangle \right), \right. \\ \left. \left\{ \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_R, \left\langle \text{tbl}(\mathbf{r}_1), \text{rrk}(\mathbf{r}_1) \right\rangle \right), \text{AccP}(\mathbf{r}_1), \right. \right. \\ \left. \left. \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_R, \left\langle \text{tbl}(\mathbf{r}_2), \text{rrk}(\mathbf{r}_2) \right\rangle \right), \text{AccP}(\mathbf{r}_2) \right\}_{(\mathbf{r}_1, \mathbf{r}_2) \in \text{DB}_{\mathbf{a}_i = X_i}} \right),$$

where $\text{AccP}(\mathbf{r})$, $\text{AccP}(\mathbf{r}_1)$, $\text{AccP}(\mathbf{r}_2)$ and $\text{AccP}(\text{EMM}_{\mathbf{a}_i})$ indicate if and when \mathbf{r} , \mathbf{r}_1 , \mathbf{r}_2 and $\text{EMM}_{\mathbf{a}_i}$ have been accessed before. In this case, $\mathcal{Z}(\mathbf{a}_i, X_i)$ captures the leakage produced when the server queries EDX to retrieve some $\text{EMM}_{\mathbf{a}_i}$ which it in turn queries to retrieve row tokens with which to query EMM_R . It reveals whether the selection term is type-2, the Σ_{DX} query leakage on the coordinates of \mathbf{a}_i , the Σ_{MM} *setup* leakage on $\text{MM}_{\mathbf{a}_i}$ and if and when $\text{EMM}_{\mathbf{a}_i}$ has been accessed in the past. In addition, it reveals the query leakage of Σ_{MM} on the coordinates of \mathbf{a}_i and X_i and, for every pair of rows $(\mathbf{r}_1, \mathbf{r}_2)$ in $\text{DB}_{\mathbf{a}_i = X_i}$, their Σ_{MM} query leakage and if and when they were accessed in the past. Again, if instantiated with standard encrypted multi-maps, this would amount to the type of the selection, if and when \mathbf{a}_i had been queried in the past, the number of columns in DB that share a domain with \mathbf{a}_i , if and when the pair (\mathbf{a}_i, X_i) has appeared in previous queries and, for every pair of rows in $\text{DB}_{\mathbf{a}_i = X_i}$, if and when these rows have been accessed in the past.

Comparison to PPE-based solutions. As mentioned in Section 1, PPE-based solutions can handle a large class of SQL queries which includes conjunctive queries. To support conjunctive queries, however, these solutions have to rely on deterministic encryption. For example, to handle a

type-1 select on a column \mathbf{c} , they will reveal a deterministic encryption of \mathbf{c} (i.e., every element of the column is encrypted under the same key). To handle a type-2 select between two columns \mathbf{c}_1 and \mathbf{c}_2 , they will reveal deterministic encryptions of both columns (under the same key). In turn, this will provide the *frequency* information on entire columns to the server. Depending on the setting, frequency patterns can be particularly dangerous, as shown in [26].

We note that, in comparison, SPX leaks considerably less. First, it does not leak any frequency information on entire columns or rows. For type-1 selects, it only leaks information about the attribute in the select and the rows that match the term. For type-2 selects, it only leaks information about the pair of attributes (a_i, X_i) in the select and the rows that match the term. Note that this leakage is only a function of the attributes in the query and of the rows that match it, whereas the leakage in PPE-based solutions is a function of entire columns. Moreover, in the case of SPX, if the underlying multi-map and dictionary schemes are instantiated with standard constructions, the information leaked about the attributes and matching rows is “repetition” type of information, i.e., if and when they have appeared in the past. Analogously, the project operations in SPX only leak information about the attributes in the project and the columns that match it and the information being leaked “repetition” type of information.

6.1 Security of Our Construction

We now prove that SPX is adaptively semantically-secure with respect to the leakage profile described in the previous sub-section.

Theorem 6.1. *If Σ_{DX} is adaptively $(\mathcal{L}_S^{\text{dx}}, \mathcal{L}_Q^{\text{dx}})$ -semantically secure and Σ_{MM} is adaptively $(\mathcal{L}_S^{\text{mm}}, \mathcal{L}_Q^{\text{mm}})$ -secure, then SPX is $(\mathcal{L}_S^{\text{spX}}, \mathcal{L}_Q^{\text{spX}})$ -semantically secure.*

The proof of Theorem 6.1 in Appendix A.

7 Handling Extended Queries

In addition to the basic *Select-From-Where* structure, SQL queries can also include additional clauses to post-process the results of a query. The most common examples are the *Group-by*, *Order-by* and various aggregate functions which include *Sum*, *Average*, *Median*, *Count*, *Mode*, *Max* and *Min*. Here, we refer to SQL queries with such additional clauses as *extended queries*. Though SPX cannot handle extended queries explicitly we note that the additional clauses can be executed at the client. In cases where the result table R is such that $\|R\|_r \ll \sum_{i=1}^n \|T_i\|_r$, outsourcing a database with SPX will still be less computationally-expensive than executing the queries locally.

Order-by. The *Order-by* clause orders the rows in the result table R according to a specified attribute in either ascending or descending order. These queries have the form:

Select attributes From tables Where conditions Order-by attribute.

Ordering R at the client can be done in $O(m \log m)$ time where $m = \|R\|_r$. Note that the client has to perform $O(m)$ work just to decrypt the result table so, asymptotically-speaking, ordering R does not add a lot. In practice, however, *Order-by* operations are often performed on integer or string columns and in the latter case this can be expensive if the strings are long. A simple optimization is for the client to compute the order of the rows during the setup phase and to add a column that

includes a CPA-secure encryption of each row's order. When the client retrieves the result-table R it can then decrypt the column and order the rows using a *numerical* sorting algorithm (as opposed to lexicographic).

Group by. Another common SQL clause is **Group-by** which is used to partition the rows in the result table by a specified attribute. **Group-by** is typically used with aggregate functions (i.e., **Sum**, **Max** etc.). Such queries have the form:

Select *attributes* *attributes* From *tables* Where *conditions* Group-by *attributes*.

Group-by clauses can be handled at the client in $O(m)$ time, where $m = \|R\|_r$.

Aggregates. Often, one needs to compute a function on the rows of the result table (e.g., computing a sum of salaries from employee table). In SQL, this is handled with an aggregate function. For example, a query with the **Sum** function has the form:

Select Sum(*attribute*) From *tables* Where *conditions*.

Some of the aggregates can be handled by the server by extending **SPX** in the natural way. For example, to handle **Count** the server can simply return the number of rows in the encrypted result table **ET**. To handle **Sum** on a given attribute **a**, it suffices to encrypt the cells of that column with an additively-homomorphic encryption scheme and have the server return the sum of the resulting column. The remaining functions like **Max**, **Min**, **Median**, **Mode**, have to be handled at the client.

