

Improved Structured Encryption for SQL Databases via Hybrid Indexing

David Cash¹, Ruth Ng², and Adam Rivkin¹

¹ University of Chicago

² University of California San Diego

{davidcash,amrivkin}@uchicago.edu, ring@eng.ucsd.edu

Abstract. We introduce a new technique for indexing joins in encrypted SQL databases called *partially precomputed joins* which achieves lower leakage and bandwidth than those used in prior constructions. These techniques are incorporated into state-of-the-art structured encryption schemes for SQL data, yielding a *hybrid indexing* scheme with both partially and fully precomputed join indexes. We then introduce the idea of *leakage-aware query planning* by giving a heuristic that helps the client decide, at query time, which index to use so as to minimize leakage and stay below a given bandwidth budget. We conclude by simulating our constructions on real datasets, showing that our heuristic is accurate and that partially-precomputed joins perform well in practice.

1 Introduction

SQL applications are often deterred from using cloud storage solutions because they do not wish to grant a third party access to their sensitive data. Yet, in-house solutions often are less convenient than these large-scale ones and are vulnerable to compromise as well. This calls for a cryptographic solution which allows data on the cloud to be end-to-end encrypted so that the server never “sees” the sensitive data. This in turn poses a challenge when the server is called upon to perform SQL operations on the data.

Most current offerings of this technology depend heavily on property-revealing encryption (PRE), making them vulnerable to leakage abuse attacks (LAAs). For example, Always Encrypted either deterministically encrypts columns or stores them with an ordered index [4]. These techniques have been shown to offer little-to-no privacy in certain practical scenarios [37,26].

A more promising approach is structured encryption (StE) which uses auxiliary encrypted data structures (e.g. encrypted multimaps) to support a subset of SQL queries [15]. This is done by translating the SQL query into tokens which can be passed to the server to query the auxiliary structures. The outputs of this are compiled, decrypted and processed to retrieve the SQL query result. Security is measured by *leakage profiles*, which characterize what information a curious server can learn. In particular, StE-based constructions leak equal or less than PRE-based constructions and resist most known LAAs [20,37,8,9,25,26,27].

OUR CONTRIBUTIONS. Our work can be grouped into three main contributions:

1. **Partially precomputed joins:** We introduce a new way to index (equi)joins which stems from the simple observation that when the server fully precomputes (FP) joins, the client has to download and decrypt a quadratic number of rows and the server learns the equality pattern of said rows. In our approach, the server partially precomputes (PP) joins: instead of indexing exactly which rows from the input table should be concatenated and returned, it just stores the set of rows from each input table that appears anywhere in the join output. At query time, the client downloads these sets and computes the join. When this is used to support SQL queries of the form “**select * from id_1 join id_2 on $at_1 = at_2$** ”, PP outperforms FP in both leakage and bandwidth at the cost of a logarithmic factor of client computation (in the worst case).
2. **Hybrid indexing:** When we incorporate PP joins into state-of-the-art StE schemes, we discover that some queries (e.g. those with a selection subquery) cannot be computed in the same way because the server does not know the equality pattern on the join columns (i.e. how the rows “match up”). So while PP joins are still the more secure choice, they sometimes incur more bandwidth than FP. To address this, we develop a hybrid StE scheme with both forms of indexing. The client chooses which to use at query time. We provide the first heuristic (that we are aware of) to enable this type of *leakage-aware client-side*

query planning, helping the client decide how to minimize leakage without exceeding a given bandwidth budget.

3. **Simulations on real data:** We quantify the effect of using FP and PP join indexing on bandwidth incurred by simulating our constructions on data from the City of Chicago’s Data Portal and MySQL’s sample Sakila database [2,3]. On simple (non-recursive) join queries, PP’s bandwidth is on average 231 times less than FP’s but more complex (recursive) queries are split down the middle as to which option used less bandwidth. We also demonstrate the accuracy of our heuristic under different client storage constraints. Assuming client storage comparable to that which is used in SQL Server, our heuristic chose a query plan with the maximal number of PP joins 79% of the time, and the optimal query plan 68% of the time.

RELATED WORK. Encrypted databases have been treated from a variety of perspectives. Structured encryption (StE) was defined by Chase and Kamara (CK) and is a special case of SSE, which was first defined by SWP [42].

We see our work as a direct extension and improvement upon SPX and OPX, two schemes which applied StE to the problem of indexing SQL databases [15,32,34]. Both our scheme and OPX address a similar query class to the one introduced in SPX, but lower leakage by using the hashset technique from OXT and primitives inspired by CJJKRS [12,13]. In particular, our FpSj scheme in Section 4.2 bears many similarities to OPX with minor leakage improvements from using a single indexing data structure. Our PpSj and HybStl schemes (in Section 4.2 and Section 5 respectively) introduce a new technique which further lowers leakage and server storage. For non-recursive queries, there are also substantial bandwidth savings.

PRE-based solutions achieves higher query support at the cost of higher leakage [39,24,1,21,44], and are particularly susceptible to leakage abuse attacks [20,37,8,9,25,26,27].

Finally, encrypted search has also been attempted using alternate models and architectures including the database-provider model [28], MPC [16,6], ORAM [23] and trusted execution environments [35,11,5].

Other works have also partially delegated computation to the client, to reduce leakage or increase query support, though none have applied it to joins [43,17,19].

2 Preliminaries

We denote the empty string with ϵ . Given positive integer n , let $[n] = \{1, 2, \dots, n\}$. Given tuples $\mathbf{t}_1 = (x_1, \dots, x_n)$ and $\mathbf{t}_2 = (y_1, \dots, y_m)$ we write $\mathbf{t}_1 \parallel \mathbf{t}_2$ as a shorthand for $(x_1, \dots, x_n, y_1, \dots, y_m)$. We extend set operations $\cap, \cup \in, \subseteq$ from sets to tuples by interpreting the tuples as sets.

Our algorithms often make use of dictionaries \mathbf{D} which map labels $\ell \in \{0, 1\}^*$ to values $\mathbf{D}[\ell] \in \{0, 1\}^* \cup \{\perp\}$. We also adopt the shorthand $\mathbf{D.Lbls} = \{\ell \in \{0, 1\}^* : \mathbf{D}[\ell] \neq \perp\}$. A multimap \mathbf{M} is an dictionary where $\mathbf{M}[\ell]$ is either a set of strings or \perp .

PSEUDOCODE. In pseudocode, we will assume that all integers, strings and sets are initialized to 0, ϵ and \emptyset respectively. For dictionaries and multimaps, they are initialized with all labels mapping to \perp . If S is a set or dictionary value, we write $S \stackrel{\cup}{\leftarrow} x$ in pseudocode as a shorthand for $S \leftarrow S \cup \{x\}$, initializing it first to \emptyset if necessary. If \mathbf{t} is a tuple, we similarly mean $\mathbf{t} \leftarrow \mathbf{t} \parallel (x)$ by writing $\mathbf{t} \leftarrow \mathbf{t} \stackrel{\cup}{\leftarrow} x$. Finally, we will write “Define $X : pred$ ” to set X (a function or constant) in such a way that the predicate $pred$ is true. If there are undefined variables in $pred$ we treat it as a random variable and expect that X is defined such that $pred$ will always be true.

GAMES. Our work uses the code-based game-playing framework of BR [7]. Let G be a game and A an adversary. Then, we write $\Pr[G(A)]$ to denote the probability that A plays G and the latter returns true. G may provide oracles to A , and if so we write A^{O_1, \dots, O_n} to denote that A is run with access to oracles O_1, \dots, O_n .

SYMMETRIC ENCRYPTION, IND\$-SECURITY. Symmetric Encryption (SE) scheme \mathbf{SE} defines key set $\mathbf{SE.KS}$, encryption algorithm $\mathbf{SE.Enc}$ and decryption algorithm $\mathbf{SE.Dec}$. Encryption is randomized, taking a key $K_e \in \mathbf{SE.KS}$ and a message $M \in \{0, 1\}^*$ and returns a ciphertext $C \in \{0, 1\}^*$. Decryption is deterministic and takes a key and ciphertext, returning a message. \mathbf{SE} also defines a ciphertext length function $\mathbf{SE.cl}$. We require that if $C \stackrel{s}{\leftarrow} \mathbf{SE.Enc}(K_e, M)$ then $|C| = \mathbf{SE.cl}(|M|)$ and $\Pr[\mathbf{SE.Dec}(K_e, C) = M] = 1$. We want our

Game $G_{SE}^{\text{ind}\$}(A)$	Game $G_F^{\text{prf}}(A)$
$b \leftarrow_{\$} \{0, 1\}$; $K_e \leftarrow_{\$} \text{SE.KS}$	$b \leftarrow_{\$} \{0, 1\}$; $K_f \leftarrow_{\$} \text{F.KS}$
$b' \leftarrow_{\$} A^{\text{ENC}}$; Return $b = b'$	$b' \leftarrow_{\$} A^{\text{FN}}$; Return $b = b'$
Alg $\text{ENC}(m)$	Alg $\text{FN}(X)$
$c_1 \leftarrow_{\$} \text{SE.Enc}(K_e, m)$	If $\mathbf{C}[X] = \perp$ then $\mathbf{C}[X] \leftarrow_{\$} \{0, 1\}^{\text{F.ol}}$
$c_0 \leftarrow_{\$} \{0, 1\}^{ c_1 }$; Return c_b	$c_1 \leftarrow_{\$} \text{F.Ev}(K_f, X)$; $c_0 \leftarrow \mathbf{C}[X]$; Return c_b

Fig. 1. Games used in defining IND\$ security of SE scheme SE (right) and PRF security of function family F (left)

SE schemes to protect the privacy of M , so ciphertexts should be indistinguishable from a random string of length $\text{SE.cl}(|M|)$. We capture this with the game $G_{SE}^{\text{ind}\$}$ in Fig. 1 and say that a scheme is IND\$-secure if $\text{Adv}_{SE}^{\text{ind}\$}(A) = 2 \Pr[G_{SE}^{\text{ind}\$}(A)] - 1$ is small for all adversaries A .

FUNCTION FAMILIES, PRF-SECURITY. A function family F defines a key set $F.\text{KS}$ and an output length $F.\text{ol}$. It defines a deterministic evaluation algorithm $F.\text{Ev}: F.\text{KS} \times \{0, 1\}^* \rightarrow \{0, 1\}^{F.\text{ol}}$. We define PRF security for function family F via the game G_F^{prf} depicted in Fig. 1. We say that F is a PRF if $\text{Adv}_F^{\text{prf}}(A) = 2 \Pr[G_F^{\text{prf}}(A)] - 1$ is small for all adversaries A .

3 Structured Indexing for SQL data types

We now generalize CK’s definition of structured encryption and provide a new framework for modeling encrypted SQL systems [15].

ABSTRACT DATA TYPES. An *abstract data type* ADT defines a domain set ADT.Dom , a query set ADT.QS , and a deterministic specification function $\text{ADT.Spec}: \text{ADT.Dom} \times \text{ADT.QS} \rightarrow \{0, 1\}^*$.

An example is the dictionary ADT DyAdt . DyAdt.Dom , DyAdt.QS contain all possible dictionaries \mathbf{D} and labels respectively (as defined in Section 2), and $\text{DyAdt.Spec}(\mathbf{D}, \ell) = \mathbf{D}[\ell]$. Multimap ADT MmAdt is defined analogously.

STRUCTURED INDEXING. We generalize Structured Encryption (StE) schemes (as defined by CK [15]) to *structured indexing* (StI) schemes. These are StE schemes without a decryption algorithm. The intuition here is that the handling of outsourced data often indexes the data in addition to encrypting it and we would like these encrypted indexes, whatever form they take, to achieve semantic security as well. Later, we show how this primitive allows us to modularize StE schemes. A StI scheme StI for ADT defines a set of keys StI.KS and the following algorithms:

- Randomized encryption algorithm StI.Enc which takes a key $K' \in \text{StI.KS}$ and an element of ADT.Dom and returns an updated key K and index $\text{IX} \in \{0, 1\}^*$. This syntax generalizes that of CK by allowing key generation to occur within or outside StI.Enc .
- Possibly randomized token generation algorithm StI.Tok which takes a key and a query from ADT.QS , and returns fixed length token $\text{tk} \in \{0, 1\}^{\text{StI.tl}}$.
- Deterministic evaluation algorithm StI.Eval which takes a token and index, and returns a ciphertext string $C \in \{0, 1\}^*$.
- Finalization algorithm StI.Fin which takes K, q and an input string, and returns an output string.

Intuitively, the client indexes his data then encrypts this index with StI.Enc , storing IX on the server. At query time, the client uses StI.Tok to generate a token and sends it to the server who runs StI.Eval , returning C to the client. StI.Fin can be used for client-side post-processing of the data. Note that the output of StI.Eval need not be the input to StI.Fin . In our indexing schemes the server will use the output of StI.Eval as “pointers” to retrieve rows of SQL data stored in a different data structure which in turn form the input to StI.Fin .

STRUCTURED ENCRYPTION. We can now define StE as a special cases of StI. Intuitively, an StE scheme is an StI scheme where the data structure is also used to store query responses (as opposed to just indexing them). The output of evaluation can be fed into finalization for decryption and should return the query result. To highlight this, StE schemes have a *decryption algorithm* StE.Dec in place of a finalization algorithm which

<p>Game $G_{\text{StE}}^{\text{COR}}(A)$</p> <p>$(DS, st) \leftarrow A(\mathbf{s}) ; K' \leftarrow \text{StE.KS}$ If $DS \notin \text{ADT.Dom}$ then return false $(K, \text{EDS}) \leftarrow \text{StE.Enc}(K', DS)$ $A^{\text{Tok}}(\mathbf{g}, \text{EDS}, st) ;$ Return win</p> <p>Oracle $\text{Tok}(q)$</p> <p>If $q \notin \text{ADT.QS}$ then win \leftarrow false $M \leftarrow \text{StE.Dec}(K, \text{StE.Eval}(\text{StE.Tok}(K, q), \text{EDS}))$ If $\text{ADT.Spec}(DS, q) \neq M$ then win \leftarrow false Return tk</p>	<p>Game $G_{\text{StI}, \mathcal{L}, \mathcal{S}}^{\text{SS}}(A)$</p> <p>$(DS, (q_1, \dots, q_n), st) \leftarrow A(\mathbf{s})$ $b \leftarrow \{0, 1\} ; K' \leftarrow \text{StI.KS}$ If $DS \notin \text{ADT.Dom}$ or $\{q_i\}_{i=1}^n \not\subseteq \text{ADT.QS}$ then return false If $b = 1$ then $(K, \text{IX}) \leftarrow \text{StI.Enc}(K', DS)$ For $i \in [n]$ do $\text{tk}_i \leftarrow \text{StI.Tok}(K, q_i)$ Else $(\text{IX}, (\text{tk}_1, \dots, \text{tk}_n)) \leftarrow \mathcal{S}(\mathcal{L}(DS, (q_1, \dots, q_n)))$ $b' \leftarrow A(\mathbf{g}, \text{IX}, (\text{tk}_1, \dots, \text{tk}_n), st) ;$ Return $(b = b')$</p>
--	--

Fig. 2. Games used in defining correctness for StE (structured encryption scheme for ADT) and semantic security for StI (structured indexing scheme for ADT) with respect to leakage algorithm \mathcal{L} and simulator \mathcal{S} .

takes as input K, q, C and returns the query result. We define correctness via game $G_{\text{StE}}^{\text{COR}}$ in Fig. 2 and say that StE is correct if the advantage of all adversaries A , defined $\text{Adv}_{\text{StE}}^{\text{COR}}(A) = \Pr[G_{\text{StE}}^{\text{COR}}(A)]$, is low. The correctness of our schemes will depend on the collision resistance of their function family primitives. Since we assume these are PRFs to prove security, we will also assume that their key-lengths are sufficient to ensure correctness.

We subdivide StE schemes into two types. We say that a scheme StE_{rr} is *response revealing* (RR) if evaluation itself returns the query result. In other words, decryption must be such that $\text{StE}_{\text{rr}}.\text{Dec}(K, q, C) = C$ for all K, q, C . An StE scheme that is not RR is *response hiding* (RH).

We refer to StE for the multimap and dictionary data types as multimap and dictionary encryption (MME/DYE) respectively. Our constructions make use of a specific dictionary encryption scheme adapted from CJJ+’s SSE scheme \prod_{bas} (2Lev in the Clusion library) [12,36]. In this scheme, the encrypted data structure is itself a dictionary \mathbf{D}' . We start by padding all values in the input dictionary to the same length, then for each label-value pair $\ell, \mathbf{D}[\ell]$, we do $\mathbf{D}'[\text{F.Ev}(K_f, \ell)] \leftarrow \text{SE.Enc}(K_e, \mathbf{D}[\ell])$ where F is a pseudorandom function family and SE is a symmetric encryption scheme. For completeness, we include the pseudocode of this dictionary encryption scheme (which we call Dye_π) in Appendix A. Our constructions also make use of a generic RR multimap encryption scheme. We adapt Dye_π to $\text{Mme}_\pi^{\text{rr}}$ (using a counter and label-dependent K_e) as an example of such a scheme in Appendix A.

SEMANTIC SECURITY. We define semantic security for StI using game $G_{\text{StI}, \mathcal{L}, \mathcal{S}}^{\text{SS}}$ depicted in Fig. 2, where StI is a StI scheme for ADT and \mathcal{L}, \mathcal{S} are algorithms we refer to as the leakage algorithm and simulator respectively. The adversary runs in a setup and guessing phase, as indicated by the first argument to it. Its advantage is $\text{Adv}_{\text{StI}, \mathcal{L}, \mathcal{S}}^{\text{SS}}(A) = 2\Pr[G_{\text{StI}, \mathcal{L}, \mathcal{S}}^{\text{SS}}(A) = 1] - 1$. Note that when StI is an StE scheme we recover CK’s non-adaptive security notion.

3.1 SQL Data Types

We now describe our notation for SQL data, queries and operations. We then define a class of ADTs we call *SQL data types* to construct StE schemes for.

SQL RELATIONS, DATABASES, SCHEMAS. SQL relation R defines a tuple of distinct attributes $\text{R.Atts} = (at_1, \dots, at_n)$. Each attribute is a bitstring $at \in \{0, 1\}^*$ and represents a “column” in the relation. R also defines a table R.T consisting of n -tuples of bitstrings representing the “rows” in the relation. Given a row $(x_1, \dots, x_n) = \mathbf{r} \in \text{R.T}$, we refer to the i -th entry of the row with $\mathbf{r}[at_i] = x_i$. We can initialize a relation with $\text{NewRltn}(\mathbf{at})$ which returns the relation with $\text{R.Atts} = \mathbf{at}$ and no rows.

We define a *database* to be a set of relations with disjoint attributes and their (distinct) identifiers, i.e. a set of the form $\text{DB} = \{(id_1, R_1), \dots, (id_N, R_N)\}$ where $i \neq j$ implies $id_i \neq id_j$ and $R_i.\text{Atts} \cap R_j.\text{Atts} = \emptyset$. We denote the identifier set of such a database as $\text{DB.IDs} = \{id_i\}_{i \in [N]}$ and retrieve relations by identifier using $\text{DB}[id_i] = R_i$. Since database attributes are non-repeating, we allow the retrieval of a table by any of its attributes using getID (i.e. if $\text{getID}(at, \text{DB}) = id$ then $at \in \text{DB}[id].\text{Atts}$). Similarly, if $\mathbf{t} \subseteq \text{DB}[id].\text{Atts}$, then $\text{getID}(\mathbf{t}, \text{DB}) = id$.

We require that each $(id, R) \in \text{DB}$ has a *unique key attribute* $\text{uk}(id) \in \text{R.Atts}$. This functions as a “row number” which uniquely identifies each row. In other words, for all distinct $\mathbf{r}, \mathbf{r}' \in \text{R.T}$, we have $\mathbf{r}[\text{uk}(id)] \neq$

R ₁ .T	
uk(id ₁)	at ₁
aa	Alice
bb	Alice
cc	Bob
dd	Charlie
ee	David

R ₂ .T		
uk(id ₂)	at ₂	at ₃
11	Alice	Math
22	Alice	Chem
33	Bob	CS
44	Eve	CS
55	Eve	Bio

(R ₁ ⋈ _{at₁,at₂} R ₂).T				
uk(id ₁)	at ₁	uk(id ₂)	at ₂	at ₃
aa	Alice	11	Alice	Math
aa	Alice	22	Alice	Chem
bb	Alice	11	Alice	Math
bb	Alice	22	Alice	Chem
cc	Bob	33	Bob	CS

(σ _{at₂,Eve} (R ₂)).T		
uk(id ₂)	at ₂	at ₃
44	Eve	CS
55	Eve	Bio

(σ _{at₃,CS} (R ₁ ⋈ _{at₁,at₂} R ₂)).T				
uk(id ₁)	at ₁	uk(id ₂)	at ₂	at ₃
cc	Bob	33	Bob	CS

Fig. 3. Examples of SQL relations R_1, R_2 and the output of join (\bowtie) and select (σ) operations on them.

