

Efficient Oblivious Data Structures for Database Services on the Cloud

Thang Hoang* Ceyhun D. Ozkaptan[†] Gabriel Hackebeil[‡] Attila A. Yavuz*

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.

Keywords— 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., [18, 19]) or searchable encryption with various security, efficiency and

*School of EECS, Oregon State University, Corvallis, OR, 97331. Email: {hoangmin, attila.yavuz}@oregonstate.edu

[†]Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210

[‡]Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, 48109

Work done when the second and third authors were employed at Oregon State University. E-mail: {ozkaptac, hackebeg}@oregonstate.edu.

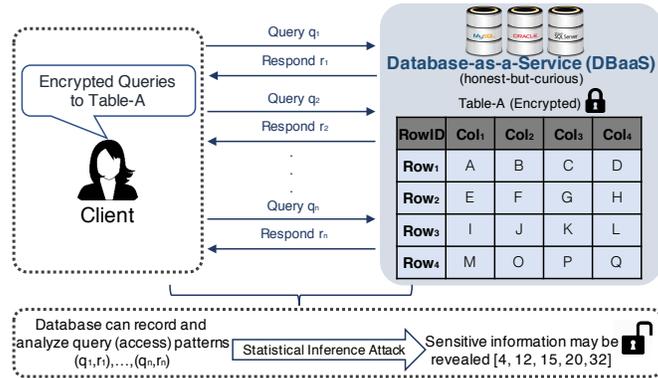


Figure 1: Information leakages through query access patterns over an encrypted database.

query functionality trade-offs (e.g., [5, 29, 13, 3, 25, 26, 30, 31]) 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 works (e.g., [4, 12, 15, 20, 32]) 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 leakages 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 a vital requirement for privacy-critical database applications to hide the access pattern.

Oblivious Random Access Machine (ORAM) [10] can be used to hide access patterns for such encrypted databases. Preliminary ORAM schemes (e.g., [10, 17]) were costly, but recent ORAM constructions (e.g., [7, 22, 23, 24, 27, 28]) showed promising results. Most efficient ORAM schemes (e.g., [9, 21, 24]) to date follow tree paradigm [22], which achieves asymptotic $\Omega(\log N)$ lower-bound [2, 10, 27] communication overhead. However, despite these improvements, there are several research gaps towards achieving efficient integration of ORAM in databases. In the following, we discuss these research gaps and limitations of the existing approaches.

1.1 Limitations of Existing Approaches

The direct application of ORAM to the structured encrypted data has been shown to be highly costly in the context of searchable encryption [1, 11]. Meanwhile, there is a very limited work on the application and integration of ORAM for encrypted database systems. To the best of our knowledge, the only work that studied ORAM schemes for real database systems was proposed by Chang et al. in [6] with a framework called SEAL-ORAM. In SEAL-ORAM, various ORAMs were implemented and compared with a MongoDB database instance where ORAM blocks are constructed in a row-oriented manner. While it shows the possibility of using ORAM for encrypted databases, the functionalities and performance offered by such a direct adaptation seem to be very limited. We outline the limitations of two direct ORAM applications in Figure 2, and further elaborate them as below:

- *Limitations of Row-Oriented Approach:* Recall that SEAL-ORAM packages each row of a database instance into an ORAM block. We refer to this approach as RowPKG. Notice that RowPKG only allows insert/delete/update queries on a row in the database table. However, to execute oblivious insert/delete/update on a column, RowPKG requires to transfer *all* blocks in ORAM, which is not only highly communication costly but also client-storage expensive. Similarly, the execution of any column-