References

- [1] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. Order preserving encryption for numeric data. In *ACM SIGMOD International Conference on Management of Data*, pages 563–574, 2004.
- [2] A. Arasu, S. Blanas, K. Eguro, R. Kaushik, D. Kossmann, R. Ramamurthy, and R. Venkatesan. Orthogonal security with cipherbase. In *CIDR*, 2013.
- [3] M. Bellare, A. Boldyreva, and A. O’Neill. Deterministic and efficiently searchable encryption. In A. Menezes, editor, *Advances in Cryptology – CRYPTO ’07*, Lecture Notes in Computer Science, pages 535–552. Springer, 2007.
- [4] A. Ben-David, N. Nisan, and B. Pinkas. Fairplaymp: a system for secure multi-party computation. In *ACM Conference on Computer and Communications Security (CCS 2008)*, pages 257–266. ACM, 2008.
- [5] A. Boldyreva, N. Chenette, Y. Lee, and A. O’neill. Order-preserving symmetric encryption. In *Advances in Cryptology - EUROCRYPT 2009*, pages 224–241, 2009.
- [6] D. Boneh, G. di Crescenzo, R. Ostrovsky, and G. Persiano. Public key encryption with keyword search. In *Advances in Cryptology – EUROCRYPT ’04*, volume 3027 of *Lecture Notes in Computer Science*, pages 506–522. Springer, 2004.

- [7] D. Cash, J. Jaeger, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Dynamic searchable encryption in very-large databases: Data structures and implementation. In *Network and Distributed System Security Symposium (NDSS '14)*, 2014.
- [8] D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *Advances in Cryptology - CRYPTO '13*. Springer, 2013.
- [9] D. Cash and S. Tessaro. The locality of searchable symmetric encryption. In *Advances in Cryptology - EUROCRYPT 2014*, 2014.
- [10] Donald D Chamberlin and Raymond F Boyce. SEQUEL: A structured english query language. In *ACM SIGFIDET workshop on Data description, access and control (SIGMOD '74)*, pages 249–264. ACM, 1974.
- [11] A. Chandra and P. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *ACM Symposium on Theory of Computing (STOC '77)*, pages 77–90. ACM, 1977.
- [12] M. Chase and S. Kamara. Structured encryption and controlled disclosure. In *Advances in Cryptology - ASIACRYPT '10*, volume 6477 of *Lecture Notes in Computer Science*, pages 577–594. Springer, 2010.
- [13] M. Chase and S. Kamara. Structured encryption and controlled disclosure. Technical Report 2011/010.pdf, IACR Cryptology ePrint Archive, 2010.
- [14] E. Codd. A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387, 1970.
- [15] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky. Searchable symmetric encryption: Improved definitions and efficient constructions. In *ACM Conference on Computer and Communications Security (CCS '06)*, pages 79–88. ACM, 2006.
- [16] S. Faber, S. Jarecki, H. Krawczyk, Q. Nguyen, M. Rosu, and M. Steiner. Rich queries on encrypted data: Beyond exact matches. In *European Symposium on Research in Computer Security (ESORICS '15)*. *Lecture Notes in Computer Science*, volume 9327, pages 123–145, 2015.
- [17] C. Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009.
- [18] O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious RAMs. *Journal of the ACM*, 43(3):431–473, 1996.
- [19] P. Grofig, M. Haerterich, I. Hang, F. Kerschbaum, M. Kohler, A. Schaad, A. Schroepfer, and W. Tighzert. Experiences and observations on the industrial implementation of a system to search over outsourced encrypted data. In *Sicherheit*, pages 115–125, 2014.
- [20] H. Hacigümüs, B. Iyer, C. Li, and S. Mehrotra. Executing sql over encrypted data in the database-service-provider model. In *ACM SIGMOD International Conference on Management of Data*, pages 216–227. ACM, 2002.

- [21] S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Outsourced symmetric private information retrieval. In *ACM Conference on Computer and Communications Security (CCS '13)*, pages 875–888, 2013.
- [22] S. Kamara and C. Papamanthou. Parallel and dynamic searchable symmetric encryption. In *Financial Cryptography and Data Security (FC '13)*, 2013.
- [23] S. Kamara, C. Papamanthou, and T. Roeder. Dynamic searchable symmetric encryption. In *ACM Conference on Computer and Communications Security (CCS '12)*. ACM Press, 2012.
- [24] K. Kurosawa and Y. Ohtaki. Uc-secure searchable symmetric encryption. In *Financial Cryptography and Data Security (FC '12)*, Lecture Notes in Computer Science, pages 285–298. Springer, 2012.
- [25] X. Meng, S. Kamara, K. Nissim, and G. Kollios. Grecs: Graph encryption for approximate shortest distance queries. In *ACM Conference on Computer and Communication Security (CCS '15)*, 2015.
- [26] M. Naveed, S. Kamara, and C. Wright. Inference attacks on property-preserving encrypted databases. In *ACM Conference on Computer and Communication Security (CCS '15)*, 2015.
- [27] M. Naveed, M. Prabhakaran, and C. Gunter. Dynamic searchable encryption via blind storage. In *IEEE Symposium on Security and Privacy (S&P '14)*, 2014.
- [28] V. Pappas, F. Krell, B. Vo, V. Kolesnikov, T. Malkin, S.-G. Choi, W. George, A. Keromytis, and S. Bellovin. Blind seer: A scalable private dbms. In *Security and Privacy (SP), 2014 IEEE Symposium on*, pages 359–374. IEEE, 2014.
- [29] R. Ada Popa, C. Redfield, N. Zeldovich, and H. Balakrishnan. CryptDB: Protecting confidentiality with encrypted query processing. In *SOSP*, pages 85–100, 2011.
- [30] D. Song, D. Wagner, and A. Perrig. Practical techniques for searching on encrypted data. In *IEEE Symposium on Research in Security and Privacy*, pages 44–55. IEEE Computer Society, 2000.
- [31] J. Van den Bussche and S. Vansummeren. Translating sql into the relational algebra. 2009. http://cs.ulb.ac.be/public/_media/teaching/infoh417/sql2alg_eng.pdf.

A Proof of Theorem 6.1

Theorem 6.1. *If Σ_{DX} is adaptively $(\mathcal{L}_S^{\text{dx}}, \mathcal{L}_Q^{\text{dx}})$ -semantically secure and Σ_{MM} is adaptively $(\mathcal{L}_S^{\text{mm}}, \mathcal{L}_Q^{\text{mm}})$ -secure, then SPX is $(\mathcal{L}_S^{\text{spx}}, \mathcal{L}_Q^{\text{spx}})$ -semantically secure.*

Proof. Let \mathcal{S}_{DX} and \mathcal{S}_{MM} be the simulators guaranteed to exist by the adaptive security of Σ_{DX} and Σ_{MM} and consider the SPX simulator \mathcal{S} that works as follows. Given $\mathcal{L}_S^{\text{spx}}(\text{DB})$, \mathcal{S} simulates EDB by computing $\text{EDX} \leftarrow \mathcal{S}_{\text{DX}}(\mathcal{L}_S^{\text{dx}}(\text{DX}))$, $\text{EMM}_R \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_R))$, $\text{EMM}_C \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_C))$ and $\text{EMM}_V \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_V))$ and outputting

$$\text{EDB} = (\text{EMM}_R, \text{EMM}_C, \text{EMM}_V, \text{EDX}).$$

Recall that SPX is response-hiding so \mathcal{S} receives $(\perp, \mathcal{L}_Q^{\text{SPX}}(\text{DB}, q))$ as input in the $\mathbf{Real}_{\text{SPX}, \mathcal{S}}(k)$ experiment. Given this input, \mathcal{S} simulates a token

$$\text{tk} = \left(t, \mathbf{s}, (e_i)_{i \in [f]}, (\text{ytk}_i)_{i \in [h]}, (\text{stk}_i)_{i \in [\ell]} \right)$$

as follows. It retrieves t and \mathbf{s} directly from $\text{XPP}(\text{DB}, q)$. It then samples $K_1 \xleftarrow{\$} \{0, 1\}^k$ and, for all $i \in [f]$, computes $e_i \leftarrow \text{Enc}_{K_1}(\mathbf{0}^{|a_i|})$, where $|a_i|$ is also from $\text{XPP}(\text{DB}, q)$.