$r[\text{uk}(id')]$. Given some $r \in \text{DB}[id]$ we refer to the tuple $(id, r[\text{uk}(id)])$ as its *coordinates* and note that it uniquely identifies that row within the database. Additionally, we refer to the values in a “column” with $\text{rng}(at, \text{DB}) = \{r[at] : r \in \text{DB}[\text{getID}(at, \text{DB})]\}$.

A database’s schema communicates all information about DB except the tables: $\text{Schema}(\text{DB}) = \{(id, R.\text{Ats}) : (id, R) \in \text{DB}\}$. As shorthand, if $\text{scma} = \text{Schema}(\text{DB})$ then $\text{scma}[id] = \text{DB}[id].\text{Ats}$ and $\text{getID}(at, \text{scma}) = \text{getID}(at, \text{DB})$. In our schemes, the client stores $\text{Schema}(\text{DB})$ as part of the key in order to appropriately format data returned by the server. This is a result of our explicit handling of schemas, coordinates and attributes, something which was left implicit in prior work.

SQL OPERATIONS. In our work, we address the secure computation of SQL (equi)joins and (equality) selections. These operations work as follows.

The selection operation is parametrized by a pair of bitstrings (at, x) , takes a relation R_1 with $at \in R_1.\text{Ats}$ as input, and returns $R = \sigma_{(at,x)}(R_1)$ where:

$$R.\text{Ats} = R_1.\text{Ats} \text{ and } R.T = \{r \in R_1.T : r[at] = x\}.$$

In Fig. 3, we provide an example of such a selection on a relation in a database.

The join infix operation is a function parametrized by two equal-length tuples of attributes $\mathbf{t}_1, \mathbf{t}_2$. It takes two relations R_1, R_2 with disjoint attribute sets where $(at_1^i, \dots, at_n^i) = \mathbf{t}_i \subseteq R_i.\text{Ats}$. It returns $R = R_1 \bowtie_{\mathbf{t}_1, \mathbf{t}_2} R_2$ where:

$$R.\text{Ats} = R_1.\text{Ats} \parallel R_2.\text{Ats} \quad \text{and} \quad R.T = \{r_1 \parallel r_2 : r_1 \in R_1.T, r_2 \in R_2.T, \forall i \in [n], r_1[at_1^i] = r_2[at_2^i]\}.$$

In the case of a join on singleton tuples, we abbreviate $\bowtie_{(at),(at')}$ as $\bowtie_{at,at'}$. In Fig. 3, we provide an example such a join. Attribute tuples can be empty in which case it returns the Cartesian product of the input rows. This is also known as the “cross” operation \times .

ADT FOR SQL DATABASES. We say that an ADT SqlIDT is a *SQL data type* if its domain elements $\text{DB} \in \text{SqlIDT}$ are *SQL databases* which take the form $\mathbf{DB} = (\text{DB}, \alpha)$ where DB is as defined in Section 3.1 and $\alpha \in \{0, 1\}^*$ is the auxiliary data. The purpose of α is to allow annotations on DB consistent with real world applications. In this work, we use α to indicate the allowed joins, and $\text{SqlIDT}.\text{Spec}$ always returns either a relation or \perp .

3.2 Constructing StE for SQL Data Types Using Encrypted Indexes

Our end goal is structurally encrypted databases supporting response-hiding SQL queries. We build these by constructing StI schemes for classes of SQL queries, then converting these into StE schemes for SQL data types via a generic transform. We now describe this conversion, then dedicate the remainder of this work to the abovementioned StI schemes.

StE, StI FOR SqlIDT. Intuitively, our StE schemes handle the indexing and storage of SQL data separately. We do the former with an StI scheme and the latter with an RH dictionary encryption scheme. This modularization simplifies pseudocode and reduces the problem of designing secure StE schemes to that of StI schemes.

<p>Alg StE.Enc(K'_i, DB) $(K_d, ED, DS) \leftarrow \text{EncRows}(\text{DB})$ $(K_i, IX) \leftarrow \text{Stl.Enc}(K'_i, DS)$ Return $((K_d, K_i), (ED, IX))$</p> <p>Subroutine EncRows((DB, α)) For $(id, R) \in \text{DB}$ do For $r \in R.T$ do $\mathbf{D}[(id, r[\text{uk}(id)])] \leftarrow r$ $(K_d, ED) \leftarrow \text{Dye}_\pi.\text{Enc}(\mathbf{D})$ For $(id, R) \in \text{DB}$ do For $r \in R.T$ do $\ell \leftarrow (id, r[\text{uk}(id)])$ $\mathbf{T}[\ell] \leftarrow \text{Dye}_\pi.\text{Tok}(K_d, \ell)$ Return $(K_d, ED, (\text{DB}, \alpha, \mathbf{T}))$</p> <p>Alg StE.Tok($(K_d, K_i), q$) $\text{tk} \leftarrow \text{Stl.Tok}(K_i, q)$; Return tk</p>	<p>Alg StE.Eval($\text{tk}, (ED, IX)$) $P \leftarrow \text{Stl.Eval}(\text{tk}, IX)$; Return $\text{EvalRows}(P, ED)$</p> <p>Subroutine EvalRows($((P_1, \dots, P_n), ED)$) For $i \in [n]$ do $C_i \leftarrow \{(c_1, \dots, c_{n'}) : (rt_1, \dots, rt_{n'}) \in P_i, c_j = \text{Dye}_\pi.\text{Eval}(rt_j, ED)\}$ $C \leftarrow (C_1, \dots, C_n)$; Return C</p> <p>Alg StE.Dec($(K_d, K_i), q, C$) Return $\text{Stl.Fin}(K_i, q, \text{DecRows}(K_d, C))$</p> <p>Subroutine DecRows($K_d, (C_1, \dots, C_n)$) For $i \in [n]$ do $M_i \leftarrow \{(m_1, \dots, m_{n'}) : (c_1, \dots, c_{n'}) \in C_i, m_j = \text{Dye}_\pi.\text{Eval}(rt_j, ED)\}$ $M \leftarrow (M_1, \dots, M_n)$; Return M</p> <p>Alg $\mathcal{L}(\text{DB}, (q_1, \dots, q_n))$ $(K, ED, DS) \leftarrow \text{EncRows}(\text{DB})$ using a random function in place of $\text{F.Ev}(K_f, \cdot)$ $lk^i \leftarrow \mathcal{L}^i(DS, (q_1, \dots, q_n))$ Return (lk^i, N, L) where L, N are the max row length and # of rows in DB</p>
--	---

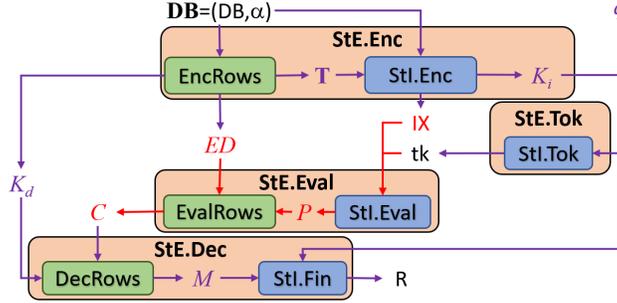


Fig. 4. Algorithms and for structured encryption scheme $\text{StE} = \text{SqlStE}[\text{Stl}, \text{SE}, \text{F}]$ expressed both in pseudocode (top) and diagrammatically (bottom left), and leakage algorithm for StE (bottom right). Here, Dye_π is the RH dictionary encryption scheme Dye_π in Appendix A (which uses SE, F as primitives) and \mathcal{L}^i is Stl 's leakage profile.

More formally, we construct an StE scheme for SQL data type SqlIDT using the transform SqlStE which takes uses an StI scheme for SqlIDT_1 (described below), symmetric encryption scheme SE and function family F . We capture the syntax and pseudocode of StE 's algorithms in Fig. 4. Note that $\text{StE.KS} = \text{Stl.KS}$ and Dye_π is the RH dictionary encryption scheme given in Appendix A which uses SE, F as primitives. It is used in $\text{EncRows}, \text{EvalRows}, \text{DecRows}$, which encrypt, retrieve and decrypt the rows of database DB . We used a specific RH dictionary encryption scheme because pathological alternatives may introduce circular security issues, preventing a more general approach.

We now describe how the algorithms in Stl and $\text{StE} = \text{SqlStE}[\text{Stl}, \text{SE}, \text{F}]$ work. During StE.Enc , algorithm EncRows will store the rows of DB in an encrypted dictionary ED using $\text{Dye}_\pi.\text{Enc}$. It also prepares a token dictionary \mathbf{T} which maps each row coordinate to a token for Dye_π . SQL data type SqlIDT_1 is the same as SqlIDT except that its domain elements now take the form $DS = (\text{DB}, \alpha, \mathbf{T})$ where $(\text{DB}, \alpha) \in \text{SqlIDT}.\text{Dom}$. The output of StE.Enc is ED and the index returned by $\text{Stl.Enc}(DS)$.

StE 's tokens are those generated by Stl . As such, the server's first step in StE.Eval is to run Stl.Eval . We require that Stl.Eval returns a *pointer tuple* $P = (P_1, \dots, P_n)$ which is a tuple of sets of tokens. The tokens in each P_i come from \mathbf{T} and point to rows from the same table. Algorithm EvalRows replaces each token with relevant (encrypted) row from ED and returns *ciphertext tuple* $C = (C_1, \dots, C_n)$, the output of StE.Eval .

During StE.Dec , algorithm DecRows decrypts each ciphertext to get *plaintext tuple* $M = (M_1, \dots, M_n)$. Stl.Fin takes these decrypted rows and performs any final client-side post-processing, returning the final output relation R .

In this work, we will define three different SQL data types, each with its own StI scheme(s). To demonstrate that all of these can be used to construct secure RH StE for their respective data type via SqlStE ,

we demonstrate that the semantic security of StE reduces to that of its primitives. The proof follows from a standard hybrid argument and is given in Appendix B.

Theorem 1. *Let $\text{StE} = \text{SqlStE}[\text{StI}, \text{SE}, \text{F}]$ be a correct StE scheme for SqlDT. Then given algorithms $\mathcal{L}^i, \mathcal{S}^i$ and adversary A we can define \mathcal{L} as in Fig. 4 and construct $\mathcal{S}, A_s, A_f, A_i$ such that:*

$$\text{Adv}_{\text{StE}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A) \leq \text{Adv}_{\text{SE}}^{\text{ind}^{\mathcal{S}}}(A_s) + \text{Adv}_{\text{F}}^{\text{prf}}(A_f) + \text{Adv}_{\text{StI}, \mathcal{L}^i, \mathcal{S}^i}^{\text{ss}}(A_i).$$

4 Partially Precomputed Joins

We demonstrate our framework from Section 3 in action with two SQL data types: JnDT and SjDT. The former only supports non-recursive join queries and is presented for the purpose of introducing partially precomputed (PP) join indexing. The latter allows recursive queries, cluster joins and equality selections, and demonstrates how OPX’s techniques can be modified to use PP joins.

4.1 Indexing of Non-Recursive Joins

JOIN DATA TYPE JnDT. We define JnDT.Dom to contain (DB, α) such that DB is a database and α is the set of join queries supported (i.e. if A is the set of attributes in DB that are not unique key attributes, then $\alpha \subseteq \{(at_1, at_2) \in A \times A : \text{getID}(at_1, \text{DB}) \neq \text{getID}(at_2, \text{DB})\}$). Our goal here is to capture SQL queries of the form “ id_1 **join** id_2 **on** $at_1 = at_2$ ” where $id_1, id_2 \in \text{DB.IDs}$ and $at_i \in \text{DB}[id_i].\text{Ats}$.

We allow queries to be any pair of attributes (i.e. $\text{JnDT.QS} = \{(at_1, at_2) : at_i \in \{0, 1\}^*\}$), but JnDT.Spec only computes the join if $(at_1, at_2) \in \alpha$:

$$\text{JnDT.Spec}((at_1, at_2), (\text{DB}, \alpha)) = \text{DB}[\text{getID}(at_1, \text{DB})] \bowtie_{at_1, at_2} \text{DB}[\text{getID}(at_2, \text{DB})]$$

and returns \perp otherwise. From here on we assume that all queries made are “non-trivial” meaning they return relations with at least one row.

FP INDEXING. FpJn is an StI scheme that “fully precomputes” joins and is modeled after SPX’s handling of “type-2 selections” and OPX’s handling of “leaf joins” [32,34]. The intuition here is that the output relation for each possible join query is precomputed and pointers to the rows therein are stored as an entry in a RR encrypted multimap. FpJn ’s detailed algorithms and leakage profile are given in Fig. 5. Note that $\text{FpJn.KS} = \text{Mme.KS}$ and that each row in the output of a particular join is indexed as a pair of pointers to rows in DB .

Since join queries are handled directly by Mme the leakage and efficiency of FpJn depends entirely on Mme . For the rest of this discussion, we will assume Mme is one of the mainstream multimap encryption schemes (e.g. [12,18,15]) with the “standard” leakage profile consisting the label space size $|\mathbf{M.Lbls}|$, multimap size $\sum_{\ell \in \mathbf{M.Lbls}} |\mathbf{M}[\ell]|$, query pattern (equality pattern of queries ℓ_1, \dots, ℓ_n) and query responses $\mathbf{M}[\ell_1], \dots, \mathbf{M}[\ell_n]$.

Notice that when a query (at_1, at_2) is made in FpJn , the query responses reveal the equality pattern of columns at_1, at_2 for rows that appear in the join output. To illustrate, if the query is made on $\text{DB} = \{(id_1, R_1), (id_2, R_2)\}$ where R_1, R_2 are as depicted in Fig. 3, the server learns that the first two rows of each R_i all have the same value in their at_1, at_2 columns, but won’t reveal anything about the last two rows of each R_i apart from the fact that they are not returned in the join. Note that in the worst case, the join returns all rows from both relations and the search pattern leakage reveals the entire equality pattern of both columns. This leakage is comparable to PRE-based techniques like deterministic encryption or adjustable joins (an observation also made by DPPS [20]). This is significant because, as discussed in Section 1, LAAs are highly effective against PRE and can be applied in this case. Beyond the worst case, FP indexing leaks strictly less than PRE-based solutions but this does not make them immune to LAAs. In particular, we believe that attacks (such as those using ℓ_p -optimization or graph matching [37,8]) can be extended to make use of partial equality patterns and cross column correlations, and be effective against FpJn ’s leakage.

We also note that FpJn achieves lower leakage than the analogous indexing in SPX or OPX because it uses a single multimap. The latter schemes had one encrypted multimap for each attribute (i.e. \mathbf{M}_{at_1} indexes all joins $(at_1, at_2) \in \alpha$) this leaks additional metadata and tells the adversary when two queries join on the same at_1 .

Algs $\boxed{\text{FpJn.Enc}(K'_m, (DB, \alpha, \mathbf{T}))}$, $\boxed{\text{PpJn.Enc}(K'_m, (DB, \alpha, \mathbf{T}))}$	
For $(at_1, at_2) \in \alpha$ do For $\mathbf{r} \in (\text{DB}[\text{getID}(at_i, DB)] \bowtie_{at_1, at_2} \text{DB}[\text{getID}(at_i, DB)]) \cdot \mathbf{T}$ do $\mathbf{rt}_1 \leftarrow \mathbf{T}[(id_1, \mathbf{r}[\text{uk}(id_1)])]$; $\mathbf{rt}_2 \leftarrow \mathbf{T}[(id_2, \mathbf{r}[\text{uk}(id_2)])]$ $\boxed{\mathbf{M}[(at_1, at_2)] \leftarrow^{\cup} (\mathbf{rt}_1, \mathbf{rt}_2)}$; $\boxed{\text{For } i \in \{1, 2\} \text{ do } \mathbf{M}[(at_1, at_2, i)] \leftarrow^{\cup} \mathbf{rt}_i}$ $(K_m, \text{IX}) \leftarrow^s \text{Mme.Enc}(K'_m, \mathbf{M})$; Return $((K_m, \text{Schema}(DB)), \text{IX})$	
Alg $\text{FpJn.Tok}((K_m, \text{scma}), q)$ Return $\text{Mme.Tok}(K_m, q)$ Alg $\text{FpJn.Eval}(\text{tk}, \text{IX})$ Return $(\text{Mme.Eval}(\text{tk}, \text{IX}))$ Alg $\text{FpJn.Fin}((K_m, \text{scma}), q, (M))$ $\mathbf{at}_1 \leftarrow \text{scma}[\text{getID}(\mathbf{at}_1, \text{scma})]$ $\mathbf{at}_2 \leftarrow \text{scma}[\text{getID}(\mathbf{at}_2, \text{scma})]$ $\mathbf{R} \leftarrow \text{NewRltn}(\mathbf{at}_1 \parallel \mathbf{at}_2)$ $\mathbf{R} \cdot \mathbf{T} \leftarrow \{\mathbf{r}_1 \parallel \mathbf{r}_2 : (\mathbf{r}_1, \mathbf{r}_2) \in M\}$ Return \mathbf{R}	Alg $\text{PpJn.Tok}((K_m, \text{scma}), (at_1, at_2))$ $\mathbf{mt}_1 \leftarrow^s \text{Mme.Tok}(K_m, (at_1, at_2, 1))$ $\mathbf{mt}_2 \leftarrow^s \text{Mme.Tok}(K_m, (at_1, at_2, 2))$ Return $(\mathbf{mt}_1, \mathbf{mt}_2)$ Alg $\text{PpJn.Eval}((\text{tk}_1, \text{tk}_2), \text{IX})$ Return $(\text{Mme.Eval}(\text{tk}_1, \text{IX}), \text{Mme.Eval}(\text{tk}_2, \text{IX}))$ Alg $\text{PpJn.Fin}((K_m, \text{scma}), (at_1, at_2), (M_1, M_2))$ For $i = 1, 2$ do $\mathbf{R}_i \leftarrow \text{NewRltn}(\text{scma}[\text{getID}(at_i, \text{scma})])$ $\mathbf{R}_i \cdot \mathbf{T} \leftarrow \{\mathbf{r} : (\mathbf{r}) \in M_i\}$ Return $\mathbf{R}_1 \bowtie_{at_1, at_2} \mathbf{R}_2$
Alg $\mathcal{L}^f(\text{DS}, (q_1, \dots, q_n))$ Construct \mathbf{M} as in $\text{FpJn.Enc}(\cdot, \text{DS})$ Return $\mathcal{L}^m(\mathbf{M}, (q_1, \dots, q_n))$	Alg $\mathcal{L}^p(\text{DS}, ((at_1, at'_1), \dots, (at_n, at'_n)))$ Construct \mathbf{M} as in $\text{PpJn.Enc}(\cdot, \text{DS})$ For $i \in [n]$ do $q_{2i-1} \leftarrow (at_i, at'_i, 1)$; $q_{2i} \leftarrow (at_i, at'_i, 2)$ Return $\mathcal{L}^m(\mathbf{M}, (q_1, \dots, q_{2n}))$

Fig. 5. Algorithms of StI schemes FpJn, PpJn (top) and their leakage algorithms (bottom) where Mme is a RR multimap encryption scheme. Note that in the encryption algorithm, boxed code belongs only to the respective algorithm.

PP INDEXING. We introduce a new StI scheme PpJn which performs “partially precomputed” indexing, whose algorithms are also depicted in Fig. 5. PpJn.Enc proceeds in the same way as FpJn.Enc but we store the rows from each input relation separately. In other words, if $\mathbf{M}_f, \mathbf{M}_p$ are the multimaps constructed in the respective setup algorithms, then $\mathbf{M}_p[(at_1, at_2, i)] = \{\mathbf{rt}_i : (\mathbf{rt}_1, \mathbf{rt}_2) \in \mathbf{M}_f[(at_1, at_2)]\}$ for $i = 1, 2$ and $(at_1, at_2) \in \alpha$. Notice that this means the client needs to reassemble the output relation from the two sets of rows in StI.Fin. We recommend that the client do so by sorting then joining the columns, avoiding the quadratic time nested loop join where rows are compared pairwise.

