

Towards Practical Oblivious Join Processing

Zhao Chang , Dong Xie, Sheng Wang, Feifei Li, and Yulong Shen 

Abstract—In cloud computing, remote accesses over the cloud data inevitably bring the issue of trust. Despite strong encryption schemes, adversaries can still learn sensitive information from encrypted data by observing data access patterns. Oblivious RAMs (ORAMs) are proposed to protect against access pattern attacks. However, directly deploying ORAM constructions in an encrypted database brings large computational overhead. In this work, we focus on oblivious joins over a cloud database. Existing studies in the literature are restricted to either primary-foreign key joins or binary equi-joins. Our major contribution is to support general band joins and multiway equi-joins. For oblivious join without ORAMs, we extend the existing binary equi-join algorithm to support general band joins obliviously. For oblivious join with ORAMs, we integrate B -tree indices into ORAMs for each input table and retrieve blocks through the indices in join processing. The key point is to avoid retrieving tuples that make no contribution to the final join result and bound the number of accesses to each B -tree index. The effectiveness and efficiency of our algorithms are demonstrated through extensive evaluations over real-world datasets. Our method shows orders of magnitude speedup for oblivious multiway equi-joins in comparison with baseline algorithms.

Index Terms—Data privacy, oblivious index, oblivious join, oblivious RAM.

I. INTRODUCTION

MANY cloud service providers offer cloud-based database systems such as Amazon RDS and Redshift, Azure SQL, and Google Cloud SQL. Data encryption is a necessary step for keeping sensitive information secure and private on a cloud. To that end, encrypted databases such as Cipherbase [1], [2], CryptDB [3], TrustedDB [4], SDB [5], and Monomi [6], as well as related query execution techniques [7], [8], [9], [10] have been developed. But query access patterns still pose a privacy threat and leak sensitive information [11], [12], [13], [14]. It is possible to analyze the importance of different areas in the database, e.g., by counting the frequency of accessing data items [15], [16], [17], [18]. With background knowledge, the server may learn a lot about user queries and/or data [11], [19], [20].

Manuscript received 25 October 2022; revised 4 May 2023; accepted 26 August 2023. Date of publication 1 September 2023; date of current version 8 March 2024. This work was supported in part by the National Key R&D Program of China under Grant 2021YFB3101100 and in part by the National Natural Science Foundation of China under Grants 62220106004, 61972308, and 62302368. This work was partially done when Zhao Chang worked as a visiting scholar at Alibaba DAMO Academy. Recommended for acceptance by M. A. Cheema. (Corresponding author: Yulong Shen.)

Zhao Chang and Yulong Shen are with the Xidian University, Xi'an 710126, China (e-mail: changzhao@xidian.edu.cn; ylshen@mail.xidian.edu.cn).

Dong Xie is with the The Pennsylvania State University, State College, PA 16801 USA (e-mail: dongx@psu.edu).

Sheng Wang and Feifei Li are with the Alibaba Group, Hangzhou 311121, China (e-mail: sh.wang@alibaba-inc.com; lifeifei@alibaba-inc.com).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TKDE.2023.3310038>, provided by the authors.

Digital Object Identifier 10.1109/TKDE.2023.3310038




Fig. 1. Strawman solution to oblivious many-to-many join.

Oblivious RAMs (ORAMs) [21], [22], [23] allow the client to access encrypted data on a server without revealing her access patterns. However, most ORAM constructions are still too expensive to be deployed in a large database [11]. Recent studies [14], [24], [25], [26], [27], [28], [29] also explore building oblivious data structures or indices over encrypted data, but none of them support complex queries (e.g., joins). The key point is that ORAM does not protect the number of block accesses inherently for a general query operator. Hence, existing solutions to integrating indices into ORAMs leak the number of accesses to any index in processing. We will address the security issue in our algorithms in Sections V and VI.

Joins are commonly used operations in relational databases. In this work, we consider the problem of computing join functions in an oblivious way. Li and Chen [30] first studies oblivious theta-joins, but their algorithms are no better than a Cartesian product. Arasu and Kaushik [13] presents oblivious algorithms for a rich class of database queries including equi-joins. However, Krastnikov et al. [31] points out that the details in [13] are incomplete, and no practical implementation is provided to show the empirical results. Opaque [12] and ObliDB [32] are efficient only for the special case of one-to-many equi-join, e.g., primary-foreign key join. Krastnikov et al. [31] proposes a novel oblivious algorithm for general binary equi-joins. However, it is non-trivial to extend the algorithm to join multiple tables obliviously. A series of oblivious binary joins will disclose the intermediate table sizes, which may leak some sensitive information, e.g., data distribution or sparseness of the intermediate join graph. ObliDB [32] offers an oblivious hash join algorithm to support general equi-joins over multiple tables, but it is equivalent to a Cartesian product. Table I shows the comparison of oblivious join algorithms.

Example 1: Fig. 1 shows that Opaque Join [12] and 0-OM Join [32] do not work for many-to-many join, due to leaking some sensitive information (e.g., join degree).

Given two input tables T_1 and T_2 , they first put tuples from both input tables into one single table T , and obliviously sort T according to the join key. Next, they perform a linear scan over the single sorted table T , and join each tuple originally from T_1 with the corresponding tuples originally from T_2 .

In the original setting, they need to ensure the invariant that after accessing every input tuple in T , they write out exactly

TABLE I
COMPARISON OF OBLIVIOUS JOIN ALGORITHMS

	Join Type ^a			Algorithm	Complexity Analysis ^b			
	BE	BD	ME		Computation Overhead ^c	Cloud Storage	Client Storage	
Li and Chen [18]	✓	✓	✓	BE	A1	$\Omega(\prod_{j=1}^{\ell} T_j)$	$O(\prod_{j=1}^{\ell} T_j)$	$O(1)$
				BD	A2		$O(T_{in} + T_{out})$	$O(M)$
				ME	A3		$\Omega(T_{in} + T_{out})$	$O(M)$
Arasu and Kaushik [14]	✓	×	✓	BE ME	Equi-Join	$O((T_{in} + T_{out}) \log(T_{in} + T_{out}))$	$O(T_{in} + T_{out})$	$O(\log(T_{in} + T_{out}))$
Opaque [5]	×	×	×	PF	Opaque Join	$O((T_{in} + T_{out}) \log^2((T_{in} + T_{out})/M))$	$O(T_{in} + T_{out})$	$O(M)$
ObliDB [6]	✓	×	✓	PF	0-OM Join	$O((T_{in} + T_{out}) \log(T_{in} + T_{out}))$	$O(T_{in} + T_{out})$	$O(1)$
				BE ME	Hash Join	$O(\prod_{j=1}^{\ell} T_j)$	$O(\prod_{j=1}^{\ell} T_j)$	$O(M)$
ODBJ [2]	✓	×	×	BE BD	Binary Join	$O((T_{in} + T_{out}) \log(T_{in} + T_{out}))$	$O(T_{in} + T_{out})$	$O(1)$
Ours ^d	✓	✓	✓	BE BD	SMJ	$O((T_{in} + T_{out}) \cdot (\log_M(T_{in} + T_{out}) + \log T_{in}))$	$O(T_{in} + T_{out})$	$O(T_{in} /B + M + \log T_{in})$
				BE BD	INLJ (+Cache)	$O((T_1 + T_{out}) \cdot (\log_M(T_1 + T_{out}) + \log T_1 + \Delta \log T_2))$	$O(T_{in} + T_{out})$	$O(T_{in} /B + M + \log T_{in})$
				ME	INLJ (+Cache)	$O((T_{in} + T_{out}) \cdot (\log_M(T_{in} + T_{out}) + \log T_1 + \Delta \sum_{j=2}^{\ell} \log T_j))$	$O(T_{in} + T_{out})$	$O(T_{in} /B + M + \sum_{j=1}^{\ell} \log T_j)$

^aWe denote binary equi-join as BE, band join as BD, acyclic multiway equi-join as ME, primary-foreign key join as PF.

^bWe denote the total size of all input tables as $|T_{in}| = \sum_{j=1}^{\ell} |T_j|$ and the real join result size as $|T_{out}|$.

^cWe assume an oblivious sorting needs $O(n \log n)$ time cost [1], and an oblivious external sorting needs $O(n \log^2(n/m))$ time cost as with Table 2 in [6] and Table 1 in [2]. We also assume an oblivious filtering needs $O(n \log_m n)$ time cost, as with Theorem 6 in [45].

^dWe denote sort-merge join as SMJ and index nested-loop join as INLJ. We denote the number of outsourced levels in each B -tree as Δ .

one real or dummy join record. But for many-to-many join, they cannot keep the invariant above. For example, after accessing tuple $T_2(2, 1)$, which can match two tuples $T_1(2, 1)$ and $T_1(2, 2)$, they must write out two join records $T_1(2, 1) \bowtie T_2(2, 1)$ and $T_1(2, 2) \bowtie T_2(2, 1)$ before the next access over T , i.e., the number of output records between two accesses over T leaks the join degree. Processing tuples $T_2(2, 2)$ and $T_2(2, 3)$ brings the same security issue.

In summary, prior studies are still unable to address the major challenge in oblivious join. They are only efficient for foreign key join [12], [32], or restricted to binary join [31], or not leading to practical implementation [13], [30].

Our major objective is to support general band joins and multiway equi-joins obliviously. Band join [34] is a binary join between tables T_1 and T_2 on numeric attributes $T_1.A$ and $T_2.B$ with the join condition $T_1.A - c_1 \leq T_2.B \leq T_1.A + c_2$, where c_1 and c_2 are numeric values satisfying $c_1 \geq 0$ and $c_2 \geq 0$. In particular, a band join will reduce to a binary equi-join, when $c_1 = c_2 = 0$. First, we extend the binary equi-join algorithm in Krastnikov et al. [31] to support general band joins obliviously. Second, we propose two band join algorithms using ORAMs: sort-merge join and index nested-loop join. We integrate B -tree indices into ORAMs for input tables and retrieve blocks through indices obliviously to perform our algorithms. The key point is to bound the number of accesses to any index. Furthermore, we extend the index nested-loop join to support multiway equi-joins obliviously. The key idea is to avoid retrieving tuples that make no contribution to the final join result and bound the total number of block accesses. Note that ORAM can be viewed as a blackbox, providing read and write interface, while hiding access patterns. We can introduce some novel ORAMs (e.g., [35], [36], [37]) to improve the performance. We can also leverage other types of indices (e.g., Obliv [26]) rather than B -tree to perform our algorithms, as long as they can support both point and range queries obliviously. Our major contributions are listed as follows.

- We extend the binary equi-join algorithm in Krastnikov et al. [31] to support general band joins obliviously in Section IV. Note that existing studies (except [30]) do not work for any non-equi joins.
- We also propose two band join algorithms using ORAMs: sort-merge join and index nested-loop join in Section V-A and V-B. The key point is to bound the number of accesses to each B -tree index.
- We support acyclic equi-joins over multiple tables obliviously using index nested-loop join in Section VI. We avoid retrieving tuples that make no contribution to the final join result and bound the total number of block accesses.
- We conduct extensive experiments on real-world datasets in Section IX. The results demonstrate a superior performance gain (orders of magnitude speedup for oblivious multiway equi-joins) over baseline algorithms.

