

Authenticating Aggregate Range Queries over Dynamic Multidimensional Dataset

Jia XU

jiaxu2001@gmail.com

National University of Singapore

Abstract. We are interested in the integrity of the query results from an outsourced database service provider. Alice passes a set \mathbf{D} of d -dimensional points, together with some authentication tag \mathbf{T} , to an untrusted service provider Bob. Later, Alice issues some query over \mathbf{D} to Bob, and Bob should produce a query result and a proof based on \mathbf{D} and \mathbf{T} . Alice wants to verify the integrity of the query result with the help of the proof, using only the private key. Xu J. *et al.* [25] proposed an authentication scheme to solve this problem for multidimensional aggregate range query, including SUM, COUNT, MIN, MAX and MEDIAN, and multidimensional range selection query, with $O(d^2)$ communication overhead. However, their scheme only applies to *static* database. This paper extends their method to support *dynamic* operations on the dataset, including inserting or deleting a point. The communication overhead of our scheme is $O(d^2 \log N)$, where N is the number of data points in the dataset.

Key words: Authentication, Multidimensional Aggregate Query, Dynamic Database, Secure Outsourced Database, Provable Remote Computing

1 Introduction

Alice has a set D of d -dimensional points. She preprocesses the dataset D using her private key to generate some authentication tag T . She sends (outsources) D and T to an untrusted service provider Bob. Then Alice deletes the original copy of dataset D and tag T from her local storage. Later Alice may issue a query over D to Bob, for example, an aggregate query conditional on a multidimensional range selection, and Bob should produce the query result and a proof based on D and T . Alice wants to authenticate the query result, using only her private key.

We are concerned about the communication cost and the storage overhead on Bob's side. Such requirements exclude the following two straightforward approaches: (1) Bob sends back the whole dataset D with its tag T ; (2) During preprocessing, Alice generates and signs answers to all possible queries.

The problem we study in this paper fits in the framework of the outsourced database applications [6, 11], which emerged in early 2000s as an example of "software-as-a-service" (SaaS). By outsourcing database management, backup services and other IT needs to a professional service provider, companies can reduce expensive cost in purchase of equipments and even more expensive cost in hiring or training qualified IT specialists to maintain the IT services [17].

Xu J. *et al.* [25] proposed a scheme, called MAIA, to authenticate aggregate range query, including SUM, COUNT, MIN, MAX and MEDIAN, over a static d -dimensional dataset, with

$O(d^2)$ communication overhead. Since MAIA indexes each point in the dataset using integers and integers are not densely ordered, MAIA cannot handle with dynamic operations, like inserting to or deleting a point from the dataset.

In this paper, we propose a method to index each point in the dataset with a rational number in the interval $[0, 1)$. Since rational number is densely ordered, i.e. there is at least one rational number in-between any two unequal rational numbers, our indexing method support dynamic operations. The new indexing method utilizes a special ternary tree, which we call *Doit tree*— Dynamic Ordered Indexing Ternary Tree.

Based on the new indexing scheme, we propose a scheme, called dynamic MAIA, to authenticate the same types of aggregate range queries as MAIA, over dynamic multidimensional dataset. The communication overhead is $O(d^2 \log N)$, where N is the number of points in the dataset.

Table 1: Performance of our scheme *Dynamic MAIA*, compared with MAIA proposed by Xu J. *et al.* [25]. Note both schemes are much more efficient in computation cost in 1D case, compared with high dimensional case (See annotation \star). We point out that the high computation cost on prover can be mitigated with horizontal partition of the dataset and parallel execution on each partition.

Scheme	Communication overhead (bits)	Key Size	Storage overhead	Computation (Verifier Alice)	Computation (Prover Bob)	Query
MAIA ([25])	$O(d^2)$	$O(d)$	$O(dN)$	$O(d \log N) \dagger \star$	$O(dN \log N) \ddagger \star$	SUM, COUNT, MIN, MAX, MEDIAN, QUANTILE, Range Selection
Dynamic MAIA (This paper)	$O(d^2 \log N)$	$O(d)$	$O(dN)$	$O(d \log N) \dagger \star$	$O(dN \log N) \ddagger \star$	INSERT, DELETE, SUM, COUNT, MIN, MAX, MEDIAN, QUANTILE, Range Selection

d : The dimension of data point in the dataset.
 \dagger : $O(d \log N)$ modular multiplications.
 \star : If the query range is 1D, the cost is $O(|S|)$.

N : The number of tuples in the dataset.
 \ddagger : $O(dN \log N)$ modular exponentiations.

1.1 Contribution

Our main contributions can be summarized as follows:

1. We propose a Dynamic Ordered Indexing Ternary Tree, called as *Doit tree*, which maps a datapoint to a rational number in $[0, 1)$ and preserves the order, and supports insertion or deletion of datapoints.
2. Based on *Doit tree*, we propose dynamic MAIA, to authenticate aggregate range query over dynamic multidimensional dataset.
3. Our scheme dynamic MAIA is efficient (See Table 1) and secure (See Theorem 1).