This small change in indexing technique has substantial impact on bandwidth and security. In the worst case, the number of rows sent with FP is quadratic while PP’s is linear. This bandwidth reduction occurs because two sets of rows are sent instead of their cross product. Notice that modulo some metadata information (i.e. the multimap sizes), the PP leakage can be derived from the FP leakage meaning that PP indexing is no worse than FP indexing. In fact, if more than one row is returned to any query PP leakage is strictly lower. To illustrate, when join query (at_1, at_2) is made to the aforementioned database in Fig. 3, the adversary sees that the first three rows of both tables were returned and can infer that each row has at least one matching value in the other column – nothing specific about their equality patterns.

In summary, PpJn is the superior indexing choice for JnDT because its leakage is strictly lower, bandwidth is no worse and efficiency is comparable.

SEMANTIC SECURITY. The security of FpJn, PpJn reduce to that of Mme. The proof follows directly from the definition of Mme’s semantic security and is given in Appendix C.

Theorem 2. Let \mathcal{L}, \mathcal{S} be the leakage algorithm and simulator for Mme. Let $\mathcal{L}^f, \mathcal{L}^p$ be the leakage algorithms given in Fig. 5. Then, given adversary A there exists adversary A_m and simulator \mathcal{S}^p such that:

$$\mathbf{Adv}_{\text{FpJn}, \mathcal{L}^f, \mathcal{S}}^{\text{SS}}(A) \leq \mathbf{Adv}_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{SS}}(A) \text{ and } \mathbf{Adv}_{\text{PpJn}, \mathcal{L}^p, \mathcal{S}^p}^{\text{SS}}(A) \leq \mathbf{Adv}_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{SS}}(A_m).$$

4.2 PP indexing for recursive queries

SjDT. We expand the query support of JnDT to include equality selections, cluster joins (joins on more than one attribute) and recursively defined queries. The resultant query class is similar to the SPJ algebra defined by CM [14] except for the omission of the projection operation which we note can be handled as a post-processing step requiring no cryptographic techniques.

We capture this via the SQL data type SjDT. Its domain is unchanged from JnDT.Dom except that α allows tuple pairs in addition to attribute pairs. Below we describe the forms, evaluation and SQL equivalent of $q \in \text{SjDT.QS}$. (Note that the $\mathbf{r}, \mathbf{s}, \mathbf{j}$ flags are included in SjDT queries for domain separation.) These are defined recursively so q_i, \mathbf{q}_i are themselves queries of the respective type.

Query Type	SjDT query q	SjDT.Spec(q, \mathbf{DB})	SQL query \mathbf{q}
Relation retrieval	(\mathbf{r}, id)	$\text{DB}[id]$ where $\mathbf{DB} = (\mathbf{DB}, \alpha)$	select * from id
(Equality) selections	(\mathbf{s}, at, x, q_1)	$\sigma_{(at,x)}(\text{SjDT.Spec}(q_1, \mathbf{DB}))$	select * from \mathbf{q}_1 where $at = c$
(Equi)joins	$(\mathbf{j}, \mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$ where $(\mathbf{t}_1, \mathbf{t}_2) \in \alpha$	$(\text{SjDT.Spec}(q_1, \mathbf{DB})) \bowtie_{\mathbf{t}_1, \mathbf{t}_2} (\text{SjDT.Spec}(q_2, \mathbf{DB}))$	select * from \mathbf{q}_1 join \mathbf{q}_2 on $\mathbf{t}_1 = \mathbf{t}_2$

We say that queries of the form (\mathbf{r}, id) , $(\mathbf{s}, at, x, (\mathbf{r}, id))$ or $(\mathbf{j}, \mathbf{t}_1, \mathbf{t}_2, (\mathbf{r}, id_1), (\mathbf{r}, id_2))$ are *non-recursive* and all others are *recursive*. We require that all attributes in each \mathbf{t}_i come from the same relation in \mathbf{DB} (i.e. $\mathbf{t}_i \subseteq \text{DB}[id].\text{Ats}$ for some $id \in \text{DB.IDs}$). While allowing cluster joins may lead to an exponential-size index, a judicious database administrator would not allow this – cluster joins are rarely used and usually known in advance.

HASHSET FILTERING. To minimize the leakage of recursive queries in our StI schemes we employ the filtering hashset technique introduced in OXT [13]. We now review this technique and establish some notation for it.

This filtering hashset is a set denoted HS containing outputs of a function family \mathbf{F} where $\mathbf{F.KS} = \{0, 1\}^{\text{F.ol}}$. In our algorithms, the hashset will be used to associate predicate bitstrings with a row tokens (from \mathbf{T}). Later, given a predicate p 's key $K = \mathbf{F.Ev}(K_f, p)$ we can filter a set of row tokens, retaining only those which satisfy the predicate. We formalize this via the following algorithms:

<p>Alg HsEnc(K_f, SET)</p> <p>For $(p, \text{rt}) \in \text{SET}$ do $\text{HS} \stackrel{\cup}{\leftarrow} \mathbf{F.Ev}(\mathbf{F.Ev}(K_f, p), \text{rt})$ Return HS</p>	<p>Alg HsFilter($K, (P_1, \dots, P_n), \text{HS}$)</p> <p>For $i \in [n]$ do For $\text{rt} \in P_i$ if $\mathbf{F.Ev}(K, \text{rt}) \in \text{HS}$ then $S \stackrel{\cup}{\leftarrow} \text{rt}$ If $S \neq \emptyset$ then $P_i \leftarrow S$ Return (P_1, \dots, P_n)</p>
---	---

For notational convenience in our pseudocodes, HsFilter takes as input a tuple set $P = (P_1, \dots, P_n)$. It then attempts to filter each P_i and retain only the tuples where at least one rt satisfies the predicate. However, if no such tuple exists, it does not perform the filtering at all.

PP INDEXING FOR SjDT. We are now ready to extend the PP indexing technique introduced in Section 4 to construct StI for SjDT. On a high level, we do so by using an inverted index (similar to those used for SSE) to handle selections and a filtering hashset to handle recursively defined queries. The result is StI scheme PpSj whose algorithms are depicted in Fig. 6.

Now we provide some intuition for PpSj's algorithms. The scheme has two server-side data structures: an encrypted multimap and a hashset. The multimap is used to index non-recursive queries by mapping a query-derived label to the relevant rows in the database. For example, the label for relation retrieval query (\mathbf{r}, id) is the query itself and its values are row tokens associated to rows in $\text{DB}[id]$ (i.e. $\{(\mathbf{T}[(id, \mathbf{r}[\text{uk}(id)]))]) : \mathbf{r} \in \text{DB}[id].\mathbf{T}\}$). Note that the latter are singleton tuples because we required that pointer tuples be made out of tuples of tokens. The hashset is used to filter the sets in a pointer tuple during a recursive query. For example, when processing the query $(\mathbf{j}, \mathbf{t}_1, \mathbf{t}_2, (\mathbf{s}, at, x, (\mathbf{r}, id_1)), (\mathbf{r}, id_2))$ (a select followed by a join), the server would use the multimap to retrieve row tokens for each of the non-recursive subqueries (i.e. $(\mathbf{s}, at, x, (\mathbf{r}, id_1))$ and (\mathbf{r}, id_2)). The token would also include two keys which can be used with HsFilter which tests if the rows being pointed to (in $\text{DB}[id_1]$ or $\text{DB}[id_2]$) are in $\text{DB}[id_1] \bowtie_{\mathbf{t}_1, \mathbf{t}_2} \text{DB}[id_2]$.

<p>Alg PpSj.Enc($K'_m, (DB, \alpha, \mathbf{T})$)</p> <p>For all $(id, R) \in DB$ and $r \in R.T$ do $rt \leftarrow \mathbf{T}[(id, r[uk(id)])]$; $\mathbf{M}[(r, id)] \stackrel{\cup}{\leftarrow} (rt)$ For $at \in R.Ats$ where $at \neq uk(id)$ do $\mathbf{M}[(s, at, r[at])] \stackrel{\cup}{\leftarrow} (rt)$; $\mathbf{SET} \stackrel{\cup}{\leftarrow} ((s, at, r[at]), rt)$</p> <p>For $(t_1, t_2) \in \alpha$ do $id_1 \leftarrow \text{getID}(t_1)$; $id_2 \leftarrow \text{getID}(t_2)$ For $r \in (DB[id_1] \bowtie_{t_1, t_2} DB[id_2]).T$ do For $i = 1, 2$ do $rt \leftarrow \mathbf{T}[(id_i, r[uk(id_i)])]$; $\mathbf{M}[(j, t_1, t_2, i)] \stackrel{\cup}{\leftarrow} (rt)$ $\mathbf{SET} \stackrel{\cup}{\leftarrow} ((j, t_1, t_2, i), rt)$</p> <p>$(K_m, EM) \leftarrow^s \text{Mme.Enc}(K'_m, \mathbf{M})$; $K_f \leftarrow^s F.KS$; $HS \leftarrow \text{HsEnc}(K_f, \mathbf{SET})$ Return $((\text{Schema}(DB), K_m, K_f), (EM, HS))$</p> <p>Alg PpSj.Tok($(scma, K_m, K_f), q$)</p> <p>If $q = (r, id)$ then return $(r, \text{Mme.Tok}(K_m, (r, id)))$ Else if $q = (s, at, x, (r, id))$ then return $(r, \text{Mme.Tok}(K_m, (s, at, x)))$ Else if $q = (s, at, x, q_1)$ then $tk_1 \leftarrow^s \text{PpSj.Tok}((scma, K_m, K_f), q_1)$ Return $(s, F.Ev(K_f, (s, at, x)), tk_1)$ Else if $q = (j, t_1, t_2, q_1, q_2)$ then For $i = 1, 2$ do If $q_i = (r, id_i)$ then $tk_i \leftarrow^s (r, \text{Mme.Tok}(K_m, (j, t_1, t_2, i)))$ Else $tk' \leftarrow^s \text{PpSj.Tok}((scma, K_m, K_f), q_i)$ $tk_i \leftarrow (s, F.Ev(K_f, (j, t_1, t_2, i), tk'))$ Return (j, tk_1, tk_2)</p>	<p>Alg PpSj.Eval(tk, IX)</p> <p>$(EM, HS) \leftarrow IX$ If $tk = (r, tk_1)$ then Return $(\text{Mme.Eval}(tk, IX))$ Else If $tk = (s, K, tk_1)$ then $P \leftarrow \text{PpSj.Eval}(tk_1, IX)$ Return $\text{HsFilter}(K, P, HS)$ Else if $tk = (j, tk_1, tk_2)$ For $i = 1, 2$ do $P^i \leftarrow \text{PpSj.Eval}(tk_i, IX)$ Return $P^1 \parallel P^2$</p> <p>Alg PpSj.Fin($K_i, q, (M_1)$)</p> <p>$(scma, K_m, K_f) \leftarrow K_i$ If $q = (r, id)$ then $R \leftarrow \text{NewRltn}(scma[id])$ $R.T \leftarrow r : (r) \in M_1$; Return R Else if $q = (s, at, x, q_1)$ then Return $\text{PpSj.Fin}(K_i, q_1, (M_1))$ Else if $q = (j, t_1, t_2, q_1, q_2)$ then Define $M^1, M^2 : M^1 \parallel M^2 = M_1$, M^1 has as many M_i as q_i has subqueries of the form (r, id) For $i = 1, 2$ do $R_i \leftarrow \text{PpSj.Fin}(K_i, q_i, M^i)$ Return $R_1 \bowtie_{t_1, t_2} R_2$</p>
--	--

Fig. 6. Algorithms for PpSj the StI scheme for SjdT using PP indexing.

FP INDEXING FOR SjdT. We analogously extend FpJn introduced in Section 4 to construct FpSj, an StI for SjdT. Just like with PpSj, non-recursive queries will be added to the encrypted multimap that is used to index the non-recursive joins while all recursive queries are filtered using the hashset. The only subtlety in this extension is the handling of “internal joins” which are queries of the form $q = (j, t_1, t_2, (r, id), q_1)$ (or $q = (j, t_1, t_2, q_1, (r, id))$) because we want to limit the row tokens leaked from id to those who join with some row returned by q_1 . Similar to OPX, we construct an index where each row token returned in the subquery will “point to” the tokens of the rows joined to in $DB[id]$. As alluded to in Section 3.2, this self-referential indexing (where Mme tokens are stored in \mathbf{M}) may introduce circular security issues if pathological Mme primitives are used. We avoid this by indexing internal joins with a specific, non-pathological primitive (as was done in OPX). To avoid the increased leakage and complexity of an additional data structure, we will assume that Mme is the $\text{Mme}_{\pi}^{\text{rr}}$ scheme recounted in Appendix A and co-locate this index with the one used for non-recursive queries. Notice that this subtlety does not come up in PpSj because we do not reveal join equality patterns so all recursive joins can be handled similarly.

The StE scheme $\text{StE} = \text{SqlStE}[\text{FpSj}, \text{SE}, F]$ is essentially the same as OPX with minor improvements in leakage (analogous to those described in our discussion of FpJn in Section 4) and a slightly revised approach to “internal joins”. For this reason, we defer a complete description of FpSj to Appendix E.

PpSj LEAKAGE PROFILE. While a pseudocode description of PpSj’s leakage profile may seem convoluted, we believe the intuition behind it enables helpful comparisons with FpSj and OPX [34]. As such, we aim to give some intuition by describing the components of PpSj’s leakage profile via a running example, deferring a full description of PpSj’s leakage algorithm and the associated security proof to Appendix D. Below, we assume that MME primitives have the “standard” leakage profile (as described in Appendix A).

Our example database contains R_1, R_2 from Fig. 3. If no queries are made, the server-side data structures reveal only *metadata leakage*. This includes the number of values in the multimap, the maximum-length of

a value in the multimap and the number of F outputs in the hashset. The leakage of FpSj is comparable but on OPX it is higher because different data structures are used to index different SQL operations.

We will refer to all other forms of leakage as “query dependent leakage”. This is where PP indexing has substantial savings over FP and OPX.

Now lets assume the client makes the following queries: $q_1 = (\mathbf{s}, at_3, \mathbf{CS}, (\mathbf{r}, id_2))$, $q_2 = (\mathbf{s}, at_2, \mathbf{Eve}, (\mathbf{r}, id_2))$, $q_3 = (\mathbf{r}, id_1)$ and $q_4 = (\mathbf{j}, at_1, at_2, (\mathbf{r}, id_1), (\mathbf{s}, at_3, \mathbf{CS}, (\mathbf{r}, id_2)))$. The server will receive four tokens, where tk_1, tk_2, tk_3 are such that $tk_i = (\mathbf{r}, mt_i)$ and $tk_4 = (\mathbf{j}, (\mathbf{r}, mt_4), (\mathbf{j}, K, mt_5))$. Here, each mt_i is a token for Mme while K is a hashset key. Just from inspecting these, the adversary learns the *recursion structure* of the queries. Specifically, he learns that the first three queries were non-recursive while q_4 was a join followed by a select. This leakage is slightly lower in FpSj, PpSj compared to OPX because the adversary cannot differentiate between non-recursive selections and relation retrievals.

The Mme tokens leak the *multimap query pattern* and *multimap responses*. The former reveals whenever the associated query or subquery is repeated. In our case, the adversary learns that mt_1, mt_5 are associated to the same query. Note that this does not extend to mt_3, mt_4 because the latter is in a join. From the multimap query responses he “sees” the row tokens that are returned by each $\text{Mme.Eval}(mt_i, \text{EM})$. This reveals the equality pattern of the rows returned by each associated query/subquery. For example, this reveals that q_1, q_2 both return two rows, one of which is shared. On join queries, we enjoy similar leakage savings as described in the non-recursive case. For example, tk_4 will reveal that three rows are returned from the left relation (i.e. id_1) but doesn’t say anything about whether they are in the final output relation or how they “match up” with rows from the right relation. In FpSj and OPX, both of the above are revealed.

Finally, the hashset keys reveal the *hashset key query pattern* and *hashset filtering results*. The former reveals when the exact same predicate is repeated and is detectable because the keys would be the same. The latter is because the adversary is free to apply hashset keys (in the tokens) to filter all the row tokens he can retrieve from EM thereby learning the *hashset filtering results*. This means that using K he can learn that one row returned by q_2 satisfies the predicate associated to K even though it is not in the output of q_4 . Similarly, he learns that two rows returned by q_1 satisfies the predicate but only one is returned by q_4 . Using FpSj the adversary would additionally learn which row returned in q_3 is “paired up” with this row in the q_4 output.

LEAKAGE COMPARISON. From the above discussion, one might expect decreasing query-dependent leakage from PpSj to FpSj to OPX. While the leakage for FpSj can always be derived from OPX, the comparison of PpSj to FpSj is not as straightforward because they sometimes do not return the same rows when recursive queries are made (which we discuss in more detail below).

However, when restricted to non-recursive queries, PpSj’s query-dependent leakage is strictly superior for the same reasons that PpJn was superior in Section 4. Extending this, we can upper bound the leakage lk_p of PpSj on queries q_1, \dots, q_n with its leakage lk'_p on the minimal set of non-recursive queries q'_1, \dots, q'_m with which the server can still deduce the pointer tuples it should return on q_1, \dots, q_n . Doing the same for FpSj, we have $lk_f \leq lk'_f$ as well. Then, via the above observation about non-recursive queries we have $lk'_p \leq lk'_f$, with the inequality being strict if at least one join query with at least two rows is made. Our being able to *bound PpSj’s query-dependent leakage lower than FpSj’s* gives credence to the intuition that PpSj is the more secure variant in practice.

EFFICIENCY DRAWBACK OF PpSj. Comparing the bandwidth of PpSj, FpSj is also not clear cut: On non-recursive queries, PpSj will perform equal or better than FpSj but on recursive queries the converse is sometimes true.

Consider the query $q = (\mathbf{s}, at_3, \mathbf{CS}, (\mathbf{j}, at_1, at_2, (\mathbf{r}, id_1), (\mathbf{r}, id_2)))$ in our toy example. With FpSj, the server returns pointers to the two rows that feature in the output relation (i.e. those with coordinates $(id_1, cc), (id_2, 33)$) but PpSj returns four (i.e. with coordinates $(id_1, aa), (id_1, bb), (id_1, cc), (id_2, 33)$) because without the equality pattern over the join columns and it cannot filter out the first and second rows of R_1 .

More generally, this overhead may occur anytime that a recursive query (involving at least one join) is made and grows with query complexity. Depending on the data and query workload, this overhead ranges from negligible to quite substantial, something we explore further in Section 6.

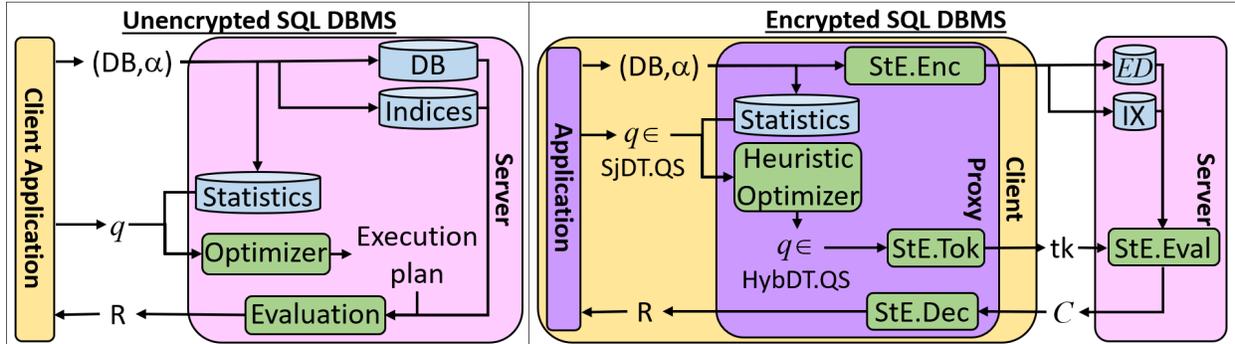


Fig. 7. Data/ query processing in unencrypted SQL databases (left) and the analogous processes using **SqlStE** with hybrid indexing (right).