II. BACKGROUND AND RELATED WORK

A. Generic ORAMs and Path-ORAM

Generic ORAMs: ORAM [21], [22], [23] allows the client to access encrypted data in the server while hiding her access patterns. ORAM is modeled similar as a key-value store and hides the access patterns with the same length of operations (i.e., $\text{get}()$ and $\text{put}()$) to make them computationally indistinguishable to the server. It consists of two components: an ORAM data structure and an ORAM query protocol. The client and server run the ORAM query protocol to read and write any data to the ORAM data structure. A few advanced ORAMs [38], [39], [40], [41], [42], [43], [44] work on file systems, multiple clients, parallelization, asynchronicity and distributed data stores. We may leverage them as our secure ORAM storage, since we treat ORAM as a blackbox.

Path-ORAM: In this work, we adopt Path-ORAM [45] due to good performance and simplicity. It organizes the ORAM data structure as a full binary tree where each node is a bucket with a

fixed number of encrypted blocks. It maintains the *invariant* that at any time, each block b is always placed in some bucket along the path to the leaf node that b is mapped to. The *stash* stores a few blocks that have not been written back to the binary tree in server. The *position map* keeps track of the mapping between blocks and leaf node IDs, which brings a linear space cost to the client. To store N blocks of size B , a basic Path-ORAM requires $O(\log N + N/B)$ client memory and $O(\log N)$ cost per query.

B. Oblivious Sorting and Filtering

Oblivious Sorting: Items can be sorted by accessing in a *fixed, predefined order*. Bitonic sort [46] needs $O(N \log^2 N)$ time cost but with small constant factor. It can be extended to an oblivious external sort with $O(N \log^2(N/M))$ time cost using client memory size M [12], [32]. A few algorithms [47], [48], [49] achieve $O(N \log N)$ time cost but may fail with a small probability [48], or lead to large constant factors [47] and non-trivial implementation [49]. Recently, Shi [28] proposes an oblivious heap sort with $O(N \log N)$ time cost, which works better in memory but is not IO-efficient.

Oblivious Filtering: Dummy records can be removed by oblivious filtering. Prior studies [12], [13], [30] and the conference version [50] adopt an oblivious sorting to filter out dummy records. Actually, it can be done by oblivious compaction. OptORAMa [37] achieves this in $O(N)$ time but needs some non-trivial techniques. In this work, we adopt a simple oblivious compaction algorithm [33] with $O(N \log_M N)$ time cost, where M is trusted memory size.

C. Oblivious Data Structure and Index

Prior studies [14], [24], [25], [26], [27], [28], [29] build oblivious tree structures or indices. For certain data structure whose access pattern exhibits some predictability, they make the structure “oblivious” to improve the performance rather than bluntly storing blocks from the structure into ORAM.

ORAM+B-Tree: B -tree indices can be introduced to speed up the oblivious query processing [32], [51]. The client *ignores* the semantic difference of (encrypted) index and data blocks and stores them into ORAM. When answering any query, the client starts with retrieving the root block (of the index) from the server and then traverse down the tree. Intuitively, the client queries the index by running the same algorithm as that over a standard B -tree. The only difference is that each index or data block is retrieved through ORAM.

Oblivious B-Tree: Oblivious B -tree [32], [51] is designed to avoid storing the position map in client. The main idea is that each index node keeps block IDs and position tags of its children nodes. When retrieving any node through ORAM, we have acquired the position tags of its children nodes simultaneously. Note that most query algorithms over tree indices traverse the tree from the root to leaf nodes. As a result, the client only needs to remember the position tag of the root node, and all other position map information can be fetched on the fly as part of the query algorithm.

Index Caching: Index caching is a popular tree-based ORAM optimization [35], [52], [53]. The client can cache *one specific level* of B -tree index to speed up the query performance. Due to large fanout in B -tree index, this overhead to the client storage is far less than storing the entire index.

Note that the techniques above do not protect *how many accesses to the data structure*. In our method, we integrate indices into ORAMs and address the security issue in the scenario of

oblivious join, as long as the indices support both point and range queries obliviously.

D. Oblivious Query Processing

Xie et al. [54] proposes ORAM solutions to shortest path computation. ZeroTrace [53] supports oblivious get/put/insert operations over set/dictionary/list interfaces. Obladi [55] provides ACID transactions while hiding access patterns. OCQ [56] performs oblivious cooperative analytics in a decentralized manner. Snoopy [57] designs an oblivious storage based on oblivious load balancer and subORAMs. Chu et al. [58] focuses on differentially oblivious join whose problem definition is different from our work.

Note that existing solutions [12], [31], [32] rely on Trusted Execution Environments (TEE) (e.g., Intel SGX [59], [60]). However, TEE is orthogonal to oblivious algorithms and has no advantage to the obliviousness.

E. Other Related Work

Secure Multi-Party Computation: Secure multi-party computation (MPC) allows multiple parties to perform data analytics over their private data, while no party learns the data from another party. Hence, MPC-based solutions [56], [61], [62], [63], [64], [65] have a different problem setting from our cloud database setting.

Differential privacy: Differential privacy (DP) protects against attacks with guaranteed probabilistic accuracy. They build index [66] and key-value data collection [67], and support general SQL queries [68], [69], [70]. However, DP-based solutions [66], [67], [68], [69], [70], [71], [72], [73] provide *differential privacy for query results*, while we provide the *obliviousness in query processing*.

III. PROBLEM DEFINITION AND OVERVIEW

The formulation includes a client and a cloud server. The client, who has a small and secure memory, stores her data into the large but untrusted cloud storage. In online processing, the client issues join queries against the server.

We follow the definition in Opaque [12] and ObliDB [32]. Let D be the relational database (where some B -tree indices may be integrated) in the cloud and Q be a join query. Let $\text{Size}(D)$ be the sizing information of database D , which includes numbers and sizes of tables, rows, columns, and attributes in D , but does not include any attribute values. Let $\text{Sch}(D)$ be the schema information of database D , which includes table and column names in D (easily hidden using encryption). Let $\text{IOSize}(D, Q)$ be the input/output sizes of running Q over D . Note that $\text{IOSize}(D, Q)$ does not include the sizes of all intermediate join tables for any join query Q over multiple tables in D , which must be protected against the adversary. Let $\text{Trace}(\cdot)$ be the trace of server location accesses and network traffic patterns in query processing. Table II lists the notations used in this paper.

Definition 1: Oblivious Join [12]. For any two relational databases D and D' and two join queries Q and Q' , where $\text{Size}(D) = \text{Size}(D')$, $\text{Sch}(D) = \text{Sch}(D')$ and $\text{IOSize}(D, Q) = \text{IOSize}(D', Q')$, we denote the access patterns produced by the join algorithm OJoin running Q and Q' over D and D' as $\text{Trace}(\text{OJoin}(D, Q))$ and $\text{Trace}(\text{OJoin}(D', Q'))$. OJoin is an oblivious join algorithm, if

- 1) OJoin ensures the confidentiality; and

TABLE II
NOTATIONS

D, D'	Relational database
Q, Q'	Join query
c_1, c_2	Non-negative parameters in band join condition $T_1.A - c_1 \leq T_2.B \leq T_1.A + c_2$
N	Number of real data blocks in the database
B	Block size in terms of number of entries
M	Number of blocks held in client memory
T_1, \dots, T_ℓ	ℓ ($\ell \geq 2$) input tables
$T_{\text{out}}(R_{\text{real}})$	Final join result table
$ T_{\text{in}} $	Total number of tuples in all input tables
$ T_{\text{out}} (R_{\text{real}})$	Number of real join records in $T_{\text{out}}(R_{\text{real}})$
Δ	Number of outsourced levels in each B -tree
SMJ	Sort-merge join
INLJ(+Cache)	Index nested-loop join (with caching)
Size(D)	Sizing information of the database D
Sch(D)	Scheme information of the database D
IOSize(D, Q)	Input/Output sizes of running Q over D
OJoin(D, Q)	Oblivious join operator of running Q over D
Trace(\cdot)	Server location accesses and network traffic patterns in query processing

2) access patterns $\text{Trace}(\text{OJoin}(D, Q))$ and $\text{Trace}(\text{OJoin}(D', Q'))$ have the same length and computationally indistinguishable for anyone but the client.

We support two oblivious join approaches including non-ORAM approach (see Section IV) and ORAM approach (see Sections V and VI). In preprocessing stage, the client partitions the data into blocks and encrypts these data blocks. In particular for ORAM approach, the client builds an ORAM data structure (e.g., Path-ORAM) over the encrypted blocks and integrates some B -tree indices into the ORAM data structure using ORAM+ B -tree or oblivious B -tree (see Section II-C). Then, the client uploads the encrypted blocks or the ORAM data structure to the cloud storage, and keeps the encryption keys and other metadata (e.g., ORAM stash and position map in Path-ORAM) at her side. In online processing, the client runs the oblivious join algorithms by performing a series of oblivious operations or ORAM operations, which reads/writes blocks from/to the server and generates the query results.

Segmenting ORAM: In ORAM approach, we separate one single ORAM into multiple smaller ORAMs (denoted as SepORAM) to reduce the cost of each ORAM access, as in OblIDB [32]. For each input table, we build an ORAM for data blocks and another smaller ORAM for index blocks. The comparison in Table I is based on this setting. We also consider one single ORAM setting (denoted as OneORAM) and make the related discussion in Section VII.

Security Model: We consider a “honest-but-curious” server. Data is encrypted, retrieved, and stored in *atomic units* (i.e., blocks). All blocks are of the same size and are indistinguishable for the server. We use N to denote the number of real data blocks in the database, and each encrypted block contains B bytes. Note that the number of entries that fit in a block is $\Theta(B)$, and the constants will vary depending on the types of entries, e.g., encrypted index entry, encrypted attribute value, and position tag in ORAM.

By default, we follow the security guarantee in Definition 1 in both non-ORAM approach and ORAM approach (including both SepORAM and OneORAM settings). We provide the security analysis and proof in Section VIII.

We also introduce a padding mode to ease the volume leakage in final output size, as in Opaque [12] and OblIDB [32].

The join result size will be padded to an upper bound size, which leaks nothing regarding the join query but the upper bound size. Besides, we may introduce some novel padding techniques. For example, explore differential privacy rather than full obliviousness to reduce the padding size [69], or pad the result size to the closest power of a constant x (e.g., 2 or 4) [74], [75], [76], leading to at most $\log_x |R_{\text{worst}}|$ distinct result sizes, where $|R_{\text{worst}}|$ is the Cartesian product size in join scenario.

Note that our approaches do not consider privacy leakage through any side-channel attack (like time taken for each operation). Prior orthogonal studies [77], [78], [79] can help to alleviate such leakage.

IV. OBLIVIOUS BAND JOIN WITHOUT ORAM

We extend the binary equi-join algorithm in Krastnikov et al. [31] to support general band join obliviously. First, we obliviously compute the degree information in the join graph. Second, we obliviously make copies for each tuple according to the join degree and perform an oblivious one-to-one mapping operation to generate the final join output.

A. Join Degree Computation

Algorithm 1 shows the details of join degree computation. For each input table $T_i(j, d)$, we denote the join key as j and the remaining attributes as d . We mainly focus on table T_1 , and the computation on T_2 goes in a similar way.