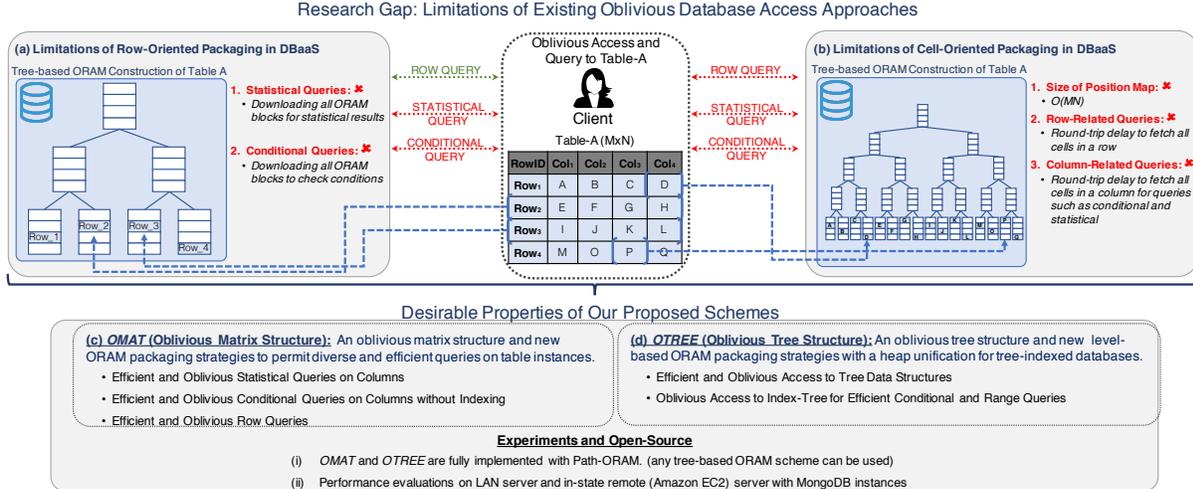


Figure 2: Research gap to be addressed and desirable properties of the proposed schemes.

related queries such as statistical or conditional queries is also very costly, as they require downloading all the ORAM blocks and may not be practical for large databases. RowPKG (i.e., row-oriented packaging) and its limitations are outlined in Figure 2-(a).

• *Limitations of Cell-Oriented Approach:* Another approach is to package each cell of the database into an ORAM block. However, it significantly increases the size of position map, which is a data structure stored on the client to enable tree-based ORAM. To eliminate position map, oblivious 2D-grid structure (referred as ODS-2D) proposed by Wang et. al [28] can be used to store database table by clustering $O(\log(N))$ cells into an ORAM block with pointers. However, this approach increases the number of requests to be sent for each query to fetch all cells in a row or column. This incurs end-to-end delay due to a large number of round-trip delays, and therefore, is not suitable for large databases. Cell-oriented packaging and its limitations are summarized in Figure 2-(b).

The above discussion indicates that there is a significant need for an efficient oblivious data structure that permits diverse types of queries on encrypted databases. Hence, in this paper, we seek answers to the following research questions:

“How can we create an efficient oblivious data structure for encrypted databases that achieve diverse types of queries with a low overhead? How can we harness asymptotically optimal ORAMs over structured data to create an oblivious data structure?”

1.2 Our Contributions

Given the availability of asymptotically efficient ORAM building blocks, our objective is to create new oblivious data structures by harnessing such ORAMs in efficient ways. Specifically, we proposed two efficient oblivious data structures that permit various types of queries on encrypted databases:

(i) Our first scheme is referred to as *Oblivious Matrix Structure* (OMAT) (Section 3.1). The main idea behind OMAT is to create an oblivious matrix structure that permits efficient queries over table objects in a database not only for rows but also columns. This is achieved via various strategies that are specifically tailored for matrix structure with a delicate balance between query diversity and ORAM overhead. This allows OMAT to perform various types of oblivious queries without downloading a large number of ORAM blocks or storing a very large position map. (ii) Our second scheme is referred to as *Oblivious Tree Structure* (OTREE) (Section 3.2), which is designed for oblivious accesses on tree-structured

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
<i>single column-related query (e.g., statistical, conditional queries)</i>					
RowPKG [6]	$Z \cdot (B_1 \cdot N) \cdot (2M - 1)$	1.00	$O(M \cdot N) \cdot w(1)$	6096 s	776 s
ODS-2D [28]	$(M/4) \cdot [Z \cdot (16 \cdot B_1) \cdot \log_2(M \cdot N/16)]$	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
<i>single row-related query (e.g., insert/delete/update queries)</i>					
RowPKG [6]	$Z \cdot (B_2 \cdot N) \cdot \log_2(M)$	1.00	$O(N \cdot \log(M)) \cdot w(1)$	567 ms	56 ms
ODS-2D [28]	$(N/4) \cdot [Z \cdot (16 \cdot B_2) \cdot \log_2(M \cdot N/16)]$	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
<i>traversal on database tree index (e.g., range queries)</i>					
<i>non-caching</i>					
ODS-Tree [28]	$2 \cdot Z_1 \cdot B \cdot (H + 1)^2$	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
<i>half-top caching</i>					
ODS-Tree [28]	$2 \cdot Z_1 \cdot B \cdot \left\lceil \frac{H+1}{2} \right\rceil \cdot (H + 1)$	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 [16], 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.

database instances. Given a column whose values are sorted into a tree data structure (i.e., database index), OTREE allows efficient oblivious conditional queries (e.g., a range query). OTREE achieves better performance than the existing oblivious data structures [28] 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 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× more communication efficient than that of RowPKG and this enables OMAT to perform queries approximately 13× faster than that of RowPKG. Compared to ODS-2D, although OMAT is only 1.6 × communication efficient, it performs approximately 5× faster in practice due to the large number of additional round-trip delays. OTREE achieves better performance than ODS for obliviously 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× 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 [24] due to its efficiency, simplicity and not requiring any server-side computation.

