

Efficient Oblivious Data Structures for Database Services on the Cloud

Thang Hoang, Ceyhun D. Ozkaptan, Gabriel Hackebeil, and Attila A. Yavuz, *Member, IEEE*

Abstract—Database-as-a-service (DBaaS) allows the client to store and manage structured data on the cloud remotely. Despite its merits, DBaaS also brings significant privacy issues. Existing encryption techniques (e.g., SQL-aware encryption) can mitigate privacy concerns, but they still leak information through access patterns which are vulnerable to statistical inference attacks. Oblivious Random Access Machine (ORAM) can seal such leakages, but the recent studies showed significant challenges on the integration of ORAM into databases. Specifically, the direct usage of ORAM on databases is not only costly but also permits very limited query functionalities. We propose new oblivious data structures called *Oblivious Matrix Structure* (OMAT) and *Oblivious Tree Structure* (OTREE), which allow tree-based ORAM to be integrated into database systems in a more efficient manner with diverse query functionalities supported. OMAT provides special ORAM packaging strategies for table structures, which not only offers a significantly better performance but also enables a broad range of query types that may not be practical in existing frameworks. OTREE allows oblivious conditional queries to be deployed on tree-indexed databases more efficient than existing techniques. We fully implemented our proposed techniques and evaluated their performance on a real cloud database with various metrics, compared with state-of-the-art counterparts.

Index Terms—Privacy-enhancing Technologies; Oblivious Data Structure; ORAM

1 INTRODUCTION

Services for outsourcing data storage and related infrastructure to the cloud have grown in the last decade due to the savings it offers to companies in terms of capital and operational costs. For instance, major cloud providers (e.g., Amazon, Microsoft) offer Database-as-a-service (DBaaS) that provides relational database management systems on the cloud. This enables a client to store and manage structured data remotely. Despite its merits, DBaaS raises privacy issues. The client may encrypt the data with standard encryption, but this then prevents searching or updating information over the cloud, thereby invalidating the effective utilization of database.

Various privacy enhancing technologies have been developed towards addressing this privacy versus data utilization dilemma. For instance, the client can use special encryption techniques such as SQL-aware encryption (e.g., [1], [2]) or searchable encryption with various security, efficiency and query functionality trade-offs ((e.g., [3], [4], [5], [6], [7], [8], [9], [10])) to achieve the data confidentiality and usability on the cloud. However, even such encryption techniques might not be sufficient for privacy-critical database applications (e.g., healthcare), as sensitive information may be revealed through access patterns during the execution of queries on the encrypted database. Recent work (e.g., [11],

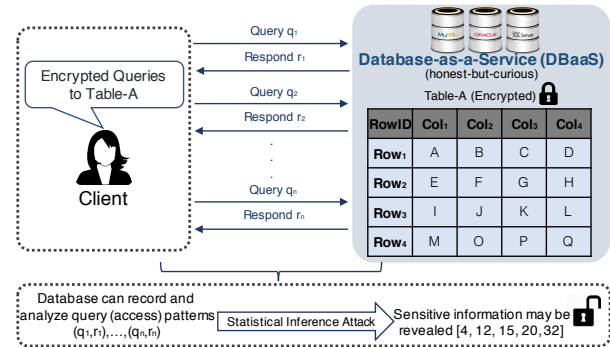


Fig. 1: Problem Statement: Information leakages through query access patterns over an encrypted database.

[12], [13], [14], [15]) showed that the information leakage through access patterns along with some prior contextual knowledge can be used to launch statistical inference attacks and reveal vital information about encrypted queries and database. For example, such information leaks may expose the prognosis of illness for a patient or types/timing of financial transactions over valuable assets based on encrypted queries. Therefore, it is an important requirement for privacy-critical database applications to hide the access pattern.

Oblivious Random Access Machine (ORAM) [16] can be used to hide access patterns for such encrypted databases. Preliminary ORAM schemes (e.g., [16], [17]) were costly, but recent ORAM constructions (e.g., [18], [19], [20], [21], [22], [23]) showed promising results. Most efficient ORAM schemes (e.g., [21], [24], [25]) to date follow tree paradigm [19], which achieves asymptotic $\Omega(\log N)$ lower-bound [16], [22], [26] communication overhead. However, despite these improvements, there are several research gaps towards achieving efficient integration of ORAM in databases. In the

- Thang Hoang is with the School of EECS, Oregon State University, Corvallis, OR, 97331. E-mail: hoangmin@oregonstate.edu
- Attila A. Yavuz is with the Department of Computer Science and Engineering, University of South Florida, Tampa, FL, 33620. E-mail: attilayavuz@usf.edu
- Gabriel Hackebeil is with the Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, 48109.
- Ceyhun D. Ozkaptan is with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210.
- Work done when the second, the third and the fourth authors were employed at Oregon State University. E-mail: {ozkaptan, hackebeg, attila.yavuz}@oregonstate.edu.



2168-7161 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

TABLE 1: Transmission cost and client storage for compared schemes.

Scheme	Communication Cost ^a	Efficiency ^b	Client Storage ^c	End-to-End Delay ^d	
				Moderate Network	High Network
single column-related query (e.g., statistical, conditional queries)					
RowPKG [29]	$Z \cdot (B_1 \cdot N) \cdot (2M - 1)$	1.00	$O(M \cdot N) \cdot w(1)$	6096 s	776 s
ODS-2D [23]	$(M/4) \cdot [Z \cdot (16 \cdot B_1) \cdot \log_2(M \cdot N/16)]$	17.04	$O(M \cdot \log(M \cdot N)) \cdot w(1)$	1245 s	292 s
OMAT	$Z^2 \cdot (B_1 \cdot M) \cdot \log_2(N)$	28.44	$O(M \cdot \log(N)) \cdot w(1)$	475 s	60 s
single row-related query (e.g., insert/delete/update queries)					
RowPKG [29]	$Z \cdot (B_2 \cdot N) \cdot \log_2(M)$	1.00	$O(N \cdot \log(M)) \cdot w(1)$	567 ms	56 ms
ODS-2D [23]	$(N/4) \cdot [Z \cdot (16 \cdot B_2) \cdot \log_2(M \cdot N/16)]$	0.19	$O(N \cdot \log(M \cdot N)) \cdot w(1)$	2380 ms	350 ms
OMAT	$Z^2 \cdot (B_2 \cdot N) \cdot \log_2(M)$	0.25	$O(N \cdot \log(M)) \cdot w(1)$	2032 ms	128 ms
traversal on database tree index (e.g., range queries)					
non-caching					
ODS-Tree [23]	$2 \cdot Z_1 \cdot B \cdot (H + 1)^2$	1.00	$O(H) \cdot w(1)$	7929 ms	1318 ms
OTREE	$Z_2 \cdot B \cdot (H + 1) \cdot (H + 2)$	1.60	$O(H) \cdot w(1)$	3762 ms	592 ms
half-top caching					
ODS-Tree [23]	$2 \cdot Z_1 \cdot B \cdot \left\lceil \frac{H+1}{2} \right\rceil \cdot (H + 1)$	1.00	$O(\sqrt{2^H}) + O(H) \cdot w(1)$	5979 ms	1008 ms
OTREE	$Z_2 \cdot B \cdot \left\lceil \frac{H+1}{2} \right\rceil \cdot \left(\left\lceil \frac{H+1}{2} \right\rceil + 1 \right)$	3.20	$O(\sqrt{2^H}) + O(H) \cdot w(1)$	1676 ms	272 ms

- *Table Notations:* M and N denote the total number of (real) rows and columns in the matrix data structure, respectively. H is the height of the tree data structure. Z and B denote the bucket size and size of each block (in bytes), respectively.
- *Settings:* We instantiate our schemes and their counterparts with underlying Path-ORAM for a fair comparison. The bottom half of the table compares OTREE and ODS-Tree when combined with tree-top caching technique proposed in [30], in which we assume the top half of tree-based ORAM is cached on the client during all access requests.
- *Server Storage:* All of the oblivious matrix structures require $O(MN)$ server storage, however, the storage of OMAT is a constant (e.g., $Z = 4$) factor larger than others. OTREE is twice more storage efficient than ODS.
- ^a Represents the total cost in terms of bytes to be processed (e.g., communication/computation depends on the underlying ORAM scheme) between the client and the server for each request. For OMAT, ODS-2D and RowPKG, the cost is for one access operation per query. For OTREE and ODS, the cost is for traversing an arbitrary path in a binary tree.
- ^b Denotes the communication cost efficiency compared to chosen baseline, where $Z = 4$, $B_1 = 64$, $B_2 = 128$, $M = 2^{15}$, $N = 2^9$ for ODS-2D, RowPKG and OMAT, and $Z_1 = 4$, $Z_2 = 5$ (for stability), $B = 4096$, $H = 20$ for ODS-Tree and OTREE.
- ^c Client storage consists of the worst-case stash size to keep fetched data. Additionally, the position map of OMAT and RowPKG are $O((M + N) \log(M + N))$ and $O(M \cdot \log(M))$, respectively. For ODS based structures and OTREE, position map requires $O(1)$ storage due to pointers and half-top cached blocks are also included in client storage.
- ^d The delays were measured with a MongoDB instance running on Amazon EC2 connected with the client on two different network settings which are described in Section 5.1.

performance than the existing oblivious data structures [23] for such settings since the structure of the data is already known.