5 Hybrid indexing

We showed that the choice between FP and PP indexing depends heavily on query load. This motivates our hybrid StI scheme that postpones this decision till query time. We first cover the technical details of supporting both indexing techniques, then give a heuristic for the client to choose between them.

HYBRID DATA PROCESSING. We give a new ADT where each join in a query is annotated with the desired indexing technique, **HybDT**. This ADT is equivalent to **SjDT** except that its join queries take the form $(op, \mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$ where $op \in \{\mathbf{fp}, \mathbf{pp}\}$. When evaluating **HybDT.Spec**, these are both functionally equivalent to the analogous **SjDT** join query's $(j, \mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$.

The hybrid system we envision makes the same assumption as in (unencrypted) SQL DBMSes – that client queries are unoptimized and have no canonical form – and therefore mirrors its data flow as depicted in Fig. 7. It also borrows its architecture (i.e. use of a client-side proxy) from existing encrypted SQL solutions [39,43]. The client's **SjDT** query is annotated using a heuristic optimizer to get a **HybDT** query. This latter query is then tokenized, evaluated and decrypted using hybrid indexing scheme **HybStI** in $\text{StE} = \mathbf{SqlStE}[\text{HybStI}, \text{SE}, \text{F}]$.

As best we know, no existing work has looked into query optimization in StE schemes. We believe this area to be of independent interest because unlike encrypted systems where optimization runs on the server (with full access to the data) and is solely interested in maximizing efficiency, optimization in encrypted SQL DBMSes should be done (at least partially) by the proxy with only precomputed statistics about the data and may additionally seek to minimize leakage. We initiate this study with our heuristic below.

HybStI DETAILS. This StI merges **FpSj**, **PpSj** by essentially storing both kinds of indexes on the server. More specifically, **HybStI.Enc** will merge the multimaps and hashsets generated by **PpSj**, **FpSj** (avoiding repetition where possible) so that it can take join tokens of either form. When a **HybDT** join query is made, the client indicates which index to use in its query with op .

We believe the intuition for how **HybStI** works is straightforward, so we defer a full discussion on its pseudocode, leakage algorithm and security to Appendix F. The only subtlety comes when a query contains both FP and PP joins. Notice that pointer tuples in this case will contain more than one P_i (unlike **FpSj**) and the tuples in at least one P_i will contain more than one rt (unlike **PpSj**). As such, after the client performs the PP joins in **HybStI.Fin** some column reordering may be necessary.

HybStI LEAKAGE. We will describe **HybStI**'s leakage profile in comparison to that of **PpSj** and **FpSj**. The metadata leakage is comparable, with each size (multimap or hashset) being the sum of respective **FpSj** and **PpSj** sizes. The recursion structure leakage is technically higher but only because we leak the join annotations that weren't present in the other two schemes.

For the same reason that **PpSj** and **FpSj**'s query-dependent leakages were not directly comparable, they also cannot be compared with that of **HybStI**. However, like we did in Section 4.2, we can upper bound **HybStI**'s query-dependent leakage on $q_1, \dots, q_n \in \text{HybDT.QS}$ with that of q'_1, \dots, q'_m , the minimal set of non-recursive queries in **HybDT.QS** (with consistent join annotation) with which the server can still compute its output on q_1, \dots, q_n . This leakage is no better than the analogous bound in **PpSj** and no worse than that

of FpSj , this confirms the intuition that hybrid indexing achieves an *intermediate level of query-dependent leakage* compared to solely using FP or PP indexing.

LEAKAGE-AWARE QUERY PLANNING. The join annotation selected by our query planning heuristic will minimize leakage without exceeding a predetermined bandwidth limit. More specifically, suppose the user supplies a query $q \in \text{SjDT}$ with J joins and a bandwidth limit L indicating the maximum number of rows from ED that can be returned in the ciphertext tuple. We estimate the bandwidth of all possible HybDT queries, then select an annotation by:

1. Eliminating options which exceed L rows. If none remain, return \perp .
2. Maximize number of PP joins
3. If multiple choices remain, minimize bandwidth.

We argue that our setup is realistic because (1) we expect the J joins made in a query to be modest enough for the client to evaluate all 2^J HybDT queries, (2) bandwidth measurement can be reduced to the number of rows from ED sent as they are padded to the same length, and (3) it is common for SQL applications to limit bandwidth to prevent the client from maxing out its memory.

<p>Alg EvalBW(q)</p> <p>If $q = (\mathbf{r}, id)$ then $\mathbf{B}[(id)] \leftarrow \mathcal{N}(id)$</p> <p>Else if $q = (\mathbf{s}, at, x, q_1)$ then</p> <p style="padding-left: 20px;">$\mathbf{B} \leftarrow \text{EvalBW}(q_1)$; $id \leftarrow \text{getID}(at, \text{scma})$</p> <p style="padding-left: 20px;">$\mathbf{B}[(id)] \leftarrow \mathbf{B}[(id)] \cdot \mathcal{H}_{at}(x)$</p> <p>Else if $q = (op, \mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$ then</p> <p style="padding-left: 20px;">$\mathbf{B} \leftarrow \text{EvalBW}(q_1) \cup \text{EvalBW}(q_2)$</p> <p style="padding-left: 20px;">For $i = 1, 2$</p> <p style="padding-left: 40px;">Define $\mathbf{i}_i : \text{getID}(\mathbf{t}_i, \text{scma}) \in \mathbf{i}_i \in \mathbf{B}.\text{LbIs}$</p> <p style="padding-left: 20px;">If $op = \text{fp}$ then</p> <p style="padding-left: 40px;">$N \leftarrow \frac{\mathcal{F}(\mathbf{t}_1, \mathbf{t}_2) \cdot \mathbf{B}[\mathbf{i}_1] \cdot \mathbf{B}[\mathbf{i}_2]}{\mathcal{N}(\text{getID}(\mathbf{t}_1, \text{scma})) \cdot \mathcal{N}(\text{getID}(\mathbf{t}_2, \text{scma}))}$</p> <p style="padding-left: 40px;">$\mathbf{B}[\mathbf{i}_1 \parallel \mathbf{i}_2] \leftarrow 2 \cdot N$</p> <p style="padding-left: 40px;">$\mathbf{B}[\mathbf{i}_1] \leftarrow \perp$; $\mathbf{B}[\mathbf{i}_2] \leftarrow \perp$</p> <p style="padding-left: 20px;">Else if $op = \text{pp}$ then</p> <p style="padding-left: 40px;">For $i = 1, 2$ do</p> <p style="padding-left: 60px;">$\mathbf{B}[\mathbf{i}_i] \leftarrow \mathcal{P}_i(\mathbf{t}_1, \mathbf{t}_2) \cdot \frac{\mathbf{B}[\mathbf{i}_i]}{\mathcal{N}(\text{getID}(\mathbf{t}_1, \text{scma}))}$</p> <p>Return \mathbf{B}</p>	<p>Schema scma</p> <p>Return Schema(DB)</p> <p>Table size $\mathcal{N}(id)$</p> <p>Return $\text{DB}[id].\mathbf{T}$</p> <p>Freq. histogram $\mathcal{H}_{at}(x)$</p> <p>$\mathbf{R} \leftarrow \sigma_{(at, x)}(\text{DB}[\text{getID}(at, \text{scma})])$</p> <p>Return $\frac{ \mathbf{R}.\mathbf{T} }{\mathcal{N}(id)}$</p> <p>FP join size $\mathcal{F}(\mathbf{t}_1, \mathbf{t}_2)$</p> <p>For $i = 1, 2$ do $id_i \leftarrow \text{getID}(\mathbf{t}_i, \text{scma})$</p> <p>$\mathbf{R} \leftarrow \text{DB}[id_1] \bowtie_{\mathbf{t}_1, \mathbf{t}_2} \text{DB}[id_2]$</p> <p>Return $\mathbf{R}.\mathbf{T}$</p> <p>PP join sizes $\mathcal{P}_j(\mathbf{t}_1, \mathbf{t}_2)$</p> <p>For $i = 1, 2$ do $id_i \leftarrow \text{getID}(\mathbf{t}_i, \text{scma})$</p> <p>$\mathbf{R} \leftarrow \text{DB}[id_1] \bowtie_{\mathbf{t}_1, \mathbf{t}_2} \text{DB}[id_2]$</p> <p>Return $\{\mathbf{r}[\text{uk}(id_j)] : \mathbf{r} \in \mathbf{R}.\mathbf{T}\}$</p>
---	---

Fig. 8. EvalBW algorithm (left) defined in terms of precomputed statistics (right) stored on the client. Our heuristic assumes that q incurs bandwidth $\sum_{\mathbf{i} \in \mathbf{B}.\text{LbIs}} \mathbf{B}[\mathbf{i}]$ where $\mathbf{B} = \text{EvalBW}(q)$.

To complete this setup, we need a way for the client to estimate the bandwidth of a query with only partial information about \mathbf{DB} computed during setup. These precomputed statistics are listed on the right of Fig. 8 and the bandwidth estimation algorithm is EvalBW. Intuitively, EvalBW will populate a dictionary \mathbf{B} with entries $\mathbf{B}[\mathbf{i}]$ representing the bandwidth for the ciphertext set containing rows from all $\text{DB}[id]$ where $id \in \mathbf{i}$. We estimate that a query $q \in \text{HybDT}.\text{QS}$ incurs bandwidth $\sum_{\mathbf{i} \in \mathbf{B}.\text{LbIs}} \mathbf{B}[\mathbf{i}]$ where $\mathbf{B} = \text{EvalBW}(q)$. We will next explore the tradeoffs involved in storing these statistics.

MEMORY TRADEOFFS. Notice that the client storage required for the precomputed statistics (as given in Fig. 8) increases with number of joins (i.e. $|\alpha|$) and size of histograms (i.e. $|\text{rng}(at, \text{DB})|$ for each at). In practice, data may be too complex or client devices may be too memory strapped (e.g. mobile devices) to store this in full. We describe two tradeoffs application designers can explore to better fit their system requirements.

When it is unfeasible to store full frequency histograms for some at , the client can partition $\text{rng}(at, \text{DB})$ into ranges and store this bucketed frequency histogram. EvalBW will approximate $\mathcal{H}_{at}(x)$ by assuming that values within a bucket are uniformly distributed. This approach is used in SQL server and the literature recommends 200 equiDepth (as opposed to equiWidth) buckets [41,10]. In the extreme case, the client uses a single bucket and needs only store $|\text{rng}(at, \text{DB})|$ and uses $\mathcal{H}_{at}(x) \approx \frac{1}{|\text{rng}(at, \text{DB})|}$. Note that this only works when the elements of $\text{rng}(at, \text{DB})$ can be closely approximated and ordered. For example, this may not work

Query Type	Indexed Data	Chicago data set					Sakila data set				
		JnDT		SjDT		HybDT	JnDT		SjDT		HybDT
		FpJn	PpJn	FpSj	PpSj	HybStl	FpJn	PpJn	FpSj	PpSj	HybStl
Non-recursive join	MM lbls	1249	2498	1249	2498	3747	631	1262	631	1262	1893
	MM vals	1.495e10	2.796e7	1.495e10	2.796e7	1.498e10	5.103e8	2.201e6	5.103e8	2.201e6	5.125e8
Recursive join	MM lbls	–	–	2.796e7	–	2.796e7	–	–	2.202e6	–	2.202e6
	MM vals	–	–	1.496e10	–	1.496e10	–	–	5.107e8	–	5.107e8
	HS vals	–	–	7.477e9	2.796e7	7.505e9	–	–	2.552e8	2.201e6	2.574e8
Relation retrieval	MM lbls	–	–	15	15	15	–	–	15	15	15
	MM vals	–	–	4.010e5	4.010e5	4.010e5	–	–	4.409e4	4.409e4	4.409e4
Select	MM lbls	–	–	1.082e6	1.082e6	1.082e6	–	–	1.190e5	1.190e5	1.190e5
	MM vals	–	–	5.749e6	5.749e6	5.749e6	–	–	2.945e5	2.945e5	2.945e5
	HS vals	–	–	5.749e6	5.749e6	5.749e6	–	–	2.945e5	2.945e5	2.945e5
Total	MM lbls	1249	2498	2.905e7	1.085e6	2.905e7	631	1262	2.321e6	1.203e5	2.322e6
	MM vals	1.495e10	2.796e7	2.991e10	3.412e7	2.994e10	5.103e8	2.201e6	1.021e9	2.540e6	1.023e9
	HS vals	–	–	7.483e9	3.371e7	7.511e9	–	–	2.555e8	2.496e6	2.577e8

Fig. 9. Simulated server storage for each data set using each of our schemes in terms of multimap (MM) labels/values and hashset (HS) values broken down by the query type being indexed (i.e. relation retrievals, non-recursive/recursive joins, or selections).

with a “name” column because the names in $\text{rng}(at, \text{DB})$ are not dense in any easily enumerated set. In general, bucketing sacrifices the accuracy of EvalBW to reduce client memory. We study this tradeoff more in Section 6.

Above, we assumed the client would pre-compute and store the join sizes. When this is infeasible due to memory constraints, the client can alternatively compute join sizes using table sizes and the $\mathcal{H}_{at}(x)$ during EvalBW whenever $\text{rng}(at)$ is enumerable. Notice that we can express each co-occurrence frequency as a function of the relevant occurrence frequencies. With a single attribute join, let $X = \text{rng}(at_1, \text{DB}) \cap \text{rng}(at_2, \text{DB})$, $N_i = \mathcal{N}(\text{getID}(at_i, \text{scma}))$ then

$$\mathcal{F}(at_1, at_2) = N_1 \cdot N_2 \cdot \sum_{x \in X} \mathcal{H}_{at_1}(x) \cdot \mathcal{H}_{at_2}(x) \quad \text{and} \quad \mathcal{P}_j(at_1, at_2) = N_j \cdot \sum_{x \in X} \mathcal{H}_{at_j}(x).$$

We can extend this to a cluster join $(\mathbf{t}_1, \mathbf{t}_2)$ where $\mathbf{t}_j = (at_1^j, \dots, at_n^j)$. We substitute the above histogram values for $\mathcal{H}_{\mathbf{t}_j}(x_1, \dots, x_n)$ and take the sum over all (x_1, \dots, x_n) where $x_i \in \text{rng}(at_i^1, \text{DB}) \cap \text{rng}(at_i^2, \text{DB})$. These frequencies are approximated by assuming that columns are independently distributed: $\mathcal{H}_{\mathbf{t}_i}(x_1, \dots, x_n) \approx \prod_{i \in [n]} \mathcal{H}_{at_i^j}(x_i)$. Note also that accuracy issues are compounded if frequency histograms are themselves estimated using bucketing. In general, approximating join sizes trades efficiency (of EvalBW) and accuracy (for cluster joins) to reduce memory.

6 Simulations on Real-World Datasets

To get some indication of how our schemes would fare in practice we simulate the storage and bandwidth they would incur in a real-world context. We show that in practice, PP indexing is likely to be more storage efficient than FP. We also confirm three claims made in this work: (1) PP indexing has equal or better bandwidth than FP on non-recursive joins (i.e. JnDT queries), (2) On recursive selects and joins (i.e. SjDT queries), the analogous choice is data and query dependent, and (3) our heuristic is accurate in finding optimal hybrid query execution plans.

We note that our goal here is not to make broad statements about all SQL data nor to perform a full system evaluation. We see our simulations more as a sanity check which might motivate large-scale implementations of our schemes. Additionally, we are not aware of any benchmarks with just join and select queries so we generate our own as described below.

SIMULATION SETUP. Our simulation dataset uses all relations from the MySQL Sakila benchmark ³ and the following fifteen frequently accessed relations from Chicago’s Open Data Portal: Bike_Racks, Census_Data,

³ We excluded the film_text relation since it is a subset of the film relation

Join category	# joins	Ratio of FpJnto PpJn BW		
		Min	Ave	Max
One-one	237	1.0	1.0	1.0
One-many	711	1.0	1.8	2.0
Many-many	932	1.5	465	8000

Fig. 10. Breakdown of all possible non-recursive join queries which returns at least one row by join types. For each type, we simulated the number of rows that would be sent using FP and PP indexing, and report the minimum, average and maximum overhead incurred.

Crimes_2019, Employee_Debt, Fire_Stations, Grafitti, Housing, IUCR_Codes, Land_Inventory, Libraries, Lobbyists, Police_Stations, Reloc_Vehicles, Street_Names, Towed_Vehicles. In total, our setup involved 30 relations, 175 attributes and 219,992 rows. In Appendix G, we present some additional summary statistics to better understand our source data. We also provide a full, annotated source code for our simulations in [40].

We include in α all single-attribute joins that return at least one row. This helps to filter out meaningless join queries (e.g. joining on “language” and “actor”). We consider joins within the Sakila relations and joins within the Chicago relations, but we do not attempt joins between the two independent sources. We generate recursive queries with J joins and S selections by selecting uniformly at random J distinct joins from α as well as S attributes and elements of their domains, discarding queries that return no rows. When $J \geq 2$ we only use input tables with less than 1000 rows to avoid very large output relations.

SERVER STORAGE. With the above setup we can get an idea of how much server-side storage would be required by each of our indexing schemes. Recall that our schemes make use of a RR multimap primitive and/or a hashset filtering primitive. Therefore, in Fig. 9 we report the number of multimap⁴ labels and values as well as the values in hashset HS for each of our StI schemes. We present our simulation results for the two datasets separately since the Chicago data set contains many more rows and would dominate the Sakila statistics. Additionally, we also show a breakdown of these statistics in terms of the queries they index to better understand the cost of each type of query support.

A number of observations can be made from this data. In our simulation we see that even though there are more selections to index (as evidenced by the number of labels), the multimap size (i.e. number of values) is dominated by join indexes. We expect this cost to be lower in a real system because a judicious database administrator can reduce the set of supported joins (α) to a smaller number than we did. Our simulation also brings forth another advantage of PP join indexing – it is more storage efficient by several orders of magnitude. This is because each row token is stored at most once per join (the same thing which causes PP to have better bandwidth) and, in the case of S_jDT, there is no need for the “internal join” indexing which essentially doubles the multimap’s labels and values. Finally, for the above reason, the storage overhead of hybrid indexing over FP indexing is very small so systems which currently use indexing schemes like FP (e.g. OPX or SPX) can upgrade its security at low cost.

JOIN CATEGORIES. We partition joins into three classes which behave quite differently: *one-one*, *one-many* and *many-many*. We say that a join $R \leftarrow R_1 \bowtie_{at_1, at_2} R_2$ is one-one if each row in R_1, R_2 occurs at most once in R . It is one-many if the above is true for one relation but not for the other. It is many-many if there exists rows in both R_1, R_2 which occur more than once in R . We record the breakdown of these classes in our datasets in Fig. 10.

STI FOR J_nDT. In Section 4 we showed that PP indexing has superior bandwidth on non-recursive join queries. We demonstrate that these savings by computing all 1880 possible joins in α and report our findings in Fig. 10. As one would expect, PP indexing always performs equal or better to FP – they perform equally for one-one joins but there are moderate and significant savings for one-many and many-many joins respectively.

STI FOR S_jDT. In Section 4.2 we noted that neither PP nor FP joins are strictly superior when it comes to recursive S_jDT queries. We demonstrate this using our datasets. For each combination of 1 to 3 joins and 0 to 2 selects, we randomly sampled 25 queries and report the results in Fig. 11. As can be seen, neither scheme can reliably achieve the optimal bandwidth. While FpS_j performed better on average, its maximum overhead exceeds that of PpS_j in about half the cases.

⁴ Note that in the case of FpS_j, HybStI, this includes the multimap for internal joins.