- *Full-fledged Implementation and Comprehensive Experimental Evaluation:* 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.

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 [10]. Existing ORAM schemes rely on IND-CPA encryption [14] and an oblivious shuffling to ensure that any data access patterns of the same length are computationally indistinguishable by anyone but the client.

Recent ORAMs (e.g., [8, 9, 21, 24]) follow the tree paradigm [22], 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 fetching and eviction. For each access operation, the client gets the path ID of accessing block from the position map and sends the path

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
$i := \text{pm}[id]$	block identified by id is currently associated with leaf node i , i.e., it resides somewhere along $\mathcal{P}(i)$ or in the stash.

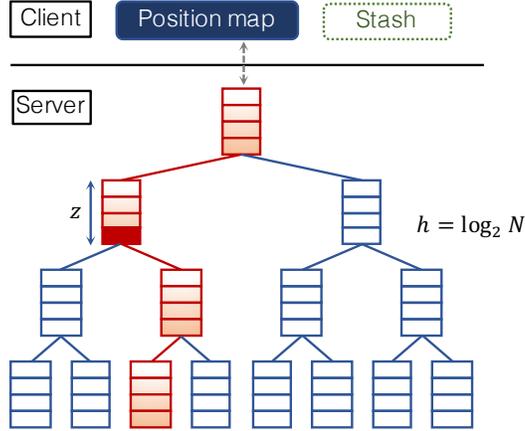


Figure 3: Tree-based ORAM structure [22].

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 call the eviction procedure which re-encrypts the block and pushes it back to the ORAM tree in such a way that the server does not know which block was being accessed. Notice that although recently proposed ORAM schemes that follow the tree paradigm (e.g., [24, 21, 9]) provide different trade-offs between communication and computation overhead, they all rely on the aforementioned basic procedures.

Path-ORAM: Path-ORAM [24] 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 [23]. A new random 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 have been shown to be highly costly in practice, we do not discuss them in this work.

2.2 Oblivious Data Structure

Oblivious Data Structure (ODS) proposed by Wang et al. [28] 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.

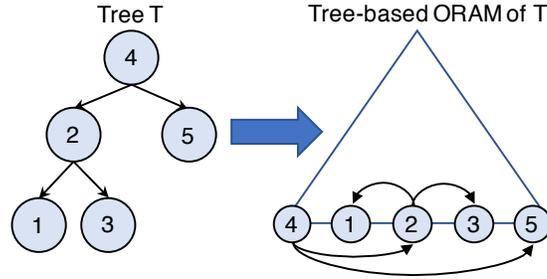


Figure 4: Oblivious Data Structure for a tree structure [28].

2.3 ORAM Implementation Framework

One of the most reliable and complete ORAM frameworks is CURIOUS [1], which gives a complete implementation of the state-of-the-art ORAM schemes (e.g., Path-ORAM [24]) in Java. In this paper, we chose CURIOUS to implement our oblivious data structures as it can be adopted with the library of various database drivers such as MongoDB or MySQL.

3 The Proposed Techniques

We now present our proposed oblivious data structures, which are specially designed for efficient operations in database settings. We propose two schemes, one is referred to as *Oblivious Matrix Structure* (OMAT) and the other is *Oblivious Tree Structure* (OTREE).

For our oblivious data structures, we choose Path-ORAM [24] as underlying ORAM for the following reasons: (i) It is simple yet achieves asymptotic efficiency. (ii) Unlike some recent ORAMs [9, 21] 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 under frameworks (e.g., CURIOUS [1]) 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 [1] and database systems in specific [6] 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 [6], 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.