We illustrate desirable properties of our schemes in Figure 2-(c,d), and further discuss them as below:

- *Highly efficient and diverse oblivious queries:* OMAT supports a diverse set of queries to be executed with ORAM. Specifically, OMAT permits oblivious statistical queries over value-based columns such as SUM, AVG, MAX and MIN. Moreover, oblivious queries on rows (e.g., insert, update) can be executed on an attribute with a similar cost. As shown in Table 1, with the given parameters and experimental setup, executing a column-related query such as statistical or conditional query with OMAT is approximately $28\times$ more communication efficient than that of RowPKG and this enables OMAT to perform queries approximately $13\times$ faster than that of RowPKG. Compared to ODS-2D, although OMAT is only $1.6\times$ communication efficient, it performs approximately $5\times$ faster in practice due to the large number of additional round-trip delays. OTREE achieves better performance than ODS for obviously accessing database index, which is constructed from the values of a column as a tree data structure. The communication cost of OTREE is $1.6\times$ less than that of ODS without any caching. This gain can be increased up to $3.2\times$ with tree-top caching strategy.
- *Generic Instantiations from Tree-based ORAM Schemes:* Any tree-based ORAM scheme can be used for both OMAT and

OTREE instantiations. This provides a flexibility in selecting the base ORAM scheme, which can be adjusted according to the performance requirements of specific applications. Note that, in this paper, we instantiated our schemes with Path-ORAM [21] due to its efficiency, simplicity and not requiring any server-side computation.

- *Comprehensive Experiments and Open-source Framework:* We implemented OTREE, OMAT, and their counterparts under the same framework. We evaluated their performance with a database instance of MongoDB running on a remote AmazonEC2 server with two different network settings: (1) moderate-speed network and (2) high-speed network. This permits us to observe the impact of real network and cloud environment. We will also release our cryptographic software for public use.

2 PRELIMINARIES

We now present cryptographic techniques and implementation frameworks that are used by or are relevant to our proposed schemes.

2.1 Tree-based ORAM

ORAM enables a client to access encrypted data on an untrusted server without exposing the access patterns (e.g., memory blocks, their access time and order) to the server [16]. Existing ORAM schemes rely on IND-CPA encryption [31] and an oblivious shuffling to ensure that any

TABLE 2: Summary of notations in tree-based ORAM.

Symbol	Description
N	Total number of nodes in the tree-based ORAM
H	Height of the ORAM tree structure
b, B	Block and Block size
z	Capacity (in blocks) of each node
$\mathcal{P}(i)$	Path from leaf node i to root bucket in the tree
$\mathcal{P}(i, \ell)$	Bucket at level ℓ along the path $\mathcal{P}(i)$
S	Client's local stash (optional)
pm	Client's local position map
$i := \text{pm}[\text{id}]$	block identified by id is currently associated with leaf node i , i.e., it resides somewhere along $\mathcal{P}(i)$ or (optional) in the stash.




Fig. 3: Tree-based ORAM structure [19].

data access patterns of the same length are computationally indistinguishable by anyone but the client.

Recent ORAMs (e.g., [21], [24], [25], [32]) follow the tree paradigm [19], which consists of two main data structures: A perfectly-balanced tree data structure stored on the server side and a position map (denoted as pm) stored at client side (Figure 3). Each node in the tree is called as a bucket (denoted as B) which can store up to Z data blocks (e.g., $Z = 4$). Each block b has a unique identifier id and all blocks have the same size B (4 KB). A tree-based ORAM with N leaf nodes can store up to N real blocks, and other empty slots are filled with dummy data. $\mathcal{P}(i)$ denotes a path from the root to leaf i of the tree. The position map pm holds the location among 2^N possible paths $\mathcal{P}(i)$ for every block with identifier id . The size of pm is $\mathcal{O}(N \log N)$ which can be reduced to $\mathcal{O}(1)$ by using recursive ORAMs to store pm on the server with the communication overhead up to $\mathcal{O}(\log N)$ for each access operation. Table 2 summarizes notations being used for tree-based ORAM scheme.

There are two basic procedures in tree-based ORAMs: Block Retrieval and Eviction. For each access operation, the client gets the path ID of accessing block from the position map and sends the path ID to the server who responds with all blocks residing in the requested path. The client decrypts and processes the received data to obtain the desired block and executes the eviction, which re-encrypts downloaded block(s) and pushes them back to the ORAM tree. Notice that although recently proposed ORAM schemes that follow the tree paradigm (e.g., [21], [24], [25]) provide different trade-offs between communication and computation overhead, they all rely on the aforementioned basic procedures.

Path-ORAM: Path-ORAM [21] follows the same block fetching procedure of tree-based ORAM schemes. After fetching, all downloaded real blocks are temporarily placed in client's stash, a component introduced in [20]. A new random




Fig. 4: Oblivious Data Structure for a tree structure [23].

address is then assigned to the accessed block and the local position map is updated. Next, the blocks in the stash are evicted to the same path in the server. Path-ORAM offers asymptotically optimal communication and computation cost of $\mathcal{O}(\log N)$ by storing $\mathcal{O}(N \log N)$ -sized position map. As the recursive ORAMs were known to be highly costly, we do not discuss them in this work.

2.2 Oblivious Data Structure

Oblivious Data Structure (ODS) proposed by Wang et al. [23] leverages “pointer techniques” to reduce the bulk storage of position map components in non-recursive ORAM schemes to $\mathcal{O}(1)$, if the data to be accessed have some specific structures (e.g., grid, tree, etc.). For instance, given a binary search-sorted array as illustrated in Figure 4, the ORAM block is augmented with $k + 1$ additional slots that hold the position of the block along with the positions and identifiers of its children as $b := (\text{id}, \text{data}, \text{pos}, \text{childmap})$, where id is the block identifier, data is the block data, pos is its position in ORAM structure, and childmap is a miniature position map with entries $(\text{id}_i, \text{pos}_i)$ for k children. To ensure that the childmap is up to date, a child block must be accessed through at most one parent at any given time. If a block does not have a parent (e.g., the root of a tree), its position will be stored in the client. A parent block should never be written back to the server without updating positions of its children blocks.

2.3 ORAM Implementation Framework

One of the most reliable and complete ORAM frameworks is CURIOUS [27], which gives a complete implementation of the state-of-the-art ORAM schemes (e.g., Path-ORAM [21]) in Java. In this paper, we chose CURIOUS to implement our oblivious data structures as it can be adopted with database drivers such as MongoDB or MySQL.

3 PROPOSED TECHNIQUES




Fig. 5: Overview of our proposed techniques.

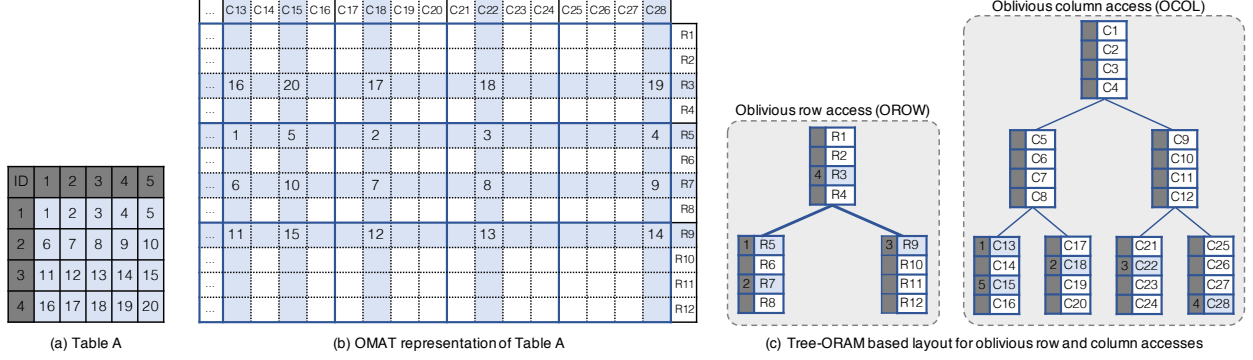


Fig. 6: OMAT structure for oblivious access on table.