First, we obliviously sort T_1 and T_2 lexicographically by (j, d) , and add a unique id to each tuple (Line 1–4). We parameterize oblivious sorting with a lexicographic ordering on chosen attributes, e.g., $\text{OSort}(T_i)\langle j \uparrow, d \uparrow \rangle$ sorts T_i by increasing j attribute, followed by increasing d attributes.

Then, we aim to generate augmented tables \tilde{T}_1 and \tilde{T}_2 with join degree α and position pos , such that each tuple $\tilde{t}_1 \in \tilde{T}_1$ matches $\tilde{t}_1.\alpha$ tuples in \tilde{T}_2 , where each matched tuple \tilde{t}_2 has a unique $\tilde{t}_2.id \in (\tilde{t}_1.pos, \tilde{t}_1.pos + \tilde{t}_1.\alpha]$ (Line 5–15).

Specifically, we generate two auxiliary tables T_R and T_S for T_1 , where $T_R.j \leftarrow T_1.j - c_1$ and $T_S.j \leftarrow T_1.j + c_2$ (Line 6–7). Suppose tuple $t_1 \in T_1$ corresponds to tuples $t_R \in T_R$ and $t_S \in T_S$. According to the join condition, any tuple $t_2 \in T_2$ with $t_2.j \in [t_R.j, t_S.j]$ will match $t_1 \in T_1$. Then, we generate augmented tables \tilde{T}_R and \tilde{T}_S with position pos , such that any tuple $t_2 \in T_2$ with $t_2.id \in (t_R.pos, t_S.pos]$ will match $t_1 \in T_1$, with the help of a union T_U of tables T_R, T_S and T_2 (Line 8–13). Finally, we generate \tilde{T}_1 from \tilde{T}_R and \tilde{T}_S , where the position $\tilde{T}_1.pos \leftarrow \tilde{T}_R.pos$ and the join degree $\alpha \leftarrow \tilde{T}_S.pos - \tilde{T}_R.pos$ (Line 14).

Algorithm 1 takes $O((|T_1| + |T_2|) \log(|T_1| + |T_2|))$ time cost, when an oblivious sort needs $O(n \log n)$ time cost [28].

Example 2: An example is given in Fig. 2. We mainly focus on table T_1 . First, we sort T_1 and T_2 lexicographically by (j, d) (Line 1–4). Then, we generate two auxiliary tables T_R and T_S , where $T_R.j \leftarrow T_1.j - 2$ and $T_S.j \leftarrow T_1.j + 1$ (Line 6–7). Note that $t_1 = (3, 4, 3) \in T_1$ corresponds to $(3, 2, 3) \in T_R$ and $(3, 5, 3) \in T_S$. According to the join condition, $(2, 2, 2)$, $(3, 2, 4)$ and $(4, 5, 3) \in T_2$ with $j \in [2, 5]$ are 3 matches of $t_1 \in T_1$.

Now, we compute a union T_U of tables T_R, T_S and T_2 , and sort T_U lexicographically by $(j, tid: T_R < T_2 < T_S)$ (Line 8–9), e.g., $(3, 2, 3, T_R) < (2, 2, 2, T_2) = (3, 2, 4, T_2) < (1, 2, 2, T_S)$. Note that $(2, 2, 2, T_2)$, $(3, 2, 4, T_2)$ and $(4, 5, 3, T_2)$ are 3 matches of $(3, 4, 3) \in T_1$, and they all rank between $(3, 2, 3, T_R)$ and




Fig. 2. Example of join degree computation.




Fig. 3. Example of table expansion and alignment.

Algorithm 1: Join Degree Computation.

Require: Input: two tables $T_1(j, d)$ and $T_2(j, d)$ with join condition $T_1.j - c_1 \leq T_2.j \leq T_1.j + c_2$.

Output: $\tilde{T}_1(id, j, d, pos, \alpha)$ and $\tilde{T}_2(id, j, d, pos, \alpha)$.

- 1: **for** $i \leftarrow 1$ to 2 **do**
 - 2: $T_i \leftarrow \text{OSort}(T_i)\langle j \uparrow, d \uparrow \rangle$;
 - 3: $T_i(id, j, d) \leftarrow T_i(id \leftarrow \text{ID}, j, d)$; \triangleright add id column
 - 4: **end for**
 - 5: **for** $i \leftarrow 1$ to 2 **do**
 - 6: $T_R(id, j, d) \leftarrow T_i(id, j - c_i, d)$;
 - 7: $T_S(id, j, d) \leftarrow T_i(id, j + c_{3-i}, d)$;
 - 8: $T_U(id, j, d, tid) \leftarrow T_R \cup T_S \cup T_{3-i}$; \triangleright add tid column
 - 9: $T_U \leftarrow \text{OSort}(T_U)\langle j \uparrow, tid : T_R < T_{3-i} < T_S \rangle$;
 - 10: $T_U(id, j, d, tid, pos) \leftarrow \text{Fill-Pos}(T_U)$;
 - 11: $T_U \leftarrow \text{OSort}(T_U)\langle tid : T_R < T_S < T_{3-i}, id \uparrow \rangle$;
 - 12: $\tilde{T}_R \leftarrow \pi_{id, j, d, pos}(T_U[1 \dots |T_i|])$;
 - 13: $\tilde{T}_S \leftarrow \pi_{id, j, d, pos}(T_U[|T_i| + 1 \dots 2|T_i|])$;
 - 14: $\tilde{T}_i \leftarrow T_i(id, j, d, pos \leftarrow \tilde{T}_R.pos, \alpha \leftarrow \tilde{T}_S.pos - \tilde{T}_R.pos)$;
 - 15: **end for**
 - 16: **return** \tilde{T}_1 and \tilde{T}_2 ;
-

$(3, 5, 3, T_S)$ in T_U . We scan T_U and assign the current number of tuples with $tid = T_2$ to $T_U.pos$ ($\text{Fill-Pos}(T_U)$) in Line 10).

After that, we extract \tilde{T}_R and \tilde{T}_S from T_U by re-sorting T_U (Line 11–13). Finally, we generate \tilde{T}_1 from \tilde{T}_R and \tilde{T}_S , where $\tilde{T}_1.pos \leftarrow \tilde{T}_R.pos$ and $\tilde{T}_1.\alpha \leftarrow \tilde{T}_S.pos - \tilde{T}_R.pos$ (Line 14). Note that $\tilde{t}_1 = (3, 4, 3, 1, 3) \in \tilde{T}_1$ matches $\tilde{t}_1.\alpha = 3$ tuples in \tilde{T}_2 , $(2, 2, 2, 0, 3)$, $(3, 2, 4, 0, 3)$ and $(4, 5, 3, 2, 1)$ with $id \in (1, 4] = (\tilde{t}_1.pos, \tilde{t}_1.pos + \tilde{t}_1.\alpha]$.

B. Table Expansion and Alignment

Algorithm 2 shows the details of table expansion and alignment. After obtaining the join degree α , we need to make copies for each tuple based on α (aka *table expansion*). We obviously expand each \tilde{T}_i into table S_i using Algorithm 4 in [31], where S_i consists of α (contiguous) copies of each tuple $(id, j, d, pos) \in \tilde{T}_i$ (Line 2). Then, we obviously align

Algorithm 2: Table Expansion and Alignment.

Require: Input: two tables $\tilde{T}_1(id, j, d, pos, \alpha)$ and

$\tilde{T}_2(id, j, d, pos, \alpha)$ with band join parameters c_1 and c_2 .

Output: join result table $T_{out}(j_1, d_1, j_2, d_2)$.

- 1: **for** $i \leftarrow 1$ to 2 **do**
 - 2: $S_i(id, j, d, pos) \leftarrow \text{OExpand}(\tilde{T}_i, \alpha)$;
 - 3: $S_i(id, j, d, pos, gid) \leftarrow S_i(id, j, d, pos, gid \leftarrow \text{ID}_{id})$;
 - 4: $S_i(id, j, d, pos) \leftarrow S_i(id, j, d, pos \leftarrow pos + gid)$;
 - 5: **end for**
 - 6: $S_2 \leftarrow \text{OSort}(pos \uparrow, id \uparrow)(S_2)$;
 - 7: $T_{out}(j_1, d_1, j_2, d_2) \leftarrow (S_1.j, S_1.d, S_2.j, S_2.d)$;
 - 8: **return** T_{out} ;
-

S_2 with S_1 (aka *table alignment*), so that each join record corresponds to a row of S_1 and a row of S_2 with matching index (Line 3–6). Finally, we generate the join output table T_{out} by concatenating (j, d) attributes in S_1 and S_2 (Line 7).

Example 3: An example is given in Fig. 3. First, we obviously expand each \tilde{T}_i into table S_i (Line 2). For example, for tuple $\tilde{t}_1 = (3, 4, 3, 1) \in \tilde{T}_1(id, j, d, pos)$, we make $\tilde{t}_1.\alpha = 3$ copies of \tilde{t}_1 to match $(2, 2, 2, 0)$, $(3, 2, 4, 0)$ and $(4, 5, 3, 2)$ in $\tilde{T}_2(id, j, d, pos)$ with $\tilde{T}_2.id \in (\tilde{t}_1.pos, \tilde{t}_1.pos + \tilde{t}_1.\alpha] = (1, 4]$.

Then, we obviously align S_2 with S_1 (Line 3–6).

1) We perform a grouping identity operation by scanning S_i (Line 3). For example, \tilde{t}_1 's 3 copies $(3, 4, 3, 1, 1)$, $(3, 4, 3, 1, 2)$ and $(3, 4, 3, 1, 3)$ belong to the same group, and each gets a different $gid = 1, 2$ and 3.

2) We update pos attribute as $pos \leftarrow pos + gid$ in table S_i (Line 4). After that, \tilde{t}_1 's 3 copies in S_1 will be $(3, 4, 3, 2)$, $(3, 4, 3, 3)$ and $(3, 4, 3, 4)$, and \tilde{t}_1 's 3 matches in S_2 will be $(2, 2, 2, 3)$, $(3, 2, 4, 3)$ and $(4, 5, 3, 3)$. Now, any tuple $s_2 \in S_2$ matches the only one tuple $s_1 \in S_1$, where $s_1.id = s_2.pos$ and $s_1.pos = s_2.id$.

3) After 2), S_1 has been permuted lexicographically by (id, pos) . Hence, we obviously sort table S_2 lexicographically by (pos, id) to achieve the table alignment (Line 6).

Finally, we generate the join output table T_{out} by simply concatenating (j, d) attributes in S_1 and S_2 (Line 7).

Algorithm 2 consists of two parts: table expansion and table alignment. We assume an oblivious sorting needs $O(n \log n)$ time cost [28]. For oblivious table expansion, the time cost is $O((|T_1| + |T_2|) \log(|T_1| + |T_2|) + |R_{\text{real}}| \log |R_{\text{real}}|)$. For oblivious table alignment, the time cost is $O(|R_{\text{real}}| \log |R_{\text{real}}|)$. Hence, the total time cost is $O((|T_1| + |T_2|) \log(|T_1| + |T_2|) + |R_{\text{real}}| \log |R_{\text{real}}|)$.

V. OBLIVIOUS BAND JOIN WITH ORAM

A. Oblivious Sort-Merge Join

Our algorithm is similar to the traditional sort merge join but with some differences. In preprocessing, we integrate non-clustered B -tree indices into ORAMs for each input table in advance, where each leaf index entry keeps a pointer to the data tuple. Leaf index entries are sorted as per the attribute. For each input table, we build an ORAM structure for data blocks and another smaller one for index blocks.