3.1.1 Oblivious Data Structure for OMAT

The main data structure we use for oblivious access on a table is a matrix-structured table (as in a binary matrix in [29]). Given an input table T of size $M \times N$, we allocate a matrix 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:

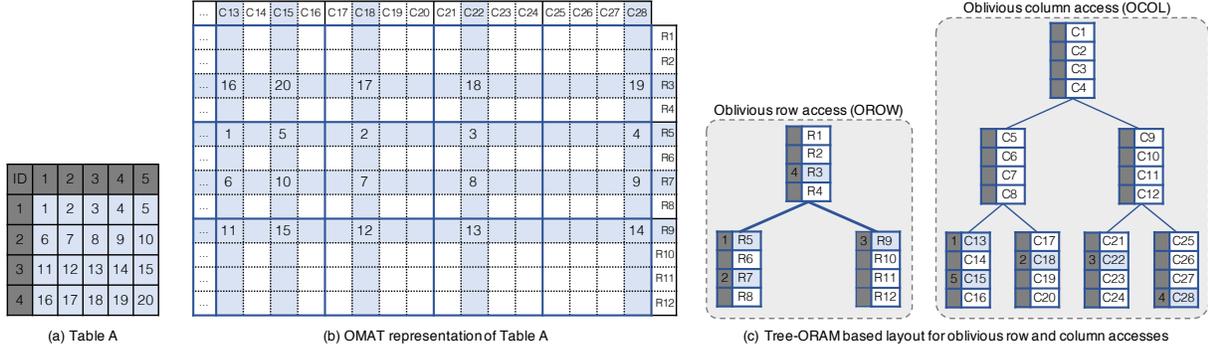


Figure 5: OMAT structure for oblivious access on matrix-like table.

The layout of OMAT matrix M can be interpreted as two logical tree-based ORAMs defined as oblivious rows (denoted as OROW) and oblivious columns (denoted as OCOL), as illustrated in Figure 5. That is, the ORAM for row access on OROW is formed by a set of blocks $b_i := (id_i, data_i)$, where id_i is either a unique identifier if b_i contains the content of a row of the table T or otherwise null, and $data_i \leftarrow M[i, *]$. We group Z subsequent rows in M to form a bucket (i.e., node) in the OROW structure. Similarly, the ORAM for column access on OCOL is formed by $b_j = (id_j, data_j)$. Each (bucket) node in OCOL is formed by grouping Z subsequent blocks.

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

3.1.2 OMAT Access Protocol

We present our OMAT scheme, which is instantiated with Path-ORAM, in Scheme 1. We use subroutines $\text{ReadBucket}(\cdot, \cdot)$ and $\text{WriteBucket}(\cdot, \cdot)$, which are tailored for OMAT purposes from Path-ORAM. These subroutines take an extra parameter dim as the input, indicating which dimension the client wishes to obliviously read/write on. For example, if $\text{dim} = \text{col}$, then $\text{ReadBucket}(\text{dim}, \text{bucket})$ subroutine will read Z columns from the bucket corresponding to the OCOL structure. Moreover, we use two stashes denoted as S_{row} and S_{col} to store the accessed rows and columns separately due to Path-ORAM operations. As row access operation will update a row data to a new position, it is required to update the order of all columns currently in the stash S_{col} and vice versa to maintain the consistency (Scheme 1, line 14). Rows/columns are IND-CPA decrypted and re-encrypted as they are read and written, 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 selected randomly) simultaneously for each access.