For all $i \in [h]$, if

$$\mathcal{P}(\mathbf{a}_i) = \left(\text{real}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_C, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right), ((c_j)_{j \in [\#c]}, \text{AccP}(\mathbf{c})) \right)$$

it sets

$$\text{ytk}_i := \text{ptk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((ct_j)_{j \in [\#c]}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_C, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right) \right),$$

where $ct_j \leftarrow \text{Enc}_{K_1}(\mathbf{0}^{|c_j|})$ if the column has never been accessed before and where ct_j is the previously used ciphertext otherwise. If, on the other hand, $\mathcal{P}(\mathbf{a}_i) = (\text{virtual}, \mathbf{a}_i)$ it sets $\text{ytk}_i := \mathbf{a}_i$.

Then, for all $i \in [\ell]$, it simulates stk_i as follows. If $\mathcal{Z}(\mathbf{a}_i)$ is type-1, it first computes for all $\mathbf{r} \in \text{DB}_{\mathbf{a}_i = X_i}$,

$$\text{rtk}_{\mathbf{r}} \leftarrow \mathcal{S}_{\text{MM}} \left((ct_j)_{j \in [\#\mathbf{r}]}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_R, \left\langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \right\rangle \right) \right)$$

where $ct_j \leftarrow \text{Enc}_{K_1}(\mathbf{0}^{|r_j|})$ if \mathbf{r} has never been accessed and ct_j is the previously used ciphertext otherwise. It then computes the token

$$\text{stk}_i := \text{vtk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((\text{rtk}_{\mathbf{r}})_{\mathbf{r} \in \text{DB}_{\mathbf{a}_i = X_i}}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_V, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right) \right).$$

If, on the other hand, $\mathcal{Z}(\mathbf{a}_i)$ is type-2, \mathcal{S} starts by computing

$$\text{dtk}_i \leftarrow \mathcal{S}_{\text{DX}} \left(\text{EMM}_{\mathbf{a}_i}, \mathcal{L}_Q^{\text{dx}} \left(\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right) \right)$$

where $\text{EMM}_{\mathbf{a}_i} \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_{\mathbf{a}_i}))$ if it has never been accessed before and where $\text{EMM}_{\mathbf{a}_i}$ is the previously-used structure otherwise. It then computes, for all $(\mathbf{r}_1, \mathbf{r}_2) \in \text{DB}_{\mathbf{a}_i = X_i}$,

$$\text{rtk}_1 \leftarrow \mathcal{S}_{\text{MM}} \left((ct_j)_{j \in [\#\mathbf{r}_1]}, \mathcal{L}_Q^{\text{MM}} \left(\text{MM}_R, \langle \text{tbl}(\mathbf{r}_1), \text{rrk}(\mathbf{r}_1) \rangle \right) \right)$$

and

$$\text{rtk}_2 \leftarrow \mathcal{S}_{\text{MM}} \left((ct'_j)_{j \in [\#\mathbf{r}_2]}, \mathcal{L}_Q^{\text{MM}} \left(\text{MM}_R, \langle \text{tbl}(\mathbf{r}_2), \text{rrk}(\mathbf{r}_2) \rangle \right) \right),$$

where $ct_j \leftarrow \text{Enc}_{K_1}(\mathbf{0}^{|r_1[j]|})$ and $ct'_j \leftarrow \text{Enc}_{K_1}(\mathbf{0}^{|r_2[j]|})$ if they have never been accessed and ct_j and ct'_j are the previously-used ciphertexts otherwise. Finally, it computes

$$\text{jtk}_i \leftarrow \mathcal{S}_{\text{MM}} \left(\left\{ \text{rtk}_{\mathbf{r}_1}, \text{rtk}_{\mathbf{r}_2} \right\}_{(\mathbf{r}_1, \mathbf{r}_2) \in \text{DB}_{\mathbf{a}_i = X_i}}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_{\mathbf{a}_i}, \left\langle \text{tbl}(\mathbf{a}_i), \text{tbl}(X_i) \right\rangle \right) \right),$$

and sets $\text{stk}_i = (\text{dtk}_i, \text{jtk}_i)$.

It remains to show that for all probabilistic polynomial-time adversaries \mathcal{A} , the probability that $\mathbf{Real}_{\text{SPX}, \mathcal{A}}(k)$ outputs 1 is negligibly-close to the probability that $\mathbf{Ideal}_{\text{SPX}, \mathcal{A}, \mathcal{S}}(k)$ outputs 1. We do this using the following sequence of games:

Game₀ : is the same as a $\mathbf{Real}_{\text{SPX},\mathcal{A}}(k)$ experiment. For ease of exposition, we define Loc_Ψ as the tuple of left-hand-side attributes in the type-2 terms of Ψ ; i.e., the \mathbf{a}_i 's in the terms of form $\mathbf{a}_i = \mathbf{x}_i$.

Game₁ : is the same as **Game₀**, except that EMM_C is replaced with the output of $\mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_C))$ and, for every normal form query q , each real y-token ytk_i is replaced with the output of

$$\mathcal{S}_{\text{MM}}\left(\left(\text{ct}_j\right)_{j \in [\#\mathbf{c}]}, \text{AccP}(\mathbf{c}), \mathcal{L}_Q^{\text{mm}}\left(\text{MM}_C, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle\right)\right).$$

Game₂ : is the same as **Game₁**, except that EMM_V is replaced with the output of $\mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_V))$ and, for every normal form query q , every type-1 s-token stk_i is replaced with the output of

$$\mathcal{S}_{\text{MM}}\left(\left(\text{rtk}_r\right)_{r \in \text{DB}_{\mathbf{a}_i = \mathbf{x}_i}}, \mathcal{L}_Q^{\text{mm}}\left(\text{MM}_V, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle\right)\right).$$

Game₃ : is the same as **Game₂**, except that EDX is replaced with the output of $\mathcal{S}_{\text{DX}}(\mathcal{L}_S^{\text{dx}}(\text{DX}))$ and, for every normal form query q , every dictionary token dtk_i in type-2 s-tokens $\text{stk}_i = (\text{dtk}_i, \text{jtk}_i)$ are replaced with the output of

$$\mathcal{S}_{\text{DX}}\left(\text{EMM}_{\mathbf{a}_i}, \text{AccP}(\text{EMM}_{\mathbf{a}_i}), \mathcal{L}_Q^{\text{dx}}\left(\left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle\right)\right).$$