Query type	Ratio of BW to ideal						Bucketed $B = 1$			Bucketed $B = 200$			Full histograms					
	FpSj			PpSj			Correct	Wrong			Correct	Wrong						
	Min	Ave	Max	Min	Ave	Max		R1	R2	R3		R1	R2	R3	R1	R2	R3	
1 \bowtie , 0 σ	1.0	9.6	37	1.0	1.0	1.0	14	11	0	0	25	0	0	0	25	0	0	0
1 \bowtie , 1 σ	1.0	1.6	4.0	1.0	60	302	6	17	0	6	12	0	12	1	16	0	8	1
1 \bowtie , 2 σ	1.0	1.3	2.0	1.0	90	500	5	0	0	20	14	0	1	10	15	0	0	10
2 \bowtie , 0 σ	1.0	3.3	57	1.3	13	54	5	0	0	20	21	3	0	1	25	0	0	0
2 \bowtie , 1 σ	1.0	15	201	1.0	41	201	15	0	1	9	18	6	1	0	24	0	1	0
2 \bowtie , 2 σ	1.0	14	121	1.0	93	535	17	8	0	0	20	0	1	4	23	0	0	2
3 \bowtie , 0 σ	1.0	7.2	48	2.4	9.1	17	5	20	0	0	8	10	7	0	21	2	2	0
3 \bowtie , 1 σ	1.0	6.5	63	2.6	23	60	2	23	0	0	19	1	5	0	25	0	0	0
3 \bowtie , 2 σ	1.0	5.0	61	2.3	30	84	7	18	0	0	16	4	5	0	24	0	1	0

Fig. 11. Left: On randomly generated queries involving the indicated number of joins (\bowtie) and selects (σ), we report the minimum, average and maximum ratios of rows sent using each indexing technique compared to the theoretical minimum possible. Right: Using the same queries, we report the accuracy of our heuristic under three different client storage settings. When a suboptimal query execution plan is returned, we report the point at which our heuristic fails (with R3 being the closest to success).

HYBRID STI. In Section 5 we provided a heuristic for client-side leakage-aware query planning. We demonstrate its efficacy when frequency histograms are estimated via three bucketing options: $B = |\text{rng}(at, DB)|$ (full histograms), $B = 200$ and $B = 1$. We use the same 225 queries as the SjdT simulations and set the bandwidth limit L for each $q \in \text{SjdT}$ to be the mean incurred by all 2^J possible HybDT queries to ensure that the optimization is non-trivial. Additionally, join sizes $\mathcal{F}, \mathcal{P}_1, \mathcal{P}_2$ are estimated using the histogram. Therefore, our simulation is conservative and we expect our heuristic to perform better in applications with a fixed L and precomputed join sizes.

In Fig. 11 we show how our heuristic performed for each query type and histogram estimation technique. When the optimal join annotation is not returned we note which “level” the heuristic failed at, where the levels are defined in relation to our definition of “optimality” given in Section 5. In particular, an R1 failure means the returned q' exceeds bandwidth limit L when StE.Eval is run, an R2 failure means q' used more FP joins than was necessary to reduce bandwidth below L and an R3 failure means q' was not the smallest bandwidth option which uses the minimal number of FP joins while meeting L .

Unsurprisingly, there is a direct tradeoff between client memory and the heuristic’s accuracy: across all 225 queries, the heuristic returned the optimal q' on 198 with full histograms but only 143 and 76 when $B = 200$ and $B = 1$ respectively. More interestingly, our heuristic seems to improve when the search space increases: when there is one join the heuristic performed slightly better averaged across all three B values than guessing (58.7% vs 50%) but when there are three it performs significantly better (56.4% vs 12.5%). This demonstrates that our heuristic works when it is most needful since we expect the bandwidth overhead from an incorrect choice to increase with query complexity.

7 Conclusion

Our work introduces partially precomputed join indexing and incorporates it into a hybrid StE scheme. While we did not explore it in this work, we believe that our schemes can be extended to support dynamic queries and adaptive security via multimap primitives of the same kind. We believe the former can be achieved in a similar way to KM’s extension of SPX to SPX⁺. To achieve the latter, our schemes can be reframed in JN’s model for adaptive compromise [30]. Future work can also extend our query support, possibly by incorporating cryptographic techniques for range queries or aggregations [22,29]. Higher query support would also enable more rigorous testing using real-world applications and query benchmarks. Stronger security can be achieved using lower-leakage indexing primitives [31,38,33].

We also introduce leakage-aware query planning which we believe to be of independent interest as it incorporates structured indexing into DBMS architecture, which may help StE become a part of commercial DBMSes. Future work could improve our heuristic’s efficiency and accuracy, or develop analogous hybrid schemes for other query classes.

8 Acknowledgements

We would like to thank the anonymous reviewers for their comments on our work. We are also grateful to Mihir Bellare and Francesca Falzon for discussion and insights. Cash was supported in part by NSF CNS 1703953. Ng was supported by DSO National Laboratories. Rivkin was supported by the Liew Family College Research Fellows Fund.

References

1. encrypted-bigquery-client. <https://github.com/google/encrypted-bigquery-client>, 2015.
2. City of chicago data portal. <https://data.cityofchicago.org/>, 2021.
3. Sakila sample database. <https://dev.mysql.com/doc/sakila/en/>, 2021.
4. P. Antonopoulos, A. Arasu, K. D. Singh, K. Eguro, N. Gupta, R. Jain, R. Kaushik, H. Kodavalla, D. Kossmann, N. Ogg, et al. Azure sql database always encrypted. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 1511–1525, 2020.
5. S. Bajaj and R. Sion. Trustseddb: A trusted hardware-based database with privacy and data confidentiality. *IEEE Transactions on Knowledge and Data Engineering*, 26(3):752–765, 2013.
6. J. Bater, G. Elliott, C. Eggen, S. Goel, A. Kho, and J. Rogers. Smcql: secure querying for federated databases. *Proceedings of the VLDB Endowment*, 10(6):673–684, 2017.
7. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 409–426. Springer, 2006.
8. V. Bindschaedler, P. Grubbs, D. Cash, T. Ristenpart, and V. Shmatikov. The tao of inference in privacy-protected databases. *Proceedings of the VLDB Endowment*, 11(11):1715–1728, 2018.
9. L. Blackstone, S. Kamara, and T. Moataz. Revisiting leakage abuse attacks. Cryptology ePrint Archive, Report 2019/1175, 2019. <https://eprint.iacr.org/2019/1175>.
10. N. Bruno and L. Gravano. *Statistics on query expressions in relational database management systems*. PhD thesis, Columbia University, 2003.
11. Y. Cao, W. Fan, Y. Wang, and K. Yi. Querying shared data with security heterogeneity. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 575–585, 2020.
12. D. Cash, J. Jaeger, S. Jarecki, C. S. Jutla, H. Krawczyk, M.-C. Rosu, and M. Steiner. Dynamic searchable encryption in very-large databases: data structures and implementation. In *NDSS*, volume 14, pages 23–26. Citeseer, 2014.
13. D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Roşu, and M. Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *Annual cryptology conference*, pages 353–373. Springer, 2013.
14. A. K. Chandra and P. M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 77–90, 1977.
15. M. Chase and S. Kamara. Structured encryption and controlled disclosure. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 577–594. Springer, 2010.
16. S. S. Chow, J.-H. Lee, and L. Subramanian. Two-party computation model for privacy-preserving queries over distributed databases. In *NDSS*. Citeseer, 2009.
17. V. Ciriani, S. D. C. Di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati. Keep a few: Outsourcing data while maintaining confidentiality. In *European Symposium on Research in Computer Security*, pages 440–455. Springer, 2009.
18. R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky. Searchable symmetric encryption: Improved definitions and efficient constructions. Cryptology ePrint Archive, Report 2006/210, 2006. <https://eprint.iacr.org/2006/210>.
19. E. Damiani, S. D. C. Vimercati, S. Jajodia, S. Paraboschi, and P. Samarati. Balancing confidentiality and efficiency in untrusted relational dbmss. In *Proceedings of the 10th ACM conference on Computer and communications security*, pages 93–102, 2003.
20. I. Demertzis, D. Papadopoulos, C. Papamanthou, and S. Shintre. Seal: Attack mitigation for encrypted databases via adjustable leakage. In *29th USENIX Security Symposium (USENIX Security 20)*, 2020.
21. S. Evdokimov and O. Günther. Encryption techniques for secure database outsourcing. In *European Symposium on Research in Computer Security*, pages 327–342. Springer, 2007.
22. S. Faber, S. Jarecki, H. Krawczyk, Q. Nguyen, M. Rosu, and M. Steiner. Rich queries on encrypted data: Beyond exact matches. Cryptology ePrint Archive, Report 2015/927, 2015. <https://eprint.iacr.org/2015/927>.
23. S. Garg, P. Mohassel, and C. Papamanthou. Tworam: efficient oblivious ram in two rounds with applications to searchable encryption. In *Annual International Cryptology Conference*, pages 563–592. Springer, 2016.

24. P. Grofig, I. Hang, M. Härterich, F. Kerschbaum, M. Kohler, A. Schaad, A. Schröpfer, and W. Tighzert. Privacy by encrypted databases. In *Privacy Technologies and Policy - Second Annual Privacy Forum, APF 2014, Athens, Greece, May 20-21, 2014. Proceedings*, pages 56–69, 2014.
25. P. Grubbs, M.-S. Lacharité, B. Minaud, and K. G. Paterson. Pump up the volume: Practical database reconstruction from volume leakage on range queries. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 315–331, 2018.
26. P. Grubbs, K. Sekniqi, V. Bindschaedler, M. Naveed, and T. Ristenpart. Leakage-abuse attacks against order-revealing encryption. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 655–672. IEEE, 2017.
27. Z. Gui, O. Johnson, and B. Warinschi. Encrypted databases: New volume attacks against range queries. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 361–378, 2019.
28. H. Hacigümüş, B. Iyer, C. Li, and S. Mehrotra. Executing sql over encrypted data in the database-service-provider model. In *Proceedings of the 2002 ACM SIGMOD international conference on Management of data*, pages 216–227, 2002.
29. T. Hackenjos, F. Hahn, and F. Kerschbaum. Sagma: Secure aggregation grouped by multiple attributes. *ACM SIGMOD Record*, 2020.
30. J. Jaeger and N. Tyagi. Handling adaptive compromise for practical encryption schemes. In *Annual International Cryptology Conference*, pages 3–32. Springer, 2020.
31. S. Kamara and T. Moataz. Encrypted multi-maps with computationally-secure leakage. *IACR Cryptol. ePrint Arch.*, 2018:978, 2018.
32. S. Kamara and T. Moataz. Sql on structurally-encrypted databases. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 149–180. Springer, 2018.
33. S. Kamara, T. Moataz, and O. Ohrimenko. Structured encryption and leakage suppression. In *Annual International Cryptology Conference*, pages 339–370. Springer, 2018.
34. S. Kamara, T. Moataz, S. Zdonik, and Z. Zhao. An optimal relational database encryption scheme. *Cryptology ePrint Archive*, Report 2020/274, 2020. <https://eprint.iacr.org/2020/274> Accessed: 2020-02-29.
35. M. Kantarcıoglu and C. Clifton. Security issues in querying encrypted data. In *IFIP Annual Conference on Data and Applications Security and Privacy*, pages 325–337. Springer, 2005.
36. E. S. Lab. The clusion library. <https://github.com/encryptedsystems/Clusion>, 2020.
37. M. Naveed, S. Kamara, and C. V. Wright. Inference attacks on property-preserving encrypted databases. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 644–655, 2015.
38. S. Patel, G. Persiano, K. Yeo, and M. Yung. Mitigating leakage in secure cloud-hosted data structures: Volume-hiding for multi-maps via hashing. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 79–93, 2019.
39. R. A. Popa, C. M. Redfield, N. Zeldovich, and H. Balakrishnan. Cryptdb: protecting confidentiality with encrypted query processing. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, pages 85–100, 2011.
40. A. Rivkin. Hybrid indexing simulations. <https://github.com/AdamRivkin/Hybrid-Indexing-Simulations>, 2021.
41. J. Sack. Optimizing your query plans with the sql server 2014 cardinality estimator, 2014.
42. D. X. Song, D. Wagner, and A. Perrig. Practical techniques for searches on encrypted data. In *Proceeding 2000 IEEE Symposium on Security and Privacy. S&P 2000*, pages 44–55. IEEE, 2000.
43. S. L. Tu, M. F. Kaashoek, S. R. Madden, and N. Zeldovich. Processing analytical queries over encrypted data. 2013.
44. Z. Yang, S. Zhong, and R. N. Wright. Privacy-preserving queries on encrypted data. In *European Symposium on Research in Computer Security*, pages 479–495. Springer, 2006.

A CJJ+’s Multimap/Dictionary Encryption Schemes

CJJ+’s RH DICTIONARY ENCRYPTION SCHEME. In our StE scheme constructed using **SqlStE** (in Section 3) we use a specific RH dictionary encryption scheme to store the rows in DB. We formalize this as Dye_π whose algorithms are in Fig. 12. The primitives (given as input to **SqlStE**) used in Dye_π are symmetric encryption scheme SE and function family F. Note that $\text{Dye}_\pi.\text{KS} = \text{F.KS} \times \text{SE.KS}$.

EXAMPLE RR MULTIMAP ENCRYPTION SCHEME. In our StI schemes such as PpJn, PpSj, HybStI we use a RR multimap Mme as a primitive. We give an example of such a scheme $\text{Mme}_\pi^{\text{rr}}$ which is also based on Π_{bas} . Its algorithms are in Fig. 12. The primitives are as in Dye_π but we require that $\text{SE.KS} = \{0, 1\}^{\text{F.ol}}$. Note that $\text{Mme}_\pi^{\text{rr}}.\text{KS} = \text{F.KS}$.

<p>Alg Dye_π.Enc((K_f, K_e), D)</p> <p>Pad all values in D to the same length</p> <p>For $\ell \in \mathbf{D}.\text{Lbls}$ do $\mathbf{D}'[\text{F.Ev}(K_f, \ell)] \leftarrow \text{SE.Enc}(K_e, \mathbf{D}[\ell])$</p> <p>Return ((K_f, K_e), D')</p> <p>Alg Dye_π.Tok((K_f, K_e), ℓ)</p> <p>$\text{tk} \leftarrow \text{F.Ev}(K_f, \ell)$; Return tk</p>	<p>Alg Dye_π.Eval(tk, D')</p> <p>Return D'[tk]</p> <p>Alg Dye_π.Dec((K_f, K_e), C)</p> <p>Unpad and return SE.Dec(K_e, C)</p>
<p>Alg Mme_π^{rr}.Enc(K_f, M)</p> <p>Pad all values in M to the same length</p> <p>For $\ell \in \mathbf{M}.\text{Lbls}$ do</p> <p style="padding-left: 20px;">$K_e \leftarrow \text{F.Ev}(K_f, \ell \ 0)$; $K \leftarrow \text{F.Ev}(K_f, \ell \ 1)$</p> <p style="padding-left: 20px;">For $v \in \mathbf{M}[\ell]$ do</p> <p style="padding-left: 40px;">$\mathbf{D}[\text{F.Ev}(K, \text{ctr})] \leftarrow \text{SE.Enc}(K_e, v)$; $\text{ctr} \leftarrow \text{ctr} + 1$</p> <p>Return (K_f, D)</p>	<p>Alg Mme_π^{rr}.Tok(K_f, ℓ)</p> <p>Return (F.Ev(K_f, $\ell \ 0$), F.Ev(K_f, $\ell \ 1$))</p> <p>Alg Mme_π^{rr}.Eval((K_e, K), D)</p> <p>While $\mathbf{D}[\text{F.Ev}(K, \text{ctr})] \neq \perp$ do</p> <p style="padding-left: 20px;">$x \leftarrow \text{SE.Dec}(K_e, \mathbf{D}[\text{F.Ev}(K, \text{ctr})])$</p> <p style="padding-left: 20px;">$\text{ctr} \leftarrow \text{ctr} + 1$; Unpad x then $S \stackrel{\cup}{\leftarrow} x$</p> <p>Return S</p>

Fig. 12. Algorithms for RH dictionary encryption scheme Dye_π and RR multimap encryption scheme Mme_π^{rr}.

B Proof of Theorem 1

Theorem 1. Let $\text{StE} = \text{SqlStE}[\text{Stl}, \text{SE}, \text{F}]$ be a correct StE scheme for SqlDT. Then given algorithms $\mathcal{L}^i, \mathcal{S}^i$ and adversary A we can define \mathcal{L} as in Section 3.2 and construct $\mathcal{S}, A_s, A_f, A_i$ such that:

$$\mathbf{Adv}_{\text{StE}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A) \leq \mathbf{Adv}_{\text{SE}}^{\text{ind}\$}(A_s) + \mathbf{Adv}_{\text{F}}^{\text{prf}}(A_f) + \mathbf{Adv}_{\text{Stl}, \mathcal{L}^i, \mathcal{S}^i}^{\text{ss}}(A_i).$$

Proof. The adversaries, simulator and games G_0, G_1, G_2, G_3 are given in Fig. 13. Notice that the EncRows algorithm used in the adversaries and games is given at the top, and uses two oracles ENC, FN which the algorithms define. Let b be the challenge bit selected in $G_{\text{StE}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A)$.

Notice that we can express $\mathbf{Adv}_{\text{StE}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A) = \Pr[G_{\text{StE}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A)|b = 1] - \Pr[G_{\text{StE}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A)|b = 0] = \Pr[G_3] - \Pr[G_0]$. In $b = 1$ case, this follows directly from the definition of A_i . In the $b = 0$ case, this follows from the definition of $\mathcal{L}^i, \mathcal{S}^i$.

The only difference between G_0 and G_1 is whether $\text{IX}, \text{tk}_1, \dots, \text{tk}_n$ are generated using Stl's algorithms or \mathcal{S} . In both cases, **D'**'s values are encrypted using SE.Enc. This is the same differentiation going on in the semantic security game so $G_{\text{Stl}, \mathcal{L}^i, \mathcal{S}^i}^{\text{ss}}(A_i) = \Pr[G_1] - \Pr[G_0]$. Similarly the difference between G_1 and G_2 is whether the values in **D'** are the output of SE.Enc or random strings which is what is going on in the IND $\$$ -security game $G_{\text{SE}}^{\text{ind}\$}(A_s)$, so $\mathbf{Adv}_{\text{SE}}^{\text{ind}\$}(A_s) = \Pr[G_2] - \Pr[G_1]$. Once again, the difference between G_2 and G_3 is whether the labels in **D'** (i.e. the tokens in Dye_π.Enc) are generated using F.Ev or a random function which is what is going on in the PRF-security game $G_{\text{F}}^{\text{prf}}(A_f)$, so $\mathbf{Adv}_{\text{F}}^{\text{prf}}(A_f) = \Pr[G_3] - \Pr[G_2]$.

Combining all the above equations gives the desired bound on $\mathbf{Adv}_{\text{StE}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A)$.

C Proof of Theorem 2

Theorem 2. Let \mathcal{L}, \mathcal{S} be the leakage algorithm and simulator for Mme. Let $\mathcal{L}^f, \mathcal{L}^p$ be the leakage algorithms given in Fig. 5. Then, for all adversaries A :

$$\mathbf{Adv}_{\text{FpJn}, \mathcal{L}^f, \mathcal{S}}^{\text{ss}}(A) \leq \mathbf{Adv}_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A).$$

Additionally, there exists adversary A_m and simulator \mathcal{S}^p such that:

$$\mathbf{Adv}_{\text{PpJn}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}(A) \leq \mathbf{Adv}_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A_m).$$

Proof. The first result follows directly from the definition of FpJn, \mathcal{L}^f . The second result requires us to define A_m, \mathcal{S}^p , which we do in Fig. 14. In both of these, tokens are just concatenated and deconcatenated as needed by the definition of PpJn. The result follows immediately.