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

```

3.1.3 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 obviously fetching 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:
 $C[*, id'] \leftarrow \text{OMAT.Access}(\text{read}, \text{col}, id')$.
2. Get IDs of rows whose value larger than k , and such IDs are in pm_{row} :
 $\mathcal{I} \leftarrow \{id \mid id \in \text{pm}_{\text{row}}.\text{id} \wedge C[id, id'] > k\}$

3. Access on OROW to get the desired result:
 $R[id, *] \leftarrow \text{OMAT.Access}(\text{read}, \text{row}, \text{id}), \text{ for each } \text{id} \in \mathcal{I}.$

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 8 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 [28]. 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.

3.2.1 Oblivious Data Structure for OTREE

Given a tree structured data T of height H as input, we first construct the OTREE structure of height H with ORAM buckets as illustrated in Figures 6(a–b). Then, each node of 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 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 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 [28]. Hence, each node of 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 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.

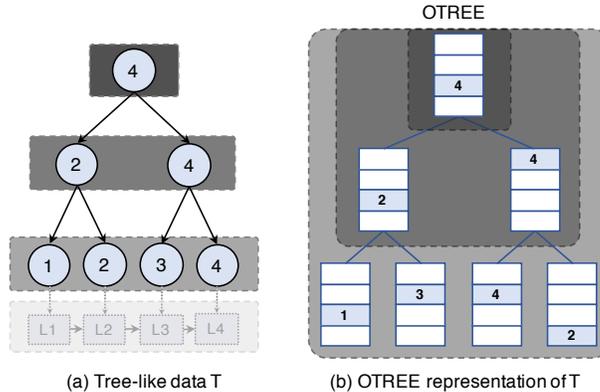


Figure 6: The OTREE layout for a tree-like input.

3.2.2 OTREE Access Protocol

We give the proposed OTREE scheme, which is instantiated with Path-ORAM, in Scheme 2. Similar to OMAT in Section 3.1, OTREE can also be instantiated with any tree-based ORAM (e.g., Ring-ORAM[21], Onion-ORAM [9]) with modifications to corresponding read/write (i.e., ReadBucket, WriteBucket) and eviction (lines 20-25) procedures but preserving the constraints of OTREE regarding the deepest level of nodes. OTREE also receives a significant benefit from caching mechanisms like top-tree caching [16], which can speed up bulk access requests. Notice that the construction and constraints of OTREE require stability analysis to ensure that tree-based ORAM scheme on OTREE behaves similarly to ODS in terms of stash overflow. We provide empirical stability analysis of OTREE with Path-ORAM as follows.

Scheme 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:  $b_0 \leftarrow$  Read block with  $\text{id}_0 = 0$  from  $S$ 
4: for each  $\ell \in \{0, \dots, H - 1\}$  do
5:   if compare(id,  $\text{id}_\ell$ ) = go_right then
6:      $(\text{id}_{\ell+1}, x_{\ell+1}) \leftarrow b_\ell.\text{child}[1]$ 
7:      $b_\ell.\text{child}[1].\text{pos} \xleftarrow{\$} \{0, \dots, 2^\ell - 1\}$ 
8:   else
9:      $(\text{id}_{\ell+1}, x_{\ell+1}) \leftarrow b_\ell.\text{child}[0]$ 
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:    $b_\ell \leftarrow$  Read block  $\text{id}_\ell$  from  $S$ 
13:   if id =  $\text{id}_\ell$  then
14:     data  $\leftarrow b_\ell.\text{data}$ 
15:     if op = write then
16:        $S \leftarrow (S \setminus \{b_\ell\}) \cup \{(\text{id}, \text{data}^*, \text{child})\}$ 
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:      $S' \leftarrow$  Select  $\min(|S'|, z)$  blocks from  $S'$ 
20:      $S \leftarrow S \setminus S'$ 
21:     WriteBucket( $\mathcal{P}_\ell(x_\ell, \ell')$ ,  $S'$ )
22: return data

```

3.2.3 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., [21, 16]) 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