Game_{3+l} for $l \in [\#\text{Loc}_\Psi]$: is the same as **Game_{2+l}**, except that $\text{EMM}_{\text{Loc}_\Psi[l]}$ is replaced with the output of $\mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_{\text{Loc}_\Psi[l]}))$ and for all type-2 terms $\mathbf{a}_i = \mathbf{x}_i$ in Ψ such that $\mathbf{a}_i = \text{Loc}_\Psi[l]$, the join tokens jtk_i in the type-2 s-token $\text{stk}_i = (\text{dtk}_i, \text{jtk}_i)$ is replaced with the output of

$$\mathcal{S}_{\text{MM}}\left(\left\{\text{rtk}_{r_1}, \text{rtk}_{r_2}\right\}_{(r_1, r_2) \in \text{DB}_{\text{Loc}_\Psi[l] = \mathbf{x}_i}}, \mathcal{L}_Q^{\text{mm}}\left(\text{MM}_{\text{Loc}_\Psi[l]}, \left\langle \text{tbl}(\text{Loc}_\Psi[l]), \text{crk}(\text{Loc}_\Psi[l]) \right\rangle\right), \left\langle \text{tbl}(\mathbf{x}_i), \text{crk}(\mathbf{x}_i) \right\rangle\right).$$

Game_{4+#Loc_Ψ} : is the same as **Game_{3+#Loc_Ψ}**, except that EMM_R is replaced with the output of $\mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_R))$ and every row token rtk_r for a row r is replaced with the output of of

$$\mathcal{S}_{\text{MM}}\left(\left\{\text{ct}_j\right\}_{j \in [\#\mathbf{r}]}, \text{AccP}(\mathbf{r}), \mathcal{L}_Q^{\text{mm}}\left(\text{MM}_R, \left\langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \right\rangle\right)\right)$$

where $\text{ct}_j \leftarrow \text{Enc}_{K_1}(r_j)$.

Game_{5+#Loc_Ψ} : is the same as **Game_{4+#Loc_Ψ}**, except that every SKE encryption ct of a message m is replaced with $\text{ct} \leftarrow \text{Enc}_{K_1}(\mathbf{0}^{|m|})$. Note that message size information is always available to the simulator \mathcal{S} through its leakage.

Game_{6+#Loc_Ψ} : is the same as **Game_{5+#Loc_Ψ}**, except that for every token tk for a normal form query q , t and \mathbf{s} are replaced with t and \mathbf{s} from $\text{XPP}(\text{DB}, q)$. Note that **Game_{6+#Loc_Ψ}** is identical to $\mathbf{Ideal}_{\text{SPX},\mathcal{A},\mathcal{S}}(k)$.

Claim 1. For all PPT adversaries \mathcal{A} ,

$$\left| \Pr[\text{Game}_0 = 1] - \Pr[\text{Game}_1 = 1] \right| \leq \text{negl}(k).$$

We show that if there exists a PPT adversary \mathcal{A} that contradicts the claim, then there exists a PPT adversary \mathcal{B} that breaks the adaptive semantic security of Σ_{MM} with respect to an arbitrary PPT simulator \mathcal{S}' . \mathcal{B} starts by running \mathcal{A} . When \mathcal{A} outputs $\text{DB} = (\mathsf{T}_1, \dots, \mathsf{T}_n)$, \mathcal{B} creates a dictionary DX , a column multi-map MM_C , a row multi-map MM_R , a value multi-map MM_V and a set of multi-maps $\{\text{MM}_{\mathbf{c}}\}_{\mathbf{c} \in \text{DB}^\top}$, from DB . Note that this can be done by executing Steps 1 through 10 of SPX.Setup . \mathcal{B} then outputs MM_C . Upon receiving EMM_C^* —from either a $\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k)$ experiment or an $\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment— \mathcal{B} sends ERD to \mathcal{A} , where

$$\text{EDB} = (\text{EMM}_R, \text{EMM}_C^*, \text{EMM}_V, \text{EDX}),$$

with $(K_D, \text{EDX}) \leftarrow \Sigma_{\text{DX}}.\text{Setup}(1^k, \text{DX})$ such that for all $\mathbf{c} \in \text{DB}^\top$, $\text{DX}[\mathbf{c}] = \text{EMM}_{\mathbf{c}}$ where $(K_{\mathbf{c}}, \text{EMM}_{\mathbf{c}}) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_{\mathbf{c}})$, $(K_R, \text{EMM}_R) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_R)$ and $(K_V, \text{EMM}_V) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_V)$.

Whenever \mathcal{A} outputs a normal form query

$$q = \pi_{\mathbf{a}_1, \dots, \mathbf{a}_h} \left([a_1] \times \dots \times [a_f] \times \sigma_{\Psi}(\mathsf{T}_{i_1} \times \dots \times \mathsf{T}_{i_t}) \right),$$

where $\Psi = \mathbf{a}_1 = X_1 \wedge \dots \wedge \mathbf{a}_\ell = X_\ell$, \mathcal{B} does the following. For all $i \in [h]$, \mathcal{B} outputs $\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle$ as its own query to EMM_C^* and receives ytk_i^* . It then sends to \mathcal{A} the token

$$\text{tk} = (t, \mathbf{s}, (\mathbf{e}_i)_{i \in [f]}, (\text{ytk}_i^*)_{i \in [h]}, (\text{stk}_i)_{i \in [\ell]});$$

where t is the number of queried tables, $s = \sum_{j=1}^t \|T_{i_j}\|_c$ and for all $i \in [f]$, $\mathbf{e}_i \leftarrow \text{Enc}_{K_1}(a_i)$. For all $i \in [\ell]$, if $\mathbf{a}_i = X_i$ is type-1, then $\text{stk}_i = \text{vtk}_i$ such that $\text{vtk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_V} \left(\left\langle X_i, \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right\rangle \right)$.

Otherwise if $\mathbf{a}_i = X_i$ is type-2, then $\text{stk}_i = (\text{dtk}_i, \text{jtk}_i)$ such that $\text{dtk}_i \leftarrow \Sigma_{\text{DX}}.\text{Token}_{K_D} \left(\left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right)$ and $\text{jtk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_{\mathbf{a}_i}} \left(\left\langle \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle, \langle \text{tbl}(X_i), \text{crk}(X_i) \rangle \right\rangle \right)$.

After answering all of \mathcal{A} 's queries, \mathcal{A} outputs a bit which \mathcal{B} returns as its own output. Note that if \mathcal{B} is executed in a $\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k)$ experiment then, by construction, \mathcal{A} 's view is exactly its view in Game_0 . On the other hand, if \mathcal{B} is executed in an $\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment, then \mathcal{A} 's view is, by construction, exactly its view in Game_1 . It follows by our initial assumption that

$$\left| \Pr[\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k) = 1] - \Pr[\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k) = 1] \right|$$

is non-negligible, which is a contradiction. □

Claim 2. For all PPT adversaries \mathcal{A} ,

$$\left| \Pr[\text{Game}_1 = 1] - \Pr[\text{Game}_2 = 1] \right| \leq \text{negl}(k).$$

We show that if there exists a PPT adversary \mathcal{A} that contradicts the claim, then there exists a PPT adversary \mathcal{B} that breaks the adaptive semantic security of Σ_{MM} with respect to an arbitrary PPT simulator \mathcal{S}' . \mathcal{B} starts by running \mathcal{A} . When \mathcal{A} outputs $\text{DB} = (\mathsf{T}_1, \dots, \mathsf{T}_n)$, \mathcal{B} creates a dictionary DX , a column multi-map MM_C , a row multi-map MM_R , a value multi-map MM_V and a set of multi-maps $\{\text{MM}_{\mathbf{c}}\}_{\mathbf{c} \in \text{DB}^\top}$, from DB . \mathcal{B} then outputs MM_V . Upon receiving EMM_V^* —from either a $\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k)$ experiment or an $\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment— \mathcal{B} sends ERD to \mathcal{A} , where

$$\text{ERD} = (\text{EMM}_R, \text{EMM}_C, \text{EMM}_V^*, \text{EDX}),$$

with $\text{EMM}_C \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_C))$, $(K_D, \text{EDX}) \leftarrow \Sigma_{\text{DX}}.\text{Setup}(1^k, \text{DX})$ such that for all $\mathbf{c} \in \text{DB}^\top$, $\text{DX}[\mathbf{c}] = \text{EMM}_{\mathbf{c}}$ where $(K_{\mathbf{c}}, \text{EMM}_{\mathbf{c}}) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_{\mathbf{c}})$ and $(K_R, \text{EMM}_R) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_R)$.