In each join step, we keep the *invariant* that we retrieve the tuple needed from each input table *alternatively*. A dummy tuple is retrieved as necessary. It ensures the full obliviousness, since

Algorithm 3: Oblivious Sort-Merge Band Join.

Require: Input: two tables $T_1(j, d)$ and $T_2(j, d)$ with join condition $T_1.j - c_1 \leq T_2.j \leq T_1.j + c_2$.

Output: join result table T_{out} .

```

1: Initialize  $T_{out} \leftarrow \emptyset$ .
2: Initialize  $t_1, t_2 \leftarrow \emptyset$ .
3: for  $i \leftarrow 1$  to 2 do
4:    $t_i \leftarrow T_i.getFirst()$ ;
5: end for
6: while  $t_1 \neq \perp$  or  $t_2 \neq \perp$  do
7:    $res \leftarrow t_2.j - t_1.j$ ;
8:   if  $-c_1 \leq res \leq c_2$  then
9:      $begin \leftarrow t_2$ ;
10:    while  $-c_1 \leq res \leq c_2$  do
11:       $T_{out}.put(Join(t_1, t_2))$ ;
12:       $T_1.getDummy()$ ;  $t_2 \leftarrow T_2.getNext()$ ;
13:       $res \leftarrow t_2.j - t_1.j$ ;
14:    end while
15:     $T_{out}.put(\perp)$ ;
16:     $t_2 \leftarrow begin$ ;
17:     $t_1 \leftarrow T_1.getNext()$ ;  $T_2.getDummy()$ ;
18:  else
19:     $T_{out}.put(\perp)$ ;
20:    if  $res > c_2$  then
21:       $t_1 \leftarrow T_1.getNext()$ ;  $T_2.getDummy()$ ;
22:    else
23:       $T_1.getDummy()$ ;  $t_2 \leftarrow T_2.getNext()$ ;
24:    end if
25:  end if
26: end while
27:  $T_{out} \leftarrow OFilter(T_{out})$ ;
28: return  $T_{out}$ ;

```

each tuple retrieval needs the same number of ORAM accesses for each input table. Then, we perform a join comparison in each step. If there is a match, we write out a join record; otherwise, we write out a dummy record as necessary.

Algorithm 3 joins two tables T_1 and T_2 . Whenever we perform a getNext() over one input table (T_1 or T_2), we also perform a dummy operation getDummy() over the other table (T_2 or T_1) to ensure the obliviousness.

First, Algorithm 3 initializes $T_{out} := \emptyset$ (Line 1). Then, we retrieve the first two tuples from T_1 and T_2 as t_1 and t_2 (Line 2–5). While either t_1 or t_2 is real, we compute the join comparison result “res” between them (Line 6–7). We keep the *invariant* above that we always pull tuples from T_1 and T_2 alternatively for either of two possible cases:

1) t_1 matches t_2 : First, we save the current t_2 to a temporary tuple “begin” (Line 9). We keep writing out the join record $Join(t_1, t_2)$ to T_{out} , and retrieving the next tuple from T_2 as t_2 , until the newly retrieved t_2 does not match t_1 (Line 10–14). Once they do not match, we write out a dummy record and assign “begin” back to t_2 (Line 15–16). Finally, we will retrieve the next tuple from T_1 (Line 17) and move to the next iteration (Line 6–7).

2) t_1 does not match t_2 : Since they do not match, we first write out a dummy record (Line 19). If $res > c_2$, we retrieve the next tuple from T_1 (Line 20–21). Otherwise, we retrieve the next tuple from T_2 (Line 22–23). Finally, we move to the next iteration (Line 6–7).

T_1		T_2		Join Comparison			
j	d	j	d	T_1	T_2	res	getNext()
1	2	1	3	1	2	1	3
3	1	2	2	1	2	1	3
4	3	2	4	1	2	2	1
8	1	5	3	1	2	2	4
T_1		T_2		$T_1 \rightarrow (3, 1)$			
3	1	1	3	1	2	5	3
3	1	2	2	1	2	4	1
3	1	2	4	1	2	4	1
3	1	5	3	1	2	5	3
3	1	5	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
4	3	1	3	1	2	5	3
T_1		T_2		$T_1 \rightarrow (4, 3)$			
4	3	1	3	1	2	5	3

Algorithm 4: Oblivious Index Nested-Loop Band Join.

Require: Input: two tables $T_1(j, d)$ and $T_2(j, d)$ with join condition $T_1.j - c_1 \leq T_2.j \leq T_1.j + c_2$.

Output: join result table T_{out} .

```

1: Initialize  $T_{\text{out}} \leftarrow \emptyset$ .
2: Initialize  $t_1, t_2 \leftarrow \emptyset$ .
3: for  $i \leftarrow 1$  to  $|T_1|$  do
4:    $t_1 \leftarrow T_1.\text{getNext}()$ ;
5:    $t_2 \leftarrow T_2.\text{getFirst}(t_1.j - c_1)$ ;
6:   while  $t_1.j - c_1 \leq t_2.j \leq t_1.j + c_2$  do
7:      $T_{\text{out}}.\text{put}(\text{Join}(t_1, t_2))$ ;
8:      $T_1.\text{getDummy}()$ ;
9:      $t_2 \leftarrow T_2.\text{getNext}()$ ;
10:  end while
11:  $T_{\text{out}}.\text{put}(\perp)$ ;
12: end for
13:  $T_{\text{out}} \leftarrow \text{OFilter}(T_{\text{out}})$ ;
14: return  $T_{\text{out}}$ ;
    
```

$\text{res} < -c_1$, and the cursor on T_2 advances (Line 23). The process above will end when both cursors reach the end of T_1 and T_2 . Hence, we will invoke $\text{getNext}() |T_1| + |T_2|$ times. Therefore, the total number of join steps in Part I is $|T_1| + |T_2| + 1$.

Part II: The process in Line 10-14 in Algorithm 3.

Note that each join step leads to one join comparison. In Part II, each join comparison leads to writing out one real join record. Since the number of real join records is $|R_{\text{real}}|$, the number of join steps in Part II is also $|R_{\text{real}}|$.

Based on Part I and II, $\text{Num}_{\text{js}} = |T_1| + |T_2| + |R_{\text{real}}| + 1$.

B. Oblivious Index Nested-Loop Join

In our index nested-loop join, we integrate B -tree indices into ORAMs for each input table and retrieve tuples by querying the indices through ORAMs. In detail, the outer loop is to scan table T_1 . While accessing each tuple in T_1 , the algorithm retrieves matched tuples from table T_2 through B -tree index. In each join step, we ensure the *invariant* that we retrieve the tuple needed from each input table *alternatively*. A dummy tuple is retrieved from table T_1 as necessary. The difference on two tables is that we retrieve tuples from T_1 one by one according to sequential block IDs, while for table T_2 we retrieve the tuple that we need by searching over a whole B -tree path. After each pair of tuple retrievals, we make a join comparison of the current two tuples. If there is a match, we write out the join record; otherwise, a dummy record is output as necessary.

Algorithm 4 joins two tables T_1 and T_2 . Algorithm 4 begins with initializing an empty output table T_{out} (Line 1). The outer loop is to iterate over each tuple in table T_1 (Line 3). Each time we retrieve a new tuple t_1 from T_1 (Line 4), we first retrieve a tuple t_2 from T_2 , which is the first tuple satisfying $t_2.j \geq t_1.j - c_1$ (Line 5). If those two tuples can match, we write the join record $\text{Join}(t_1, t_2)$ to the output table T_{out} (Line 7) and retrieve the next tuple from T_2 as t_2 (Line 9). To ensure the obliviousness, we also perform a dummy retrieval from T_1 (Line 8). We repeat the process above until the newly retrieved t_2 does not match the current t_1 . Once they do not match, we write out a dummy record (Line 11) and step into the next iteration. The final step is to obliviously filter out dummy records from T_{out} and only keep real join records (Line 13).

Example 5: An example is given in Fig. 5. When we retrieve tuple $t_1 \leftarrow T_1(1, 2)$ from T_1 (Line 4), we first retrieve tuple t_2

j	d	j	d	Join Comparison				T_1	T_2	match		
T_1	T_2	T_1	T_2	j	d	j	d	j	d	match		
1	2	1	3	3	1	2	4	3	1	2	4	✓
3	1	2	2	1	2	1	3	3	1	5	3	×
4	3	2	4	1	2	2	2	4	3	2	2	✓
8	1	5	3	1	2	2	4	4	3	2	4	✓
				1	2	5	3	4	3	5	3	✓
				3	1	1	3	4	3	⊥	⊥	×
				3	1	2	2	8	1	⊥	⊥	×

Fig. 5. Example of index nested-loop join with ORAMs.

Table T_1		Table T_2		Join Comparison											
A	B	A	C	T_1	T_2	T_3	T_4	T_1	T_2	T_3	T_4	match			
1	1	1	1	A	B	A	C	A	B	A	C	D	E		
2	1	2	1	1	1	1	1	1	1	1	1	⊥	false		
2	2	2	2	2	1	2	1	2	1	2	1	⊥	disable		
2	3	3	1	2	1	2	1	2	1	1	1	⊥	false		
				2	1	2	1	2	1	2	1	1	true		
				2	1	2	3	2	1	2	3	⊥	false		
				⊥	⊥	⊥	⊥	⊥	⊥	⊥	⊥	⊥	disable		
				2	2	2	1	2	2	2	1	1	1	2	true
				2	3	2	1	2	1	⊥	⊥	⊥	false		

padding dummy operations

Fig. 6. Example of oblivious multiway equi-join.

$\leftarrow T_2(1, 3)$ from T_2 , which is the first tuple satisfying $t_2.j \geq t_1.j - c_1$ (Line 5). While t_1 can match t_2 , we keep writing out the join record $\text{Join}(t_1, t_2)$ (Line 7) and retrieving the succeeding tuples $T_2(2, 2)$ and $T_2(2, 4)$ from T_2 as the new t_2 (Line 9). Once the newly retrieved $t_2 \leftarrow T_2(5, 3)$ does not match $t_1 = T_1(1, 2)$, we step into the next iteration and process the next tuple $t_1 \leftarrow T_1(3, 1)$ from T_1 . In particular, once we cannot find any tuple needed from T_2 , we retrieve a dummy tuple \perp from T_2 and logically let the matching result be false (e.g., the last two rows in Join Comparison in Fig. 5). The final step is to obliviously filter out dummy records from T_{out} (Line 13).

Theorem 2: For any two input tables T_1 and T_2 and the real join result R_{real} , let Num_{js} be the number of join steps. It is a function of $|T_1|$, $|T_2|$ and $|R_{\text{real}}|$. Specifically, we have

$$\text{Num}_{\text{js}} = f(|T_1|, |T_2|, |R_{\text{real}}|) = |T_1| + |R_{\text{real}}|.$$

VI. OBLIVIOUS MULTIWAY EQUI-JOIN

We extend our Algorithm 4 to support *acyclic multiway equi-joins* obliviously. The key idea is to avoid retrieving tuples that make no contribution to the final join result to bound the total number of block accesses.