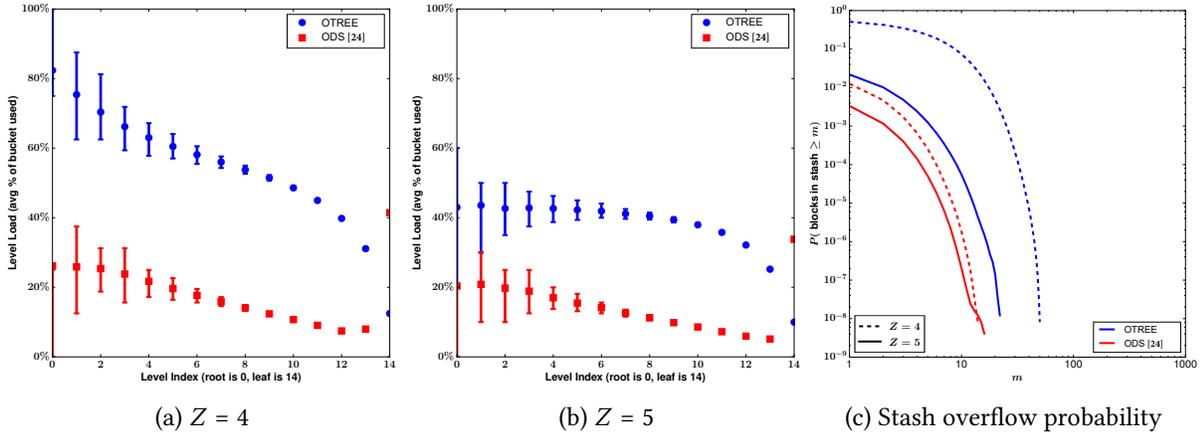


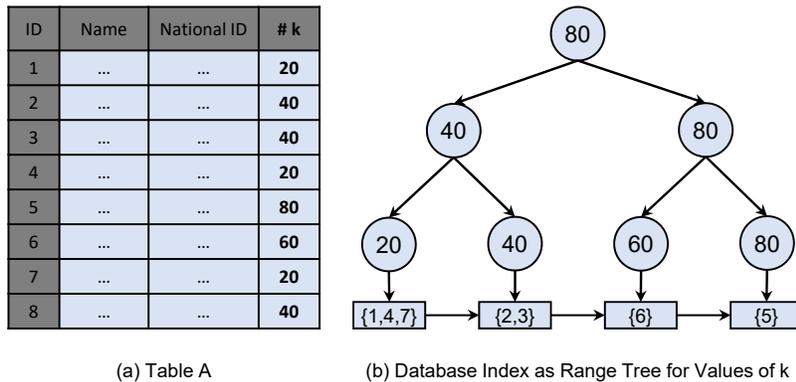
Figure 7: (a,b) Average bucket load within each level of the ORAM tree for different bucket sizes, (c) the probability of stash size exceeding the threshold.

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 7 - 7c 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 7 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 7c 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.

3.2.4 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 8. 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



(a) Table A

(b) Database Index as Range Tree for Values of k

Figure 8: Tree-indexed column values, and a linked list to retrieve matched IDs for conditional query.

ODS to store it in another ORAM structure (see [28] 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 obliviously as follows:

1. Traverse a path with OTREE to get a leaf node: $b_1 \leftarrow \text{OTREE.Access}(k)$.
2. Get ID and position map of linked-list node which node points to: $(\text{id}, \text{pos}) \leftarrow b.\text{childmap}$
3. Access on ODS to get the desired result:
 $\mathcal{R} \leftarrow \text{ODS.Access}(\text{id}, \text{pos}, \text{null})$

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 [24], is concise as the security of our proposed schemes are evident from their base ORAM.

Definition 1 (ORAM security [24]). *Let $\vec{y} := ((\text{op}_1, \text{id}_1, \text{data}_1), \dots, (\text{op}_M, \text{id}_M, \text{data}_M))$ 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.*

Claim 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 (sketch). Let M be an OMAT structure comprised of two logical tree-based ORAM structures OROW and OCOL as described in Section 3.1.1 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 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 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.

Claim 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 (sketch). Let T be a tree data structure of height H . Let 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 $u \in T_\ell$ being accessed in the tree beyond that it is from 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

5 Performance Evaluation

5.1 Configurations

- *Implementation:* We implemented our schemes and their counterparts on CURIOUS framework [1]. 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.

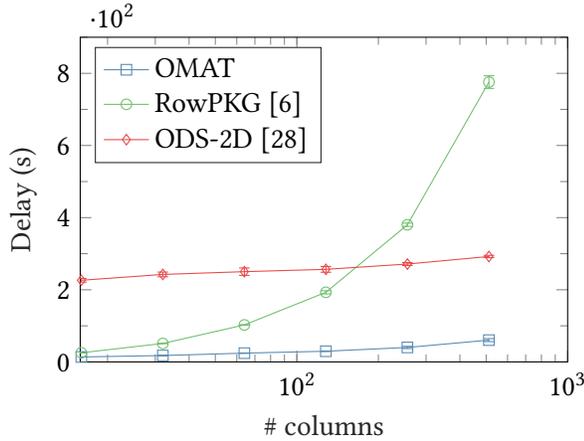
- *Experimental Setup and Configurations:* For our experiments, we used two different client machines on two different network settings: (i) A desktop computer that runs CentOS 7.2 and is equipped with Intel Xeon CPU E3-1230, 16 GB RAM; (ii) A laptop computer that runs Ubuntu 16.04 and is equipped with 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) that includes decryption, re-encryption and transmission times to perform a query; (ii) Client storage that includes the size of stash and position map; (iii) Server storage. We compare 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 compare the response times of traversing an arbitrary path on the tree-indexed database.

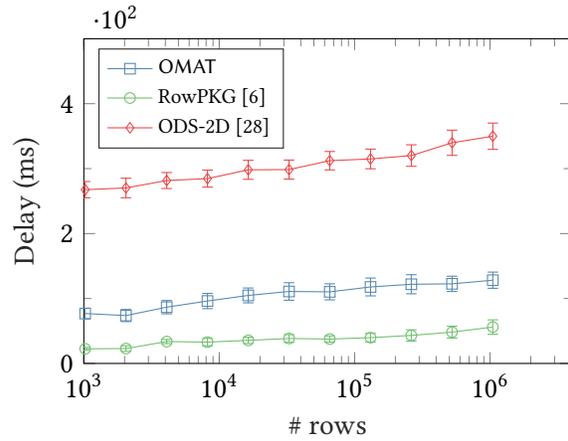
We now describe our experimental results and compare our schemes with their counterparts.