<p>Alg $\mathcal{S}(lk^i, N, L)$</p> <p>$(IX, (tk_1, \dots, tk_n)) \leftarrow \mathcal{S}^i(lk^i)$ $P \leftarrow \bigcup_{i \in [n]} \text{Stl.Eval}(tk_i, IX)$ For $rt \in \bigcup_{rt \in P} rt$ do $\mathbf{D}'[rt] \leftarrow_{\\$} \{0, 1\}^{\text{SE.cl}(L)}$ While $\mathbf{D}'.\text{Lbls} < N$ do $rt \leftarrow_{\\$} \{0, 1\}^{\text{F.ol}} ; \mathbf{D}'[rt] \leftarrow_{\\$} \{0, 1\}^{\text{SE.cl}(L)}$ Return $((IX, \mathbf{D}'), (tk_1, \dots, tk_n))$</p>	<p>Subroutine $\text{EncRows}^{\text{ENC, FN}}((DB, \alpha))$</p> <p>For $(id, R) \in DB$ do For $r \in R.T$ do $\mathbf{D}[(id, r[\text{uk}(id)])] \leftarrow r$ Pad all values in \mathbf{D} to the same length For $\ell \in \mathbf{D}.\text{Lbls}$ do $\mathbf{T}[\ell] \leftarrow_{\\$} \text{FN}(\ell) ; \mathbf{D}'[\mathbf{T}[\ell]] \leftarrow_{\\$} \text{ENC}(\mathbf{D}[\ell])$ Return $(\mathbf{D}', \alpha, \mathbf{T})$</p>
<p>Adversary $A_i(\mathbf{s})$</p> <p>$(DB, \mathbf{q}, st) \leftarrow_{\\$} A(\mathbf{s})$ $K_e \leftarrow_{\\$} \text{SE.KS} ; K_f \leftarrow_{\\$} \text{F.KS}$ Define $\text{ENC} : \text{ENC}(x) = \text{SE.Enc}(K_e, \cdot)$ Define $\text{FN} : \text{FN}(x) = \text{F.Ev}(K_f, \cdot)$ $(\mathbf{D}', \alpha, \mathbf{T}) \leftarrow_{\\$} \text{EncRows}^{\text{ENC, FN}}(DB)$ Return $((\mathbf{D}', \alpha, \mathbf{T}), \mathbf{q}, (\mathbf{D}', st))$</p> <p>Adversary $A_i(\mathbf{g}, IX, \mathbf{tk}, (\mathbf{D}', st))$</p> <p>$b' \leftarrow_{\\$} A(\mathbf{g}, (IX, \mathbf{D}'), \mathbf{tk}, st)$ Return b'</p>	<p>Adversaries $\boxed{A_s^{\text{ENC}}} \boxed{A_f^{\text{FN}}}$</p> <p>$(DB, \mathbf{q}, st) \leftarrow_{\\$} A(\mathbf{s}) ; K_f \leftarrow_{\\$} \text{F.KS}$ Define $\text{FN} : \text{FN}(x) = \text{F.Ev}(K_f, \cdot)$</p> <p style="border: 1px dashed black; padding: 2px;">Let ENC be a random function from $\{0, 1\}^L$ to $\{0, 1\}^{\text{SE.cl}(L)}$</p> <p>$(\mathbf{D}', \alpha, \mathbf{T}) \leftarrow_{\\$} \text{EncRows}^{\text{ENC, FN}}(DB)$ $lk^i \leftarrow_{\\$} \mathcal{L}^i(DS, \mathbf{q})$ $(IX, (tk_1, \dots, tk_n)) \leftarrow \mathcal{S}^i(lk^i)$ $b' \leftarrow_{\\$} A(\mathbf{g}, (IX, \mathbf{D}'), (tk_1, \dots, tk_n), st)$ Return b'</p>
<p>Games $\boxed{G_0(A)}$ $\boxed{G_1(A)}$</p> <p>$(DB, \mathbf{q}, st) \leftarrow_{\\$} A(\mathbf{s})$ $K_e \leftarrow_{\\$} \text{SE.KS} ; K_f \leftarrow_{\\$} \text{F.KS}$ Define $\text{ENC} : \text{ENC}(x) = \text{SE.Enc}(K_e, \cdot)$ Define $\text{FN} : \text{FN}(x) = \text{F.Ev}(K_f, \cdot)$ $(\mathbf{D}', \alpha, \mathbf{T}) \leftarrow_{\\$} \text{EncRows}^{\text{ENC, FN}}(DB)$</p> <div style="border: 1px solid black; padding: 2px; margin: 5px 0;"> $K'_i \leftarrow_{\\$} \text{Stl.KS}$ $(K_i, IX) \leftarrow_{\\$} \text{Stl.Enc}(K'_i, DS)$ For $i \in [n]$ do $tk_i \leftarrow_{\\$} \text{Stl.Tok}(K_i, q_i)$ </div> <p style="border: 1px dashed black; padding: 2px; margin: 5px 0;">$lk^i \leftarrow_{\\$} \mathcal{L}^i(DS, (q_1, \dots, q_n))$</p> <p style="border: 1px dashed black; padding: 2px; margin: 5px 0;">$(IX, (tk_1, \dots, tk_n)) \leftarrow \mathcal{S}^i(lk^i)$</p> <p>$b' \leftarrow_{\\$} A(\mathbf{g}, (IX, \mathbf{D}'), (tk_1, \dots, tk_n), st)$ Return $b' = 1$</p>	<p>Games $\boxed{G_2(A)}$ $\boxed{G_3(A)}$</p> <p>$(DB, \mathbf{q}, st) \leftarrow_{\\$} A(\mathbf{s}) ; K_f \leftarrow_{\\$} \text{F.KS}$ Let ENC be a random function from $\{0, 1\}^L$ to $\{0, 1\}^{\text{SE.cl}(L)}$</p> <div style="border: 1px solid black; padding: 2px; margin: 5px 0;">Define $\text{FN} : \text{FN}(x) = \text{F.Ev}(K_f, \cdot)$</div> <p style="border: 1px dashed black; padding: 2px; margin: 5px 0;">Let $\text{FN}(\cdot)$ be a random function from $\{0, 1\}^*$ to $\{0, 1\}^{\text{F.ol}}$</p> <p>$(\mathbf{D}', \alpha, \mathbf{T}) \leftarrow_{\\$} \text{EncRows}^{\text{ENC, FN}}(DB)$ $lk^i \leftarrow_{\\$} \mathcal{L}^i(DS, (q_1, \dots, q_n))$ $(IX, (tk_1, \dots, tk_n)) \leftarrow \mathcal{S}^i(lk^i)$ $b' \leftarrow_{\\$} A(\mathbf{g}, (IX, \mathbf{D}'), (tk_1, \dots, tk_n), st)$ Return $b' = 1$</p>

Fig. 13. Simulator, adversaries and games used in the proof of Theorem 1.

D Leakage Profile and Security Proof for PpSj

LEAKAGE. In Section 4.2 we overviewed the different forms of leakage that make up PpSj's leakage profile. Here, we provide a full pseudocode of \mathcal{L}^P for completeness and give some intuition to aid in reading it.

\mathcal{L}^P 's pseudocode is given in the top left of Fig. 15. It calls the three subroutines which compute the query-dependent leakage. RS is first called on each of the q_1, \dots, q_n . Through this, counters c_q, c_p are maintained which count the number of accesses to \mathbf{M} (to retrieve a value) and HS (to filter based on a predicate). The labels or predicates associated to each of these subqueries are logged in the vectors \mathbf{q}, \mathbf{p} . The r_1, \dots, r_n returned during these calls are part of the leakage. It reveals the “structure” of each query that was made.

Vector \mathbf{q} are the queries made to the multimap primitive and is therefore an input to \mathcal{L} . The output of this makes up the multimap leakage (e.g. multimap query equality pattern) and will be returned by \mathcal{L}^P . Vector \mathbf{p} is given as input to QP and HF which compute the equality pattern and filtering results on the hashset predicates respectively.

Theorem 3. *Let \mathcal{L}, \mathcal{S} be the leakage algorithm and simulator for Mme respectively. Let $\mathcal{L}^P, \mathcal{S}^P$ be as defined in Fig. 15 and let F be the function family used. Then for all adversaries A there exists adversaries A_m, A_f such that:*

$$\text{Adv}_{\text{PpSj}, \mathcal{L}^P, \mathcal{S}^P}^{\text{SS}}(A) \leq \text{Adv}_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{SS}}(A_m) + (p + 1) \cdot \text{Adv}_F^{\text{prf}}(A_f).$$

Adversary $A_m(s)$ Return $A(s)$ Adversary $A_m(g, IX, (tk_1, \dots, tk_{2n}), st)$ $tk \leftarrow ((tk_1, tk_2), \dots, (tk_{2n-1}, tk_{2n}))$ $b' \leftarrow_s A(g, IX, tk, st)$; Return b'	Alg $S^P(lk^m)$ $(EM, (tk_1, \dots, tk_{2n})) \leftarrow S^m(lk^m)$ $tk \leftarrow ((tk_1, tk_2), \dots, (tk_{2n-1}, tk_{2n}))$ Return (EM, tk)
--	---

Fig. 14. Simulators (right) and adversaries (left) used in the proof of Theorem 2.

Here, p is the number of distinct predicates used in constructing HS.

Proof. Adversary A_m is given in Fig. 15. In the same diagram, we see A_1, A_2 which are both PRF adversaries playing G_F^{prf} . We define A_f to randomly pick one at run time and use it.

Now we can proceed via a standard hybrid argument. Let b_p, b_f, b_m be the challenge bits in $G_{\text{PpSj}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}$, G_F^{prf} and $G_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}$ respectively.

From the various advantage definitions, we have that $\mathbf{Adv}_{\text{PpSj}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}(A) = \Pr[G_{\text{PpSj}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}(A)|b_p = 1] - \Pr[G_{\text{PpSj}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}(A)|b_p = 0]$, $\mathbf{Adv}_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A_m) = \Pr[G_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A_m)|b_m = 1] - \Pr[G_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A_m)|b_m = 0]$, and $\mathbf{Adv}_F^{\text{prf}}(A_1) = \Pr[G_F^{\text{prf}}(A_1)|b_f = 1] - \Pr[G_F^{\text{prf}}(A_1)|b_f = 0]$. Notice also that $\Pr[G_F^{\text{prf}}(A_2)|b_f = 0, c = i] = \Pr[G_F^{\text{prf}}(A_2)|b_f = 1, c = i + 1]$ for $i \in [p - 1]$ and $\Pr[G_F^{\text{prf}}(A_2)|b_f = 1, c = j] - \Pr[G_F^{\text{prf}}(A_2)|b_f = 1, c = j] \leq \mathbf{Adv}_F^{\text{prf}}(A_2)$ for $j \in [p]$. This means that

$$p \cdot \mathbf{Adv}_F^{\text{prf}}(A_2) \geq \Pr[G_F^{\text{prf}}(A_2)|b_f = 1, c = p] - \Pr[G_F^{\text{prf}}(A_1)|b_f = 0, c = 1].$$

Notice that A_m in $G_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}$ uses the game to simulate multimap encryption and performs the rest itself as it happens in the “real world” of $G_{\text{PpSj}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}(A)$. This gives $\Pr[G_{\text{PpSj}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}(A)|b_p = 1] = \Pr[G_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A_m)|b_m = 1]$. Similarly, A_1 simulates multimap encryption as in the “ideal world” of $G_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}$ and defers the filtering key production to FN which gives us $\Pr[G_{\text{Mme}, \mathcal{L}, \mathcal{S}}^{\text{ss}}(A_m)|b_m = 0] = \Pr[G_F^{\text{prf}}(A_1)|b_f = 1]$. When A_2 plays $G_F^{\text{prf}}(A_2)$, if $c = p$ then all the K_i will be randomly selected. This means $\Pr[G_F^{\text{prf}}(A_1)|b_f = 0] = \Pr[G_F^{\text{prf}}(A_2)|b_f = 1, c = p]$. Over p hybrids, we get to the version where all the $\text{F.Ev}(K_i, \cdot)$ (where K_i is not revealed to the adversary) are simulated with random functions, giving us $\Pr[G_F^{\text{prf}}(A_1)|b_f = 0, c = 1] = \Pr[G_{\text{PpSj}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}(A)|b_p = 0]$ because this selects all of HS elements as S^P does.

E FpSj Details

PSEUDOCODE. The pseudocode for FpSj is given in Fig. 16.

Notice that $\text{FpSj.KS} = \text{Mme}_\pi^{\text{rr}}.\text{KS} = \text{F.KS}$. The flags ii_j, ij used come from the terms used for query classification by KMZZ in [34] where recursive joins are split into “internal joins” (i.e. queries of the form $(j, \mathbf{t}_1, \mathbf{t}_2, (\mathbf{r}, id_1), q_2)$ or $(j, \mathbf{t}_1, \mathbf{t}_2, (\mathbf{r}, id_1), q_2)$) and “intermediate internal joins” (i.e. those of form $(j, \mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$).

As alluded to in Section 4.2, the biggest difference between FP and PP indexing of SjDT queries is the handling of internal joins. As was done in OPX, this necessitates the use of a specific RR multimap encryption primitive (namely $\text{Mme}_\pi^{\text{rr}}$ from Appendix A) to minimize leakage. In the FpSj’s encryption algorithm, we manually add entries to the server-side data structure of $\text{Mme}_\pi^{\text{rr}}$ (i.e. dictionary \mathbf{D}) to index these joins.

LEAKAGE. The leakage algorithm \mathcal{L}^f for FpSj is given in Fig. 17.

As mentioned in Section 4.2, the differences between this and \mathcal{L}^p all stem from their different handling of joins. As depicted in Fig. 15, we break down the latter’s query-dependent leakage into the recursion structure (\mathbf{r} computed by RS), leakage due to queries to the underlying multimap encryption scheme (lk computed using \mathcal{L}), hashset filtering results (SET' computed using HF), hashset query patterns (computed using $\text{QP}(\mathbf{p})$) and the total number of hashset predicates made (c_p). For examples and intuition of each of these forms of leakage, see the examples given in Section 4.2.

With \mathcal{L}^f , we must compute leakage due to three different types of join queries (leaf, internal or internal intermediate). For leaf joins, the difference in leakage for the two SjDT StI schemes is exactly that of the two JnDT schemes given in Section 4. For internal intermediate joins, these are handled entirely using hashset filtering, much like the recursive joins in PpSj. As such, the leakage is comparable (in that we reveal the

<p>Alg $\mathcal{L}^P(\text{DS}, (q_1, \dots, q_n))$</p> <p>Construct \mathbf{M}, SET as in $\text{PpSj.Enc}(\cdot, \text{DS})$</p> <p>For $i = 1, \dots, n$ do</p> <p style="padding-left: 20px;">$(r_i, \mathbf{q}, \mathbf{p}, c_q, c_p) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, c_q, c_p)$</p> <p>$\mathbf{r} \leftarrow (r_1, \dots, r_n)$; $lk \leftarrow \mathcal{L}(\mathbf{M}, \mathbf{q})$</p> <p>$\text{SET}' \leftarrow \text{HF}(\mathbf{p}, \bigcup_{q \in \mathbf{q}} \mathbf{M}[q], \text{SET})$</p> <p>Return $(\mathbf{r}, lk, \text{QP}(\mathbf{p}), c_p, \text{SET}', \text{SET}')$</p> <p>Subroutine $\text{HF}((p_1, \dots, p_n), S, \text{SET})$</p> <p>For all $i \in [n]$ and $\mathbf{rt} \in S$ do</p> <p style="padding-left: 20px;">If $(p_i, \mathbf{rt}) \in \text{SET}$ then $\text{SET}' \leftarrow \bigcup (i, \mathbf{rt})$</p> <p>Return SET'</p> <p>Subroutine $\text{QP}((p_1, \dots, p_n))$</p> <p>For all $i, j \in [n]$ if $p_i = p_j$ then</p> <p style="padding-left: 20px;">$\mathbf{P}[i, j] \leftarrow 1$ else $\mathbf{P}[i, j] \leftarrow 0$</p> <p>Return \mathbf{P}</p>	<p>Subroutine $\text{RS}(q, \mathbf{q}, \mathbf{p}, c_q, c_p)$</p> <p>If $q = (\mathbf{r}, id)$ then $\mathbf{q} \leftarrow \bigcup (\mathbf{r}, id)$; $r \leftarrow (\mathbf{m}, c_q)$; $c_q \leftarrow c_q + 1$</p> <p>Else if $q = (\mathbf{s}, at, x, (\mathbf{r}, id))$ then</p> <p style="padding-left: 20px;">$\mathbf{q} \leftarrow \bigcup (\mathbf{s}, at, x)$; $r \leftarrow (\mathbf{m}, c_q)$; $c_q \leftarrow c_q + 1$</p> <p>Else if $q = (\mathbf{s}, at, x, q_1)$ then</p> <p style="padding-left: 20px;">$(r_1, \mathbf{q}, \mathbf{p}, c_q, c_p) \leftarrow \text{RS}(q_1, \mathbf{q}, \mathbf{p}, c_q, c_p)$</p> <p style="padding-left: 20px;">$\mathbf{p} \leftarrow \bigcup (\mathbf{s}, at, x)$; $r \leftarrow (\mathbf{p}, c_p, r_1)$; $c_p \leftarrow c_p + 1$</p> <p>Else if $q = (\mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$ then</p> <p style="padding-left: 20px;">For $i = 1, 2$ do</p> <p style="padding-left: 40px;">If $q_i = (\mathbf{r}, id)$ then</p> <p style="padding-left: 60px;">$\mathbf{q} \leftarrow \bigcup (\mathbf{j}, \mathbf{t}_1, \mathbf{t}_2, i)$; $r_i \leftarrow (\mathbf{m}, c_q)$; $c_q \leftarrow c_q + 1$</p> <p style="padding-left: 40px;">Else</p> <p style="padding-left: 60px;">$(r'_i, \mathbf{q}, \mathbf{p}, c_q, c_p) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, c_q, c_p)$</p> <p style="padding-left: 60px;">$\mathbf{p} \leftarrow \bigcup (\mathbf{j}, \mathbf{t}_1, \mathbf{t}_2, i)$; $r_i \leftarrow (\mathbf{p}, c_p, r'_i)$; $c_p \leftarrow c_p + 1$</p> <p style="padding-left: 40px;">$r \leftarrow (\mathbf{j}, r_1, r_2)$</p> <p>Return $(r, \mathbf{q}, \mathbf{p}, c_q, c_p)$</p>
<p>Alg $\mathcal{S}^P((r_1, \dots, r_n), lk, \mathbf{P}, c_p, \text{SET}', N)$</p> <p>$(EM, \mathbf{mt}) \leftarrow \mathcal{S}(lk)$</p> <p>For $i = 1, \dots, c_p$ do</p> <p style="padding-left: 20px;">If $\exists c \in [i]$ where $\mathbf{P}[c, i] = 1$ then $K_i \leftarrow K_c$</p> <p style="padding-left: 40px;">else $K_i \leftarrow \mathcal{F.KS}$</p> <p>For $(i, \mathbf{rt}) \in \text{SET}'$ do $\text{HS} \leftarrow \bigcup \mathcal{F.Ev}(K_i, \mathbf{rt})$</p> <p>While $\text{HS} < N$ do $x \leftarrow \{0, 1\}^{\mathcal{F.ol}}$; $\text{HS} \leftarrow \bigcup x$</p> <p>For $i = 1, \dots, n$ do</p> <p style="padding-left: 20px;">$\mathbf{tk}_i \leftarrow \text{QuerySim}(r_i, \mathbf{mt}, (K_1, \dots, K_{c_p}))$</p> <p>Return $((EM, \text{HS}), (\mathbf{tk}_1, \dots, \mathbf{tk}_n))$</p>	<p>Subroutine $\text{QuerySim}(r, (\mathbf{mt}_1, \dots, \mathbf{mt}_{c_q}), (K_1, \dots, K_{c_p}))$</p> <p>If $r = (\mathbf{m}, i)$ then return $(\mathbf{r}, \mathbf{mt}_i)$</p> <p>Else if $r = (\mathbf{p}, i, r_1)$ then</p> <p style="padding-left: 20px;">Return $(\mathbf{s}, K_i, \text{QuerySim}(r_1, \mathbf{mt}, \mathbf{k}))$</p> <p>Else if $r = (\mathbf{j}, r_1, r_2)$</p> <p style="padding-left: 20px;">For $i = 1, 2$ do</p> <p style="padding-left: 40px;">If $r_i = (\mathbf{m}, i)$ then $\mathbf{tk}_i \leftarrow (\mathbf{r}, \mathbf{mt}_i)$</p> <p style="padding-left: 40px;">Else if $r_i = (\mathbf{p}, j, r')$ then</p> <p style="padding-left: 60px;">$\mathbf{tk}' \leftarrow \text{QuerySim}(r', \mathbf{mt}, \mathbf{k})$; $\mathbf{tk}_i \leftarrow (\mathbf{s}, \mathbf{tk}', K_j)$</p> <p>Return $(\mathbf{j}, \mathbf{tk}_1, \mathbf{tk}_2)$</p>
<p>Alg $A_m(\mathbf{s})$</p> <p>$(\text{DS}, (q_1, \dots, q_n), st) \leftarrow \mathcal{A}(\mathbf{s})$</p> <p>Construct \mathbf{M}, SET as in $\text{PpSj.Enc}(\cdot, \text{DS})$</p> <p>For $i = 1, \dots, n$ do</p> <p style="padding-left: 20px;">$(r_i, \mathbf{q}, \mathbf{p}, c_q, c_p) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, c_q, c_p)$</p> <p>Return $(\mathbf{M}, \mathbf{q}, (\text{SET}, \mathbf{p}, (r_1, \dots, r_n), st))$</p> <p>Alg $A_m(\mathbf{g}, EM, \mathbf{mt}, (\text{SET}, \mathbf{p}, \mathbf{r}, st))$</p> <p>$(p_1, \dots, p_{c_p}) \leftarrow \mathbf{p}$</p> <p>$(r_1, \dots, r_n) \leftarrow \mathbf{r}$</p> <p>$K_f \leftarrow \mathcal{F.KS}$</p> <p>For $i = 1, \dots, c_p$ do</p> <p style="padding-left: 20px;">$K_i \leftarrow \mathcal{F.Ev}(K_f, p_i)$</p> <p>$\mathbf{k} \leftarrow (K_1, \dots, K_{c_p})$</p> <p>For $i \in [n]$ do</p> <p style="padding-left: 20px;">$\mathbf{tk} \leftarrow \bigcup \text{QuerySim}(r_i, \mathbf{mt}, \mathbf{k})$</p> <p>Return $A(\mathbf{g}, (EM, \text{HS}), \mathbf{tk}, st)$</p>	<p>Alg $\boxed{A_1^{\text{FN}}}, \boxed{A_2^{\text{FN}}}$</p> <p>$(\text{DS}, (q_1, \dots, q_n), st) \leftarrow \mathcal{A}(\mathbf{s})$</p> <p>Construct \mathbf{M}, SET as in $\text{PpSj.Enc}(\cdot, \text{DS})$</p> <p>For $i = 1, \dots, n$ do $(r_i, \mathbf{q}, \mathbf{p}, c_q, c_p) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, c_q, c_p)$</p> <p>$(EM, \mathbf{mt}) \leftarrow \mathcal{S}(\mathcal{L}(\mathbf{M}, \mathbf{q}))$</p> <p>$(p_1, \dots, p_{c_p}) \leftarrow \mathbf{p}$; $c \leftarrow \mathcal{C}[p]$; $\text{ctr} \leftarrow 1$</p> <p>For $(p, \mathbf{rt}) \in \text{SET}$ do</p> <div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p>If $K_p = \epsilon$ then $K_p \leftarrow \mathcal{F.N}(p)$</p> <p>$\text{HS} \leftarrow \bigcup \mathcal{F.Ev}(K_p, \mathbf{rt})$</p> </div> <div style="border: 1px dashed black; padding: 5px; margin: 5px 0;"> <p>If $K_p = \epsilon$ then</p> <p style="padding-left: 20px;">If $\text{ctr} < c$ or $p \in \mathbf{p}$ then $K_p \leftarrow \mathcal{F.KS}$; $\text{ctr} \leftarrow \text{ctr} + 1$</p> <p style="padding-left: 20px;">Else if $\text{ctr} = c$ then $K_p \leftarrow \perp$; $\text{ctr} \leftarrow \text{ctr} + 1$</p> <p style="padding-left: 20px;">If $K_p = \epsilon$ then $x \leftarrow \{0, 1\}^{\mathcal{F.ol}}$; $\text{HS} \leftarrow \bigcup x$</p> <p style="padding-left: 20px;">Else if $K_p = \perp$ then $\text{HS} \leftarrow \mathcal{F.N}(\mathbf{rt})$</p> <p style="padding-left: 20px;">Else $\text{HS} \leftarrow \bigcup \mathcal{F.Ev}(K_p, \mathbf{rt})$</p> </div> <p>For $i \in [n]$ do $\mathbf{tk} \leftarrow \bigcup \text{QuerySim}(r_i, \mathbf{mt}, (K_{p_1}, \dots, K_{p_{c_p}}))$</p> <p>Return $A(\mathbf{g}, (EM, \text{HS}), \mathbf{tk}, st)$</p>