1.2 Organization

The rest of this paper is organized as follows: Section 2 briefs the related works, Section 3 reviews the scheme MAIA proposed in [25], Section 4 proposes the Doit tree, and Section 5 constructs our authentication scheme dynamic MAIA.

2 Related work

Researches [6, 11] in secure outsourced database emerged and quickly developed since early 2000's. Most papers focus on privacy-preserving SQL query execution over encrypted datasets [11, 12, 18, 7], and authentication of query results [6, 15, 5, 19, 17, 20, 23, 4, 14, 24, 1, 26, 16, 13, 9].

There are roughly four categories of approaches for outsourced database authentication in the literatures [6, 15, 5, 19, 17, 20, 23, 4, 14, 24, 1, 26, 16, 13, 9]. (1) (Homomorphic, or aggregatable) Cryptographic primitives, like collision-resistant hash, digital signature, commitment [17, 10, 3]. (2) Merkle Hash Tree and variants [16, 14, 4]. (3) Computational geometry approach [15, 4, 1]. (4) Inserting and auditing fake tuples [24].

In the “integrity authentication” track, most papers are working on authentication of simple range selection queries [15, 19, 17, 20, 1, 4]. [19, 20] authenticated 1D range selection queries, with linear (in the number of tuples selected by the query condition) communication cost and storage overhead. [17] verified range selection queries using aggregated signatures (like RSA [22], BLS [2]). [4] proposed a linear (or superlinear) scheme, using chained signatures over a “verification R-Tree” with partitions of tuples as tree nodes, to authenticate windows query, range query, kNN query, and RNN query. To the best of our knowledge, the current most efficient authentication scheme for range selection queries is [1], which proposed an efficient authentication scheme for 1D and 2D range selection queries over a grid dataset (e.g. GIS or image data) with $O(1)$ communication cost and linear storage overhead. [23] claimed to authenticate arbitrary queries, but their security model is too weak: a playful adversary can easily break their scheme.

3 Review of MAIA

In this section, we review the scheme MAIA proposed by Xu J. *et al.* [25]. For the sake of presentation of this paper, we will re-organize the presentation of [25].

Let $\mathbf{D}_x = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ be a set of N d -dimensional points. In the setup phase, Alice *normalizes* [21] dataset \mathbf{D}_x , to obtain dataset $\mathbf{D} = \{\mathbf{i}_1, \mathbf{i}_2, \dots, \mathbf{i}_N\}$. Next, Alice runs key generating algorithm KGen to generate a private key \mathcal{K} , and encodes the dataset \mathbf{D} using data encoding algorithm DEnc with key \mathcal{K} , to obtain tags \mathbf{T} . At the end of setup phase, Alice deletes $\mathbf{D}_x, \mathbf{D}, \mathbf{T}$ from her storage after sending them to Bob, and keeps key \mathcal{K} private.

The query phase consists of multiple query sessions. In each session, Alice issues an aggregate range query to Bob. Bob generates the query result, and produces a proof, with help of Alice, by running the interactive algorithm ProVer with Alice. After receiving Bob's query result and proof, Alice verifies authenticity of the query result using the private key \mathcal{K} . The type of queries include SUM, COUNT, MIN, MAX and MEDIAN.

Specially, among all components of MAIA, we take out all algorithms related to precess of index and encapsulate them as an indexing scheme (Idx , ITag , Translate , Compress , Uncompress). Idx maps a datapoint $\mathbf{x} \in \mathbf{D}_x$ to an integer vector \mathbf{i} in $[N]^d$ and preverses order for each dimension. ITag produces an authentication tag for index \mathbf{i} . Translate converts a query w.r.t. datapoint \mathbf{x} 's to a query w.r.t. indices \mathbf{i} 's, which has the same query result as the previous query. For any query range \mathbf{R} , Compress can convert the set $\{\text{ITag}(\mathbf{i}) : \mathbf{i} \in \mathbf{R}\}$ (called as *Help-Info* in [25]) into a d -dimensional vector of small size, and Uncompress can reverse the process.

In this paper, we will propose an alternative indexing scheme to replace the one in MAIA [25], and keep the other parts of MAIA unchanged. The resulting scheme is called dynamic MAIA, which can support dynamic operations, like inserting or deleting a datapoint.

4 Dynamic Ordered Indexing Ternary Tree

We first propose a *Dynamic Ordered Indexing Ternary* tree, called as *Doit* tree in Section 4.1, and then present an indexing scheme based on the *Doit* tree in Section 4.2.

4.1 Doit Tree




Fig. 1: Doit Tree: Dynamic Ordered Indexing Ternary Tree for dataset $\{x_1, x_2, x_3, x_4\}$.

Let $\text{S2R} : \{0, 1, 2\}^* \rightarrow [0, 1)$ be function that converts a string in $\{0, 1, 2\}^*$ to a real value in $[0, 1)$. On input string s in $\{0, 1, 2\}^*$, S2R treats the string “0.” $\|s$ as a rational number in ternary numeral system, and converts it to $y \in [0, 1)$. As a result, $\text{S2R}(s) = y$.