Whenever \mathcal{A} outputs a normal form query q , \mathcal{B} does the following. For all $i \in [\ell]$ such that $\mathbf{a}_i = X_i$ is type-2, \mathcal{B} outputs

$$\left\langle X_i, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right\rangle$$

as its own query to EMM_V^* and receives vtk_i^* . It then sends to \mathcal{A} the token

$$\text{tk} = (t, \mathbf{s}, (\mathbf{e}_i)_{i \in [f]}, (\mathbf{ytk})_{i \in [h]}, (\mathbf{stk}_i)_{i \in [\ell]}^*);$$

where t is the number of queried tables, $s = \sum_{j=1}^t \|T_{i_j}\|_c$ and for all $i \in [f]$, $\mathbf{e}_i \leftarrow \text{Enc}_{K_1}(a_i)$. If $\mathbf{a}_i \in \mathbb{S}(\text{DB})$,

$$\mathbf{ytk}_i := \mathbf{ptk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((\text{ct}_j)_{j \in [\#\mathbf{c}]}, \text{AccP}(\mathbf{a}_i), \mathcal{L}_{\mathbb{Q}}^{\text{mm}} \left(\text{MM}_C, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right) \right).$$

Given $\text{AccP}(\mathbf{a}_i)$, ct_j is either equal to $\text{Enc}_{K_1}(c_j)$ if the column has never been accessed before, or the previously used ciphertext otherwise. If, on the other hand, $\mathbf{a}_i \in \mathbb{S}(\text{VDB})$, $\mathbf{ytk}_i := \mathbf{a}_i$. For all $i \in [\ell]$, if $\mathbf{a}_i = X_i$ is type-2, then $\mathbf{stk}_i = (\text{dtk}_i, \text{jtk}_i)$ such that $\text{dtk}_i \leftarrow \Sigma_{\text{DX}}.\text{Token}_{K_D} \left(\left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle \right)$ and $\text{jtk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_{\mathbf{a}_i}} \left(\left\langle \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle, \left\langle \text{tbl}(X_i), \text{crk}(X_i) \right\rangle \right) \right)$.

After answering all of \mathcal{A} 's queries, \mathcal{A} outputs a bit which \mathcal{B} returns as its own output. Note that if \mathcal{B} is executed in a $\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k)$ experiment then, by construction, \mathcal{A} 's view is exactly its view in Game_1 . On the other hand, if \mathcal{B} is executed in an $\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment, then \mathcal{A} 's view is, by construction, exactly its view in Game_2 . It follows by our initial assumption that

$$\left| \Pr[\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k) = 1] - \Pr[\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k) = 1] \right|$$

is non-negligible, which is a contradiction. □

Claim 3. For all PPT adversaries \mathcal{A} ,

$$\left| \Pr[\text{Game}_2 = 1] - \Pr[\text{Game}_3 = 1] \right| \leq \text{negl}(k).$$

We show that if there exists a PPT adversary \mathcal{A} that contradicts the claim, then there exists a PPT adversary \mathcal{B} that breaks the adaptive semantic security of Σ_{DX} with respect to an arbitrary PPT simulator \mathcal{S}' . \mathcal{B} starts by running \mathcal{A} . When \mathcal{A} outputs $\text{DB} = (\mathbb{T}_1, \dots, \mathbb{T}_n)$, \mathcal{B} creates a dictionary DX , a column multi-map MM_C , a row multi-map MM_R , a value multi-map MM_V and a set of multi-maps $\{\text{MM}_c\}_{c \in \text{DB}^\top}$, from DB . \mathcal{B} then outputs DX . Upon receiving EDX^* —from either a $\mathbf{Real}_{\Sigma_{\text{DX}}, \mathcal{B}}(k)$ experiment or an $\mathbf{Ideal}_{\Sigma_{\text{DX}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment— \mathcal{B} sends ERD to \mathcal{A} , where

$$\text{EDB} = (\text{EMM}_R, \text{EMM}_C, \text{EMM}_V, \text{EDX}^*),$$

with $\text{EMM}_C \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_C))$, $\text{EMM}_V \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_V))$ and $(K_R, \text{EMM}_R) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_R)$.

Whenever \mathcal{A} outputs a normal form query q , \mathcal{B} does the following. For all $i \in [\ell]$ such that $\mathbf{a}_i = X_i$ is type-2, \mathcal{B} outputs $\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle$ as its own query to EDX^* and receives dtk_i^* . It then sends to \mathcal{A} the token

$$\text{tk} = (t, \mathbf{s}, (e_i)_{i \in [f]}, (\text{ytk})_{i \in [h]}, (\text{stk}_i)_{i \in [\ell]}^*);$$

where t is the number of queried tables, $s = \sum_{j=1}^t \|T_{i_j}\|_c$ and for all $i \in [f]$, $e_i \leftarrow \text{Enc}_{K_1}(\mathbf{a}_i)$. If $\mathbf{a}_i \in \mathbb{S}(\text{DB})$,

$$\text{ytk}_i := \text{ptk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((ct_j)_{j \in [\#c]}, \text{AccP}(\mathbf{a}_i), \mathcal{L}_{\mathbb{Q}}^{\text{mm}} \left(\text{MM}_C, \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right) \right).$$

Given $\text{AccP}(\mathbf{a}_i)$, ct_j is either equal to $\text{Enc}_{K_1}(c_j)$ if the column has never been accessed before, or the previously used ciphertext otherwise. If, on the other hand, $\mathbf{a}_i \in \mathbb{S}(\text{VDB})$, $\text{ytk}_i := \mathbf{a}_i$. For all $i \in [\ell]$, if $\mathbf{a}_i = X_i$ is type-1, $\text{stk}_i := \text{vtk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((\text{rtk}_{\mathbf{r}})_{\mathbf{r} \in \text{DB}_{\mathbf{a}_i = X_i}}, \mathcal{L}_{\mathbb{Q}}^{\text{mm}} \left(\text{MM}_V, \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right) \right)$, with $\text{rtk}_{\mathbf{r}} \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_R} \left(\langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \rangle \right)$ for $\mathbf{r} \in \text{DB}_{\mathbf{a}_i = X_i}$. If $\mathbf{a}_i = X_i$ is type-2, then $\text{stk}_i = (\text{dtk}_i^*, \text{jtk}_i)$ such that $\text{jtk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_{\mathbf{a}_i}} \left(\left\langle \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle, \langle \text{tbl}(X_i), \text{crk}(X_i) \rangle \right\rangle \right)$.