We now present our proposed oblivious data structures, which are specially designed for efficient operations in database settings. We propose two schemes including *Oblivious Matrix Structure* (OMAT) and *Oblivious Tree Structure* (OTREE). OMAT supports efficient oblivious statistical queries on generic table instances, while OTREE supports range and conditional queries on tree-indexed instances. Figure 5 outlines the overview of our proposed techniques. For our oblivious data structures, we choose Path-ORAM [21] as underlying ORAM for the following reasons: (i) It is simple yet achieves asymptotic efficiency. (ii) Unlike some recent ORAMs [24], [25] that require computations at the server side, it requires only read/write operations. This is useful since such advanced cryptographic operations might not be readily offered by well-known database instances (e.g., MongoDB, MySQL). (iii) The availability of Path-ORAM implementations on existing frameworks (e.g., CURIOUS [27]) enables a fair experimental comparison of the proposed techniques with the state-of-the-art.

3.1 Oblivious Access on Table Structures

The direct application of tree-based ORAMs to access encrypted tables in general [27] and database systems in specific [29] have been shown to be highly inefficient for large datasets. Specifically, if each row in the table is packaged into an ORAM block as in [29], then performing queries to fetch a column in such a table (e.g., statistics) would require the client to download all blocks in the ORAM structure, which is impractical. Similarly, applying ORAM to each cell in the table incurs a high network delay and client storage overhead due to a large position map. Thus, we investigate on how to translate the table into an oblivious data structure so that each row and column of it can be both accessed efficiently by a given ORAM scheme. Below, we first describe our oblivious data structure and then present our OMAT access scheme on top of it.

- *Oblivious Data Structure for OMAT.* The main data structure that we use for oblivious access on a table is a matrix. Given an input table \mathbf{T} of size $M \times N$, we allocate a matrix \mathbf{M} of size $Z \cdot 2^{\lceil \log_2(M) \rceil - 1} \times Z \cdot 2^{\lceil \log_2(N) \rceil - 1}$. We arrange tree-based ORAM building blocks for oblivious access as follows:

The layout of OMAT matrix \mathbf{M} can be interpreted as two logical tree-based ORAMs defined as oblivious rows

Algorithm 1 $\text{data} \leftarrow \text{OMAT.Access}(\text{op}, \text{dim}, \text{id})$

```

1:  $b \leftarrow \text{pm}_{\text{dim}}[\text{id}].\text{pathID}$ 
2: if  $\text{dim} = \text{col}$  then
3:    $\text{pm}_{\text{dim}}[\text{id}].\text{pathID} \xleftarrow{\$} \{1, \dots, 2^{\lceil \log_2(N) \rceil - 1}\}$ 
4:    $H \leftarrow \lceil \log_2(N) \rceil$ 
5: else
6:    $\text{pm}_{\text{dim}}[\text{id}].\text{pathID} \xleftarrow{\$} \{1, \dots, 2^{\lceil \log_2(M) \rceil - 1}\}$ 
7:    $H \leftarrow \lceil \log_2(M) \rceil$ 
8:    $\triangleright$  Read all rows/columns on the path  $\mathcal{P}(b)$ 
9:   for each  $\ell \in \{0, \dots, H\}$  do
10:     $\mathcal{S}_{\text{dim}} \leftarrow \mathcal{S}_{\text{dim}} \cup \text{ReadBucket}(\text{dim}, \mathcal{P}(b, \ell))$ 
11:    $\text{data} \leftarrow \text{Read row/column with id from } \mathcal{S}_{\text{dim}}$ 
12:    $\text{data} \leftarrow \text{FilterDummy}(\text{data}, \mathcal{S}_{-\text{dim}})$ 
13:    $\mathcal{S}_{-\text{dim}} \leftarrow \text{Update}(\mathcal{S}_{-\text{dim}}, \text{pm}_{\text{dim}})$ 
14:   if  $\text{op} = \text{write}$  then
15:     $\mathcal{S}_{\text{dim}} \leftarrow (\mathcal{S}_{\text{dim}} \setminus \{(\text{id}, \text{data})\}) \cup \{(\text{id}, \text{data}^*)\}$ 
16:     $\triangleright$  Evict blocks from the stash
17:    for each  $\ell \in \{H, \dots, 0\}$  do
18:       $\mathcal{S}'_{\text{dim}} \leftarrow \{(\text{id}', \text{data}') \in \mathcal{S}_{\text{dim}} \mid \mathcal{P}(b, \ell) = \mathcal{P}(\text{pm}_{\text{dim}}[\text{id}']. \text{pathID}, \ell)\}$ 
19:       $\mathcal{S}'_{\text{dim}} \leftarrow \text{Select min}(|\mathcal{S}'_{\text{dim}}|, Z) \text{ blocks from } \mathcal{S}'_{\text{dim}}$ 
20:       $\mathcal{S}_{\text{dim}} \leftarrow \mathcal{S}_{\text{dim}} \setminus \mathcal{S}'_{\text{dim}}$ 
21:       $o \leftarrow 1$ 
22:      for each  $(\text{id}', \text{data}') \in \mathcal{S}'_{\text{dim}}$  do
23:         $\text{pm}[\text{id}']. \text{level} \leftarrow \ell$ 
24:         $\text{pm}[\text{id}']. \text{order} \leftarrow o, o \leftarrow o + 1$ 
25:         $\text{WriteBucket}(\text{dim}, \mathcal{P}(b, \ell), \mathcal{S}'_{\text{dim}})$ 
26: return data

```

(denoted as OROW) and oblivious columns (denoted as OCOL), as illustrated in Figure 6. That is, the ORAM for row access on OROW is formed by a set of blocks $b_i := (\text{id}_i, \text{data}_i)$, where id_i is either a unique identifier if b_i contains the content of a row of the table \mathbf{T} or otherwise null, and $\text{data}_i \leftarrow \mathbf{M}[i, *]$. We group Z subsequent rows in \mathbf{M} to form a bucket (i.e., node) in the OROW structure. Similarly, the ORAM for column access on OCOL is formed by $b_j = (\text{id}_j, \text{data}_j)$. Each (bucket) node in OCOL is formed by grouping Z subsequent blocks.

We assign each row $\mathbf{T}[i', *]$ ($i' = 1, \dots, M$) and each column $\mathbf{T}[*', j']$ ($j' = 1, \dots, N$) to a random leaf node IDs $u_{i'}$ and $v_{j'}$ in OROW and OCOL, respectively. That is, the data of $\mathbf{T}[i', *]$ and $\mathbf{T}[*', j']$ reside in some rows and columns of \mathbf{M} along the assigned paths $\mathcal{P}(u_{i'})$ in OROW and $\mathcal{P}(v_{j'})$ in OCOL, respectively. In other words, $\mathbf{M}[i', j] \leftarrow \mathbf{T}[i', j']$, where $\mathbf{M}[i', *] \in \mathcal{P}(u_{i'})$ in OROW and $\mathbf{M}[*', j] \in \mathcal{P}(v_{j'})$ in OCOL. Our construction requires two position maps pm_{row}

and pm_{col} to store assigned paths for each row $\mathbf{T}[i', *]$ and each column $\mathbf{T}[:, j']$ of table \mathbf{T} in OROW and OCOL, respectively. Our position maps store all necessary information to locate the exact position of a row/column data in the tree-based ORAM structures as $\text{pm} := (\text{id}, (\text{pathID}, \text{level}, \text{order}))$, where $0 \leq \text{level} \leq \log_2(N)$ indicates the level of the bucket, in which the row/column with id resides, and $1 \leq \text{order} \leq Z$ indicates its order in the bucket.