We build a complete ternary tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ with vertex set \mathcal{V} and edge set \mathcal{E} . Let $\mathcal{L} \subset \mathcal{V}$ be the set of all leaf nodes. We associate each edge $\mathbf{e} \in \mathcal{E}$ with attributes $\mathbf{e.trit}$ and $\mathbf{e.p}$, and each vertex $\mathbf{v} \in \mathcal{V}$ with attributes $\mathbf{v.lab}$, $\mathbf{v.idx}$, $\mathbf{v.range}$ and $\mathbf{v.tag}$. These attributes are defined as follows.

1. Edge

- (a) Each internal node has three child nodes: *left child*, *middle child*, *right child*. For any edge $\mathbf{e}_l \in \mathcal{E}$ pointing to a left child, $\mathbf{e}_l.\text{trit} \leftarrow 0$; for any edge $\mathbf{e}_m \in \mathcal{E}$ pointing to a middle child, $\mathbf{e}_m.\text{trit} \leftarrow 1$; for any edge $\mathbf{e}_r \in \mathcal{E}$ pointing to a right child, $\mathbf{e}_r.\text{trit} \leftarrow 2$.
- (b) Each tree edge $\mathbf{e} \in \mathcal{E}$ is associated with a public random prime, denoted as $\mathbf{e}.\text{prime}$.

2. Vertex

- (a) For any node $\mathbf{v} \in \mathcal{V}$, let $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \dots, \mathbf{e}_w)$ be the unique simple path from the root \mathbf{r} to node \mathbf{v} , where \mathbf{e}_1 is adjacent to the root and \mathbf{e}_w is adjacent to node \mathbf{v} . Then $\mathbf{v}.\text{lab} \leftarrow \mathbf{e}_1.\text{trit} \parallel \mathbf{e}_2.\text{trit} \dots \parallel \mathbf{e}_w.\text{trit}$, where \parallel is the string concatenation operator. Specially, for the root \mathbf{r} , $\mathbf{r}.\text{lab} \leftarrow \epsilon$, where ϵ is the empty string.
- (b) For each node $\mathbf{v} \in \mathcal{V}$, $\mathbf{v}.\text{idx} \leftarrow \text{S2R}(\mathbf{v}.\text{lab})$.
- (c) For each node $\mathbf{v} \in \mathcal{V}$, $\mathbf{v}.\text{range} \leftarrow [\mathbf{v}.\text{idx}, \mathbf{v}.\text{idx} + 3^{-|\mathbf{v}.\text{lab}|}) \subseteq [0, 1)$, where $|\mathbf{v}.\text{lab}|$ denotes the length of trit-string $\mathbf{v}.\text{lab}$. Particularly, for root \mathbf{r} , $\mathbf{r}.\text{range} = [0, 1)$, since $\mathbf{r}.\text{lab} = \epsilon$, and $\text{S2R}(\epsilon) = 0$, $|\epsilon| = 0$.
- (d) For any node $\mathbf{v} \in \mathcal{V}$, let $(\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_u)$ be the unique simple path from the root to \mathbf{v} . Then $\mathbf{v}.\text{tag} = \mathbf{r}.\text{tag} \prod_{1 \leq j \leq u} \mathbf{e}_j.\text{prime} \pmod n$, where \mathbf{r} denotes the root of the tree.

4.1.1 Indexing a dataset. Now, we use the newly constructed ternary tree to index a 1D dataset.

Let dataset $\mathbf{D}_x = \{x_1, x_2, \dots, x_N\} \subset \mathbb{N}$ and $x_1 < x_2 < x_3 \dots < x_N$. For simplicity of presentation, assume $N = (3^H - 1)/2$ for some integer H . Build a complete ternary tree with $(2N + 1)$ leaves as above.

For the convenience of presentation, we name each vertex and edge.

1. For any $\mathbf{v} \in \mathcal{V}$ with $\mathbf{v}.\text{lab} = s$, we name \mathbf{v} with V_s . For example, the root node $\mathbf{r} = V_\epsilon$.
2. For any edge $\mathbf{e} = (\mathbf{u}, \mathbf{v}) \in \mathcal{E}$, where \mathbf{u} is the parent node of \mathbf{v} and $\mathbf{v}.\text{lab} = s$, we name \mathbf{e} with E_s , and name the prime $\mathbf{e}.\text{prime}$ with p_s (See Figure 1).

Next, we associate an attribute val to some leaf nodes and define the mapping $\text{idx} : \mathbf{D}_x \rightarrow [0, 1)$ as follows.

1. From left to right, number all leaf nodes with $1, 2, 3, \dots, (2N + 1)$. So the leftmost leaf node is the 1st leaf, and the rightmost leaf node is the $(2N + 1)$ -th leaf.
2. For each leaf node \mathbf{v} , associate an attribute, denoted as $\mathbf{v}.\text{val}$. For the $(2j)$ -th leaf \mathbf{v} , $j \in [N]$, $\mathbf{v}.\text{val} \leftarrow x_j$ and $\text{idx}(x_j) \leftarrow \mathbf{v}.\text{idx}$; for any other leaf \mathbf{v} , $\mathbf{v}.\text{val} \leftarrow \perp$. We call a leaf \mathbf{v} as *unassigned leaf*, if $\mathbf{v}.\text{val} = \perp$. Otherwise, we call it *assigned leaf*. Let \mathcal{L}^- be the set of all assigned leaf nodes and \mathcal{L}^+ be the set of all unassigned leaf nodes. $(\mathcal{L}^+, \mathcal{L}^-)$ forms a partition of the set \mathcal{L} of all leaf nodes.