Fig. 15. Leakage profile (top), simulator (middle) and adversaries (bottom) used in the proof of Theorem 3. In \mathcal{L}^P , $\text{RS}, \mathcal{L}, \text{HF}, \text{QP}$ compute the recursion structure leakage, Mme's leakage profile, hashset filtering results and hashset query pattern respectively, as discussed in Section 4.2. In \mathcal{S}^P , \mathcal{S} is a simulator for Mme. Note that when A_f (from Theorem 3) is run it randomly selects one of A_1, A_2 and runs it.

<p>Alg FpSj.Enc($K_m, (DB, \alpha, \mathbf{T})$)</p> <p>For all $(id, R) \in DB$ and $r \in R.T$ do $rt \leftarrow \mathbf{T}[(id, r[\text{uk}(id)])]$; $\mathbf{M}[(r, id)] \stackrel{\cup}{\leftarrow} (rt)$ For $at \in R.Ats$ where $at \neq \text{uk}(id)$ do $\mathbf{M}[(s, at, r[at])] \stackrel{\cup}{\leftarrow} (rt)$; $\mathbf{SET} \stackrel{\cup}{\leftarrow} ((s, at, r[at]), rt)$ For $(t_1, t_2) \in \alpha$ do $id_1 \leftarrow \text{getID}(t_1)$; $id_2 \leftarrow \text{getID}(t_2)$ For $r \in (DB[id_1] \bowtie_{t_1, t_2} DB[id_2]).T$ do For $i = 1, 2$ do $rt_i \leftarrow \mathbf{T}[(id_i, r[\text{uk}(id_i)])]$ $\mathbf{M}[(j, t_1, t_2)] \stackrel{\cup}{\leftarrow} (rt_1, rt_2)$; $\mathbf{M}_1[(t_1, t_2, rt_1, 1)] \stackrel{\cup}{\leftarrow} rt_2$ $\mathbf{M}_1[(t_1, t_2, rt_2, 2)] \stackrel{\cup}{\leftarrow} rt_1$; $\mathbf{SET} \stackrel{\cup}{\leftarrow} ((ij, t_1, t_2), (rt_1, rt_2))$ $(K_m, \mathbf{D}) \leftarrow \text{Mme}_{\pi}^{\text{rt}}.\text{Enc}(K_m, \mathbf{M})$ For $(t_1, t_2, rt, i) \in \mathbf{M}_1.Lbls$ do For $j = 0, 1$ do $K_j \leftarrow \mathbf{F.Ev}(\mathbf{F.Ev}(K_m, (ij, t_1, t_2, i)), rt j)$ $\{rt_1, \dots, rt_n\} \leftarrow \mathbf{M}_1[(t_1, t_2, rt, i)]$ For $k \in [n]$ do Pad rt_k to \mathbf{M}'s max. value length $\mathbf{D}[\mathbf{F.Ev}(K_0, k)] \leftarrow \text{SE.Enc}(K_1, rt_k)$ $K_f \leftarrow \text{F.KS}$; $\mathbf{HS} \leftarrow \text{HsEnc}(K_f, \mathbf{SET})$ Return $((\text{Schema}(DB), K_m, K_f), (\mathbf{D}, \mathbf{HS}))$</p> <p>Alg FpSj.Tok(K_i, q)</p> <p>$(\text{scma}, K_m, K_f) \leftarrow K_i$ If $q = (r, id)$ then return $(r, \text{Mme}_{\pi}^{\text{rt}}.\text{Tok}(K_m, (r, id)))$ Else if $q = (s, at, x, (r, id))$ then Return $(r, \text{Mme}_{\pi}^{\text{rt}}.\text{Tok}(K_m, (s, at, x)))$ Else if $q = (s, at, x, q_1)$ then Return $(s, \mathbf{F.Ev}(K_f, (s, at, x)), \text{FpSj.Tok}(K_i, q_1))$ Else if $q = (j, t_1, t_2, (r, id_1), (r, id_2))$ then Return $(r, \text{Mme}_{\pi}^{\text{rt}}.\text{Tok}(K_m, (j, t_1, t_2)))$ Else if $q = (j, t_1, t_2, q_1, (r, id))$ then Return $(ij, \mathbf{F.Ev}(K_m, (ij, t_1, t_2, 1)), 1, \text{FpSj.Tok}(K_i, q_1))$ Else if $q = (j, t_1, t_2, (r, id), q_1)$ then Return $(ij, \mathbf{F.Ev}(K_m, (ij, t_1, t_2, 2)), 2, \text{FpSj.Tok}(K_i, q_1))$ Else if $q = (j, t_1, t_2, q_1, q_2)$ then For $i = 1, 2$ do $tk_i \leftarrow \text{FpSj.Tok}(K_i, q_i)$ Return $(ij, \mathbf{F.Ev}(K_f, (ij, t_1, t_2)), tk_1, tk_2)$</p>	<p>Alg FpSj.Eval($tk, (\mathbf{D}, \mathbf{HS})$)</p> <p>If $tk = (r, mt)$ then Return $(\text{Mme}_{\pi}^{\text{rt}}.\text{Eval}(mt, \mathbf{D}))$ Else if $tk = (s, K, tk_1)$ then $P \leftarrow \text{FpSj.Eval}(tk_1, (\mathbf{D}, \mathbf{HS}))$ Return $\text{HsFilter}(K, P, \mathbf{HS})$ Else if $tk = (ij, K, i, tk_1)$ then $(P_1) \leftarrow \text{FpSj.Eval}(tk_1, (\mathbf{D}, \mathbf{HS}))$ For $rt \in \mathbf{rt} \in P_1$ do For $j = 0, 1$ do $K_j \leftarrow \mathbf{F.Ev}(K, rt j)$ $x \leftarrow \mathbf{D}[\mathbf{F.Ev}(K_0, c_{rt})]$ While $x \neq \perp$ do $S \stackrel{\cup}{\leftarrow} (\text{SE.Dec}(K_1, x), \mathbf{rt})$ $c_{rt} \leftarrow c_{rt} + 1$ $x \leftarrow \mathbf{D}[\mathbf{F.Ev}(K_0, c_{rt})]$ If $i = 1$ then $P_1 \leftarrow \{(rt) \mathbf{rt} : (rt, \mathbf{rt}) \in S\}$ Else $P_1 \leftarrow \{\mathbf{rt} \mathbf{rt} : (rt, \mathbf{rt}) \in S\}$ Return (P_1) Else if $tk = (ij, K, tk_1, tk_2)$ then For $i = 1, 2$ do $(P_i) \leftarrow \text{FpSj.Eval}(tk_i, (\mathbf{D}, \mathbf{HS}))$ For $\mathbf{rt}_1 \in P_1$ and $\mathbf{rt}_2 \in P_2$ do For $rt_1 \in \mathbf{rt}_1$ and $rt_2 \in \mathbf{rt}_2$ do If $\mathbf{F.Ev}(K, (rt_1, rt_2)) \in \mathbf{HS}$ then $P_0 \stackrel{\cup}{\leftarrow} \mathbf{rt}_1 \mathbf{rt}_2$ Return (P_0)</p> <p>Alg FpSj.Fin($K_i, q, (M_1)$)</p> <p>$(\text{scma}, K_m, K_f) \leftarrow K_i$ Using scma and q, compute the attributes \mathbf{at} in $\text{SjDT.Spec}(q, \mathbf{DS})$ $\mathbf{R} \leftarrow \text{NewRltn}(\mathbf{at})$ For $(m_1, \dots, m_n) \in M_1$ do $\mathbf{R.T} \stackrel{\cup}{\leftarrow} m_1 \dots m_n$ Return \mathbf{R}</p>
---	--

Fig. 16. Algorithms for FpSj, the StI scheme for SjDT using FP indexing.

equality pattern and filtering results of the predicates) with the only subtlety coming from the fact that FpSj associates the join predicate with a pair of row tokens, thereby leakage the equality pattern of the join (but restricted to the rows that have been retrieved by the subqueries).

This leaves the leakage from internal intermediate joins. Recall that the indexing of such joins involved manual additions to the output of $\text{Mme}_{\pi}^{\text{rt}}.\text{Enc}$ (i.e. \mathbf{D}). As such, the leakage algorithm must include information to simulate these entries. This includes the final number of values in \mathbf{D} (i.e. M), the length of these values (i.e. ℓ), the query pattern of such joins (i.e. $\text{QP}(\mathbf{j})$), query responses to these (i.e. \mathbf{I}) and the number of such joins (i.e. c_j).

SECURITY. The security of FpSj Its security is given in Theorem 4 below, which uses the simulator \mathcal{S}^f given in Fig. 17.

Theorem 4. Let \mathcal{L}, \mathcal{S} be the leakage algorithm and simulator for $\text{Mme}_{\pi}^{\text{rt}}$ respectively (given in [12]). Let \mathbf{F}, SE be the primitives used in $\text{Mme}_{\pi}^{\text{rt}}$ and FpSj's algorithms. Let $\mathcal{L}^f, \mathcal{S}^f$ be as defined in Fig. 17 and Fig. 17

respectively. Then for all adversaries A there exists adversaries A_f, A_s such that:

$$\mathbf{Adv}_{\text{PpSj}, \mathcal{L}^p, \mathcal{S}^p}^{\text{ss}}(A) \leq (m + m_1) \mathbf{Adv}_{\text{SE}}^{\text{ind}^s}(A_s) + (m + m_1 + p + 1).$$

Here, m, m_1 are the number of labels in \mathbf{M}, \mathbf{M}_1 respectively and p is the number of distinct predicates used in constructing HS.

Proof. This proof is quite similar to Theorem 3, except that we can now reduce security straight to SE, F because we assumed the use of the $\text{Mme}_{\pi}^{\text{rr}}$ multimap encryption scheme. We therefore omit a full description of the adversaries and proof except to say that the multiplicative factors in the bound come from the number of SE and F keys that are used in FpSj.Enc (including those which are an output of F.Ev).

F HybStl Details

PSEUDOCODE. HybStl 's algorithms merge the techniques used in FpSj, PpSj to support hybrid queries from HybDT . The pseudocode is given in Fig. 18. Note that HybFinalize is a subroutine recursively called by HybStl.Fin to perform client-side (i.e. PP) joins.

As mentioned in Section 5 the only subtlety in this handling is that the ordering of attributes in the tuple sets passed into HybStl.Fin may not be consistent with the desired output relation. This edge case may occur when a recursive join query is made which makes use of both types of indexing. This can be easily remedied because HybStl.Fin has access to schema scma and query q . Therefore, it can compute the desired order of attributes in the output relation.

LEAKAGE. For completeness, the leakage profile of HybStl is described via pseudocode in Fig. 19. This leakage algorithm merges our two previous ones (i.e. \mathcal{L}^f and \mathcal{L}^p) in the straightforward way. In particular, the only difference between \mathcal{L}^f (Fig. 17) and \mathcal{L}^h is the recursion structure of partially precomputed joins which are handled in the style of \mathcal{L}^p (Fig. 15).

SECURITY. As alluded to above, the proof of HybStl 's security (with respect to \mathcal{L}^h) is very similar to the result for Theorem 4. As such, we state the security bound below but omit the proof for brevity.

Theorem 5. *Let \mathcal{L} be the leakage algorithm and simulator for $\text{Mme}_{\pi}^{\text{rr}}$ (given in [12]) and \mathcal{L}^h be as defined in Fig. 19. Let F, SE be the primitives used in $\text{Mme}_{\pi}^{\text{rr}}$ and HybStl 's algorithms. Then for all adversaries A there exists A_f, A_s, \mathcal{S}^h such that:*

$$\mathbf{Adv}_{\text{HybStl}, \mathcal{L}^h, \mathcal{S}^h}^{\text{ss}}(A) \leq (m + m_1) \mathbf{Adv}_{\text{SE}}^{\text{ind}^s}(A_s) + (m + m_1 + p + 1).$$

Here, m, m_1 are the number of labels in \mathbf{M}, \mathbf{M}_1 respectively and p is the number of distinct predicates used in constructing HS.

G Supplementary Content for Simulations

The Data Portal stores each Chicago relation separately and intends each relation to be useful on its own – independent from the other relations. The Sakila database also has 15 relations, with a total of 46,238 rows and 88 attributes. Unlike the Chicago database, the Sakila relations have a clear logical structure in the schema such that each relation has a role defined relative to the other relations. Details about the Sakila schema can be found on the MySQL website. By including one example database without a structured schema and one with, we hope to model two different use cases – one where the DBA treats each relation as existing independently and one where the DBA carefully pre-plans the entire organization.

To give an idea of the data distribution in our data sets, we give some summary statistics about each in Table 20. We report each relation's name, number of attributes, number of rows, and minimum, mean, and maximum attribute densities. The density of an attribute is the average occurrence frequency of the values in that column. In other words, for relation $\text{DB}[id]$ at's attribute density is $\frac{|\text{rng}(at)|}{|\text{DB}[id].\bar{T}|}$.

