

Dynamic Searchable Symmetric Encryption

DSSE

INTRODUCTION TO CYBER SECURITY

ESTR 4306 | LIU YICUN

Motivating Problem

The clients want a encrypted database which can update dynamically.

The clients want to do the updates without re-encrypting the whole data base.

The updates should affect the previous DB as less as possible.

Queries and updates are supposed to reveal as less information as possible.

Efficiency vs. Functionality

Abstract Idea of “Dynamic”

Additions:

- Edge Addition: Simple. Just add edge (x, y) into the DB, noted that node x and y have already existed.
- Node Addition: Based on edge addition. Add a set of new edges (\cdot, y) connecting node y . The client need to keep all edges connecting node y in mind.

Deletions:

- Edge Deletion: Just remove (x, y) . Simple, but unrealistic for massive deletion.
- Node Deletion: Client give the name of the node y to the server. The server do the rest of the job. Server traverses (\cdot, y) and deletes all these edges.

Notations

1. Edge Addition: $u = (\text{Add}, x, y, w)$ adds the edge (x, y, w) to G , where x and y are originally not connected to each other. However, x and y can be originally connected to other nodes respectively.
2. Node Addition: $u = (\text{Add}, \mathbf{x}, y, \mathbf{w})$ (or $u = (\text{Add}, x, \mathbf{y}, \mathbf{w})$) adds the edges (x_i, y, w_i) where $x_i \in \mathbf{x}$ (or (x, y_i, w_i) where $y_i \in \mathbf{y}$) and $w_i \in \mathbf{w}$ to G , where node y (or x) originally does not exist in G .
3. Edge Deletion: $u = (\text{Del}, x, y, *)$ deletes the edge of the form (x, y, \cdot) from G .
4. Node Deletion: $u = (\text{Del}, x, *, *)$ (or $u = (\text{Del}, *, y, *)$) deletes the set $(x, *, *)$ (or $(*, y, *)$) of edges from G .

General DSSE

Bipartite graphs defined by the spaces \mathcal{X} , \mathcal{Y} , and \mathcal{W} .

- $(K, \text{EDB}) \leftarrow \text{Setup}(1^\lambda, M, N)$: *In the setup algorithm, the user inputs the security parameter λ , the size M and N of the spaces \mathcal{X} and \mathcal{Y} respectively. It outputs a key K , and an (initially empty) encrypted database EDB , which is outsourced to the server.*
- $((K', R), (\text{EDB}', R)) \leftarrow \text{Query}_e((K, q), \text{EDB})$:
 - ✓ $(K', \tau_q) \leftarrow \text{QueryToken}(K, q)$: Generate a query token
 - ✓ $(\text{EDB}', R) \leftarrow \text{Query}_e(\tau_q, \text{EDB})$: Return the result R of the query. Update EDB .
- $(K', \text{EDB}') \leftarrow \text{Update}_e((K, u), \text{EDB})$. $u \in \{\text{Add}, \text{Del}\} \times \mathcal{X} \times \mathcal{Y} \times \mathcal{W}$.
 $(K', \tau_u) \leftarrow \text{UpdateToken}(K, u)$ $\text{EDB}' \leftarrow \text{Update}_e(\tau_u, \text{EDB})$

Forward Privacy

Example: The client adds a new edge (x, y, w) . The server might be able to tell that this new edge is connecting to node x , using the previous query result on x .

“Forward Privacy” was previously discussed in [SPS14, RG15], however it is captured by the leakage function when query.

Captured by \mathcal{L}_q could be problematic: in our example, the criterion to determine whether forward privacy has been achieved happens in updates, but not queries!

Our definition: based on update leakage function

Forward Privacy

- Half-forward Privacy:

Addition of edge (x, y) doesn't reveal both x and y . Hide one end.