We call the constructed ternary tree **Doit Tree** for dataset \mathbf{D}_x , which stands for “*Dynamic Ordered Indexing Ternary*” Tree. Figure 1 shows an example **Doit** tree for a dataset $\{x_1, x_2, x_3, x_4\}$.

A **Doit** tree has the following properties.

Property 1 (Doit Tree)

1. (Ternary Tree) Every internal node has exactly three children, called left child, middle child and right child.
2. (Partition) For an internal node v with left child t_1 , middle child v_m and right child t_r , $(v_1.\text{range}, v_m.\text{range}, v_r.\text{range})$ form a partition of $v.\text{range}$. Additionally, for root node r , $r.\text{range} = [0, 1)$.
3. (Ordered) For any two nodes $v_i, v_j \in \mathcal{V}$, if v_i appears before v_j in the preorder traversal, then $v_i.\text{idx} \leq v_j.\text{idx}$.
4. (Dynamic) For any $x_j, x_{j+1} \in \mathbf{D}$, let $v_j, v_{j+1} \in \mathcal{L}^-$ be their corresponding leaf nodes, i.e. $v_j.\text{idx} = \text{idx}(x_j)$ and $v_{j+1} = \text{idx}(x_{j+1})$, there exists one and only one leaf node between v_j and v_{j+1} and this leaf node is unassigned. Additionally, both the leftmost leaf and the rightmost leaf node are unassigned. Those unassigned leaf nodes in set \mathcal{L}^+ allow the insertions of new values. After insertions of new nodes, these four properties still hold.

4.2 Indexing Scheme for 1D Dataset

Now we show an indexing scheme based on the Doit tree. As mentioned previously, an indexing scheme consists of `Idx`, `Translate`, `ITag`, `Compress` and `Uncompress`.

4.2.1 itag. At first, we define function $\text{R2S} : [0, 1) \rightarrow \{0, 1, 2\}^*$ to convert a real value in $[0, 1)$ to a trit-string, which is the inverse function of S2R . Given $i \in [0, 1)$, let $s_i \in \{0, 1, 2\}^*$ be the expression of i in ternary numeral system. If $|s_i| < H + 2$, then append $(H + 2 - |s_i|)$ 0's to s_i : $s_i \leftarrow s_i \parallel \underbrace{00\dots0}_{H+2-|s_i| \text{ 0's}}$. Remove the first two symbols from s_i . Output s_i .

Given $i \in [0, 1)$, we convert it into a trit-string $\text{R2S}(i) \in \{0, 1, 2\}^*$. We define

$$\text{itag}(i) \stackrel{\text{def}}{=} V_{\text{R2S}(i)}.\text{tag}.$$

Let $\text{prefix}(s, j)$ be the prefix of length j of string s and let $h : \{0, 1, 2\}^* \rightarrow \text{Prime}$ be a hash function [8]. In another equivalent expression,

$$\text{itag}(i) \stackrel{\text{def}}{=} g_p \prod_{1 \leq j \leq |s|}^{s = \text{R2S}(i)} h(\text{prefix}(s, j)) \pmod n.$$

4.2.2 insert. The input is (\mathcal{T}, x) , where \mathcal{T} is a Doit tree and $x \in \mathbb{N}$.

1. Number all leaf nodes in \mathcal{L} from left to right with $1, 2, 3, \dots, |\mathcal{L}|$. Denote the i -th leaf node with ℓ_i . Find ζ such that $\ell_\zeta, \ell_{\zeta+2} \in \mathcal{L}^-$ and $\ell_\zeta.\text{val} \leq x < \ell_{\zeta+2}.\text{val}$.
2. Expand node $\ell_{\zeta+1}$ and let t_1, t_m and t_r denote its left child, middle child and right child, respectively. *Comment: As a result, $\ell_{\zeta+1}$ becomes an internal node and t_1, t_m and t_r are three leaf nodes.*
3. Update the tree \mathcal{T} :
 - (a) Let $t_m.\text{val} \leftarrow x$;
 - (b) Remove $\ell_{\zeta+1}$ from \mathcal{L} and \mathcal{L}^+ ;
 - (c) Add t_1, t_m and t_r into \mathcal{L} ;
 - (d) Add t_1 and t_r into \mathcal{L}^+ ;

(e) Add τ_m into \mathcal{L}^- , etc.

Figure 2 shows an example of insertion. Note that the tree \mathcal{T} , set \mathcal{L} of leaves, set \mathcal{L}^+ of unassigned leaves, set \mathcal{L}^- of assigned leaves are dynamic, and will be updated during each insertion. However, for each node v and each edge e , the attributes $v.\text{lab}$, $v.\text{idx}$, $v.\text{range}$, $e.\text{trit}$, $e.\text{prime}$ are static.