After answering all of \mathcal{A} 's queries, \mathcal{A} outputs a bit which \mathcal{B} returns as its own output. Note that if \mathcal{B} is executed in a $\mathbf{Real}_{\Sigma_{\text{DX}}, \mathcal{B}}(k)$ experiment then, by construction, \mathcal{A} 's view is exactly its view in Game_2 . On the other hand, if \mathcal{B} is executed in an $\mathbf{Ideal}_{\Sigma_{\text{DX}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment, then \mathcal{A} 's view is, by construction, exactly its view in Game_3 . It follows by our initial assumption that

$$\left| \Pr [\mathbf{Real}_{\Sigma_{\text{DX}}, \mathcal{B}}(k) = 1] - \Pr [\mathbf{Ideal}_{\Sigma_{\text{DX}}, \mathcal{B}, \mathcal{S}'}(k) = 1] \right|$$

is non-negligible, which is a contradiction. □

Claim 4. For all $l \in [\#\text{Loc}_{\Psi}]$, all PPT adversaries \mathcal{A} ,

$$\left| \Pr [\text{Game}_{2+l} = 1] - \Pr [\text{Game}_{3+l} = 1] \right| \leq \text{negl}(k).$$

We show that if there exists a PPT adversary \mathcal{A} that contradicts the claim, then there exists a PPT adversary \mathcal{B} that breaks the adaptive semantic security of Σ_{MM} with respect to an arbitrary PPT simulator \mathcal{S}' . \mathcal{B} starts by running \mathcal{A} . When \mathcal{A} outputs $\text{DB} = (\mathbb{T}_1, \dots, \mathbb{T}_n)$, \mathcal{B} creates a dictionary

DX, a column multi-map MM_C , a row multi-map MM_R , a value multi-map MM_V and a set of multi-maps $\{\text{MM}_c\}_{c \in \text{DB}^\top}$, from DB. \mathcal{B} then outputs $\text{MM}_{\text{Loc}_\Psi[l]}$. Upon receiving $\text{EMM}_{\text{Loc}_\Psi[l]}^*$ —from either a $\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k)$ experiment or an $\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment— \mathcal{B} sends ERD to \mathcal{A} , where

$$\text{EDB} = (\text{EMM}_R, \text{EMM}_C, \text{EMM}_V, \text{EDX}),$$

with $\text{EMM}_C \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_C))$, $\text{EMM}_V \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_V))$, $\text{EDX} \leftarrow \mathcal{S}_{\text{DX}}(\mathcal{L}_S^{\text{dx}}(\text{DX}))$ and $(K_R, \text{EMM}_R) \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_R)$.

Whenever \mathcal{A} outputs a normal form query q , \mathcal{B} does the following. For all $i \in [\ell]$ such that $\mathbf{a}_i = X_i$ with $\mathbf{a}_i = \text{Loc}_\Psi[l]$, \mathcal{B} outputs

$$\left\langle \langle \text{tbl}(\text{Loc}_\Psi[l]), \text{crk}(\text{Loc}_\Psi[l]) \rangle, \langle \text{tbl}(X_i), \text{crk}(X_i) \rangle \right\rangle$$

as its own query to $\text{EMM}_{\text{Loc}_\Psi[l]}^*$ and receives jtk_i^* . It then sends to \mathcal{A} the token

$$\text{tk} = (t, \mathbf{s}, (\mathbf{e}_i)_{i \in [f]}, (\text{ytk})_{i \in [h]}, (\text{stk}_i)_{i \in [\ell]}^*);$$

where t is the number of queried tables, $s = \sum_{j=1}^t \|T_{i_j}\|_c$ and for all $i \in [f]$, $\mathbf{e}_i \leftarrow \text{Enc}_{K_1}(\mathbf{a}_i)$. If $\mathbf{a}_i \in \mathbb{S}(\text{DB})$,

$$\text{ytk}_i := \text{ptk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((\text{ct}_j)_{j \in [\#c]}, \text{AccP}(\mathbf{a}_i), \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_C, \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right) \right).$$

Given $\text{AccP}(\mathbf{a}_i)$, ct_j is either equal to $\text{Enc}_{K_1}(c_j)$ if the column has never been accessed before, or the previously used ciphertext otherwise. If, on the other hand, $\mathbf{a}_i \in \mathbb{S}(\text{VDB})$, $\text{ytk}_i := \mathbf{a}_i$. For all $i \in [\ell]$, if $\mathbf{a}_i = X_i$ is type-1, $\text{stk}_i := \text{vtk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((\text{rtk}_r)_{r \in \text{DB}_{\mathbf{a}_i = X_i}}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_V, \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right) \right)$, with $\text{rtk}_r \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_R} \left(\langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \rangle \right)$ for $\mathbf{r} \in \text{DB}_{\mathbf{a}_i = X_i}$. If $\mathbf{a}_i = X_i$ is type-2, then $\text{stk}_i = (\text{dtk}_i, \text{jtk}_i^*)$ equals:

- if $z = l$, $\text{dtk}_i \leftarrow \mathcal{S}_{\text{DX}} \left(\text{EMM}_{\text{Loc}_\Psi[l]}^*, \mathcal{L}_Q^{\text{dx}} \left(\langle \text{tbl}(\text{Loc}_\Psi[l]), \text{crk}(\text{Loc}_\Psi[l]) \rangle \right) \right)$ and $\text{jtk}_i = \text{jtk}_i^*$
- for all $z < l$, $\text{dtk}_i \leftarrow \mathcal{S}_{\text{DX}} \left(\text{EMM}_{\text{Loc}_\Psi[z]}, \text{AccP}(\text{EMM}_{\text{Loc}_\Psi[z]}), \mathcal{L}_Q^{\text{dx}} \left(\langle \text{tbl}(\text{Loc}_\Psi[z]), \text{crk}(\text{Loc}_\Psi[z]) \rangle \right) \right)$ with $\text{AccP}(\text{EMM}_{\text{Loc}_\Psi[z]})$ is a leakage that captures when and where $\text{EMM}_{\text{Loc}_\Psi[z]}$ was generated. Thus, it can be either equal to $\text{EMM}_{\text{Loc}_\Psi[z]} \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_S^{\text{mm}}(\text{MM}_{\text{Loc}_\Psi[z]}))$, or, a previously simulated one, and

$$\text{jtk}_i \leftarrow \mathcal{S}_{\text{MM}} \left(\left\{ \text{rtk}_{\mathbf{r}_1}, \text{rtk}_{\mathbf{r}_2} \right\}_{(\mathbf{r}_1, \mathbf{r}_2) \in \text{DB}_{\text{Loc}_\Psi[z] = X_i}}, \mathcal{L}_Q^{\text{mm}} \left(\text{MM}_{\text{Loc}_\Psi[z]}, \langle \text{tbl}(\text{Loc}_\Psi[z]), \text{tbl}(X_i) \rangle \right) \right),$$

with $\text{rtk}_{\mathbf{r}_1} \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_R} \left(\langle \text{tbl}(\mathbf{r}_1), \text{rrk}(\mathbf{r}_1) \rangle \right)$ and $\text{rtk}_{\mathbf{r}_2} \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_R} \left(\langle \text{tbl}(\mathbf{r}_2), \text{rrk}(\mathbf{r}_2) \rangle \right)$.