- *half-forward private, if for any $u_b = (\text{Add}, x_b, y_b, w)$ where $b = 0, 1$, and $x_0 = x_1$ or $y_0 = y_1$, then for any PPT distinguisher \mathcal{D} , it holds that*

$$|\Pr[\mathcal{D}(\mathcal{L}_u(u_0)) = 1] - \Pr[\mathcal{D}(\mathcal{L}_u(u_1)) = 1]| \leq \text{negl}(\lambda).$$

- Fully-forward Privacy:

Addition of edge (x, y) doesn't reveal either x and y . Hide both ends.

- *fully-forward private, if for any $u_b = (\text{Add}, x_b, y_b, w)$ where $b = 0, 1$, then for any PPT distinguisher \mathcal{D} , it holds that*

$$|\Pr[\mathcal{D}(\mathcal{L}_u(u_0)) = 1] - \Pr[\mathcal{D}(\mathcal{L}_u(u_1)) = 1]| \leq \text{negl}(\lambda).$$

Forward Privacy

- A more comprehensive definition of Forward Privacy:

Based on the update leakage function \mathcal{L}_u

$$|\Pr[\mathcal{D}(\mathcal{L}_u(u_0)) = 1] - \Pr[\mathcal{D}(\mathcal{L}_u(u_1)) = 1]| \leq \text{negl}(\lambda).$$

- Previously defined half forward privacy can be represented as $(\mathcal{U}_1 \cup \mathcal{U}_4, p_1 + p_4)$
- Previously defined fully forward privacy can be represented as $(\mathcal{U}_1 \cup \mathcal{U}_4, 1)$

| i | Sets of updates \mathcal{U}_i | Possible conditions such that $p_i = 1$ (If there are no restrictions, $p_i \equiv 1$) |
|-----|---|--|
| 1 | $\{(\text{Add}, x, y, w)\}$ | $x_0 = x_1$ or $y_0 = y_1$ |
| 2 | $\{(\text{Add}, \mathbf{x}, y, \mathbf{w})\}$ | $y_0 = y_1$ |
| 3 | $\{(\text{Add}, x, \mathbf{y}, \mathbf{w})\}$ | $x_0 = x_1$ |
| 4 | $\{(\text{Del}, x, y, *)\}$ | $x_0 = x_1$ or $y_0 = y_1$ |
| 5 | $\{(\text{Del}, x, *, *)\}$ | $ (x_0, *, *) = (x_1, *, *) $ |
| 6 | $\{(\text{Del}, *, y, *)\}$ | $ (*, y_0, *) = (*, y_1, *) $ |

Forward Privacy for Any DSSE

- DSSE scheme \mathcal{E} . (implemented before)
- Each query has a table of PRF (deterministic) keys K_x and counters c_x .
- Plaintext bipartite graph \hat{G} and ciphertext bipartite graph G
(to be achieved by an efficient data structure).
- Dictionary (or table) $\gamma^{\mathcal{Q}}$:
Maps a query $q = x \in \mathcal{X}$ to a PRF key K_x and counter c_x .
- Idea: locally maintain the table $\gamma^{\mathcal{Q}}$.

Setup

1. The dictionary and the plaintext graph are set to be empty.

2. The returned K contains the key and the dictionary. Both the EDB and the plaintext graph are returned.

$$(K, \text{EDB}) \leftarrow \text{Setup}(1^\lambda)$$

1 : $(\tilde{K}, \tilde{\text{EDB}}) \leftarrow \mathcal{E}.\text{Setup}(1^\lambda)$

2 : $\gamma^{\mathcal{Q}} = \phi, \hat{G} = \phi$

3 : **return** $K = (\tilde{K}, \gamma^{\mathcal{Q}}), \text{EDB} = (\tilde{\text{EDB}}, \hat{G})$

Addition: UpdateToken

$u = (\text{Add}, x, y, w)$ sent by the client, means adding a new edge (x, y) .

Case 1: if node x is a singleton, generate the PRF key and set counter as 1.

Case 2: if node x is not a singleton, counter + 1.

Then generate the UpdateToken for addition.