Fig. 2: Doit Tree: Insert a new value y , where $x_1 < y < x_2$.

4.2.3 MinCover. The input of algorithm **MinCover** is (v_1, v_2) , where v_1 and v_2 are two assigned leaf nodes of a Doit tree \mathcal{T} , and $v_1.\text{id} \leq v_2.\text{id}$. The equality is for special case $v_1 = v_2$. The output is a collection of tree nodes $\{r_i\}$ such that $\{r_i.\text{range}\}$ forms a partition of interval $[v_1.\text{id}, v_2.\text{id}]$.

1. From \mathcal{T} , remove all tree nodes v and their adjacent edges, such that $v.\text{id} < v_1.\text{id}$.
2. From \mathcal{T} , remove all tree nodes v and their adjacent edges, such that $v.\text{id} > v_2.\text{id}$.
3. From \mathcal{T} , remove all tree nodes, except v_1 , along the unique simple path from the root of \mathcal{T} to leaf v_1 and their adjacent edges.
4. From \mathcal{T} , remove all tree nodes, except v_2 , along the unique simple path from the root of \mathcal{T} to leaf v_2 and their adjacent edges.
5. Let $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_m\}$ denote the collection of remaining subtrees and each \mathcal{T}_i is a subtree. Let r_i be the root node of tree \mathcal{T}_i . Note $\{\mathcal{T}_i\}$ is the minimum cover for interval $[v_1.\text{id}, v_2.\text{id}]$.
6. Output (r_1, r_2, \dots, r_m) .

Figure 3 shows an example with input (V_{021}, V_{21}) .

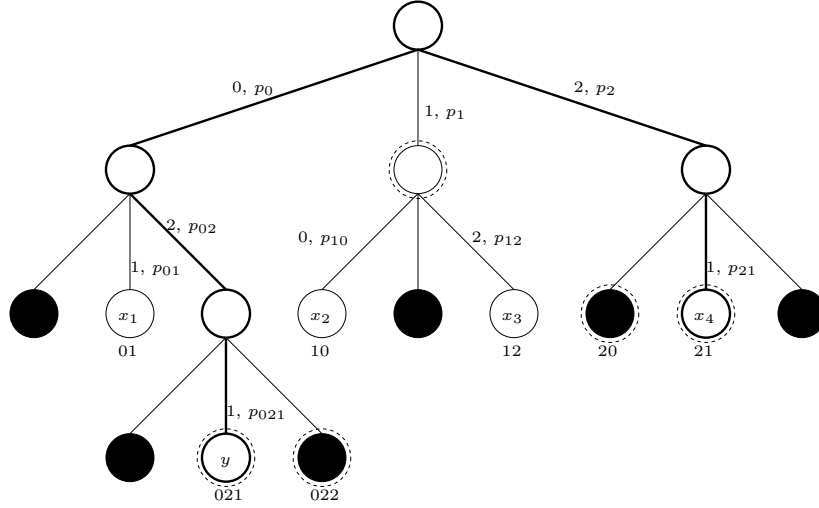


Fig. 3: Doit Tree: Compression of $[0.021, 0.21] = [0.021, 0.022] \cup [0.022, 0.1] \cup [0.1, 0.2] \cup [0.20, 0.21] \cup [0.21, 0.22]$.

4.2.4 compress. The input is $(\mathbf{v}_1, \mathbf{v}_2, \alpha)$.

1. $(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_m) \leftarrow \mathbf{MinCover}(\mathbf{v}_1, \mathbf{v}_2)$.
2. For each $i \in [m]$, $v_i \xleftarrow{\$} G_q, w_i \xleftarrow{\$} G_r$ and $\delta_i \leftarrow (\mathbf{r}_i.\text{tag})^\alpha \times v_i \times w_i \pmod n$.
3. Output $\{\delta_1, \delta_2, \dots, \delta_m\}$.

4.2.5 uncompress. The input is $(\delta_1, \delta_2, \dots, \delta_m, \mathbf{v})$.

1. Let $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_m\}$ denote the corresponding minimum cover and \mathbf{r}_i be the root node of tree \mathcal{T}_i .
2. Find $j \in [m]$, such that $\mathbf{v}.\text{id} \in \mathbf{r}_j.\text{range}$. Note \mathbf{v} should be a leaf node of the subtree \mathcal{T}_j .
3. Set $\mathbf{r}_j.\text{tag} \leftarrow \delta_j$ and compute $\mathbf{v}.\text{tag}$ as described previously.
4. Output $\mathbf{v}.\text{tag}$.

4.3 Indexing Scheme for d -Dimensional Dataset

Let $\mathbf{D}_x = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subset \mathbb{N}^d$ be a set of N d -dimensional data points, where $N = (3^H - 1)/2$ for some integer H .