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

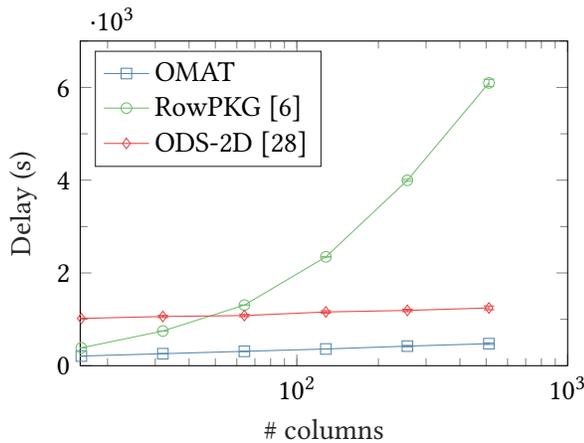


(a) Column-related queries with 2^{15} rows

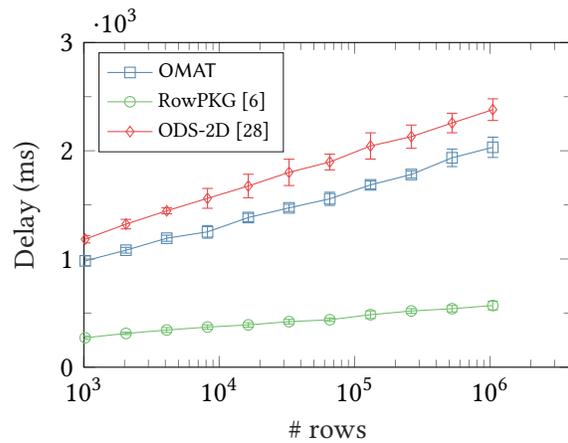


(b) Row-related queries with 2^5 columns

Figure 9: 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

Figure 10: End-to-end delay of queries for OMAT and counterparts with moderate network speed.

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 $M = 2^{15}$. Figure 9a and Figure 10a 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

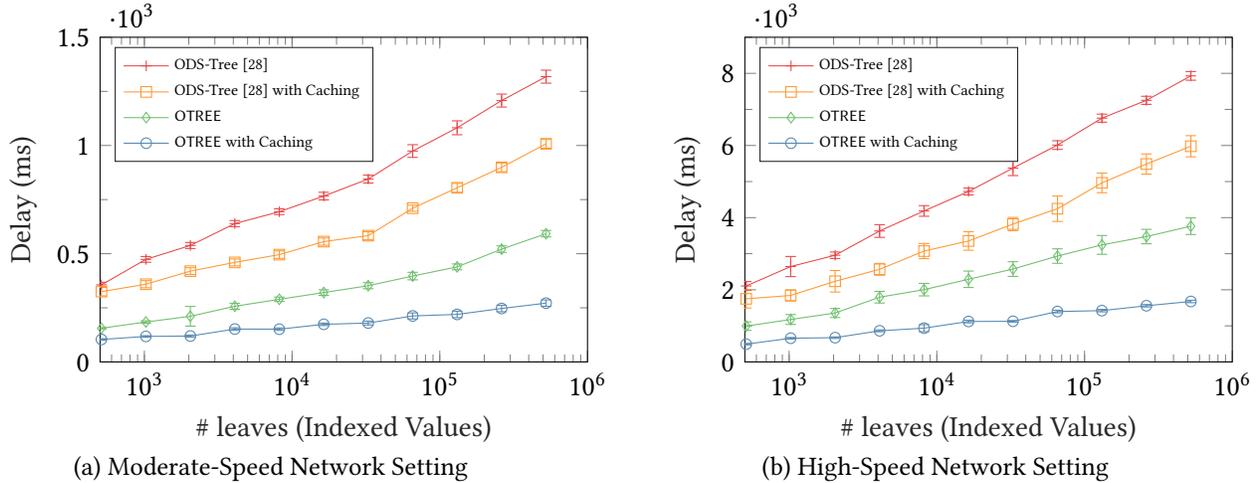


Figure 11: End-to-end delay of traversal on tree-indexed database for OTREE and ODS-Tree.