$(K', \tau_u^+) \leftarrow \text{UpdateToken}(K, u = (\text{Add}, x, y, w))$

```
1: if  $\gamma^Q[x] = \perp$  then
2:    $K_x \leftarrow \{0, 1\}^\lambda$ 
3:    $c_x \leftarrow 1$ 
4: else
5:    $(K_x, c_x) \leftarrow \gamma^Q[x]$ 
6:    $c_x \leftarrow c_x + 1$ 
7: endif
8:  $\gamma^Q[x] \leftarrow (K_x, c_x)$ 
9:  $\tau_u^+ \leftarrow \mathcal{E}.\text{UpdateToken}(\tilde{K}, (\text{Add}, F(K_x, c_x), y, w))$ 
10: return  $(K, \tau_u^+)$ 
```

Addition: Update

Simple, using the already generated updated token, and update the EDB.

1 : $\tilde{\text{EDB}} \leftarrow \mathcal{E}.\text{Update}(\tilde{\tau}_u^+, \tilde{\text{EDB}})$

2 : **return** EDB

Deletion: UpdateToken

$u = (\text{Del}, x, *, *)$ sent by the client, means deleting all edges connecting node x .

More complex than addition, cause more edges to deal with.

Case 1: if node x is a singleton

1. Get the PRF key, know the counter. Then set $\text{dict}[x]$ empty.
2. Generate DeleteToken for each edges
3. Encapsulate them into an UpdateToken

Case 2: if node x is not a singleton, UpdateToken = empty

```
1 : if  $u = (\text{Del}, x, *, *)$  then
2 :   if  $\gamma^Q[x] \neq \perp$  then
3 :      $(K_x, c_x) \leftarrow \gamma^Q[x]$ 
4 :      $\gamma^Q[x] \leftarrow \perp$ 
5 :     for  $i = 1, \dots, c_x$  do
6 :        $\tilde{\tau}_i^- \leftarrow \mathcal{E}.\text{UpdateToken}(\tilde{K}, (\text{Del}, F(K_x, i), *, *))$ 
7 :     endfor
8 :      $\tau_u^- \leftarrow (x, \{\tilde{\tau}_i^-\}_{i=1}^{c_x})$ 
9 :   else
10 :     $\tau_u^- \leftarrow x$ 
11 :   endif
```

Deletion: Update

Here we already have a update token, which contains all deletion tokens of edges $(x, *)$.

1. Delete all the edges in plaintext graph.
2. Delete all the edges in EDB.

Here we only discuss half of the situation, the other half is deleting node y , whose token generation and update procedure is very similar to the previous situation.

$\text{EDB}' \leftarrow \text{Update}_e(\tau_u^-, \text{EDB})$ for $u = (\text{Del}, \cdot, \cdot, \cdot)$

- 1 : if $\tau_u^- = (x \in \mathcal{X}, \{\tilde{\tau}_i^-\}_{i=1}^{c_x})$ then
- 2 : $\hat{G} \leftarrow \text{Update}((\text{Del}, x, *, *), \hat{G})$
- 3 : for $i = 1, \dots, c_x$ do
- 4 : $\tilde{\text{EDB}} \leftarrow \mathcal{E}.\text{Update}_e(\tau_i^-, \tilde{\text{EDB}})$
- 5 : endfor

Query: QueryToken

When $q = x$, query on a node x .

Case 1: If node x is a singleton

1. Get the PRF key and the counter in dict.
2. Generate the QueryToken and UpdateToken for each counter.
3. Encapsulate these two kinds of token into one token named QueryToken.

Case 2: x is not a singleton.

QueryToken = x .