For each dimension $j \in [d]$, let $\mathbf{D}_x^{(j)} = \{\mathbf{x}[j] : \mathbf{x} \in \mathbf{D}_x\}$, build a Doit tree $\mathcal{T}^{(j)}$ for 1D dataset $\mathbf{D}_x^{(j)}$ and define $\text{idx}^{(j)}$, $\text{compress}^{(j)}$, $\text{uncompress}^{(j)}$ correspondingly. Now we define

corresponding functions in vector forms.

$$\text{Idx}(\mathbf{x}) = (\text{idx}^{(1)}(\mathbf{x}[1]), \text{idx}^{(2)}(\mathbf{x}[2]), \dots, \text{idx}^{(d)}(\mathbf{x}[d]))$$

$$\text{ITag}(\mathbf{x}) = ((\text{itag}^{(1)}(\mathbf{x}[1]))^{s_1[1]}, (\text{itag}^{(2)}(\mathbf{x}[2]))^{s_1[2]}, \dots, (\text{itag}^{(d)}(\mathbf{x}[d]))^{s_1[d]})$$

$$\text{Compress}(\mathbf{v}_1, \mathbf{v}_2, \boldsymbol{\alpha}) = (\text{compress}^{(1)}(\mathbf{v}_1[1], \mathbf{v}_2[1], \boldsymbol{\alpha}[1]), \dots, \text{compress}^{(d)}(\mathbf{v}_1[d], \mathbf{v}_2[d], \boldsymbol{\alpha}[d]));$$

$$\text{Uncompress}(\boldsymbol{\delta}_1, \dots, \boldsymbol{\delta}_d, \mathbf{v}) = (\text{uncompress}^{(1)}(\boldsymbol{\delta}_1, \mathbf{v}[1]), \dots, \text{uncompress}^{(d)}(\boldsymbol{\delta}_d, \mathbf{v}[d])).$$

4.3.1 Translate. Let $\mathbf{D}_i = \{\text{Idx}(\mathbf{x}) : \mathbf{x} \in \mathbf{D}_x\}$ be the set of indices of points in dataset \mathbf{D}_x . Given a query with range $\mathbf{B} = [a_1, b_1] \times [a_2, b_2] \times \dots \times [a_d, b_d] \subset \mathbb{N}^d$, which asks for the sum $\bigoplus_{\mathbf{x} \in \mathbf{D}_x \cap \mathbf{B}} \mathbf{x}$, algorithm **Translate** output a query with range $\mathbf{R} = [\beta_1, \gamma_1] \times [\beta_2, \gamma_2] \times \dots \times [\beta_d, \gamma_d] \subset [0, 1]^d$, which asks for the sum $\bigoplus_{\mathbf{i} \in \mathbf{D}_i \cap \mathbf{R}} \text{Att}(\mathbf{i})$, such that (1) the two queries are equivalent, that is, $\bigoplus_{\mathbf{x} \in \mathbf{D}_x \cap \mathbf{B}} \mathbf{x} = \bigoplus_{\mathbf{i} \in \mathbf{D}_i \cap \mathbf{R}} \text{Att}(\mathbf{i})$; (2) Alice can verify the correctness of the translation.

Translate.

1. Alice sends to Bob $\{(a_i, b_i) : i \in [d]\}$.
 2. Bob returns to Alice $\{(\bar{\beta}_i, \bar{x}_i, t_{\bar{\beta}_i}; \beta_i, x_i, t_{\beta_i}; \gamma_i, y_i, t_{\gamma_i}; \bar{\gamma}_i, \bar{y}_i, t_{\bar{\gamma}_i}) : i \in [d]\}$.
 3. Alice outputs $\{(\beta_i, \gamma_i) : i \in [d]\}$, if
 - (a) for each $i \in [d]$, $\bar{x}_i < a_i \leq x_i$ and $y_i \leq b_i < \bar{y}_i$;
 - (b) $t_{\bar{\beta}_i}, t_{\beta_i}, t_{\gamma_i}$ and $t_{\bar{\gamma}_i}$, are valid **itag** values of $(\bar{\beta}_i, \bar{x}_i)$, (β_i, x_i) , (γ_i, y_i) and $(\bar{\gamma}_i, \bar{y}_i)$, respectively.
 Otherwise, Alice output \perp .
-

5 Dynamic MAIA

The dynamic MAIA is a \mathcal{RC} protocol [25] between Alice and Bob, and consists of setup phase and query phase. In setup phase, Alice runs the key generating algorithm **KGen** to generate a private key, then encodes the dataset using data encoding algorithm **DEnc** with the private key. The query phase consists of multiple query sessions. In each query session, Alice either issue a query or update the dataset by inserting or deleting a point from the dataset.

5.1 Setup phase

5.1.1 KGen. At first, we define a subroutine **KeyGen**. **KeyGen**(1^κ):

Alice

1. runs algorithm $\mathcal{G}(1^\kappa)$ to generate $(n, p, q, r, G_p, G_q, G_r, g_p, g_q, g_r)$.
2. chooses $(s_{\ell,1}, s_{\ell,2}, s_{\ell,3})$ at random from $\mathbb{Z}_p^* \times \mathbb{Z}_q^* \times \{0, 1\}^\kappa$ for each $\ell \in [d]$.