- *The Proposed OMAT Access Scheme.* We present our OMAT scheme, which is instantiated with Path-ORAM, in Algorithm 1. Specifically, given a *column* (resp. *row*) identifier (id) to be accessed¹, the client retrieves its location from the *column* (resp. *row*) position map (step 1). The client then assigns the *column* (resp. *row*) to a new location selected uniformly at random (steps 2–7). The client reads all *columns* (resp. *rows*) residing on the same path according to tree-based ORAM layout (as depicted in Figure 6-(c) to the stash (steps 8–9). In this case, we modify the original ReadBucket subroutine of Path-ORAM, where it now takes an extra parameter (dim) that indicates the dimension to be read, and outputs the corresponding Z columns/rows in the bucket. The client retrieves the *column* (resp. *row*) with id from the stash (step 10). One might observe that according to OMAT structure, the retrieved *column* (resp. *row*) will contain data from dummy rows (resp. columns) as depicted by empty blue cells in Figure 6-(b). Therefore, to obtain only the real data of the requested *column* (resp. *row*), the client filters all data from dummy rows (resp. columns) (step 11). Moreover, since the position of the retrieved column (resp. row) is moved to a new random position (steps 2–7), it is required to update all rows (resp. columns) that are currently stored in the stash at this column (resp. row) position to achieve the consistency (step 12). If the access is to update, the client then updates the *column* (resp. *row*) with new data (steps 13–14). Finally, the client performs eviction as described in Path-ORAM to flush columns (resp. rows) from the stash back to the OMAT structure in the server (steps 15–23).

Notice that all columns/rows are IND-CPA decrypted and re-encrypted as they are read and written to/from the server, respectively. We assume that it is not required to hide the information whether a column or a row is being accessed. However, this can be achieved with the cost of performing oblivious accesses on both row and column (one of them is dummy selected randomly) for each access.

- *Use Case: Statistical and Conditional Queries.* Recall that, in row-oriented packaging, implementing secure statistical queries on a column requires downloading the entire ORAM blocks from the database. In contrast, OMAT structure allows queries such as `add`, `delete`, `update` not only on its row but also on its column. Thus, we can implement statistical queries (e.g., `MAX`, `MIN`, `AVG`, `SUM`, `COUNT`, etc.) over a column in an efficient manner via OMAT. Note that OMAT also permits conditional query on rows with `WHERE` statement. Similar to statistical queries, the query can be implemented by reading the attribute column on which the `WHERE` clause looks up OCOL first to determine appropriate records that satisfy the condition, and then obliviously fetch-

ing such records on OROW structure. For example, assume that we have the following SQL-like conditional search:

```
SELECT * FROM A WHERE C > k;
```

It can be implemented by:

- 1) Read the column C with id' on OCOL:
 $C[:, \text{id}'] \leftarrow \text{OMAT.Access}(\text{read}, \text{col}, \text{id}')$.
- 2) Get IDs of rows whose value larger than k , and such IDs are in pm_{row} :
 $\mathcal{I} \leftarrow \{\text{id} \mid \text{id} \in \text{pm}_{\text{row}}, \text{id} \wedge C[\text{id}, \text{id}'] > k\}$
- 3) Access on OROW to get the desired result:
 $\mathbf{R}[\text{id}, *] \leftarrow \text{OMAT.Access}(\text{read}, \text{row}, \text{id})$, for each $\text{id} \in \mathcal{I}$.

The aforementioned approach can work with any unindexed columns. In the next section, we propose an alternative approach that can offer a better performance if the columns are indexed with certain restrictions.

3.2 Oblivious Access on Tree Structures

In an unencrypted database setting, conditional queries can be performed more efficiently, if column values are indexed by a search-efficient tree data structure (e.g., Range tree, B+ tree, AVL tree). Figure 10 illustrates an example of a column indexed by Range tree for (non)-equality/range queries, in which each leaf node points to a node in another linked-list structure that stores the list of matching IDs. We propose an oblivious tree structure called OTREE, in which indexed data for such queries are translated into a balanced tree structure. As in OMAT, OTREE can be instantiated from any tree-based ORAM scheme. Notice that oblivious access on a tree has been studied by Wang et. al. in [23]. However, our method requires less amount of data to be transmitted and processed, since the structure of indexed values (i.e., the tree data structure) is not required to be hidden, and the client is merely required to traverse an arbitrary path of the tree. We present the construction of OTREE as follows.

- *Oblivious Data Structure for OTREE:* Given a tree structured data \mathbf{T} of height H as input, we first construct the OTREE structure of height H with ORAM buckets as illustrated in Figures 7-(a,b). Then, each node of \mathbf{T} at level ℓ is assigned to a random path and placed into a bucket of OTREE which resides on the assigned path at level ℓ' where $\ell' \leq \ell$. In other words, any node of \mathbf{T} at level $0 \leq \ell \leq H$ will reside in a bucket at level ℓ or lower in OTREE. If there is no empty slot in the path, the node will be stored in the stash if OTREE is instantiated with stash-required ORAM schemes (e.g., Path-ORAM).

We assume \mathbf{T} is sorted by nodes' id and the position of nodes at level ℓ is stored in its parent node at level $\ell-1$ using the pointer technique proposed in [23]. Hence, each node of \mathbf{T} is considered as a separate block in OTREE structure as: $b := (\text{id}, \text{data}, \text{childmap})$, where id is the node identifier sorted in \mathbf{T} (e.g., indexed column value), data indicates the node data, and childmap is of structure $\langle \text{id}, \text{pos} \rangle$ that stores the position information of node's children.

- *The Proposed OTREE Access Scheme:* OTREE can be instantiated with any tree-based ORAM schemes (e.g., Ring-ORAM [25], Circuit-ORAM [22]), as similar to OMAT in Section 3.1, with modifications to corresponding retrieval/eviction procedures, but preserving the constraints of OTREE regarding

1. The access can be any types of operation such as read/add/delete/modify.

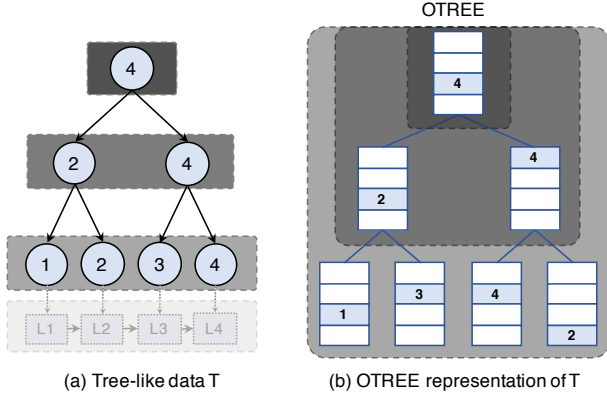


Fig. 7: The OTREE layout for a tree data.

the deepest level of nodes. OTREE also receives a significant benefit from caching mechanisms like top-tree caching [30], which can speed up bulk access requests.

We give the proposed OTREE scheme instantiated with Path-ORAM in Algorithm 2. Specifically, given the node identifier id to be accessed in the id -sorted tree structure, the client first reads the root bucket of Path-ORAM structure to obtain the root node of the tree (steps 1–3). The client then compares the requested id with the root id to decide which child of the root node should be accessed in the next step. The client accesses this child by reading its path in the Path-ORAM structure from level 0 to level 1. We notice that for each node at level l in the tree to be accessed, the client only accesses the path in the Path-ORAM structure up to level l . The process repeats until the desired id is found (steps 4–16). Finally, the client performs eviction to flush read nodes back to the Path-ORAM structure, wherein nodes at level l in the tree must reside somewhere in the Path-ORAM structure from level 0 to level l (steps 17–21).