- for all $z > l$, $\text{dtk}_i \leftarrow \mathcal{S}_{\text{DX}} \left(\text{EMM}_{\text{Loc}\Psi[z]}, \text{AccP}(\text{EMM}_{\text{Loc}\Psi[z]}), \mathcal{L}_{\mathbb{Q}}^{\text{dx}} \left(\langle \text{tbl}(\text{Loc}\Psi[z]), \text{crk}(\text{Loc}\Psi[z]) \rangle \right) \right)$
with either $\text{EMM}_{\text{Loc}\Psi[z]} \leftarrow \Sigma_{\text{MM}}.\text{Setup}(1^k, \text{MM}_{\text{Loc}\Psi[z]}, \text{ or a previously generated one, and$

$$\text{jtk}_i \leftarrow \Sigma_{\text{MM}}.\text{Token}_{K_{a_i}} \left(\left\langle \left\langle \text{tbl}(\text{Loc}\Psi[z]), \text{crk}(\text{Loc}\Psi[z]) \right\rangle, \langle \text{tbl}(X_i), \text{crk}(X_i) \rangle \right\rangle \right).$$

After answering all of \mathcal{A} 's queries, \mathcal{A} outputs a bit which \mathcal{B} returns as its own output. Note that if \mathcal{B} is executed in a $\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k)$ experiment then, by construction, \mathcal{A} 's view is exactly its view in Game_{2+l} . On the other hand, if \mathcal{B} is executed in an $\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment, then \mathcal{A} 's view is, by construction, exactly its view in Game_{3+l} . It follows by our initial assumption that

$$\left| \Pr [\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k) = 1] - \Pr [\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k) = 1] \right|$$

is non-negligible, which is a contradiction. □

Claim 5. For all PPT adversaries \mathcal{A} ,

$$\left| \Pr [\text{Game}_{3+\#\text{Loc}\Psi} = 1] - \Pr [\text{Game}_{4+\#\text{Loc}\Psi} = 1] \right| \leq \text{negl}(k).$$

We show that if there exists a PPT adversary \mathcal{A} that contradicts the claim, then there exists a PPT adversary \mathcal{B} that breaks the adaptive semantic security of Σ_{MM} with respect to an arbitrary PPT simulator \mathcal{S}' . \mathcal{B} starts by running \mathcal{A} . When \mathcal{A} outputs $\text{DB} = (\mathbb{T}_1, \dots, \mathbb{T}_n)$, \mathcal{B} creates a dictionary DX , a column multi-map MM_C , a row multi-map MM_R , a value multi-map MM_V and a set of multi-maps $\{\text{MM}_c\}_{c \in \text{DB}^\top}$, from DB . \mathcal{B} then outputs MM_R . Upon receiving EMM_R^* —from either a $\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k)$ experiment or an $\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment— \mathcal{B} sends ERD to \mathcal{A} , where

$$\text{ERD} = (\text{EMM}_R^*, \text{EMM}_C, \text{EMM}_V, \text{EDX}),$$

with $\text{EMM}_C \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_C))$, $\text{EMM}_V \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_V))$ and $\text{EDX} \leftarrow \mathcal{S}_{\text{DX}}(\mathcal{L}_{\mathbb{S}}^{\text{dx}}(\text{DX}))$.

Whenever \mathcal{A} outputs a normal form query q , \mathcal{B} does the following. \mathcal{B} outputs $\langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \rangle$ as its own query to EMM_R^* and receives rtk_r^* . It then sends to \mathcal{A} the token

$$\text{tk} = (t, \mathbf{s}, (\mathbf{e}_i)_{i \in [f]}, (\mathbf{ytk})_{i \in [h]}, (\mathbf{stk}_i)_{i \in [\ell]});$$

where t is the number of queried tables, $s = \sum_{j=1}^t \|T_{i_j}\|_c$ and for all $i \in [f]$, $\mathbf{e}_i \leftarrow \text{Enc}_{K_1}(a_i)$. If $\mathbf{a}_i \in \mathbb{S}(\text{DB})$,

$$\text{ytk}_i := \text{ptk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((\text{ct}_j)_{j \in [\#\mathbf{c}]}, \text{AccP}(\mathbf{a}_i), \mathcal{L}_{\mathbb{Q}}^{\text{mm}} \left(\text{MM}_C, \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right) \right).$$

Given $\text{AccP}(\mathbf{a}_i)$, ct_j is either equal to $\text{Enc}_{K_1}(c_j)$ if the column has never been accessed before, or the previously used ciphertext otherwise. If, on the other hand, $\mathbf{a}_i \in \mathbb{S}(\text{VDB})$, $\text{ytk}_i := \mathbf{a}_i$. For all $i \in [\ell]$, if $\mathbf{a}_i = X_i$ is type-1, $\text{stk}_i := \text{vtk}_i \leftarrow \mathcal{S}_{\text{MM}} \left((\text{rtk}_r^*)_{r \in \text{DB}_{\mathbf{a}_i = X_i}}, \mathcal{L}_{\mathbb{Q}}^{\text{mm}} \left(\text{MM}_V, \langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right) \right)$. If

$\mathbf{a}_i = X_i$ is type-2, then $\mathbf{stk}_i = (\mathbf{dtk}_i, \mathbf{jtk}_i)$ such that $\mathbf{dtk}_i \leftarrow \mathcal{S}_{\text{DX}}\left(\text{EMM}_{\mathbf{a}_i}, \text{AccP}(\text{EMM}_{\mathbf{a}_i}), \mathcal{L}_{\mathbb{Q}}^{\text{dx}}\left(\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle\right)\right)$ with $\text{AccP}(\text{EMM}_{\mathbf{a}_i})$ is a leakage that captures when and where $\text{EMM}_{\mathbf{a}_i}$ was generated. Thus, it can be either equal to $\text{EMM}_{\mathbf{a}_i} \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_{\mathbf{a}_i}))$, or, a previously simulated one, and

$$\mathbf{jtk}_i \leftarrow \mathcal{S}_{\text{MM}}\left(\left\{\text{rtk}_{\mathbf{r}_1}^*, \text{rtk}_{\mathbf{r}_2}^*\right\}_{(\mathbf{r}_1, \mathbf{r}_2) \in \text{DB}_{\mathbf{a}_i = X_i}}, \mathcal{L}_{\mathbb{Q}}^{\text{mm}}\left(\text{MM}_{\mathbf{a}_i}, \left\langle \text{tbl}(\mathbf{a}_i), \text{tbl}(X_i) \right\rangle\right)\right).$$

After answering all of \mathcal{A} 's queries, \mathcal{A} outputs a bit which \mathcal{B} returns as its own output. Note that if \mathcal{B} is executed in a $\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k)$ experiment then, by construction, \mathcal{A} 's view is exactly its view in $\mathbf{Game}_{3+\#\text{Loc}_{\Psi}}$. On the other hand, if \mathcal{B} is executed in an $\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k)$ experiment, then \mathcal{A} 's view is, by construction, exactly its view in $\mathbf{Game}_{4+\#\text{Loc}_{\Psi}}$. It follows by our initial assumption that

$$\left| \Pr[\mathbf{Real}_{\Sigma_{\text{MM}}, \mathcal{B}}(k) = 1] - \Pr[\mathbf{Ideal}_{\Sigma_{\text{MM}}, \mathcal{B}, \mathcal{S}'}(k) = 1] \right|$$

is non-negligible, which is a contradiction. □

Claim 6. For all PPT adversaries \mathcal{A} ,

$$\left| \Pr[\mathbf{Game}_{4+\#\text{Loc}_{\Psi}} = 1] - \Pr[\mathbf{Game}_{5+\#\text{Loc}_{\Psi}} = 1] \right| \leq \text{negl}(k).$$