Let $\mathbf{s}_1 = (s_{1,1}, s_{2,1}, \dots, s_{d,1})$, $\mathbf{s}_2 = (s_{1,2}, s_{2,2}, \dots, s_{d,2})$ and $\mathbf{s}_3 = (s_{1,3}, s_{2,3}, \dots, s_{d,3})$. Let $\mathcal{K} = (\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3, n, p, q, r, g_p, g_q, g_r)$. Output \mathcal{K} .

Next, we define the algorithm **KGen** based on **KeyGen**.

KGen:

1. Run **KeyGen**(1^κ) and obtains output \mathcal{K} .
2. Run **KeyGen**(1^κ) and obtains output $\bar{\mathcal{K}}$.

Output $(\mathcal{K}, \bar{\mathcal{K}})$.

5.2 DEnc.

1. For each dimension $j \in [d]$, construct a Doit tree $\mathcal{T}^{(j)}$ for $\mathbf{D}_x^{(j)} = \{\mathbf{x}[j] : \mathbf{x} \in \mathbf{D}\}$ and define **Att**, **Idx**, **ITag**, **Compress** and **Uncompress**;
2. For each $\mathbf{i} \in \mathbf{D}_i$, Alice
 - (a) chooses a number ξ_i at random from \mathbb{Z}_q^* ;
 - (b) computes a tag \mathbf{t}_i using the private key \mathcal{K} :

$$\mathbf{t}_i \leftarrow \text{Tag}_{\mathcal{K}}(\text{Att}(\mathbf{i}), \mathbf{i}; \xi_i). \quad (1)$$

3. Let $\mathbf{T} = \{\mathbf{t}_i : \mathbf{i} \in \mathbf{D}_i\}$, and $\bar{\mathbf{D}} = \bar{\mathbf{T}} = \emptyset$. Alice sends $\mathbf{D}, \bar{\mathbf{D}}, \mathbf{T}$ and $\bar{\mathbf{T}}$ to Bob, removes local copy of $\mathbf{D}, \bar{\mathbf{D}}$ and \mathbf{T} from her storage, and keeps $(\mathcal{K}, \bar{\mathcal{K}})$.
-

5.3 Query phase

5.3.1 Query. $\langle \text{Eval}, \text{Ext} \rangle(\text{SUM}(\mathbf{R}))$. Let **ProVer** be as in scheme **MAIA** [25].

1. Simulate **ProVer** on input $\text{SUM}(\mathbf{R})$ w.r.t. dataset \mathbf{D} using private key \mathcal{K} , and obtain result \mathbf{X} .
 2. Simulate **ProVer** on input $\text{SUM}(\mathbf{R})$ w.r.t. dataset $\bar{\mathbf{D}}$ using private key $\bar{\mathcal{K}}$, and obtain result $\bar{\mathbf{X}}$.
 3. Output $\mathbf{X} - \bar{\mathbf{X}}$.
-

5.3.2 Insert.

Insert($\mathbf{x}, \mathbf{D}, \mathbf{T}, \mathcal{K}, \pi^*$): Insert \mathbf{x} into \mathbf{D} .

1. For each $j \in [d]$, $\text{insert}(\mathcal{T}^{(j)}, \mathbf{x}[j])$.
 2. Let $\mathbf{i} \leftarrow \text{Idx}(\mathbf{x})$.
 3. Choose ξ from \mathbb{Z}_q^* at random, and compute tag: $\mathbf{t}_i \leftarrow \text{Tag}_{\mathcal{K}}(\mathbf{x}, \mathbf{i}; \xi)$.
 4. Add \mathbf{x} into set \mathbf{D} , and add \mathbf{t}_i into set \mathbf{T} .
 5. Update π^* : $\pi^* \leftarrow \pi^* \times \text{CTag}_{\mathcal{K}}(\mathbf{i})[1] \pmod n$.
-

5.3.3 Delete.

Delete(\mathbf{x}): Delete \mathbf{x} from the dataset $\mathbf{D}_{\mathbf{x}}$.

1. If $\mathbf{x} \in \mathbf{D}_{\mathbf{x}}$, then Insert($\mathbf{x}, \bar{\mathbf{D}}, \bar{\mathbf{T}}, \bar{\mathcal{K}}, \bar{\pi}^*$). Otherwise, do nothing.
-

Editing can be implemented with combination of deleting and inserting. We save the details.

5.4 Complexity Analysis

Without surprise, due the use of tree, the communication complexity of our scheme is $O(d^2 \log N)$. The computation and storage overhead remain the same as MAIA in Xu J. *et al.* [25]. The details are showed in Table 1 in Section 1.

5.5 Security

Xu J. *et al.* [25] gave a security model for outsourced computing, called *Provable Remote Computing (PRC)* protocol.

Theorem 1 *Dynamic MAIA as constructed above is a PRC w.r.t. aggregate SUM query, and insertion/deletion operations, if the Assumption 1 and Assumption 2 in Xu J. et al. [25] hold, and RSA signature is a secure multiplicative homomorphic signature scheme.*

6 Conclusion

In this paper, we proposed a dynamic ordered indexing ternary tree, and based it we proposed a method to authenticate aggregate range query over dynamic multidimensional dataset. The supported query types include SUM, COUNT, MIN, MAX and MEDIAN. The communication overhead is $O(d^2 \log N)$, where d is the dimension and N is the number of points in the dataset.