The construction and constraints of OTREE require sta-

Algorithm 2 (data) \leftarrow OTREE.Access($op, id, data^*$)

```

1:  $x_0 \leftarrow \text{RootPos}$ 
2:  $S \leftarrow S \cup \text{ReadBucket}(\mathcal{P}(x_0, 0), 0)$ 
3:  $b_0 \leftarrow \text{Read block with } id_0 = 0 \text{ from } S$ 
4: for each  $\ell \in \{0, \dots, H-1\}$  do
5:   if compare( $id, id_\ell$ ) = go_right then
6:      $(id_{\ell+1}, x_{\ell+1}) \leftarrow b_\ell.\text{child}[1]$ 
7:      $b_\ell.\text{child}[1].\text{pos} \xleftarrow{\$} \{0, \dots, 2^\ell - 1\}$ 
8:   else
9:      $(id_{\ell+1}, x_{\ell+1}) \leftarrow b_\ell.\text{child}[0]$ 
10:     $b_\ell.\text{child}[0].\text{pos} \xleftarrow{\$} \{0, \dots, 2^\ell\}$ 
11:   $S \leftarrow S \cup \text{ReadBucket}(\mathcal{P}_{\ell+1}(x_{\ell+1}, \ell+1))$ 
12:   $b_\ell \leftarrow \text{Read block } id_\ell \text{ from } S$ 
13:  if  $id = id_\ell$  then
14:     $data \leftarrow b_\ell.data$ 
15:    if  $op = \text{write}$  then
16:       $S \leftarrow (S \setminus \{b_\ell\}) \cup \{(id, data^*, \text{child})\}$ 
17:  for each  $\ell' \in \{\ell, \dots, 0\}$  do
18:     $S' \leftarrow \{b' \in S : \mathcal{P}_\ell(b'.\text{pos}, \ell') = \mathcal{P}_\ell(b_\ell.\text{pos}, \ell') \wedge b'.\text{level} = \ell'\}$ 
19:     $S' \leftarrow \text{Select min}(|S'|, z) \text{ blocks from } S'$ 
20:     $S \leftarrow S \setminus S'$ 
21:    WriteBucket( $\mathcal{P}_\ell(x_\ell, \ell'), S'$ )
22: return data

```

bility analysis to ensure that tree-based ORAM scheme on OTREE behaves similarly to ODS in terms of stash overflow. In the following, we provide empirical stability analysis of OTREE with Path-ORAM.

- *Stability Analysis of OTREE:* We analyze the stability of OTREE in terms of average bucket load within levels of the ORAM tree. Intuitively, one would expect an increase in average bucket load near the top of the ORAM tree, and a possible increase in the average client stash size if a Path-ORAM variant (e.g., [25], [30]) is used. We show empirically by our simulations, that OTREE behaves almost similar to ODS with a bucket size of $Z \geq 4$ with Path-ORAM. With $Z = 5$, bucket usage with OTREE structure approaches that of the stationary distribution when using an infinitely large bucket size.

Our empirical study considers experiments with an ORAM tree of height $H = 14$, storing $N = 2^{15} - 1$ blocks. We run the experiments with different bucket sizes to observe its effect on the stash size and bucket usage. We treat ORAM blocks as nodes in a complete binary tree of $H = 14$. We insert nodes into storage in breadth-first order via access functions, which is followed by a series of $(H+1)$ -length access requests, each of which consists of accessing a path of nodes from the root to a random leaf node in the binary tree. A single-round experiment is the execution of 2^{14} random root-to-leaf access sequences as described.

Figures 8 - 9 show the results of these experiments for ODS and OTREE with different bucket sizes. The results were generated by first running 1000 warm-up rounds after initialization, and then collecting statistics over 1000 test rounds. Figure 8 depicts that with a bucket size $Z = 5$, buckets near the root of the OTREE structure contain roughly two non-empty blocks (one more than the average number of blocks assigned to them). Figure 9 illustrates that with $Z \leq 4$, the probability of the stash size exceeding $O(H)$ for OTREE diminishes quickly. These results suggest that using $Z = 5$ for OTREE in order to make underlying ORAM scheme in OTREE behaves similarly to that with $Z = 4$ on ODS.

- *Use Case: Conditional Query on Columns:* We exemplify an implementation of a database index structured as OTREE for conditional queries as follows: Consider a column whose values are indexed by a sorted tree T of height h by putting distinct values as keys on leaf nodes as depicted in Figure 10. The leaf nodes of T points to a node ID in a linked-list structure that contains a list of matching IDs with the key. We translate T to OTREE, where each node at level $\ell < H$ stores the position maps of its children. We store a list of IDs in each linked-list node using an inverted index with compression. As the data structure for the linked-list, we employ ODS to store it in another ORAM structure (see [23] for details). Hence, each leaf node of T stores the position map of a linked-list node in ODS it points to. An example of a given conditional query:

SELECT * FROM A WHERE C = k,

where the column C is indexed into OTREE. It can be executed obviously as follows:

- 1) Traverse a path with OTREE to get a leaf node: $b_1 \leftarrow \text{OTREE.Access}(k)$.

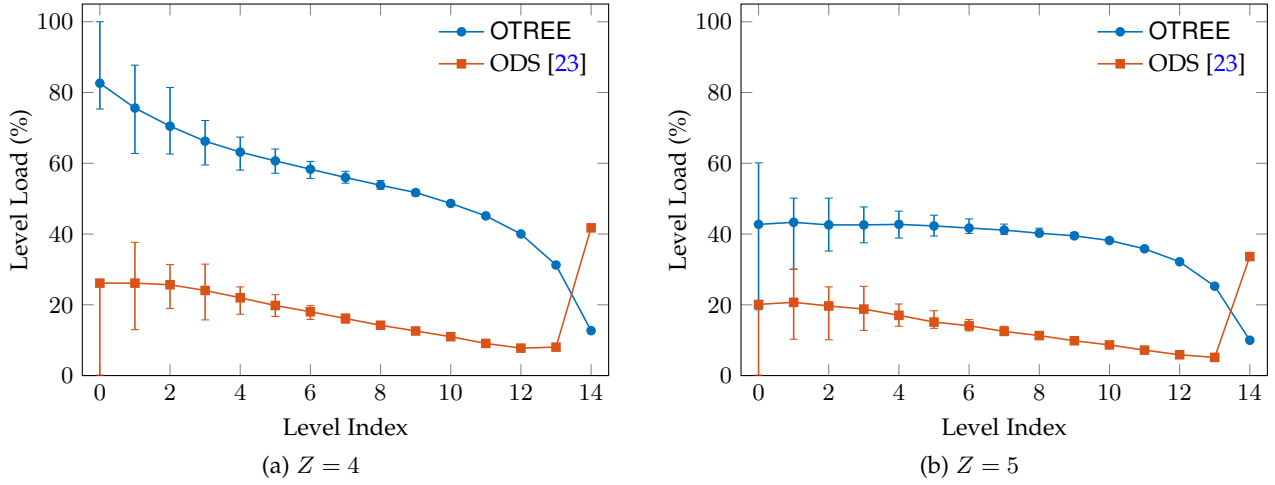


Fig. 8: Average bucket load within each level of the ORAM tree for different bucket sizes, where y -axis shows the average percentage of bucket being used and x -axis shows the bucket levels from 0 (root) to 14 (leaf).




Fig. 9: Probability of stash size exceeding the threshold.

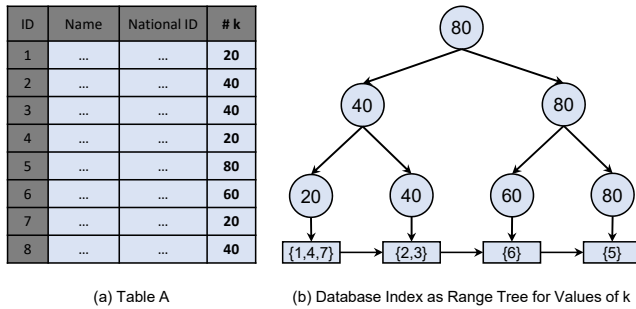


Fig. 10: Values in a column indexed as a tree, and a linked list to retrieve matching IDs for conditional queries.

- 2) Get ID and position map of linked-list node which node points to: $(id, pos) \leftarrow b.childmap$
- 3) Access on ODS to get the desired result:
 $\mathcal{R} \leftarrow ODS.Access(id, pos, null)$

The overall cost for this approach is: $O(\log(N)^2) + k \cdot O(\log(N))$, where k is the distance from the first element of the linked-list. First part is the overhead of OTREE and second part is the overhead of ODS without padding.

4 SECURITY ANALYSIS

Our security analysis, as in Path-ORAM [21], is concise as the security of our proposed schemes are evident from their base ORAM.

Definition 1 (ORAM security [21]). Let $\vec{y} := ((op_1, id_1, data_1), \dots, (op_M, id_M, data_M))$ be a data request sequence of length M , where each op_i denotes a read(id_i) or a write($id_i, data$) operation. Let $A(\vec{y})$ denote the sequence of accesses made to the server that satisfies the user data request sequence \vec{y} . An ORAM construction is secure if: (1) For any two data request sequences \vec{y} and \vec{z} of the same length, the access patterns $A(\vec{y})$ and $A(\vec{z})$ are computationally indistinguishable to an observer, and (2) it returns the data that is consistent with the input \vec{y} with probability $\geq 1 - \text{negl}(|\vec{y}|)$. That is, the ORAM fails with only a negligible probability.