Case 3: x doesn't exist in previous graph. QueryToken = empty.

$$(K', \tau_q) \leftarrow \text{QueryToken}(K, q)$$

```
1 : if  $q = x \in \mathcal{X} \wedge \gamma^Q[x] \neq \perp$  then
2 :    $(K_x, c_x) \leftarrow \gamma^Q[x]$ 
3 :    $\gamma^Q[x] \leftarrow \perp$ 
4 :   for  $i = 1, \dots, c_x$  do
5 :      $\tilde{\tau}_i \leftarrow \mathcal{E}.\text{QueryToken}(\tilde{K}, F(K_x, i))$ 
6 :      $\tilde{\tau}_i^- \leftarrow \mathcal{E}.\text{UpdateToken}(\tilde{K}, (\text{Del}, F(K_x, i), *, *))$ 
7 :   endfor
8 :    $\tau_q \leftarrow (x, \{\tilde{\tau}_i, \tilde{\tau}_i^-\}_{i=1}^{c_x})$ 
```

$$\tau_q \leftarrow x$$
$$\tau_q \leftarrow \perp$$

Query: Do Query

Why do we need the plaintext graph?

Once the query on x is done, we think that $(x, *)$ is not secure anymore. For efficiency, we throw these edges to a plaintext graph.

1. Try to get the result directly from plaintext graph
2. Get the result from EDB, delete these edges from EDB
3. Merge the result from step 1 and 2
4. Adding the query result to the plaintext graph

```
1 : Parse  $\tau_q$  as  $(x, \{\tilde{\tau}_i, \tilde{\tau}_i^-\}_{i=1}^{c_x})$ 
2 :  $R \leftarrow \text{Query}(x, \hat{G})$ 
3 : for  $i = 1, \dots, c_x$  do
4 :    $\tilde{R} \leftarrow \mathcal{E}.\text{Query}_e(\tilde{\tau}_i, \text{EDB})$ 
5 :    $\text{EDB} \leftarrow \mathcal{E}.\text{Update}_e(\tilde{\tau}_i^-, \text{EDB})$ 
6 :    $R \leftarrow R \cup \tilde{R}$ 
7 : endfor
8 : foreach  $(\tilde{x}, y, w) \in R$  do
9 :    $R \leftarrow R \setminus \{(\tilde{x}, y, w)\} \cup \{(x, y, w)\}$ 
10 :   $\hat{G} \leftarrow \text{Update}((\text{Add}, x, y, w), \hat{G})$ 
11 : endfor
12 : return  $(\text{EDB}, R)$ 
```


Security Proof

- Take the proof of previously constructed (\mathcal{U}_1, p'_1) for illustration:

$$p'_1(u_0, u_1) = 1 \text{ if and only if } y_0 = y_1 \text{ and } w_0 = w_1$$

Simulator $\mathcal{S}_{\mathcal{E}}$: which simulates the queries and updates of \mathcal{E} , which is a $(\tilde{\mathcal{L}}_q, \tilde{\mathcal{L}}_u)$ -CQA2 secured scheme.

Simulator \mathcal{S} : which receive the leakage for random x when update.

$$\text{for instance: } \mathcal{L}_u(\text{Add}, x, y, w) = (\tilde{x}, \tilde{\mathcal{L}}_u(\text{Add}, \tilde{x}, y, w))$$

Then pass $\tilde{\mathcal{L}}_u(\text{Add}, \tilde{x}, y, w)$ to $\mathcal{S}_{\mathcal{E}}$, which outputs a simulated token τ_u^+ .

If there exists an environment which distinguish the simulation and the real scheme, then \mathcal{S} can be constructed by this environment.


Cascaded Triangles

Motivations: The updates should affect the previous DB as less as possible.

Previously we used linked list, trees and others.

Idea: we separate the big binary tree into many small perfect binary trees, whose size is kept in cascaded order. We call these perfect binary trees “Cascaded Triangles”.

So the update will just affect a few triangles!



Cascaded Triangles

Consists of three dictionaries γ^Q, γ^D , and γ^δ .

Why do we need three dictionaries instead of one? Think when traverse.

Firstly, we need to store the root address of every triangles, that is the first dictionary.

Secondly, given the root, we should know what are the left and right child (if any) of the root. So the second dictionary to store the detailed information of every node is needed.

Lastly, clients are not always traverse by X, so another dictionary to store the dual root address is needed.