References

1. M. J. Atallah, Y. Cho, and A. Kundu. Efficient data authentication in an environment of untrusted third-party distributors. In *ICDE '08: Proceedings of the 2008 IEEE 24th International Conference on Data Engineering*, pages 696–704, Washington, DC, USA, 2008. IEEE Computer Society.
2. D. Boneh, B. Lynn, and H. Shacham. Short signatures from the weil pairing. *J. Cryptol.*, 17(4):297–319, 2004.
3. S. H. W. G. H. Brian Thompson, Danfeng Yao and T. Sander. Privacy-preserving computation and verification of aggregate queries on outsourced databases. In *Privacy Enhancing Technologies Symposium (PETS)*, Aug 2009.
4. W. Cheng and K.-L. Tan. Query assurance verification for outsourced multi-dimensional databases. *J. Comput. Secur.*, 17(1):101–126, 2009.
5. P. Devanbu, M. Gertz, C. Martel, and S. G. Stubblebine. Authentic data publication over the internet. *J. Comput. Secur.*, 11(3):291–314, 2003.

6. P. T. Devanbu, M. Gertz, C. U. Martel, and S. G. Stubblebine. Authentic third-party data publication. In *Proceedings of the IFIP TC11/ WG11.3 Fourteenth Annual Working Conference on Database Security*, pages 101–112, Deventer, The Netherlands, The Netherlands, 2001. Kluwer, B.V.
7. T. Ge and S. B. Zdonik. Answering aggregation queries in a secure system model. In *VLDB*, pages 519–530, 2007.
8. R. Gennaro, S. Halevi, and T. Rabin. Secure hash-and-sign signatures without the random oracle. In *EUROCRYPT*, pages 123–139, 1999.
9. M. T. Goodrich, R. Tamassia, and N. Triandopoulos. Super-efficient verification of dynamic outsourced databases. In *CT-RSA*, pages 407–424, 2008.
10. S. Haber, W. Horne, T. Sander, and D. Yao. Privacy-preserving verification of aggregate queries on outsourced databases. Technical report, HP Laboratories, 2006. HPL-2006-128.
11. H. Hacigümüş, B. Iyer, C. Li, and S. Mehrotra. Executing sql over encrypted data in the database-service-provider model. In *SIGMOD '02: Proceedings of the 2002 ACM SIGMOD international conference on Management of data*, pages 216–227, New York, NY, USA, 2002. ACM.
12. H. Hacigümüş, B. R. Iyer, and S. Mehrotra. Efficient execution of aggregation queries over encrypted relational databases. In *DASFAA*, pages 125–136, 2004.
13. K. M. HweeHwa PANG, Jilian ZHANG. Scalable verification for outsourced dynamic databases. In *Will appear in 35th International Conference on Very Large Data Bases (VLDB09)*.
14. F. Li, M. Hadjieleftheriou, G. Kollios, and L. Reyzin. Dynamic authenticated index structures for outsourced databases. In *SIGMOD Conference*, pages 121–132, 2006.
15. C. Martel, G. Nuckolls, P. Devanbu, M. Gertz, A. Kwong, and S. G. Stubblebine. A general model for authenticated data structures. *Algorithmica*, 39(1):21–41, 2004.
16. K. Mouratidis, D. Sacharidis, and H. Pang. Partially materialized digest scheme: an efficient verification method for outsourced databases. *The VLDB Journal*, 18(1):363–381, 2009.
17. E. Mykletun, M. Narasimha, and G. Tsudik. Authentication and integrity in outsourced databases. *Trans. Storage*, 2(2):107–138, 2006.
18. E. Mykletun and G. Tsudik. Aggregation queries in the database-as-a-service model. In *IFIP WG 11.3 Working Conference on Data and Applications Security*, pages 89–103, July 2006.
19. H. Pang, A. Jain, K. Ramamritham, and K.-L. Tan. Verifying completeness of relational query results in data publishing. In *SIGMOD '05: Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pages 407–418, New York, NY, USA, 2005. ACM.
20. H. Pang and K.-L. Tan. Verifying completeness of relational query answers from online servers. *ACM Trans. Inf. Syst. Secur.*, 11(2):1–50, 2008.
21. F. P. Preparata and M. I. Shamos. *Computational geometry: an introduction*. Springer-Verlag New York, Inc., New York, NY, USA, 1985.
22. R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.
23. R. Sion. Query execution assurance for outsourced databases. In *VLDB*, pages 601–612, 2005.
24. M. Xie, H. Wang, J. Yin, and X. Meng. Integrity auditing of outsourced data. In *VLDB '07: Proceedings of the 33rd international conference on Very large data bases*, pages 782–793. VLDB Endowment, 2007.
25. J. XU and E.-C. CHANG. Authenticating aggregate range queries over multidimensional dataset. Cryptology ePrint Archive, Report 2010/050, 2010. <http://eprint.iacr.org/>.
26. Y. Yang, D. Papadias, S. Papadopoulos, and P. Kalnis. Authenticated join processing in outsourced databases. In *SIGMOD Conference*, pages 5–18, 2009.