Example 6: Fig. 6 shows an example of acyclic multiway equi-join over four tables T_1 - T_4 . Due to the acyclicity, each input table can be arranged as a node in a join tree. In this tree, for any different tables T_i, T_j, T_k , if T_k is on the path from T_i to T_j , we must have $\text{Attr}(T_i) \cap \text{Attr}(T_j) \subseteq \text{Attr}(T_k)$ for their attribute sets. The algorithm of building a join tree is presented in [81]. We number input tables in a pre-order traversal of the join tree. It ensures $i < j$, if T_i is an ancestor table of T_j . We also denote the parent table of T_i in the join tree as $T_{p(i)}$.

In our index nested-loop join algorithm, the outer loop is to iterate over each tuple in root table T_1 . Each time we retrieve a new tuple (e.g., $T_1(1, 1)$) from T_1 , we search matched tuples (e.g., $T_2(1, 1)$, $T_3(1, 4)$, ...) from T_2, \dots, T_ℓ . To ensure the obliviousness, we retrieve the tuple needed from each input table in a *round-robin* way and add dummy retrievals as necessary (e.g., retrieve \perp from T_4 , due to no tuple with join key $D \geq 4$ for matching $T_3(1, 4)$, as highlighted in yellow in Fig. 6). In

each step, if there is a match (e.g., in 4th and 7th join step), we output the join record; otherwise, we output a dummy record.

To bound the total number of join steps, we make the following observations to avoid retrieving unnecessary tuples.

Observation 1: For any non-root table T_j and its parent table $T_{p(j)}$, tuple $[p(j)]$ in $T_{p(j)}$ makes no contribution to the final join result, if no tuple in T_j matches tuple $[p(j)]$. Then, tuple $[p(j)]$ can be safely *disabled* (i.e., will not be accessed in the future).

For example, for table T_4 and tuple $T_3(1, 4)$ in parent table T_3 , we find no tuple in T_4 matches $T_3(1, 4)$ (in 1st join step). Hence, $T_3(1, 4)$ makes no contribution to the final join result, and can be safely *disabled* in an additional dummy join step (in 2nd join step). In this dummy step, we perform a dummy tuple retrieval from each input table except T_3 . For T_3 , we perform a tuple disabling operation, which is indistinguishable from a tuple retrieval based on the access patterns.

When disabling any tuple, we mark its leaf entry as disabled using an additional boolean tag. If all entries in any B -tree leaf block have been marked as disabled, the parent entry in the B -tree parent block will also be marked as disabled. This can recursively go up to B -tree root block. Since the recursion goes up along a B -tree path, we can still finish each disabling operation using some additional B -tree path access through ORAM (i.e., adding some dummy join step). When retrieving a new tuple from any input table, we skip disabled entries during searching over B -tree index.

Observation 2: For any non-root table T_j and its parent table $T_{p(j)}$, tuple $[p(j)]$ in $T_{p(j)}$ makes no contribution to the final join result, if each tuple in T_j that matches tuple $[p(j)]$ has been *disabled*. Then, tuple $[p(j)]$ can also be safely *disabled*.

For example, for table T_3 and tuple $T_1(1, 1)$ in parent table T_1 , $T_3(1, 4)$ is the only tuple in T_3 that matches $T_1(1, 1)$. However, since $T_3(1, 4)$ has been disabled (in 2nd join step), we know that $T_1(1, 1)$ makes no contribution to the final join result. If the parent tuple is in a non-root table, we will disable it by adding some dummy join step as above. Otherwise, we do not physically disable any tuple in root table T_1 , since the outer loop in our algorithm iterates over each tuple in root table T_1 , and will not access any previous tuple in T_1 in the future.

Observation 3: For any non-root table T_j and its parent table $T_{p(j)}$, tuple $[p(j)]$ in $T_{p(j)}$ will have no more matches, if the current tuple tuple $[j]$ in T_j matches tuple $[p(j)]$ but the succeeding tuple in T_j has a different join key from tuple $[j]$'s.

Observation 3 is based on the property of equi-joins. For example, for table T_3 and tuple $T_1(1, 1)$ in parent table T_1 , we find that $T_3(1, 4)$ can match $T_1(1, 1)$ (in 1st join step). But since the succeeding tuple $T_3(2, 1)$ has a different join key from $T_3(1, 4)$, we can conclude that $T_3(2, 1)$ does not match $T_1(1, 1)$ in equi-join scenario. Hence, $T_1(1, 1)$ will have no more matches.

To perform this optimization, we attach another boolean tag to each leaf entry, which indicates whether the next leaf entry in T_j has the same key with the current entry in T_j . If not, we do not retrieve the next tuple from the child table T_j .

After the normal join process, we pad the number of join steps to the upper bound (e.g., the last step in Fig. 6) in Theorem 3 to ensure the obliviousness. Finally, we obviously filter out dummy records and only keep real join records. The last step is to go over all index blocks and reset boolean tags in each entry.

In brief, tuple disabling operations will introduce some additional dummy join steps, but we can still bound the total number of join steps in Theorem 3. Besides, tuple disabling

operations also bring the overhead of resetting the boolean tags after answering each join query. However, the total time complexity is dominated by regular join steps and final oblivious filtering. Hence, the time cost of resetting the boolean tags is relatively small in oblivious join processing.

Theorem 3: For any ℓ ($\ell \geq 2$) input tables T_1, \dots, T_ℓ and the real join result R_{real} , let Num_{js} be the number of join steps. It is a function of $|T_1|, \dots, |T_\ell|$ and $|R_{\text{real}}|$. Specifically, we have

$$\text{Num}_{\text{js}} = f(|T_1|, \dots, |T_\ell|, |R_{\text{real}}|) = |T_1| + 2 \sum_{j=2}^{\ell} |T_j| + |R_{\text{real}}|.$$

VII. DISCUSSION ON ONE ORAM SETTING

In this work, we separate one single ORAM into multiple smaller ORAMs (aka SepORAM setting). Now, we reconsider the optimization in OneORAM setting.

Since we retrieve all the tuples through one single ORAM, an optimization in OneORAM is to safely remove some dummy tuple retrievals to speed up join processing. To ensure the obliviousness, we must write out a real or dummy join record after each tuple retrieval in OneORAM (rather than after each join step in SepORAM). Then, we must pay the same number of ORAM accesses between writing out any two join records. In other words, we must pad the number of ORAM accesses to the maximum height of B -tree indices in OneORAM. Note that each tuple retrieval from any input table will be indistinguishable for the adversary, although he knows the total number of tuple retrievals. We can bound the total number of tuple retrievals in OneORAM, as long as it only pertains to the input and output sizes, and no additional information will be leaked.

However, there is a major drawback in OneORAM setting. Suppose there are multiple tables in the whole dataset, but only a few binary joins will be processed online. In this scenario, we must put all input tables into one single ORAM in advance, since we do not know the online workload. Hence, we have to pay much larger cost for accessing the large single ORAM rather than smaller separate ORAMs.

VIII. SECURITY ANALYSIS

We formalize our security guarantee in Theorem 4 with the same notations in Definition 1. As with Opaque [12] and OblIDB [32], our security is guaranteed by the existence of simulator SIM such that for any probabilistic polynomial-time (PPT) adversary \mathcal{A} , \mathcal{A} cannot distinguish between the real server location trace from our method and the simulated trace from simulator SIM. Since SIM only sees what we want to leak, \mathcal{A} cannot learn any additional information. A brief description on specifics of simulated traces from SIM is given in the proof of Theorem 4.

In our setting, SIM only has the access to the schema and sizing information of input and output tables, the oblivious join operator, and some specific public constants (e.g., the number of outsourced levels in each B -tree index, denoted as Δ). Note that SIM has no access to the sizes of all intermediate join tables, since we protect this sensitive information against the adversary.

Theorem 4: For any relational database D , schema $\text{Sch}(D)$, join query Q , oblivious join algorithm OJoin, and security parameter λ , there is a polynomial-time simulator SIM such that for any PPT adversary \mathcal{A} ,

$$|\Pr[\mathcal{A}(\text{SIM}(\text{Size}(D), \text{Sch}(D), \text{IOSize}(D), Q),$$

$$\begin{aligned} & \text{OJoin}(D, Q)) \Rightarrow 1] \\ & - \Pr[\mathcal{A}(\text{Trace}(\text{OJoin}(D, Q))) \Rightarrow 1] \leq \text{negl}(\lambda). \end{aligned}$$

Proof (Informal Sketch): In this proof, we show the existence of simulator SIM, and argue that access pattern of SIM is distributed indistinguishable from $\text{Trace}(\text{OJoin}(D, Q))$ (generated from algorithm $\text{OJoin}(D, Q)$). SIM reads algorithm $\text{OJoin}(D, Q)$ to determine which operations to simulate.

For Oblivious Join without ORAMs:

The security proof is similar to that of Krastnikov et al. [31] and Arasu and Kaushik [13]. First, the process of our join algorithm $\text{OJoin}(D, Q)$ guarantees that each intermediate table size only pertains to the input and output sizes $\text{IOSize}(D, Q)$. Then, we consider how SIM simulates the access patterns for the operations in $\text{OJoin}(D, Q)$ as follows.

- **Oblivious Sorting and Linear Scan:** These two operations access the blocks in a fixed, predefined order. Hence, SIM can simulate the access patterns, given the access to $\text{Sch}(D)$, $\text{Size}(D)$ and $\text{IOSize}(D, Q)$.
- **Table Augmentation:** For each iteration in Table Augmentation, we read an input tuple, compute and add derived attributes, and write out the output tuple. As with linear scan, SIM can simulate the access patterns, given $\text{Sch}(D)$, $\text{Size}(D)$ and $\text{IOSize}(D, Q)$.
- **Union of Tables:** For each iteration in Union of Tables, we read a tuple from one input table and write it out to the output table. As with linear scan, SIM can simulate the access patterns, given $\text{Sch}(D)$, $\text{Size}(D)$ and $\text{IOSize}(D, Q)$.
- **Filling Position:** Filling Position ($\text{Fill-Pos}(\cdot)$ in Algorithm 1) operation scans the input tuples while maintaining a counter in private client. The counter will be incremented once we meet specific tuples. For each input tuple, we assign the counter to a new attribute pos and write out the updated tuple. Hence, the access pattern simulation can be reduced to Table Augmentation.
- **Table Expansion:** We adopt Algorithm 4 in [31] to support Table Expansion operation. SIM can simulate the access patterns as in [31].
- **Table Alignment:** Table Alignment sorts and scans the expanded tables. For each iteration in the scan, we read two input tuples, concatenate their (j, d) attributes and write out one join record. Hence, the access pattern simulation can be reduced to Oblivious Sorting, Linear Scan, and Table Augmentation.

For Oblivious Join with ORAMs:

SIM needs to simulate access patterns for ORAM operations and oblivious filter operations (including oblivious compaction/sorting, and a few linear scans) in $\text{OJoin}(D, Q)$.

This proof is covered by Arguments 1–4. We mainly focus on separate ORAMs setting (denoted as SepORAM) in Arguments 1–3. For one ORAM setting (denoted as OneORAM), the proof relies on Argument 4: OneORAM does not introduce any more privacy leakage than SepORAM.

Argument 1: We ensure the obliviousness in each join step.