We show that if there exists a PPT adversary \mathcal{A} that contradicts the claim, then there exists a PPT adversary \mathcal{B} that breaks the adaptive semantic security of $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ with respect to an arbitrary PPT simulator \mathcal{S}' . \mathcal{B} starts by running \mathcal{A} . When \mathcal{A} outputs $\text{DB} = (\mathbb{T}_1, \dots, \mathbb{T}_n)$, \mathcal{B} creates a dictionary DX , a column multi-map MM_C , a row multi-map MM_R , a value multi-map MM_V and a set of multi-maps $\{\text{MM}_{\mathbf{c}}\}_{\mathbf{c} \in \text{DB}^{\uparrow}}$, from DB . \mathcal{B} then outputs all message m for all cells. Upon receiving ct^* for all cells in DB —from either a $\mathbf{Real}_{\text{SKE}, \mathcal{B}}(k)$ experiment or an $\mathbf{Ideal}_{\text{SKE}, \mathcal{B}, \mathcal{S}'}(k)$ experiment— \mathcal{B} sends ERD to \mathcal{A} , where

$$\text{EDB} = (\text{EMM}_R, \text{EMM}_C, \text{EMM}_V, \text{EDX}),$$

with $\text{EMM}_C \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_C))$, $\text{EMM}_V \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_V))$, $\text{EDX} \leftarrow \mathcal{S}_{\text{DX}}(\mathcal{L}_{\mathbb{S}}^{\text{dx}}(\text{DX}))$ and $\text{EMM}_R \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_R))$.

Whenever \mathcal{A} outputs a query χ , \mathcal{B} does the following. For all $i \in [f]$, \mathcal{B} outputs \mathbf{a}_i and receives \mathbf{e}_i^* . It then sends to \mathcal{A} the token

$$\mathbf{tk} = (t, \mathbf{s}, (\mathbf{e}_i)_{i \in [f]}^*, (\mathbf{ytk})_{i \in [h]}, (\mathbf{stk}_i)_{i \in [\ell]});$$

where t is the number of queried tables and $s = \sum_{j=1}^t \|T_{i_j}\|_{\mathbf{c}}$. If $\mathbf{a}_i \in \mathbb{S}(\text{DB})$,

$$\mathbf{ytk}_i := \mathbf{ptk}_i \leftarrow \mathcal{S}_{\text{MM}}\left(\left(\text{ct}_j^*\right)_{j \in [\#\mathbf{c}]}, \text{AccP}(\mathbf{a}_i), \mathcal{L}_{\mathbb{Q}}^{\text{mm}}\left(\text{MM}_C, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle\right)\right).$$

Given $\text{AccP}(\mathbf{a}_i)$, ct_j^* is either equal to a simulated ciphertext if the column has never been accessed before, or the previously used ciphertext otherwise. If, on the other hand, $\mathbf{a}_i \in \mathbb{S}(\text{VDB})$, $\mathbf{ytk}_i := \mathbf{a}_i$. For all $i \in [\ell]$, if $\mathbf{a}_i = X_i$ is type-1, $\mathbf{stk}_i := \mathbf{vtk}_i \leftarrow \mathcal{S}_{\text{MM}}\left(\left(\text{rtk}_{\mathbf{r}}\right)_{\mathbf{r} \in \text{DB}_{\mathbf{a}_i = X_i}}, \mathcal{L}_{\mathbb{Q}}^{\text{mm}}\left(\text{MM}_V, \left\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \right\rangle\right)\right)$.

If $\mathbf{a}_i = X_i$ is type-2, then $\text{stk}_i = (\text{dtk}_i, \text{jtk}_i) \text{dtk}_i \leftarrow \mathcal{S}_{\text{DX}} \left(\text{EMM}_{\mathbf{a}_i}, \text{AccP}(\text{EMM}_{\mathbf{a}_i}), \mathcal{L}_{\mathbb{Q}}^{\text{dx}} \left(\langle \text{tbl}(\mathbf{a}_i), \text{crk}(\mathbf{a}_i) \rangle \right) \right)$ with $\text{AccP}(\text{EMM}_{\mathbf{a}_i})$ is a leakage that captures when and where $\text{EMM}_{\mathbf{a}_i}$ was generated. Thus, it can be either equal to $\text{EMM}_{\mathbf{a}_i} \leftarrow \mathcal{S}_{\text{MM}}(\mathcal{L}_{\mathbb{S}}^{\text{mm}}(\text{MM}_{\mathbf{a}_i}))$, or, a previously simulated one, and

$$\text{jtk}_i \leftarrow \mathcal{S}_{\text{MM}} \left(\left\{ \text{rtk}_{\mathbf{r}_1}, \text{rtk}_{\mathbf{r}_2} \right\}_{(\mathbf{r}_1, \mathbf{r}_2) \in \text{DB}_{\mathbf{a}_i = X_i}}, \mathcal{L}_{\mathbb{Q}}^{\text{mm}} \left(\text{MM}_{\mathbf{a}_i}, \langle \text{tbl}(\mathbf{a}_i), \text{tbl}(X_i) \rangle \right) \right),$$

with $\text{rtk}_{\mathbf{r}}$ in both cases above equals

$$\text{rtk}_{\mathbf{r}} \leftarrow \mathcal{S}_{\text{MM}} \left(\left\{ \text{ct}_j^* \right\}_{j \in [\#\mathbf{r}]}, \text{AccP}(\mathbf{r}), \mathcal{L}_{\mathbb{Q}}^{\text{mm}} \left(\text{MM}_R, \langle \text{tbl}(\mathbf{r}), \text{rrk}(\mathbf{r}) \rangle \right) \right),$$

where $\text{AccP}(\mathbf{r})$ captures when and where \mathbf{r} has been queried. Thus, ct_j^* can be either freshly generated, or, re-used from a previously generated one.

After answering all of \mathcal{A} 's queries, \mathcal{A} outputs a bit which \mathcal{B} returns as its own output. Note that if \mathcal{B} is executed in a $\mathbf{Real}_{\text{SKE}, \mathcal{B}}(k)$ experiment then, by construction, \mathcal{A} 's view is exactly its view in $\text{Game}_{4+\#\text{Loc}_{\Psi}}$. On the other hand, if \mathcal{B} is executed in an $\mathbf{Ideal}_{\text{SKE}, \mathcal{B}, \mathcal{S}'}(k)$ experiment, then \mathcal{A} 's view is, by construction, exactly its view in $\text{Game}_{5+\#\text{Loc}_{\Psi}}$. It follows by our initial assumption that

$$\left| \Pr [\mathbf{Real}_{\text{SKE}, \mathcal{B}}(k) = 1] - \Pr [\mathbf{Ideal}_{\text{SKE}, \mathcal{B}, \mathcal{S}'}(k) = 1] \right|$$

is non-negligible, which is a contradiction. □

Claim 7. For all PPT adversaries \mathcal{A} ,

$$\left| \Pr [\text{Game}_{5+\#\text{Loc}_{\Psi}} = 1] - \Pr [\text{Game}_{6+\#\text{Loc}_{\Psi}} = 1] \right|.$$

It is clear that the view of both games is the same. ■