Corollary 1. Accessing OMAT leaks no information beyond (i) the size of rows and columns, (ii) whether the row or column dimension being accessed, given that the ORAM scheme being used on top is secure by Definition 1.

Proof. Let \mathbf{M} be an OMAT structure comprised of two logical tree-based ORAM structures OROW and OCOL as described in Section 3.1 with dimensions M and N , respectively. Let the bit $B = 0$ if the query is on OROW and $B = 1$, otherwise. A construction providing OMAT leaks no information about the location of a node u being accessed in \mathbf{M} beyond the bit B and dimensions (M, N) . This is due to the fact that OMAT uses a secure ORAM that satisfies Definition 1 to access each block of OROW and OCOL in \mathbf{M} . Thus, as long as the node accessed within OROW or OCOL is not distinguishable from any other node within that OROW and OCOL through the number of access requests, it is indistinguishable by Definition 1. \square

Note that the information on whether the row or column was accessed can be hidden by performing a simultaneous row and column access on both dimensions for each query. This poses a security-performance trade-off. One can also hide the size of row and column by setting OMAT matrix with equal dimensions, but this may introduce extreme costs for certain applications.

Corollary 2. *Accessing OTREE leaks no information about the actual path being traversed, given that the ORAM scheme being used on top is secure by Definition 1.*

Proof. Let \mathbf{T} be a tree data structure of height H . Let \mathbf{T}_ℓ be the set of nodes at level $0 \leq \ell \leq H$ in the tree. A construction providing OTREE leaks no information about the location of a node $u \in \mathbf{T}_\ell$ being accessed in the tree beyond that it is from \mathbf{T}_ℓ . This is due to OTREE uses a secure ORAM that satisfies Definition 1 to access each level of the tree. Thus, as long as a node accessed within level ℓ is not distinguishable from any other node within that level through the number of access requests, it will be indistinguishable according to Definition 1. \square

Side-channel leakages in Path-ORAM. There are several side-channel attacks on Path-ORAM (e.g., [33], [34]) when it is executed by the secure CPU playing on behalf of the ORAM client. In this context, since the secure CPU resides in the untrusted party, the adversary has a partial view on it to exploit the timing leakage (e.g., [33]). In our model, we assume that the client is fully trusted and it is totally apart from the adversary view (i.e., untrusted database server). Therefore, we do not consider these side-channel leakages due to the difference between our model and the secure CPU context.

5 PERFORMANCE EVALUATION

5.1 Configurations

- **Implementation:** We implemented our schemes and their counterparts on CURIOUS framework [27]. We integrated additional functionalities into the framework to perform batch read/write operations to prevent unnecessary round-trip delays, and also to communicate with MongoDB instance via MongoDB Java Driver. We chose MongoDB instance as our database and storage engine. We preferred MongoDB since its Java Driver library is well-documented and easy to use. Moreover, it supports batch updates without restrictions, which is important for consistent performance analysis.

- **Data Formatting:** We created our table structure from randomly generated data with a different number of rows, columns, and field sizes. We then used the table to construct tree-based ORAMs for compared schemes. For instance, in OMAT, we created OROW and OCOL structures from this table, while an oblivious tree structure is created for OTREE, as described in Section 3.

- **Experimental Setup and Configurations:** For our experiments, we used two different client machines on two different network settings: (i) A desktop computer that runs CentOS 7.2 and is equipped with Intel Xeon CPU E3-1230, 16 GB RAM; (ii) A laptop computer that runs Ubuntu 16.04 and is equipped with Intel i7-6700HQ, 16 GB RAM. For our remote server, we used AmazonEC2 with instance type of t2.large that runs Ubuntu Server 16.04. While the connection between the desktop and the server was a *high-speed network* with download/upload speeds of 500/400 Mbps and an average latency of 11 ms, the connection between the laptop and the server was a *moderate-speed network* with download/upload speeds of 80/6 Mbps and an average

latency of 30 ms.


- **Evaluation Metrics:** We evaluated the performance of our schemes and their counterparts based on the following metrics: (i) Response time (i.e., end-to-end delay) including decryption, re-encryption and transmission times to perform a query; (ii) Client storage including the size of stash and position map; (iii) Server storage including the size of OMAT or OTREE. We compared the response times of OMAT and its counterparts for both row- and column-related queries (e.g., statistical, conditional). For OTREE and ODS-Tree, we compared the response times of traversing an arbitrary path on the tree-indexed database. To measure the end-to-end delay, we used the `std::chrono` C++ library to get the actual duration at the client side, from the time the client sends the first command until he receives the last response from the Amazon server. For each experiment, we ran 50 times and took the average number as the final response time reported in this section. We now describe our experimental evaluation results and compare our schemes with their counterparts.

5.2 Experimental Results


- **Statistical and Conditional Queries (Column-Related):** We first analyze the response time of column-related queries for OMAT, ODS-2D and RowPKG. With these queries, the client can fetch a column from the encrypted database for statistical analysis or a conditional search. Given a column-related query, the total number of bytes to be transmitted and processed by each scheme are shown in Table 1. RowPKG's transmission cost is the size of all ORAM buckets, where $Z \cdot (B \cdot N)$ and $(2M - 1)$ denote the bucket size and the total number of buckets, respectively. As for OMAT, its oblivious data structure OCOL allows efficient queries on column dimension with $O(\log(N))$ communication overhead, which outperforms the linear overhead of $O(N)$ of RowPKG. While OMAT and RowPKG can fetch the whole column with one request, it requires $M/4$ synchronous requests for ODS-2D where each request costs $Z \cdot (16 \cdot B_1) \cdot \log_2(M \cdot N/16)$ bytes due to 4×4 clustering of the cells.

We measured the performance of OMAT and its counterparts with arbitrary column queries. In this experiment, we set parameters as $B = 64$ bytes and $Z = 4$. The number of columns N varies from 2^4 to 2^9 , where the number of rows is *fixed* to be $M = 2^{15}$. Figures 11a and 12a illustrate the performance of the schemes on two different network settings with two different client machines as described in Section 5.1. For a database table with 2^{10} rows and 2^9 columns, OMAT's average query times are 60 s and 475 s compared to RowPKG's 775 s and 6100 s, and ODS-2D's 292 s and 1245 s on high- and moderate-speed networks, respectively. This makes OMAT about $13\times$ faster than RowPKG. While OMAT performs $2.6\times$ faster than ODS-2D on the moderate-speed network, it becomes $4.9\times$ on high-speed network since the latency starts to dominate the response time of ODS-2D with $M/4$ requests due to its construction with pointers.

- **Single Row-Related Queries:** We now analyze the response time of row-related queries for OMAT and its counterparts. Given a row-related query, the total number of bytes to be transmitted and processed by OMAT and its counterparts

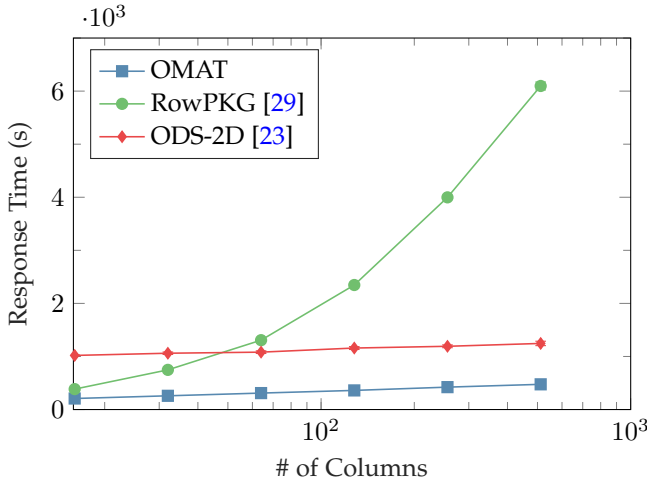


(a) Column-related queries with 2^{15} rows

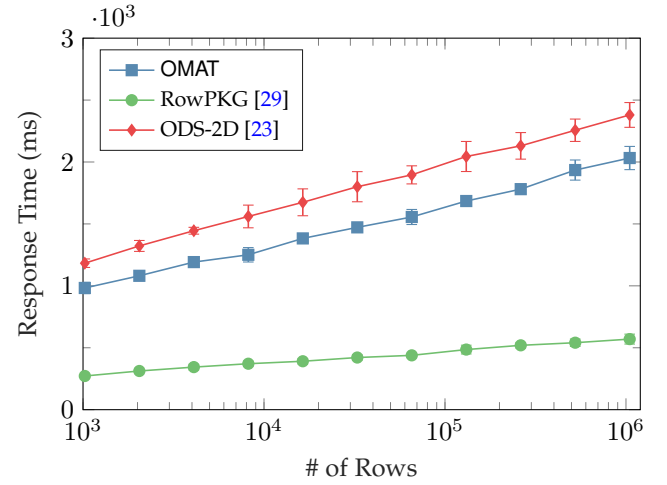


(b) Row-related queries with 2^5 columns

Fig. 11: End-to-end delay of queries for OMAT and counterparts with high-speed network setting.