First, we argue that SIM can simulate each ORAM or oblivious filtering operation. Since SIM has the access to schema $\text{Sch}(D)$ and sizing information $\text{Size}(D)$, the access pattern simulation for each of such operations is the same as that in the original ORAM scheme, or that for original oblivious compaction/sorting and linear scan operations.

Then, we argue that SIM can simulate each join step. In SepORAM, we keep the *invariant* that we always retrieve the tuples needed from each input table in a round-robin way in each

join step. Even if we do not need to retrieve any new tuple, we still retrieve a dummy tuple to ensure the obliviousness. At the end of each join step, if there is a match, we write out a join record to the output table; otherwise, we write out a dummy record as necessary. Specifically, each tuple retrieval for any input table leads to the same number of ORAM accesses, which only pertains to the height of the outsourced B -tree index. In each join step, since SIM has the access to specific public constants (e.g., the number of outsourced levels in each B -tree index), SIM can perform the corresponding number of ORAM operation simulations for each input table in a round-robin way and output a (randomized encrypted) join record.

Argument 2: We ensure the number of join steps only pertains to the input and output sizes.

In SepORAM, Theorems 1–3 guarantee that the number of join steps in algorithm $\text{OJoin}(D, Q)$ only pertains to the input and output sizes $\text{IOSize}(D, Q)$. Since SIM has the access to $\text{IOSize}(D, Q)$, SIM will know the number of join steps based on $\text{IOSize}(D, Q)$, and perform the corresponding number of join step simulations.

Argument 3: Arguments 1 and 2 ensure the simulated access pattern is indistinguishable from $\text{Trace}(\text{OJoin}(D, Q))$ in the whole process (i.e., the obliviousness in SepORAM).

Argument 4: OneORAM does not introduce any more privacy leakage than SepORAM.

For each step in OneORAM, algorithm $\text{OJoin}(D, Q)$ retrieves the tuple needed from an input table through the single ORAM, and pads the number of ORAM accesses to the maximum length of all retrieved B -tree paths. It ensures that each tuple retrieval from any input table will be indistinguishable for the adversary. Note that $\text{OJoin}(D, Q)$ may remove some dummy tuple retrievals, as long as total number of tuple retrievals only pertains to the input and output sizes $\text{IOSize}(D, Q)$. Then, after each tuple retrieval in OneORAM (rather than after each join step in SepORAM), we ensure to write out a real or dummy join record to the output table, to protect the join degree information and ensure the full obliviousness. The simulation is similar to that in SepORAM, since SIM still has the access to the background knowledge.

Theorem 4 guarantees our security in the sense of Definition 1. For binary joins, our security guarantee is the same as Krastnikov et al. [31] and oblivious mode in Opaque [12] and OblIDB [32]. For multiway joins, our security guarantee is the same as Arasu and Kaushik [13].

The simulator SIM' for padded mode behaves analogously to SIM. In padded mode, the security theorem replaces the final join output size with an upper bound size as a public parameter in simulator SIM, which indicates the padded output size.

IX. EXPERIMENTAL RESULTS

A. Experimental Setup and Datasets

We make the evaluation for OblIDB [32], ODBJ [31] and our ORAM approach. For ODBJ, we extend its implementation [82] to support general band joins. We adopt two oblivious sorting algorithms including oblivious external bitonic sorting [46] (denoted as ODBJ (Bitonic)) and oblivious heap sorting [28] (denoted as ODBJ (Heap)). For our ORAM approach, we have two settings: SepORAM and OneORAM. Each setting includes three algorithms: SMJ, INLJ and INLJ+Cache (see Table I). In “+Cache” mode, the client caches all index blocks above the leaf level, i.e., $\Delta = 1$ (see Table II).




Fig. 7. Storage cost against raw data size on TPC-H.




Fig. 8. Storage cost against raw data size on social graph.




Fig. 9. Performance of binary equi-join on TPC-H.

We also compare our method with an insecure baseline (Raw Index(+Cache)). It builds B -tree indices over data blocks and stores them in the cloud without encryption.

Setup: The client is an Ubuntu 18.04 machine with 18 GB memory. The server is an Ubuntu 18.04 machine with 256 GB memory and 2 TB hard disk. The bandwidth is 1 Gbps.

Default parameter values: We set block size $B = 4$ KB, as in [11], [40], [54]. We set trusted memory size $M = 2B$ (B is block size) in ODBJ and our method, but set $M = 50 \log N$ in OblIDB to make it finish in a reasonable period.

We evaluate the methods on the following two datasets.

TPC-H: We set default data size to 100 MB and vary data sizes from 10 MB to 1 GB in TPC-H benchmark. Query TE1-TE3 and Query TM1-TM3 come from the conference version [50]. Appendix A, available online shows Query TB1-TB2 in SQL.

- Query TE1-TE3: general equi-joins over 2 tables.
- Query TB1-TB2: general band joins over 2 tables.
- Query TM1-TM3: general multiway joins over 3-5 tables.

Social graph: Social graph [83], [84] contains twitter friendship links. We set default user number to 20,000 (with raw data size 4.5 MB) and vary user numbers from 5,000 to 200,000 (with raw data size from 1.3 MB to 58 MB). The following queries come from the conference version [50].

- Query SE1-SE3: general equi-joins over 2 tables.
- Query SM1-SM3: general multiway joins over 3-4 tables.

Remarks: The query cost for each method should be roughly proportional to the communication cost. It is confirmed by our experimental results (see Figs. 9–16). For simplicity, we mainly focus on experimental results for query cost.




Fig. 10. Performance of binary equi-join on social graph.




Fig. 11. Performance of Query TE2 against raw data size.




Fig. 12. Performance of Query SE2 against raw data size.




Fig. 13. Performance of band join on TPC-H.

B. Cloud and Client Storage Costs

Figs. 7(a) and 8(a) show cloud storage cost on two datasets. OblIDB and ODBJ achieve the minimum cloud storage cost, since they only store encrypted data blocks. Raw Index(+Cache) needs a little more cost for storing index blocks. ORAM based method has roughly 10X larger cost than Raw Index(+Cache), due to building ORAM data structure.

Figs. 7(b) and 8(b) show client memory size on two datasets. ODBJ achieves the minimum cost, since the client always keeps a constant number of blocks. For Raw Index(+Cache), the client also keeps a few more blocks along retrieved B -tree paths and may cache some index blocks. For OblIDB, we set trusted memory size $M = 50 \log N$ and make it finish as soon as possible. For ORAM based method, the client memory cost grows (roughly) linearly with raw data size, due to $O(N/B)$ blocks in the position map.




Fig. 14. Performance of Query TB1 against raw data size.




Fig. 15. Performance of multiway equi-join on TPC-H.




Fig. 16. Performance of multiway equi-join on social graph.




Fig. 17. Padded versus non-padded mode (binary equi-join).




Fig. 18. Padded versus non-padded mode (band join).

C. Performance of Binary Equi-Join

1) *Default Setting*: Figs. 9(a) and 10(a) show query cost for binary equi-join in default setting. Our SepORAM(+Cache) achieves 2X-3X and 50X-3000X better performances than ObliDB on TPC-H and social graph, since our query cost depends on input and output sizes *linearly*. The speedup difference




Fig. 19. Padded versus non-padded mode (multiway equi-join).

is mainly due to the join result size, which grows with *square of input size* on TPC-H but is *comparable with input size* on social graph.

Our SepORAM(+Cache) brings 90X-450X larger query cost than Raw Index(+Cache) except for Query SE1, and also brings 7X-15X and 40X-160X larger query cost on TPC-H and social graph than ODBJ (Bitonic) except for Query SE1. The major reason is that data tuple size is much less than block size. For index based methods (Raw Index(+Cache) and ours), only one index entry or data tuple in each retrieved block contributes to the join processing. For ODBJ method, ODBJ (Heap) brings 4X-9X and 25X-37X larger query cost on TPC-H and social graph than ODBJ (Bitonic), since oblivious heap sort [28] works better in memory but does not achieve good IO performance. In contrast, oblivious external bitonic sort [46] is more IO-efficient. Note that data tuple size is much less than block size on both datasets, even if the trusted memory contains only two blocks, the trusted memory actually holds decades of tuples.

In particular, Query SE1 joins a small table with a large one but generates few join records. Sep SMJ and Sep INLJ(+Cache) bring 2400X and 30X larger cost than Raw Index(+Cache) algorithms. Sep INLJ(+Cache) even achieves 1.7X-2.7X better performance than ODBJ (Bitonic). The reason is that query cost of Sep INLJ(+Cache) increases with large table size *logarithmically*, while that of Sep SMJ and ODBJ increases with large table size *linearly* (see Table I).

For our ORAM based method, Sep INLJ achieves 1.2X-2.6X better performance than One INLJ. As explained in Section VII, One INLJ(+Cache) has to pad the number of ORAM accesses for each tuple retrieval to the maximum length of outsourced *B*-tree paths, although this problem can be alleviated by index caching. One SMJ does not need padding, since the client always accesses an index block and then a data block for each tuple retrieval through ORAM. One SMJ even achieves 1.6X better performance than Sep SMJ on Query SE2 and SE3, due to less number of tuple retrievals based on the optimization in Section VII. Last, the index caching brings 1.2X-1.6X speedup ratio.

2) *Scalability*: Figs. 11(a) and 12(a) show query cost for Query TE2 and SE2 against raw data size. Our SepORAM(+Cache) achieves 2X-4X and 1600X-16000X better performances than ObliDB for Query TE2 and SE2, when raw data size increases from the minimum to the maximum. The speedup difference for two queries is still on account of the join result size, as explained in Section IX-C1. Compared with Raw Index(+Cache), SepORAM(+Cache) brings 75X-157X and 161X-409X larger query cost on Query TE2 and SE2. Compared with ODBJ (Bitonic), SepORAM(+Cache) brings 10X-20X and 30X-140X larger query cost on Query TE2 and SE2. The major




Fig. 20. Access pattern of our oblivious band join algorithm in ODBJ (Bitonic).




Fig. 21. Access pattern of our oblivious multiway equi-join algorithm.

reason is still that data tuple size is much less than the block size, as explained in Section IX-C1. For ODBJ method, ODBJ (Heap) brings 5X-15X and 16X-37X larger query cost on Query TE2 and SE2 than ODBJ (Bitonic), since oblivious heap sort [28] is suitable in memory but not IO-efficient. For our method, Sep INLJ achieves 1.1X-3.4X better performance than One INLJ, as explained in Section IX-C1. As in Section IX-C1, One SMJ achieves 1.4X-1.7X better performance than Sep SMJ on Query SE2 due to less number of tuple retrievals. Last, the index caching brings 1.2X-2.0X speedup ratio.

D. Performance of Band Join

Fig. 13(a) shows query cost for band join on TPC-H in default setting. Compared with Raw INLJ(+Cache), our extended ODBJ (Bitonic) brings 30X and 58X larger query cost on Query TB1 and TB2, and Sep INLJ(+Cache) brings 164X and 288X larger query cost on Query TB1 and TB2. For extended ODBJ method, ODBJ (Heap) brings 9X and 5X larger query cost on Query TB1 and TB2 than ODBJ (Bitonic), as explained in Section IX-C. For ORAM based method, Sep INLJ achieves 1.4X-2.5X better performance than One INLJ, as explained in Section IX-C. The index caching brings 1.2X-1.5X better performance. Fig. 14(a) shows query cost on Query TB1 against raw data size. When raw data size varies from 20 MB to 1 GB, our extended ODBJ (Bitonic) and Sep INLJ(+Cache) bring 15X-66X and 73X-264X larger query cost on Query TB1 compared with Raw INLJ(+Cache). For extended ODBJ method, ODBJ (Heap) brings 5X-15X larger query cost on Query TB1 than ODBJ (Bitonic). For ORAM based method, Sep INLJ achieves 1.9X-2.9X better performance than One INLJ, and index caching achieves 1.2X-1.5X better performance.