<p>Alg $\mathcal{L}^f(DS, (q_1, \dots, q_n))$ Construct $\mathbf{M}, \mathbf{M}_1, \text{SET}$ as in $\text{FpSj.Enc}(\cdot, DS)$ For $i = 1, \dots, n$ do $(r_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$ $\mathbf{r} \leftarrow (r_1, \dots, r_n)$; $lk \leftarrow \mathcal{L}(\mathbf{M}, \mathbf{q})$; $S \leftarrow \bigcup_{q \in \mathbf{q}} \mathbf{M}[q]$ While $S \neq S'$ do $S \leftarrow S'$; $(\text{SET}, S', \mathbf{I}) \leftarrow \text{IJ}(\mathbf{j}, S', \mathbf{I})$ $\text{SET}' \leftarrow \text{HF}(\mathbf{p}, S, \text{SET})$ Define M : # of vals in \mathbf{M} and \mathbf{M}_1 Define ℓ : max. length val in \mathbf{M}, \mathbf{M}_1 Return $(\mathbf{r}, lk, \text{QP}(\mathbf{p}), c_p, \text{SET}', \text{SET}' , \mathbf{I}, M, \ell, \text{QP}(\mathbf{j}), c_j)$</p> <p>Subroutine $\text{IJ}((j_1, \dots, j_n), S, \mathbf{I})$ For $i \in [n]$ $\text{rt} \in S$ do $(\mathbf{t}_1, \mathbf{t}_2, k) \leftarrow j_i$ If $\mathbf{M}_1[(\mathbf{t}_1, \mathbf{t}_2, \text{rt}, i)] \neq \perp$ then $S' \leftarrow S' \cup \mathbf{M}_1[(\mathbf{t}_1, \mathbf{t}_2, \text{rt}, i)]$; $\mathbf{I}[(\text{rt}, i)] \leftarrow \mathbf{M}_1[(\mathbf{t}_1, \mathbf{t}_2, \text{rt}, i)]$ Return (S', \mathbf{I})</p> <p>Subroutine $\text{RS}(q, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$ If $q = (\mathbf{r}, id)$ then $\mathbf{q} \leftarrow (\mathbf{r}, id)$; $r \leftarrow (m, c_q)$; $c_q \leftarrow c_q + 1$ Else if $q = (\mathbf{s}, at, x, (\mathbf{r}, id))$ then $\mathbf{q} \leftarrow (\mathbf{s}, at, x)$; $r \leftarrow (m, c_q)$; $c_q \leftarrow c_q + 1$ Else if $q = (\mathbf{s}, at, x, q_1)$ then $(r_1, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_1, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$; $\mathbf{p} \leftarrow (\mathbf{s}, at, x)$; $r \leftarrow (\mathbf{s}, c_p, r_1)$; $c_p \leftarrow c_p + 1$ Else if $q = (\mathbf{t}_1, \mathbf{t}_2, (\mathbf{r}, id_1), (\mathbf{r}, id_2))$ then $\mathbf{q} \leftarrow (\mathbf{j}, \mathbf{t}_1, \mathbf{t}_2)$; $r \leftarrow (m, c_q)$; $c_q \leftarrow c_q + 1$ Else if $q = (\mathbf{t}_1, \mathbf{t}_2, q_1, (\mathbf{r}, id))$ or $q = (\mathbf{t}_1, \mathbf{t}_2, (\mathbf{r}, id), q_1)$ then $(r_1, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_1, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$; If $q = (\mathbf{t}_1, \mathbf{t}_2, q_1, (\mathbf{r}, id))$ then $i = 1$ else $i = 2$ $\mathbf{j} \leftarrow (\mathbf{t}_1, \mathbf{t}_2, i)$; $r \leftarrow (ij, c_j, i, r_1)$; $c_j \leftarrow c_j + 1$ Else if $q = (\mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$ then For $i = 1, 2$ do $(r_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$ $\mathbf{p} \leftarrow (ij, \mathbf{t}_1, \mathbf{t}_2)$; $r \leftarrow (ij, c_p, r_1, r_2)$; $c_p \leftarrow c_p + 1$ Return $(r, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p)$</p>	<p>Subroutine $\text{HF}(\mathbf{p}, S, \text{SET})$ $(p_1, \dots, p_n) \leftarrow \mathbf{p}$ For all $i \in [n]$ If $p_i = (ij, \mathbf{t}_1, \mathbf{t}_2)$ then For $\text{rt} \in S \times S$ do If $(p_i, \text{rt}) \in \text{SET}$ then $\text{SET}' \leftarrow \bigcup (i, \text{rt})$ Else For $\text{rt} \in S$ do If $(p_i, \text{rt}) \in \text{SET}$ then $\text{SET}' \leftarrow \bigcup (i, \text{rt})$ Return SET'</p> <p>Subroutine $\text{QP}((t_1, \dots, t_n))$ For all $i, j \in [n]$ if $t_i = t_j$ then $\mathbf{T}[i, j] \leftarrow 1$ else $\mathbf{T}[i, j] \leftarrow 0$ Return \mathbf{T}</p>
<p>Alg $\mathcal{S}^f(\mathbf{r}, lk, \mathbf{P}, c_p, \text{SET}', N, \mathbf{I}, M, \ell, \mathbf{J}, c_j)$ $(r_1, \dots, r_n) \leftarrow \mathbf{r}$; $(\mathbf{D}, \mathbf{mt}) \leftarrow \mathcal{S}(lk)$ For $i = 1, \dots, c_p$ do If $\exists c \in [i]$ where $\mathbf{P}[c, i] = 1$ then $K_i \leftarrow K_c$ else $K_i \leftarrow \mathbf{F.KS}$ For $(i, x) \in \text{SET}'$ do $\text{HS} \leftarrow \mathbf{F.Ev}(K_i, x)$ While $\text{HS} < N$ do $x \leftarrow \{0, 1\}^{\text{F.ol}}$; $\text{HS} \leftarrow \bigcup x$ For $(\text{rt}, i) \in \mathbf{I.Lbls}$ do If $\exists c \in [i]$ where $\mathbf{J}[c, i] = 1$ then $K'_i \leftarrow K'_c$ else $K'_i \leftarrow \mathbf{F.KS}$ $\{\text{rt}'_1, \dots, \text{rt}'_n\} \leftarrow \mathbf{I}[(\text{rt}, i)]$; For $k = 0, 1$ do $K''_k \leftarrow \mathbf{F.Ev}(K'_i, \text{rt} k)$ For $k \in [n]$ do Pad rt'_k to length ℓ then $\mathbf{D}[\mathbf{F.Ev}(K''_0, k)] \leftarrow \text{SE.Enc}(K''_1, \text{rt}'_k)$ While $\mathbf{D.Lbls} < M$ do $x \leftarrow \{0, 1\}^{\text{F.ol}}$; $\mathbf{D}[x] \leftarrow \{0, 1\}^\ell$ For $i = 1, \dots, n$ do $\text{tk}_i \leftarrow \text{QuerySim}(r_i, \mathbf{mt}, (K_1, \dots, K_{c_p}), (K'_1, \dots, K'_{c_j}))$ Return $(\mathbf{D}, \text{HS}), (\text{tk}_1, \dots, \text{tk}_n)$</p>	<p>Subroutine $\text{QuerySim}(r, \mathbf{mt}, \mathbf{k}, \mathbf{k}')$ $(\text{mt}_1, \dots, \text{mt}_{c_q}) \leftarrow \mathbf{mt}$ $(K_1, \dots, K_{c_p}) \leftarrow \mathbf{k}$; $(K'_1, \dots, K'_{c_j}) \leftarrow \mathbf{k}'$ If $r = (m, i)$ then return $(\mathbf{r}, \text{mt}_i)$ Else if $r = (\mathbf{p}, i, r_1)$ then $\text{tk} \leftarrow \text{QuerySim}(r_1, \mathbf{mt}, \mathbf{k}, \mathbf{k}')$ Return $(\mathbf{s}, K_i, \text{tk})$ Else if $r = (ij, i, j, r_1)$ then $\text{tk} \leftarrow \text{QuerySim}(r_1, \mathbf{mt}, \mathbf{k}, \mathbf{k}')$ Return (ij, K'_i, j, tk) Else if $r = (ij, i, r_1, r_2)$ then For $j = 1, 2$ do $\text{tk}_j \leftarrow \text{QuerySim}(r_j, \mathbf{mt}, \mathbf{k}, \mathbf{k}')$ Return $(ij, K_i, \text{tk}_1, \text{tk}_2)$</p>

Fig. 17. Top: Leakage profile for FpSj. Here, \mathcal{L} is the leakage algorithm for $\text{Mme}_{\pi}^{\text{rf}}$ and subroutines IJ, HF, QP, RS compute the leakage associated to internal joins, hashset filtering, hashset query patterns and query recursion structures, as discussed in Section 4.2 and Appendix E. Bottom: Simulator used in the proof of Theorem 4 where \mathcal{S} is a simulator for $\text{Mme}_{\pi}^{\text{rf}}$.

Alg HybStl.Enc($K_m, (DB, \alpha, T)$)

For all $(id, R) \in DB$ and $r \in R.T$ do
 $rt \leftarrow T[(id, r[uk(id)])]$; $M[(r, id)] \stackrel{\leftarrow}{\leftarrow} (rt)$
For $at \in R.Ats$ where $at \neq uk(id)$ do
 $M[(s, at, r[at])] \stackrel{\leftarrow}{\leftarrow} (rt)$; $SET \stackrel{\leftarrow}{\leftarrow} ((s, at, r[at]), rt)$
For $(t_1, t_2) \in \alpha$ do
 $id_1 \leftarrow getID(t_1)$; $id_2 \leftarrow getID(t_2)$
For $r \in (DB[id_1] \bowtie_{t_1, t_2} DB[id_2]).T$ do
For $i = 1, 2$ do
 $rt_i \leftarrow T[(id_i, r[uk(id_i)])]$
 $M[(pp, t_1, t_2, i)] \stackrel{\leftarrow}{\leftarrow} (rt_i)$
 $SET \stackrel{\leftarrow}{\leftarrow} ((pp, t_1, t_2, i), rt_i)$
 $M[(fp, t_1, t_2)] \stackrel{\leftarrow}{\leftarrow} (rt_1, rt_2)$
 $M_1[(t_1, t_2, rt_1, 1)] \stackrel{\leftarrow}{\leftarrow} rt_1$; $M_1[(t_1, t_2, rt_2, 2)] \stackrel{\leftarrow}{\leftarrow} rt_2$
 $SET \stackrel{\leftarrow}{\leftarrow} ((fp, t_1, t_2), (rt_1, rt_2))$
 $(K_m, D) \leftarrow s Mme_{\pi}^{rr}.Enc(K_m, M)$
For $(t_1, t_2, rt, i) \in M_1.Lbls$ do
For $j = 0, 1$ do
 $K_j \leftarrow F.Ev(F.Ev(K_m, (ij, t_1, t_2, i)), rt[j])$
 $\{rt_1, \dots, rt_n\} \leftarrow M_1[(t_1, t_2, rt, i)]$
For $k \in [n]$ do
Pad rt_k to M 's max. value length
 $D[F.Ev(K_0, k)] \leftarrow s SE.Enc(K_1, rt_k)$
 $K_f \leftarrow s F.KS$; $HS \leftarrow HsEnc(K_f, SET)$
Return $((Schema(DB), K_m, K_f), (D, HS))$

Alg HybStl.Tok(K_i, q)

$(scma, K_m, K_f) \leftarrow K_i$
If $q = (r, id)$ then return $(r, Mme_{\pi}^{rr}.Tok(K_m, (r, id)))$
Else if $q = (s, at, x, (r, id))$ then
Return $(r, Mme_{\pi}^{rr}.Tok(K_m, (s, at, x)))$
Else if $q = (s, at, x, q_1)$ then
Return $(s, F.Ev(K_f, (s, at, x)), HybStl.Tok(K_i, q_1))$
Else if $q = (pp, t_1, t_2, q_1, q_2)$ then
For $i = 1, 2$ do
If $q_i = (r, id_i)$ then
 $tk_i \leftarrow s (r, Mme_{\pi}^{rr}.Tok(K_m, (pp, t_1, t_2, i)))$
Else
 $tk' \leftarrow s HybStl.Tok(K_i, q_i)$
 $tk_i \leftarrow (s, F.Ev(K_f, (pp, t_1, t_2, i)), tk')$
Return (pp, tk_1, tk_2)
Else if $q = (fp, t_1, t_2, (r, id_1), (r, id_2))$ then
Return $(r, Mme_{\pi}^{rr}.Tok(K_m, (fp, t_1, t_2)))$
Else if $q = (fp, t_1, t_2, q_1, (r, id))$ then
Return $(ij, F.Ev(K_m, (ij, t_1, t_2, 1)), HybStl.Tok(K_i, q_1))$
Else if $q = (fp, t_1, t_2, (r, id), q_1)$ then
Return $(ij, F.Ev(K_m, (ij, t_1, t_2, 2)), HybStl.Tok(K_i, q_1))$
Else if $q = (fp, t_1, t_2, q_1, q_2)$ then
For $i = 1, 2$ do $tk_i \leftarrow s HybStl.Tok(K_i, q_i)$
Return $(ij, F.Ev(K_f, (fp, t_1, t_2)), tk_1, tk_2)$

Alg HybStl.Eval(tk, IX)

$(D, HS) \leftarrow IX$
If $tk = (r, tk_1)$ then return $(Mme_{\pi}^{rr}.Eval(tk, IX))$
Else If $tk = (s, K, tk_1)$ then
Return $HsFilter(K, HybStl.Eval(tk_1, IX), HS)$
Else if $tk = (pp, tk_1, tk_2)$
Return $HybStl.Eval(tk_1, IX) \parallel HybStl.Eval(tk_2, IX)$
Else if $tk = (ij, K, tk_1)$ then
 $(P_1, \dots, P_n) \leftarrow HybStl.Eval(tk_1, IX)$
Define $j : \exists rt \in \mathbf{rt} \in P_j$
where $D[F.Ev(F.Ev(K, rt|0), 0)] \neq \perp$
For $rt \in \mathbf{rt} \in P_j$ do
For $i = 1, 2$ do $K_i \leftarrow F.Ev(K, rt|i)$
While $D[F.Ev(K_0, c_{rt})] \neq \perp$ do
 $rt \stackrel{\leftarrow}{\leftarrow} SE.Dec(K_1, D[F.Ev(K_0, c_{rt})])$
 $P' \stackrel{\leftarrow}{\leftarrow} \mathbf{rt}$; $c_{rt} \leftarrow c_{rt} + 1$
Return $P \setminus \{P_j\} \parallel (P')$
Else if $tk = (ijj, K, tk_1, tk_2)$ then
For $i = 1, 2$ do
 $(P_1^i, \dots, P_{n_i}^i) \leftarrow HybStl.Eval(tk_i, (D, HS))$
Define $j_1, j_2 : \exists rt_i \in \mathbf{rt}_i \in P_{j_i}^i$
where $F.Ev(K, (rt_1, rt_2)) \in HS$
For $rt_1 \in \mathbf{rt}_1 \in P_{j_1}^1$ and $rt_2 \in \mathbf{rt}_2 \in P_{j_2}^2$ do
If $F.Ev(K, (rt_1, rt_2)) \in HS$ then $P' \stackrel{\leftarrow}{\leftarrow} \mathbf{rt}_1 \parallel \mathbf{rt}_2$
If $P' \neq \emptyset$ then return $P_1 \setminus \{P_{j_1}^1\} \parallel P_2 \setminus \{P_{j_2}^2\} \parallel (P')$

Alg HybStl.Fin($(scma, K_m, K_f), q, (M_1, \dots, M_n)$)

Using $scma$ and q , compute at_1, \dots, at_n, at , the attributes in $M_1, \dots, M_n, HybDT.Spec(q, DS)$ respectively
For $i \in [n]$ do
 $R_i \leftarrow NewRltn(at_i)$
 $R_i.T \leftarrow \{m_1 \parallel \dots \parallel m_{n'} : (m_1, \dots, m_{n'}) \in M_i\}$
 $R \leftarrow HybFinalize(q, (R_1, \dots, R_n))$
If $R.Ats \neq at$ then reorder attributes in R accordingly
Return R

Subroutine HybFinalize(q, R)

If $q = (r, id)$ then $(R) \leftarrow R$; Return R
Else if $q = (s, at, x, q_1)$ then return $HybFinalize(q_1, R)$
Else if $q = (fp, t_1, t_2, q_1, q_2)$ then
 $(R_1, \dots, R_n) \leftarrow R$
Define $j : t_1 \cup t_2 \subseteq R_j.Ats$
Partition $R \setminus \{R_j\}$ into R_1, R_2 where R_i contains all the attributes in $HybDT(q_i, DS)$
 $R \leftarrow HybFinalize(q_1, R_1 \parallel (R_j))$
 $R \leftarrow HybFinalize(q_2, R_2 \parallel (R))$
Else if $q = (pp, t_1, t_2, q_1, q_2)$ then
Partition R into R_1, R_2 where R_i contains all attributes in $HybDT(q_i, DS)$
Return $HybFinalize(q_1, R_1) \bowtie_{t_1, t_2} HybFinalize(q_2, R_2)$

Fig. 18. Algorithms for HybStl, the StI scheme for HybDT using hybrid indexing. HybFinalize is a recursively called subroutine used in HybStl.Fin.

Alg $\mathcal{L}^h(\text{DS}, (q_1, \dots, q_n))$
Construct $\mathbf{M}, \mathbf{M}_1, \text{SET}$ as in $\text{FpSj.Enc}(\cdot, \text{DS})$
For $i = 1, \dots, n$ do $(r_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$
 $\mathbf{r} \leftarrow (r_1, \dots, r_n)$; $lk \leftarrow_s \mathcal{L}(\mathbf{M}, \mathbf{q})$; $S \leftarrow \bigcup_{q \in \mathbf{q}} \mathbf{M}[q]$; While $S \neq S'$ do $S \leftarrow S'$; $(\text{SET}, S', \mathbf{I}) \leftarrow \text{IJ}(\mathbf{j}, S', \mathbf{I})$
Define M : # of vals in \mathbf{M} and \mathbf{M}_1
Define ℓ : max. length val in \mathbf{M}, \mathbf{M}_1
Return $(\mathbf{r}, lk, \text{QP}(\mathbf{p}), c_p, \text{HF}(\mathbf{p}, S, \text{SET}), |\text{SET}'|, \mathbf{I}, M, \ell, \text{QP}(\mathbf{j}), c_j)$

Subroutine $\text{RS}(q, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$
If $q = (\mathbf{r}, id)$ then $\mathbf{q} \stackrel{\cup}{\leftarrow} (\mathbf{r}, id)$; $r \leftarrow (m, c_q)$; $c_q \leftarrow c_q + 1$
Else if $q = (\mathbf{s}, at, x, (\mathbf{r}, id))$ then $\mathbf{q} \stackrel{\cup}{\leftarrow} (\mathbf{s}, at, x)$; $r \leftarrow (m, c_q)$; $c_q \leftarrow c_q + 1$
Else if $q = (\mathbf{s}, at, x, q_1)$ then
 $(r_1, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_1, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$; $\mathbf{p} \stackrel{\cup}{\leftarrow} (\mathbf{s}, at, x)$; $r \leftarrow (\mathbf{s}, c_p, r_1)$; $c_p \leftarrow c_p + 1$
Else if $q = (\mathbf{pp}, \mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$ then
For $i = 1, 2$ do
If $q_i = (\mathbf{r}, id)$ then
 $\mathbf{q} \stackrel{\cup}{\leftarrow} (\mathbf{pp}, \mathbf{t}_1, \mathbf{t}_2, i)$; $r_i \leftarrow (m, c_q)$; $c_q \leftarrow c_q + 1$
Else
 $(r'_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$; $\mathbf{p} \stackrel{\cup}{\leftarrow} (\mathbf{pp}, \mathbf{t}_1, \mathbf{t}_2, i)$; $r_i \leftarrow (\mathbf{p}, c_p, r'_i)$; $c_p \leftarrow c_p + 1$
 $r \leftarrow (\mathbf{pp}, r_1, r_2)$
Else if $q = (\mathbf{fp}, \mathbf{t}_1, \mathbf{t}_2, (\mathbf{r}, id_1), (\mathbf{r}, id_2))$ then $\mathbf{q} \stackrel{\cup}{\leftarrow} (\mathbf{fp}, \mathbf{t}_1, \mathbf{t}_2)$; $r \leftarrow (m, c_q)$; $c_q \leftarrow c_q + 1$
Else if $q = (\mathbf{fp}, \mathbf{t}_1, \mathbf{t}_2, q_1, (\mathbf{r}, id))$ or $q = (\mathbf{fp}, \mathbf{t}_1, \mathbf{t}_2, (\mathbf{r}, id), q_1)$ then
 $(r_1, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_1, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$; If $q = (\mathbf{fp}, \mathbf{t}_1, \mathbf{t}_2, q_1, (\mathbf{r}, id))$ then $i = 1$ else $i = 2$
 $\mathbf{j} \stackrel{\cup}{\leftarrow} (\mathbf{t}_1, \mathbf{t}_2, i)$; $r \leftarrow (\mathbf{fp}, c_j, i, r_1)$; $c_j \leftarrow c_j + 1$
Else if $q = (\mathbf{fp}, \mathbf{t}_1, \mathbf{t}_2, q_1, q_2)$ then
For $i = 1, 2$ do $(r_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j) \leftarrow \text{RS}(q_i, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$
 $\mathbf{p} \stackrel{\cup}{\leftarrow} (\mathbf{fp}, \mathbf{t}_1, \mathbf{t}_2)$; $r \leftarrow (\mathbf{ii}, c_p, r_1, r_2)$; $c_p \leftarrow c_p + 1$
Return $(r, \mathbf{q}, \mathbf{p}, \mathbf{j}, c_q, c_p, c_j)$

Fig. 19. Leakage algorithm used in Theorem 5, the proof of security for hybrid StI scheme HybStI. The subroutines HF, IJ, QP are given in Fig. 17.

Chicago R name	R.Ats	R.T	at densities		
			Min	Ave	Max
Bike_Racks	12	5164	4e-4	0.53	1.0
Census_Data	9	78	0.72	0.91	1.0
Crimes_2019	30	1.7e5	6e-6	0.17	1.0
Employee_Debt	7	1.4e4	3e-3	0.10	0.46
Fire_Stations	7	92	0.01	0.65	1.0
Graffiti	5	67	0.09	0.68	1.0
Housing	14	915	0.03	0.32	0.56
IUCR_Codes	4	404	5e-3	0.51	1.0
Land_Inventory	19	2.0e5	3e-4	0.41	1.0
Libraries	9	81	0.01	0.63	1.0
Lobbyists	7	1537	0.03	0.22	0.66
Police_Stations	15	23	0.04	0.87	1.0
Reloc_Vehicles	20	2672	2e-3	0.51	1.0
Street_Names	7	2582	2e-3	0.31	1.0
Towed_Vehicles	10	3339	1e-3	0.19	0.99

Sakila R name	R.Ats	R.T	at densities		
			Min	Ave	Max
Address	8	603	2e-3	0.93	1.0
Actor	4	208	0.02	0.56	1.0
Category	3	16	0.06	0.69	1.0
City	4	600	2e-3	0.55	1.0
Country	3	109	0.01	0.67	1.0
Customer	10	599	2e-3	0.5	1.0
Film	14	1002	1e-3	0.37	1.0
Film_Actor	3	5462	1e-4	0.07	0.18
Film_Cat	3	1000	1e-3	0.34	1.0
Inventory	4	5481	2e-4	0.30	1.0
Language	3	6	0.17	0.72	1.0
Payment	7	1.6e4	1e-4	0.44	1.0
Staff	11	2	0.5	0.86	1.0
Store	4	2	0.5	0.88	1.0
Rental	7	1.6e4	1e-4	0.47	1.0

Fig. 20. Summary statistics for the Chicago (left) and Sakila (right) data used in our simulations.