(a) Column-related queries with 2^{15} rows



(b) Row-related queries with 2^5 columns

Fig. 12: End-to-end delay of queries for OMAT and counterparts with moderate-speed network setting.

are summarized in Table 1. For OMAT and RowPKG, $(B \cdot N)$ and $Z \cdot \log_2(M)$ denote the total row size and the overhead of Path-ORAM, respectively. Due to OMAT's OCOL and OROW structures, OMAT is always a constant factor of $Z = 4$ more costly than RowPKG. Clustering strategy of ODS-2D also introduces more cost and makes ODS-2D $4.2\times$ more costly than RowPKG when $N = 32$.

We measured the performance of OMAT and its counterparts with arbitrary row queries, where the number of rows M varies from 2^{10} to 2^{20} . The block size is $B = 128$ bytes and the number of columns is fixed as $N = 32$. By this setting, the total row/record size is $B \cdot N = 4096$ KB. Figures 11b and 12b illustrate the performance of the compared schemes for both network settings. We can see that OMAT performs slower than RowPKG by a constant factor of approximately $2.3\times$ and $3.6\times$ on high and moderate-speed network, respectively. As for ODS-2D, Figure 11b explicitly shows the effect of the round-trip delay introduced by network latency on ODS-2D due to $N/4$ synchronous requests. Although ODS-2D has similar cost with OMAT, it performs approximately 220 ms and 380 ms slower than OMAT.

• *Traversal on Tree-indexed Database:* We analyze the response

time of oblivious traversal on database index that is constructed as a range tree by putting distinct values of a column to the leaf of the tree. Figure 10 exemplifies the constructed range tree, and this structure is used along with its linked list to perform conditional queries (e.g., equality, range) on an indexed column, and fetch matching IDs. We compare our proposed OTREE and ODS-Tree with no caching and half-top caching strategies.

Given a database index tree constructed with values of the column, the total number of bytes to be transmitted and processed by OTREE and ODS-Tree without caching are $Z_2 \cdot B \cdot (H + 1) \cdot (H + 2)$ and $2 \cdot Z_1 \cdot B \cdot (H + 1)^2$, respectively, where H is the height of tree data structure. While ODS traverses the tree with $O(H)$, the additional overhead of Path-ORAM makes the total overhead to be $O(H^2)$. As for OTREE, its level restriction on ORAM storage reduces the transmission overhead by $1.6\times$. With half-top caching strategy, overheads of both schemes reduce as shown in Table 1, however, OTREE's construction benefits more from caching by performing traversal $3.2\times$ less costly than ODS-Tree.

For this experiment, we set the block size $B = 4$ KB, the

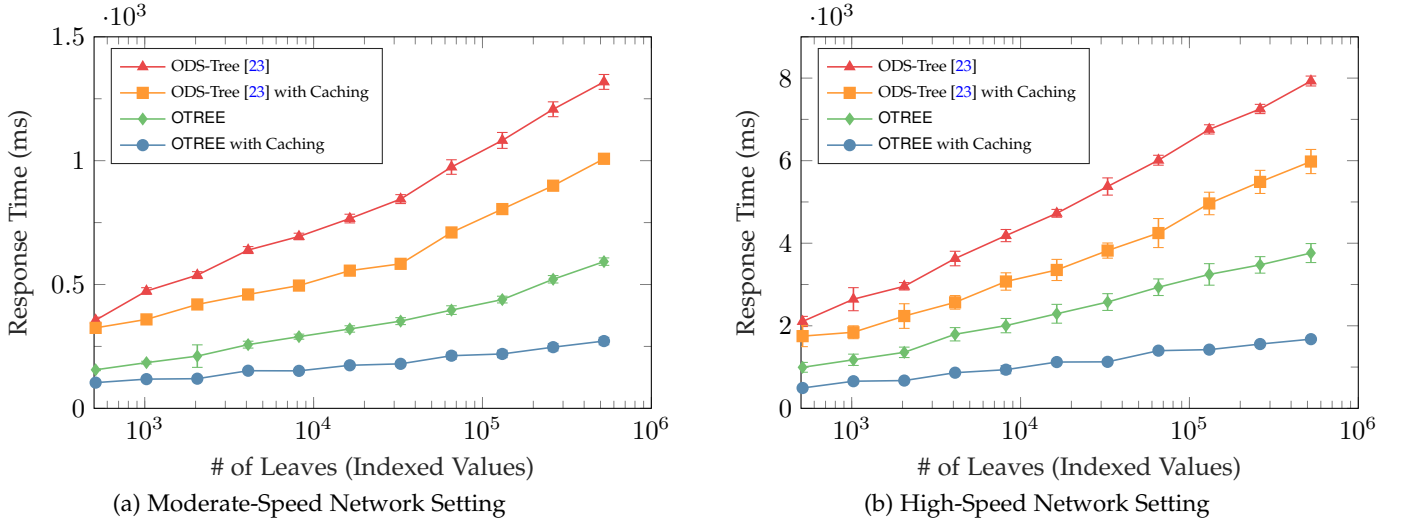


Fig. 13: End-to-end delay of traversal on tree-indexed database for OTREE and ODS-Tree.

number of blocks inside a bucket for ODS-Tree is $Z_1 = 4$, and the number of blocks inside a bucket for OTREE is $Z_2 = 5$ (see Subsection 3.2 for stability analysis). We benchmarked OTREE and ODS-Tree with arbitrary equality queries when the number of indexed values varies from 2^9 to 2^{19} . The number of indexed values is set to 2^{19} for large database setting. For both network settings, Figure 13 demonstrates the effect of half-top caching strategy and how the structure of OTREE gives more leverage in response time. While OTREE without caching performs around $2\times$ faster than its counterpart, caching allows OTREE to perform $3.6\times$ faster than ODS-Tree with caching for both network settings.

5.3 Client and Server Storage

We now analyze the client storage overhead of our schemes and their counterparts. The position map of OMAT requires $O((M + N) \cdot \log(M + N))$ storage, while RowPKG requires $O(M \cdot \log(M))$, since only the position map of rows are stored. However, the dominating factor is M , since large databases have more rows than columns. ODS's pointer technique allows it to operate with $O(1)$ storage for position map. Moreover, the worst-case stash size changes with the query type, because stash is also used to store currently fetched data and the worst-case storage costs are summarized in Table 1. For row-related queries, the worst-case stash storage is the same for both OMAT and RowPKG but ODS-2D requires more storage due to clustering. For column-related queries, RowPKG requires storing $O(M \cdot N)$ that corresponds to all ORAM buckets. Besides the query performance issues, this also makes RowPKG infeasible for very large databases to perform column-related queries. In addition, ODS-2D also requires $O(\log(M))$ times more client storage compared to OMAT. While RowPKG and ODS-2D have the same server storage size, OMAT requires constant $z\times$ more storage due to additional dummy blocks.

Since OTREE and ODS do not require position map to operate, the client storage consists of stash and additionally cached block according to the caching strategy used. For the worst-case, both schemes have the same client storage with the same caching strategy, however, OTREE's stash may be more loaded than ODS as shown in Figure 9 due to its level

restriction. Moreover, server storage of OTREE is $2\times$ times less than ODS, since Path-ORAM of ODS requires one more level than OTREE.


6 CONCLUSIONS

In this paper, we introduce two new oblivious data structures called OMAT and OTREE. The proposed techniques can be instantiated with any tree-based ORAM scheme to enable efficient private queries on database instances. OMAT enables various statistical and conditional queries on generic database tables, which may be highly inefficient for its counterparts relying on row-oriented packaging. On the other hand, OTREE provides more efficient range queries on tree-indexed database than existing ODS techniques, and also receives more benefit from caching optimizations. These properties allow OMAT and OTREE to be ideal data structures to construct oblivious database services on the cloud, which offers high security and privacy guarantee for the users.