counterparts. Given a row-related query, the total number of bytes to be transmitted and processed by OMAT and its counterparts 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. Figure 9b and 10b 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 9b 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 8 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 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 3.2.3 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 11 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 7c 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 developed two new oblivious data structure techniques that we call as OMAT and OTREE, which achieve diverse queries with high efficiency on table and tree structured database instances, respectively. We introduce strategies such as oblivious data structures tailored for encrypted database access, which are supported with special ORAM packaging, leveled access and heap unification techniques. Our schemes can be instantiated from any tree-based ORAM scheme. OMAT enables various query types that may be highly inefficient for its counterparts that only rely on row-oriented packaging. Specifically, OMAT offers significantly more efficient column-related queries such as statistical (e.g., SUM, AVG, MAX) and conditional queries (e.g., range queries) than that of its counterparts. OTREE permits an efficient private access on tree structured database instances. OTREE provides more efficient conditional queries (e.g., range query) than that of the existing ODS techniques, and also receives more benefit from caching optimizations. To the best of our knowledge, these properties make OMAT and OTREE the most efficient oblivious data structure techniques for encrypted databases known to date. Finally, our cryptographic framework will be released as an open-source library to provide opportunities for researchers to develop more efficient techniques.

References

- [1] 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 *Proceedings of the 22nd ACM SIGSAC Conference on*

- Computer and Communications Security*, pages 837–849. ACM, 2015.
- [2] 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*, pages 357–368. ACM, 2016.
 - [3] 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*, 25(1):222–233, 2014.
 - [4] 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*, pages 668–679. ACM, 2015.
 - [5] 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*, 2014:853, 2014.
 - [6] Z. Chang, D. Xie, and F. Li. Oblivious ram: a dissection and experimental evaluation. *Proceedings of the VLDB Endowment*, 9(12):1113–1124, 2016.
 - [7] B. Chen, H. Lin, and S. Tessaro. Oblivious parallel ram: Improved efficiency and generic constructions. In *Theory of Cryptography Conference*, pages 205–234. Springer, 2016.
 - [8] J. Dautrich and C. Ravishankar. Combining oram with pir to minimize bandwidth costs. In *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy*, pages 289–296. ACM, 2015.
 - [9] 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*, pages 145–174. Springer, 2016.
 - [10] O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious rams. *Journal of the ACM (JACM)*, 43(3):431–473, 1996.
 - [11] 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.
 - [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*, volume 20, page 12, 2012.
 - [13] S. Kamara, C. Papamanthou, and T. Roeder. Dynamic searchable symmetric encryption. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, pages 965–976. ACM, 2012.
 - [14] J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC, 2007.
 - [15] C. Liu, L. Zhu, M. Wang, and Y.-a. Tan. Search pattern leakage in searchable encryption: Attacks and new construction. *Information Sciences*, 265:176–188, 2014.
 - [16] 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*, pages 311–324. ACM, 2013.
 - [17] B. Pinkas and T. Reinman. Oblivious ram revisited. In *Advances in Cryptology—CRYPTO 2010*, pages 502–519. Springer, 2010.
 - [18] 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*, pages 85–100. ACM, 2011.
 - [19] R. A. Popa, C. Redfield, N. Zeldovich, and H. Balakrishnan. Cryptdb: processing queries on an encrypted database. *Communications of the ACM*, 55(9):103–111, 2012.

- [20] 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.
- [21] 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*, 2014:997, 2014.
- [22] 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*, pages 197–214. Springer, 2011.
- [23] E. Stefanov, E. Shi, and D. Song. Towards practical oblivious ram. *arXiv preprint arXiv:1106.3652*, 2011.
- [24] E. Stefanov, M. Van Dijk, 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*, pages 299–310. ACM, 2013.
- [25] 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*, pages 111–122. ACM, 2014.
- [26] B. Wang, M. Li, and H. Wang. Geometric range search on encrypted spatial data. *IEEE Transactions on Information Forensics and Security*, 11(4):704–719, 2016.
- [27] 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*, pages 850–861. ACM, 2015.
- [28] 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*, pages 215–226. ACM, 2014.
- [29] 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*, Lecture Notes in Computer Science. Springer International Publishing, August 2015.
- [30] 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.
- [31] 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*, pages 269–276. IEEE, 2017.
- [32] 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)*, pages 707–720, Austin, TX, 2016.