Cascaded Triangles

Consists of three dictionaries $\gamma^{\mathcal{Q}}$, $\gamma^{\mathcal{D}}$, and γ^{δ} .

$\gamma^{\mathcal{Q}}$ stores the root address and heights of the triangles at the client.

$$\gamma^{\mathcal{Q}}[x] = (\text{addr}_1, \dots, \text{addr}_k, h)$$

$\gamma^{\mathcal{D}}$ stores the address of the dual tuples, for dual retrieving.

γ^{δ} maps the address to a tuple (a, b).

$$a = \bar{a} = (x, y, w) \qquad b = (\text{chd}_0, \text{chd}_1)$$

$\gamma^{\mathcal{D}}$ and γ^{δ} are stored at the servers' side.

Requirement of Heights: $h_1 \leq h_2 < h_3 < \dots < h_k, \quad h \in \{0, 1, 2\}^{\lceil \lg n_q \rceil}$.

Adding a Node

First add it as a root node of a new tree.


If two shortest trees have the same height, combine them.

Carry($h + 1$) = 1. , add (a, b) where $a = (x, y, w)$ and $b = (\text{addr}_1, \text{addr}_2)$ to a random address addr

$$\gamma^Q[x] \leftarrow (\text{addr}, \text{addr}_3, \dots, \text{addr}_k, h + 1),$$

Carry($h + 1$) = 0. $a = (x, y, w)$ and $b = (\perp, \perp)$

$$\gamma^Q[x] \leftarrow (\text{addr}, \text{addr}_1, \dots, \text{addr}_k, h + 1).$$




Deleting a Node

Firstly, replace that node with the root of the shortest tree.

If it is a singleton, simply remove the corresponding linked list node.

Otherwise, update the linked list so it points to one of its child.



Deleting a Node

If we want to delete node x :

Firstly, we traverse the set $(x, *, *)$ using the traverse algorithm.

Delete all the traverse nodes, based on the address of the root node.

Then, use traverse result of x to find the duals in triangles based on y .
Replace these nodes (in the middle of some triangles) with the smallest triangles.

- Lookup $\gamma^{\mathcal{D}}[\overline{\text{addr}}_1] = \text{addr}_1$.
- Delete $\gamma^{\mathcal{D}}[\text{addr}]$ and $\gamma^{\mathcal{D}}[\overline{\text{addr}}_1]$, and update $\gamma^{\mathcal{D}}[\text{addr}_1] \leftarrow \overline{\text{addr}}$ and $\gamma^{\mathcal{D}}[\overline{\text{addr}}] \leftarrow \text{addr}_1$.
- Lookup $\gamma^{\delta}[\overline{\text{addr}}_1] = (\bar{a}_1, \bar{b}_1)$, where $\bar{b}_1 = (\overline{\text{addr}}'_0, \overline{\text{addr}}'_1)$.
- Update $\gamma^{\delta}[\overline{\text{addr}}] \leftarrow (\bar{a}_1, \bar{b})$ and delete $\gamma^{\delta}[\overline{\text{addr}}_1]$.
- Update $\gamma^{\mathcal{Q}}[y] = (\overline{\text{addr}}'_0, \overline{\text{addr}}'_1, \overline{\text{addr}}_2, \dots, \overline{\text{addr}}_k, \bar{h} - 1)$, where $h - 1$ is the reverse process of $h + 1$ defined above.

Encrypted Version

Main idea:

Split G into two graphs, one is the ciphertext graph and plaintext graph.

1. Now we only stored the ciphertext in tuples of $\gamma^Q[x]$.
2. Using two secret keys to access $a = (x, y, w)$ and $b = (ch0, ch1)$. The keys are independently generated for each x .
3. The $\gamma^Q[x]$ additionally stores the two secret keys associated to x .

Encrypted Version: Queries

Differences form non-encrypted version:

1. When retrieving of the root address, the two secret keys are also retrieved.
2. The clients sends the result of the retrieval back to the server. The server used the first key to get part of the query result. And use the second key to get the subtrees.
3. The server also returns the previous result in plaintext graph.