REFERENCES

- [1] R. A. Popa, C. 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*. ACM, 2011, pp. 85–100.
- [2] —, "Cryptdb: processing queries on an encrypted database," *Communications of the ACM*, vol. 55, no. 9, pp. 103–111, 2012.
- [3] 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," *IACR Cryptology ePrint Archive*, vol. 2014, p. 853, 2014.
- [4] A. A. Yavuz and J. Guajardo, "Dynamic searchable symmetric encryption with minimal leakage and efficient updates on commodity hardware," in *Selected Areas in Cryptography – SAC 2015*, ser. Lecture Notes in Computer Science. Springer International Publishing, August 2015.
- [5] S. Kamara, C. Papamanthou, and T. Roeder, "Dynamic searchable symmetric encryption," in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*. ACM, 2012, pp. 965–976.
- [6] N. Cao, C. Wang, M. Li, K. Ren, and W. Lou, "Privacy-preserving multi-keyword ranked search over encrypted cloud data," *IEEE Transactions on parallel and distributed systems*, vol. 25, no. 1, pp. 222–233, 2014.

- [7] B. Wang, Y. Hou, M. Li, H. Wang, and H. Li, "Maple: scalable multi-dimensional range search over encrypted cloud data with tree-based index," in *Proceedings of the 9th ACM symposium on Information, computer and communications security*. ACM, 2014, pp. 111–122.
- [8] B. Wang, M. Li, and H. Wang, "Geometric range search on encrypted spatial data," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 4, pp. 704–719, 2016.
- [9] H. Yin, Z. Qin, J. Zhang, L. Ou, and K. Li, "Achieving secure, universal, and fine-grained query results verification for secure search scheme over encrypted cloud data," *IEEE Transactions on Cloud Computing*, 2017.
- [10] R. Zhang, R. Xue, L. Liu, and L. Zheng, "Oblivious multi-keyword search for secure cloud storage service," in *Web Services (ICWS), 2017 IEEE International Conference on*. IEEE, 2017, pp. 269–276.
- [11] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart, "Leakage-abuse attacks against searchable encryption," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 668–679.
- [12] M. S. Islam, M. Kuzu, and M. Kantarcioglu, "Access pattern disclosure on searchable encryption: Ramification, attack and mitigation," in *Annual Network and Distributed System Security Symposium – NDSS*, vol. 20, 2012, p. 12.
- [13] C. Liu, L. Zhu, M. Wang, and Y.-a. Tan, "Search pattern leakage in searchable encryption: Attacks and new construction," *Information Sciences*, vol. 265, pp. 176–188, 2014.
- [14] D. Pouliot and C. V. Wright, "The shadow nemesis: Inference attacks on efficiently deployable, efficiently searchable encryption," in *Proceedings of the 2016 ACM Conference on Computer and Communications Security*. ACM, 2016.
- [15] Y. Zhang, J. Katz, and C. Papamanthou, "All your queries are belong to us: The power of file-injection attacks on searchable encryption," in *25th USENIX Security Symposium (USENIX Security 16)*, Austin, TX, 2016, pp. 707–720.
- [16] O. Goldreich and R. Ostrovsky, "Software protection and simulation on oblivious rams," *Journal of the ACM (JACM)*, vol. 43, no. 3, pp. 431–473, 1996.
- [17] B. Pinkas and T. Reinman, "Oblivious ram revisited," in *Advances in Cryptology—CRYPTO 2010*. Springer, 2010, pp. 502–519.
- [18] B. Chen, H. Lin, and S. Tessaro, "Oblivious parallel ram: Improved efficiency and generic constructions," in *Theory of Cryptography Conference*. Springer, 2016, pp. 205–234.
- [19] E. Shi, T.-H. H. Chan, E. Stefanov, and M. Li, "Oblivious ram with $o(\log n)$ 3) worst-case cost," in *Advances in Cryptology—ASIACRYPT 2011*. Springer, 2011, pp. 197–214.
- [20] E. Stefanov, E. Shi, and D. Song, "Towards practical oblivious ram," *arXiv preprint arXiv:1106.3652*, 2011.
- [21] E. Stefanov, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas, "Path oram: an extremely simple oblivious ram protocol," in *Proceedings of the 2013 ACM SIGSAC conference on Computer and Communications security*. ACM, 2013, pp. 299–310.
- [22] X. Wang, H. Chan, and E. Shi, "Circuit oram: On tightness of the goldreich-ostrovsky lower bound," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 850–861.
- [23] X. S. Wang, K. Nayak, C. Liu, T. Chan, E. Shi, E. Stefanov, and Y. Huang, "Oblivious data structures," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 215–226.
- [24] S. Devadas, M. van Dijk, C. W. Fletcher, L. Ren, E. Shi, and D. Wichs, "Onion oram: A constant bandwidth blowup oblivious ram," in *Theory of Cryptography Conference*. Springer, 2016, pp. 145–174.
- [25] L. Ren, C. W. Fletcher, A. Kwon, E. Stefanov, E. Shi, M. van Dijk, and S. Devadas, "Ring oram: Closing the gap between small and large client storage oblivious ram," *IACR Cryptology ePrint Archive*, vol. 2014, p. 997, 2014.
- [26] E. Boyle and M. Naor, "Is there an oblivious ram lower bound?" in *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*. ACM, 2016, pp. 357–368.
- [27] V. Bindaschadler, 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 *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 837–849.
- [28] T. Hoang, A. Yavuz, and J. Guajardo, "Practical and secure dynamic searchable encryption via oblivious access on distributed data structure," in *Proceedings of the 32nd Annual Computer Security Applications Conference (ACSAC)*. ACM, 2016.
- [29] Z. Chang, D. Xie, and F. Li, "Oblivious ram: a dissection and experimental evaluation," *Proceedings of the VLDB Endowment*, vol. 9, no. 12, pp. 1113–1124, 2016.
- [30] M. Maas, E. Love, E. Stefanov, M. Tiwari, E. Shi, K. Asanovic, J. Kubiawicz, and D. Song, "Phantom: Practical oblivious computation in a secure processor," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 311–324.
- [31] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*. Chapman & Hall/CRC, 2007.
- [32] J. Dautrich and C. Ravishanker, "Combining oram with pir to minimize bandwidth costs," in *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy*. ACM, 2015, pp. 289–296.
- [33] C. Bao and A. Srivastava, "Exploring timing side-channel attacks on path-orams," in *2017 IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*. IEEE, 2017, pp. 68–73.
- [34] 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 *High Performance Computer Architecture (HPCA), 2014 IEEE 20th International Symposium on*. IEEE, 2014, pp. 213–224.



Thang Hoang is currently a PhD student in the School of Electrical Engineering and Computer Science, Oregon State University (September 2015). He received his ME degree in Computer Science from Chonnam National University, Gwangju, South Korea in February, 2014, and BS degree in Computer Science from University of Natural Sciences, Saigon, Vietnam in September, 2010. His research interest currently focuses on privacy-enhancing technologies (e.g., searchable encryption, ORAM) and authentication mechanisms for mobile devices.



Ceyhun D. Ozkaptan is currently a PhD student in the Department of Electrical and Computer Engineering, The Ohio State University (August 2017). He received his BE degree from Bilkent University in Ankara, Turkey (June 2016). His research interest spans from signal processing to security and applied cryptography.



Gabriel Hackebeil is currently a PhD student in the Department of Industrial & Operations Engineering, University of Michigan (August 2017). He received his MS degree from Oregon State University in 2016, and BS degree from Texas A&M University in 2012. He is interested in optimization and cryptography.



Attila Altay Yavuz is an Assistant Professor in the Department of Computer Science and Engineering, University of South Florida (August 2018). He was an Assistant Professor in the School of Electrical Engineering and Computer Science, Oregon State University (09/2014-07/2018). He was a member of the security and privacy research group at the Robert Bosch Research and Technology Center North America (2011-2014). He received his PhD degree in Computer Science from North Carolina State

University in August 2011. He received his MS degree in Computer Science from Bogazici University (2006) in Istanbul, Turkey. He is broadly interested in design, analysis and application of cryptographic tools and protocols to enhance the security of computer networks and systems. Attila Altay Yavuz is a recipient of NSF CAREER Award (2017). His research on privacy enhancing technologies (searchable encryption) and intra-vehicular network security are in the process of technology transfer with potential world-wide deployments. He has authored more than 40 research articles in top conferences and journals along with several patents. He is a member of IEEE and ACM.