E. Performance of Multiway Equi-Join

1) *Default Setting:* Figs. 15(a) and 16(a) show query cost for multiway equi-join on two datasets in default setting. Our Sep INLJ(+Cache) achieves 10^6 X- 10^{11} X better performance than OblIDB on all queries except Query TM2. The reason is that our query cost is roughly *linear* with input and output sizes, but

OblIDB has to perform a Cartesian product. For Query TM2, the speedup ratio goes down to 280X, since the join result size is roughly proportional to Cartesian product size. Compared with Raw INLJ(+Cache), Sep INLJ(+Cache) brings 185X-985X and 37000X-70000X larger query cost on TPC-H and social graph, due to ensuring the obliviousness. For our method, Sep INLJ achieves 1.6X-5.5X better performance than One INLJ, since One INLJ has to access the large single ORAM. Last, index caching brings 1.1X-1.5X speedup ratio.

F. Padded Mode Versus Non-Padded Mode

We also make the comparison between padded mode and non-padded mode for all secured methods. We discuss three padding strategies for join result size: (1) no padding (denoted as Real Size); (2) padding to closest power of 2 (denoted as Closest Power) [74], [75], [76]; (3) padding to Cartesian product (denoted as Cartesian Product).

Figs. 17–19 show query cost against different padding strategies in default setting. For OblIDB, Cartesian Product even achieves 5X less query cost than Real Size and Closest Power, since Real Size and Closest Power need an additional oblivious filtering over the join output with Cartesian product size. For ODBJ and ORAM based method, query cost is roughly proportional to different ratios of padded join sizes to real join sizes, e.g., Closest Power introduces within 2X larger query cost than Real Size, due to padding to closest power of 2. In Cartesian Product, ODBJ (Bitonic) needs 40X-50X larger query cost than OblIDB, and ORAM based method brings 500X-1700X and 900X-5300X larger query cost on binary and multiway equi-joins than OblIDB. The first reason is that ODBJ and ORAM based method incur $\Omega(\log N)$ bandwidth overhead to ensure the obliviousness. The second reason is that we set trusted memory size $M = 50 \log N$ in OblIDB. In Cartesian Product, ODBJ (Heap) even achieves 2X and 4X better performance on Query TE2 and SE2 than ODBJ (Bitonic), since oblivious heap sort [28] achieves lower time complexity than external bitonic sort ($O(N \log N)$ versus $O(N \log^2(N/M))$), especially when the join output size is huge.

G. Access Pattern Logs

For security analysis, we verify the obliviousness of our method by comparing the logs of access patterns for different inputs, as with “Experiments: Memory Access Logs” paragraph in Section VI.A in [31]. We also visualize the access patterns in Figs. 20 and 21, as with Fig. 7 in [31]. Fig. 20 shows the access pattern of our extended oblivious band join algorithm in ODBJ (Bitonic). It joins T_1 and T_2 of size 4 into T_{out} of size 9, as with Examples 2 and 3. Specifically, horizontal axis means the discretized time, and vertical axis means the tuple index. Each light bar means a tuple read, and each dark bar means a tuple write. Fig. 21 shows the access pattern of oblivious multiway equi-join algorithm (without index caching). It joins 4 tables with $|T_1| = |T_2| = 4$ and $|T_3| = |T_4| = 3$ into T_{out} of size 2, as with Example 6. Each B -tree index for T_1 - T_4 has 3 levels: root node level, leaf entry level, and data tuple level. Specifically, horizontal axis means the discretized time. For T_1 - T_4 , vertical axis means the index level in B -tree for T_1 - T_4 ; each light bar means an ORAM read, and each dark bar means an ORAM write. For T_{out} , vertical axis means the record index; each light bar means a record read, and each dark bar means a record write. We have verified that given the specific input and output sizes ranging from 10 to 10,000, the tests for different input tuples produce the same logs of access patterns.

X. CONCLUSION

This work supports general band joins and multiway equi-joins obliviously based on non-ORAM approach [31] and ORAM approach [50]. Non-ORAM approach stores input tables in flat storage and achieves better performance in join processing, but needs some delicate design of oblivious operations. ORAM approach builds oblivious indices over input tables directly, but usually brings larger computation overhead in join processing. As with OblIDB [32], accessing a few rows in any table should use the indexed storage, while the flat storage performs better for accessing large segments. Hence, to design the query optimizer for different approaches is a crucial point in building encrypted or oblivious databases. Note that our current design does not address challenges associated with ad-hoc updates, which is a future direction. Last, how to support query concurrency in an efficient manner using ORAM is still a major challenge.

REFERENCES

- [1] A. Arasu, K. Eguro, M. Joglekar, R. Kaushik, D. Kossmann, and R. Ramamurthy, “Transaction processing on confidential data using Cipherbase,” in *Proc. IEEE Int. Conf. Data Eng.*, 2015, pp. 435–446.
- [2] A. Arasu et al., “Secure database-as-a-service with Cipherbase,” in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2013, pp. 1033–1036.
- [3] R. A. Popa, C. M. S. Redfield, N. Zeldovich, and H. Balakrishnan, “CryptDB: Protecting confidentiality with encrypted query processing,” in *Proc. Symp. Operating Syst. Princ.*, 2011, pp. 85–100.
- [4] S. Bajaj and R. Sion, “TrustedDB: A trusted hardware-based database with privacy and data confidentiality,” *IEEE Trans. Knowl. Data Eng.*, vol. 26, no. 3, pp. 752–765, Mar. 2014.
- [5] Z. He et al., “SDB: A secure query processing system with data interoperability,” in *Proc. VLDB Endowment*, vol. 8, no. 12, pp. 1876–1879, 2015.
- [6] S. Tu, M. F. Kaashoek, S. Madden, and N. Zeldovich, “Processing analytical queries over encrypted data,” in *Proc. VLDB Endowment*, vol. 6, no. 5, pp. 289–300, 2013.
- [7] A. Arasu, K. Eguro, R. Kaushik, and R. Ramamurthy, “Querying encrypted data,” in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2014, pp. 1259–1261.
- [8] H. Hacigümüs, B. R. Iyer, C. Li, and S. Mehrotra, “Executing SQL over encrypted data in the database-service-provider model,” in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2002, pp. 216–227.
- [9] B. Yao, F. Li, and X. Xiao, “Secure nearest neighbor revisited,” in *Proc. IEEE Int. Conf. Data Eng.*, 2013, pp. 733–744.
- [10] W. K. Wong, B. Kao, D. W. Cheung, R. Li, and S. Yiu, “Secure query processing with data interoperability in a cloud database environment,” in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2014, pp. 1395–1406.
- [11] Z. Chang, D. Xie, and F. Li, “Oblivious RAM: A dissection and experimental evaluation,” in *Proc. VLDB Endowment*, vol. 9, no. 12, pp. 1113–1124, 2016.
- [12] W. Zheng, A. Dave, J. G. Beekman, R. A. Popa, J. E. Gonzalez, and I. Stoica, “Opaque: An oblivious and encrypted distributed analytics platform,” in *Proc. 14th USENIX Conf. Netw. Syst. Des. Implementation*, 2017, pp. 283–298.
- [13] A. Arasu and R. Kaushik, “Oblivious query processing,” in *Proc. 17th Int. Conf. Database Theory*, 2014, pp. 26–37.
- [14] T. Hoang, C. D. Ozkaptan, G. Hackeheil, and A. A. Yavuz, “Efficient oblivious data structures for database services on the cloud,” *IEEE Trans. Cloud Comput.*, vol. 9, no. 2, pp. 598–609, Second Quarter 2018.
- [15] M. Naveed, S. Kamara, and C. V. Wright, “Inference attacks on property-preserving encrypted databases,” in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2015, pp. 644–655.
- [16] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart, “Leakage-abuse attacks against searchable encryption,” in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2015, pp. 668–679.
- [17] G. Kellaris, G. Kollios, K. Nissim, and A. O’Neill, “Generic attacks on secure outsourced databases,” in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2016, pp. 1329–1340.
- [18] V. Bindshaedler, P. Grubbs, D. Cash, T. Ristenpart, and V. Shmatikov, “The tao of inference in privacy-protected databases,” in *Proc. VLDB Endowment*, vol. 11, no. 11, pp. 1715–1728, 2018.
- [19] B. Pinkas and T. Reinman, “Oblivious RAM revisited,” in *Proc. 30th Annu. Cryptol. Conf.*, 2010, pp. 502–519.
- [20] M. S. Islam, M. Kuzu, and M. Kantarcioglu, “Access pattern disclosure on searchable encryption: Ramification, attack and mitigation,” in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2012.
- [21] O. Goldreich, “Towards a theory of software protection and simulation by oblivious RAMs,” in *Proc. 19th Annu. ACM Symp. Theory Comput.*, 1987, pp. 182–194.
- [22] R. Ostrovsky, “Efficient computation on oblivious RAMs,” in *Proc. Annu. ACM Symp. Theory Comput.*, 1990, pp. 514–523.
- [23] O. Goldreich and R. Ostrovsky, “Software protection and simulation on oblivious RAMs,” *J. ACM*, vol. 43, no. 3, pp. 431–473, 1996.
- [24] M. Keller and P. Scholl, “Efficient, oblivious data structures for MPC,” in *Proc. 20th Int. Conf. Theory Appl. Cryptol. Inf. Secur., Part II*, 2014, pp. 506–525.
- [25] X. S. Wang et al., “Oblivious data structures,” in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2014, pp. 215–226.
- [26] P. Mishra, R. Poddar, J. Chen, A. Chiesa, and R. A. Popa, “Obliv: An efficient oblivious search index,” in *Proc. IEEE Symp. Secur. Privacy*, 2018, pp. 279–296.
- [27] T. Hoang, M. O. Ozmen, Y. Jang, and A. A. Yavuz, “Hardware-supported ORAM in effect: Practical oblivious search and update on very large dataset,” in *Proc. Priv. Enhancing Technol.*, vol. 2019, no. 1, pp. 172–191, 2019.
- [28] E. Shi, “Path oblivious heap: Optimal and practical oblivious priority queue,” in *Proc. IEEE Symp. Secur. Privacy*, 2020, pp. 842–858.
- [29] Z. Jafarholi, K. G. Larsen, and M. Simkin, “Optimal oblivious priority queues,” in *Proc. ACM-SIAM Symp. Discrete Algorithms*, 2021, pp. 2366–2383.
- [30] Y. Li and M. Chen, “Privacy preserving joins,” in *Proc. IEEE Int. Conf. Data Eng.*, 2008, pp. 1352–1354.
- [31] S. Krastnikov, F. Kerschbaum, and D. Stebila, “Efficient oblivious database joins,” in *Proc. VLDB Endowment*, vol. 13, no. 11, pp. 2132–2145, 2020.
- [32] S. Eskandarian and M. Zaharia, “OblIDB: Oblivious query processing for secure databases,” in *Proc. VLDB Endowment*, vol. 13, no. 2, pp. 169–183, 2019.
- [33] M. T. Goodrich, “Data-oblivious external-memory algorithms for the compaction, selection, and sorting of outsourced data,” in *Proc. Annu. ACM Symp. Parallelism Algorithms Architectures*, 2011, pp. 379–388.

- [34] D. J. DeWitt, J. F. Naughton, and D. A. Schneider, "An evaluation of non-equijoin algorithms," in *Proc. 17th Int. Conf. Very Large Data Bases*, 1991, pp. 443–452.
- [35] L. Ren et al., "Constants count: Practical improvements to oblivious RAM," in *Proc. USENIX Secur. Symp.*, 2015, pp. 415–430.
- [36] H. Chen, I. Chillotti, and L. Ren, "Onion ring ORAM: Efficient constant bandwidth oblivious RAM from (leveled) TFHE," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2019, pp. 345–360.
- [37] G. Asharov, I. Komargodski, W. Lin, K. Nayak, E. Peseiro, and E. Shi, "OptORAMa: Optimal oblivious RAM," in *Proc. 39th Annu. Int. Conf. Theory Appl. Cryptogr. Techn., Part II*, 2020, pp. 403–432.
- [38] P. Williams, R. Sion, and A. Tomescu, "PrivateFS: A parallel oblivious file system," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2012, pp. 977–988.
- [39] J. R. Lorch, B. Parno, J. W. Mickens, M. Raykova, and J. Schiffman, "Shroud: Ensuring private access to large-scale data in the data center," in *Proc. 11th USENIX Conf. File Storage Technol.*, 2013, pp. 199–214.
- [40] E. Stefanov and E. Shi, "ObliviStore: High performance oblivious cloud storage," in *Proc. IEEE Symp. Secur. Privacy*, 2013, pp. 253–267.
- [41] V. Bindschaedler, M. Naveed, X. Pan, X. Wang, and Y. Huang, "Practicing oblivious access on cloud storage: The gap, the fallacy, and the new way forward," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2015, pp. 837–849.
- [42] C. Sahin, V. Zakhary, A. E. Abbadi, H. Lin, and S. Tessaro, "TaoStore: Overcoming asynchronicity in oblivious data storage," in *Proc. IEEE Symp. Secur. Privacy*, 2016, pp. 198–217.
- [43] A. Chakraborti and R. Sion, "ConcurORAM: High-throughput stateless parallel multi-client ORAM," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2019.
- [44] S. Maiyya et al., "QuORAM: A quorum-replicated fault tolerant ORAM datastore," in *Proc. USENIX Secur. Symp.*, 2022, pp. 3665–3682.
- [45] E. Stefanov et al., "Path ORAM: An extremely simple oblivious RAM protocol," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2013, pp. 299–310.
- [46] K. E. Batchier, "Sorting networks and their applications," in *Proc. AFIPS Spring Joint Comput. Conf.*, 1968, pp. 307–314.
- [47] M. Ajtai, J. Komlós, and E. Szemerédi, "An $O(n \log n)$ sorting network," in *Proc. 15th Annu. ACM Symp. Theory Comput.*, 1983, pp. 1–9.
- [48] M. T. Goodrich, "Randomized Shellsort: A simple oblivious sorting algorithm," in *Proc. 21st Annu. ACM-SIAM Symp. Discrete Algorithms*, 2010, pp. 1262–1277.
- [49] M. T. Goodrich, "Zig-zag sort: A simple deterministic data-oblivious sorting algorithm running in $O(n)$ time," in *Proc. Annu. ACM Symp. Theory Comput.*, 2014, pp. 684–693.
- [50] Z. Chang, D. Xie, S. Wang, and F. Li, "Towards practical oblivious join," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2022, pp. 803–817.
- [51] Z. Chang, D. Xie, F. Li, J. M. Phillips, and R. Balasubramanian, "Efficient oblivious query processing for range and kNN queries," *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 12, pp. 5741–5754, Dec. 2022.
- [52] M. Maas et al., "PHANTOM: Practical oblivious computation in a secure processor," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2013, pp. 311–324.
- [53] S. Sasy, S. Gorbunov, and C. W. Fletcher, "ZeroTrace: Oblivious memory primitives from intel SGX," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2018.
- [54] D. Xie et al., "Practical private shortest path computation based on oblivious storage," in *Proc. IEEE Int. Conf. Data Eng.*, 2016, pp. 361–372.
- [55] N. Crooks, M. Burke, E. Cecchetti, S. Harel, R. Agarwal, and L. Alvisi, "Obladi: Oblivious serializable transactions in the cloud," in *Proc. USENIX Conf. Operating Syst. Des. Implementation*, 2018, pp. 727–743.
- [56] A. Dave, C. Leung, R. A. Popa, J. E. Gonzalez, and I. Stoica, "Oblivious cooperative analytics using hardware enclaves," in *Proc. Eur. Conf. Comput. Syst.*, 2020, pp. 39:1–39:17.
- [57] E. Dauterman, V. Fang, I. Demertzis, N. Crooks, and R. A. Popa, "Snoopy: Surpassing the scalability bottleneck of oblivious storage," in *Proc. Symp. Operating Syst. Princ.*, 2021, pp. 655–671.
- [58] S. Chu, D. Zhuo, E. Shi, and T. H. Chan, "Differentially oblivious database joins: Overcoming the worst-case curse of fully oblivious algorithms," in *Proc. 2nd Conf. Inf.-Theoretic Cryptogr.*, 2021, pp. 19:1–19:24.
- [59] V. Costan and S. Devadas, "Intel SGX explained," *IACR Cryptol. ePrint Arch.*, vol. 2016, 2016, Art. no. 86.
- [60] T. Kim, Z. Lin, and C. Tsai, "CCS'17 tutorial abstract: SGX security and privacy," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2017, pp. 2613–2614.
- [61] X. S. Wang, Y. Huang, T.-H. H. Chan, A. Shelat, and E. Shi, "SCORAM: Oblivious RAM for secure computation," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2014, pp. 191–202.
- [62] C. Liu, X. S. Wang, K. Nayak, Y. Huang, and E. Shi, "ObliVM: A programming framework for secure computation," in *Proc. IEEE Symp. Secur. Privacy*, 2015, pp. 359–376.
- [63] J. Bater, G. Elliott, C. Egeen, S. Goel, A. N. Kho, and J. Rogers, "SM-CQL: Secure query processing for private data networks," in *Proc. VLDB Endowment*, vol. 10, no. 6, pp. 673–684, 2017.
- [64] N. Volgushev, M. Schwarzkopf, B. Getchell, M. Varia, A. Lapets, and A. Bestavros, "Conclave: Secure multi-party computation on big data," in *Proc. Eur. Conf. Comput. Syst.*, 2019, pp. 3:1–3:18.
- [65] Y. Wang and K. Yi, "Secure yannakakis: Join-aggregate queries over private data," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2021, pp. 1969–1981.
- [66] C. Sahin, T. Allard, R. Akbarinia, A. E. Abbadi, and E. Pacitti, "A differentially private index for range query processing in clouds," in *Proc. IEEE Int. Conf. Data Eng.*, 2018, pp. 857–868.
- [67] Q. Ye, H. Hu, X. Meng, and H. Zheng, "PrivKV: Key-value data collection with local differential privacy," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 317–331.
- [68] N. M. Johnson, J. P. Near, and D. Song, "Towards practical differential privacy for SQL queries," in *Proc. VLDB Endowment*, vol. 11, no. 5, pp. 526–539, 2018.
- [69] J. Bater, X. He, W. Ehrlich, A. Machanavajjhala, and J. Rogers, "Shrinkwrap: Efficient SQL query processing in differentially private data federations," in *Proc. VLDB Endowment*, vol. 12, no. 3, pp. 307–320, 2018.
- [70] I. Kotsogiannis et al., "PrivateSQL: A differentially private SQL query engine," in *Proc. VLDB Endowment*, vol. 12, no. 11, pp. 1371–1384, 2019.
- [71] N. Wang et al., "Collecting and analyzing multidimensional data with local differential privacy," in *Proc. IEEE Int. Conf. Data Eng.*, 2019, pp. 638–649.
- [72] T. Wang et al., "Answering multi-dimensional analytical queries under local differential privacy," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2019, pp. 159–176.
- [73] J. Yang, T. Wang, N. Li, X. Cheng, and S. Su, "Answering multi-dimensional range queries under local differential privacy," in *Proc. VLDB Endowment*, vol. 14, no. 3, pp. 378–390, 2020.
- [74] A. Chakraborti, A. J. Aviv, S. G. Choi, T. Mayberry, D. S. Roche, and R. Sion, "rORAM: Efficient range ORAM with $O(\log^2 N)$ locality," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2019.
- [75] G. Asharov, T. H. Chan, K. Nayak, R. Pass, L. Ren, and E. Shi, "Locality-preserving oblivious RAM," in *Proc. 8th Annu. Int. Conf. Theory Appl. Cryptogr. Techn., Part II*, 2019, pp. 214–243.
- [76] I. Demertzis, D. Papadopoulos, C. Papamanthou, and S. Shintre, "SEAL: Attack mitigation for encrypted databases via adjustable leakage," in *Proc. USENIX Secur. Symp.*, 2020.
- [77] C. W. Fletcher, L. Ren, X. Yu, M. van Dijk, O. Khan, and S. Devadas, "Suppressing the oblivious RAM timing channel while making information leakage and program efficiency trade-offs," in *Proc. Int. Symp. High Perform. Comput. Architecture*, 2014, pp. 213–224.
- [78] S. Chen, X. Zhang, M. K. Reiter, and Y. Zhang, "Detecting privileged side-channel attacks in shielded execution with Déjà Vu," in *Proc. ACM Asia Conf. Comput. Commun.*, 2017, pp. 7–18.
- [79] D. Gruss, J. Lettner, F. Schuster, O. Ohrimenko, I. Haller, and M. Costa, "Strong and efficient cache side-channel protection using hardware transactional memory," in *Proc. USENIX Secur. Symp.*, 2017, pp. 217–233.
- [80] Z. Chang, D. Xie, S. Wang, F. Li, and Y. Shen, "Towards practical oblivious join processing," 2022. [Online]. Available: <https://zhao-chang.github.io/paper/ojoin.pdf>
- [81] C. T. Yu and M. Z. Ozsoyoglu, "An algorithm for tree-query membership of a distributed query," in *Proc. IEEE Comput. Soc. 3rd Int. Comput. Softw. Appl. Conf.*, 1979, pp. 306–312.
- [82] "Oblivious database join algorithm," 2020. [Online]. Available: <https://git.uwaterloo.ca/skrastni/obliv-join-impl>
- [83] M. Cha, H. Haddadi, F. Benevenuto, and P. K. Gummadi, "Measuring user influence in twitter: The million follower fallacy," in *Proc. 4th Int. Conf. Weblogs Soc. Media*, 2010.
- [84] Z. Zhao, R. Christensen, F. Li, X. Hu, and K. Yi, "Random sampling over joins revisited," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2018, pp. 1525–1539.