




Injection-Secure Structured Encryption and Searchable Symmetric Encryption

Ghous Amjad, Seny Kamara, Tarik Moataz

Encrypted Search Algorithms

Trusted client




Untrusted server



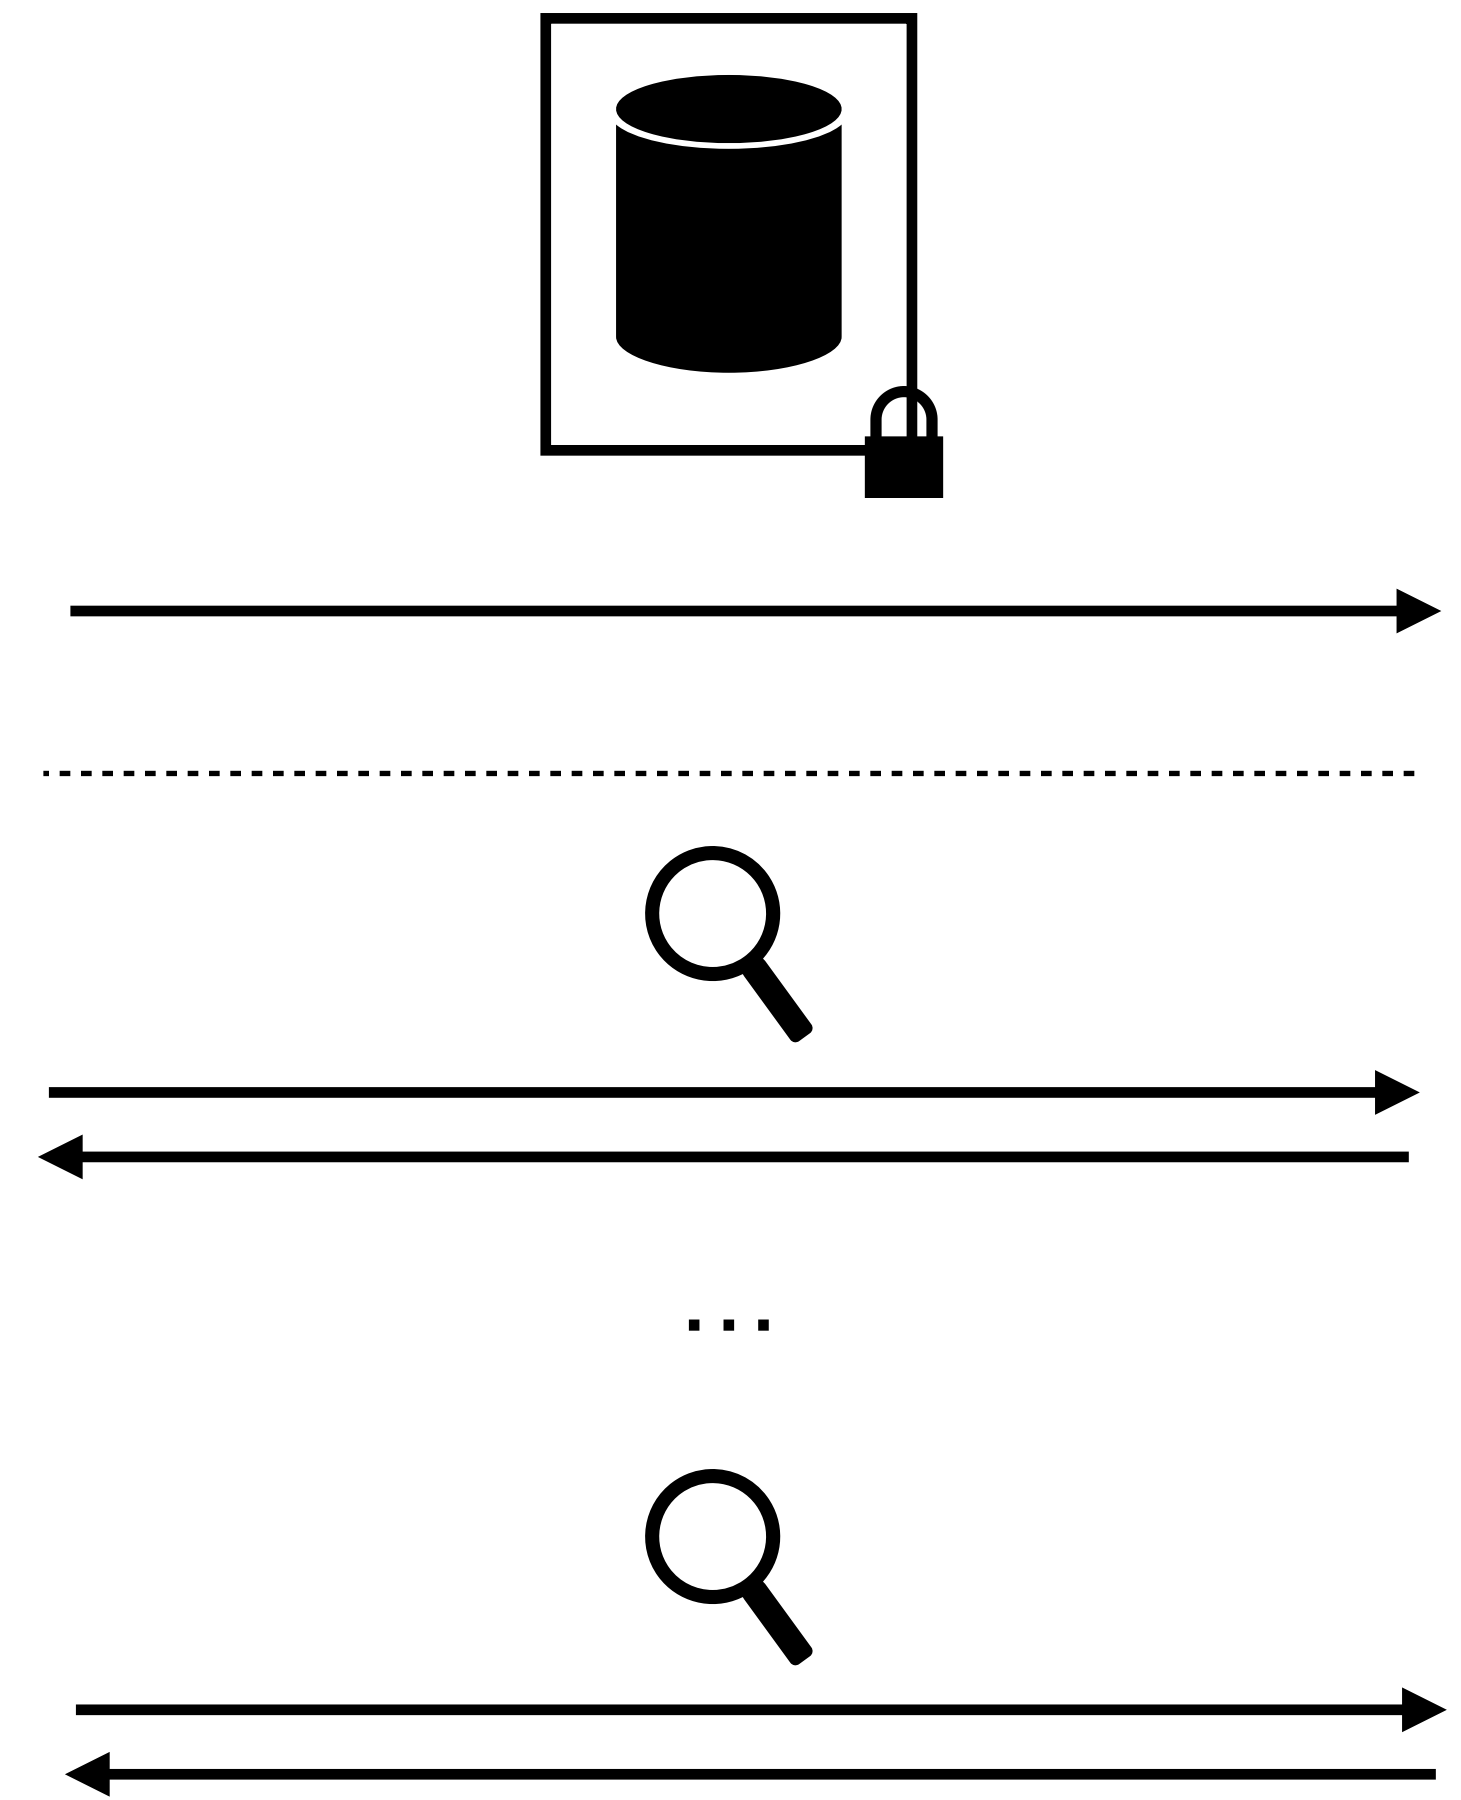
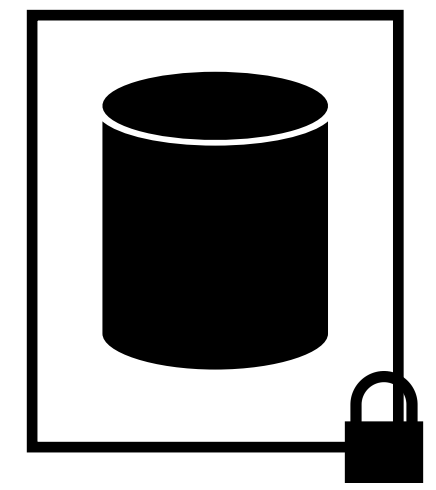
Encrypted Search Algorithms

Overview

Trusted client



Untrusted server



Encrypted Search Algorithms

Encrypted Search Algorithms

Fully-Homomorphic Encryption (FHE)

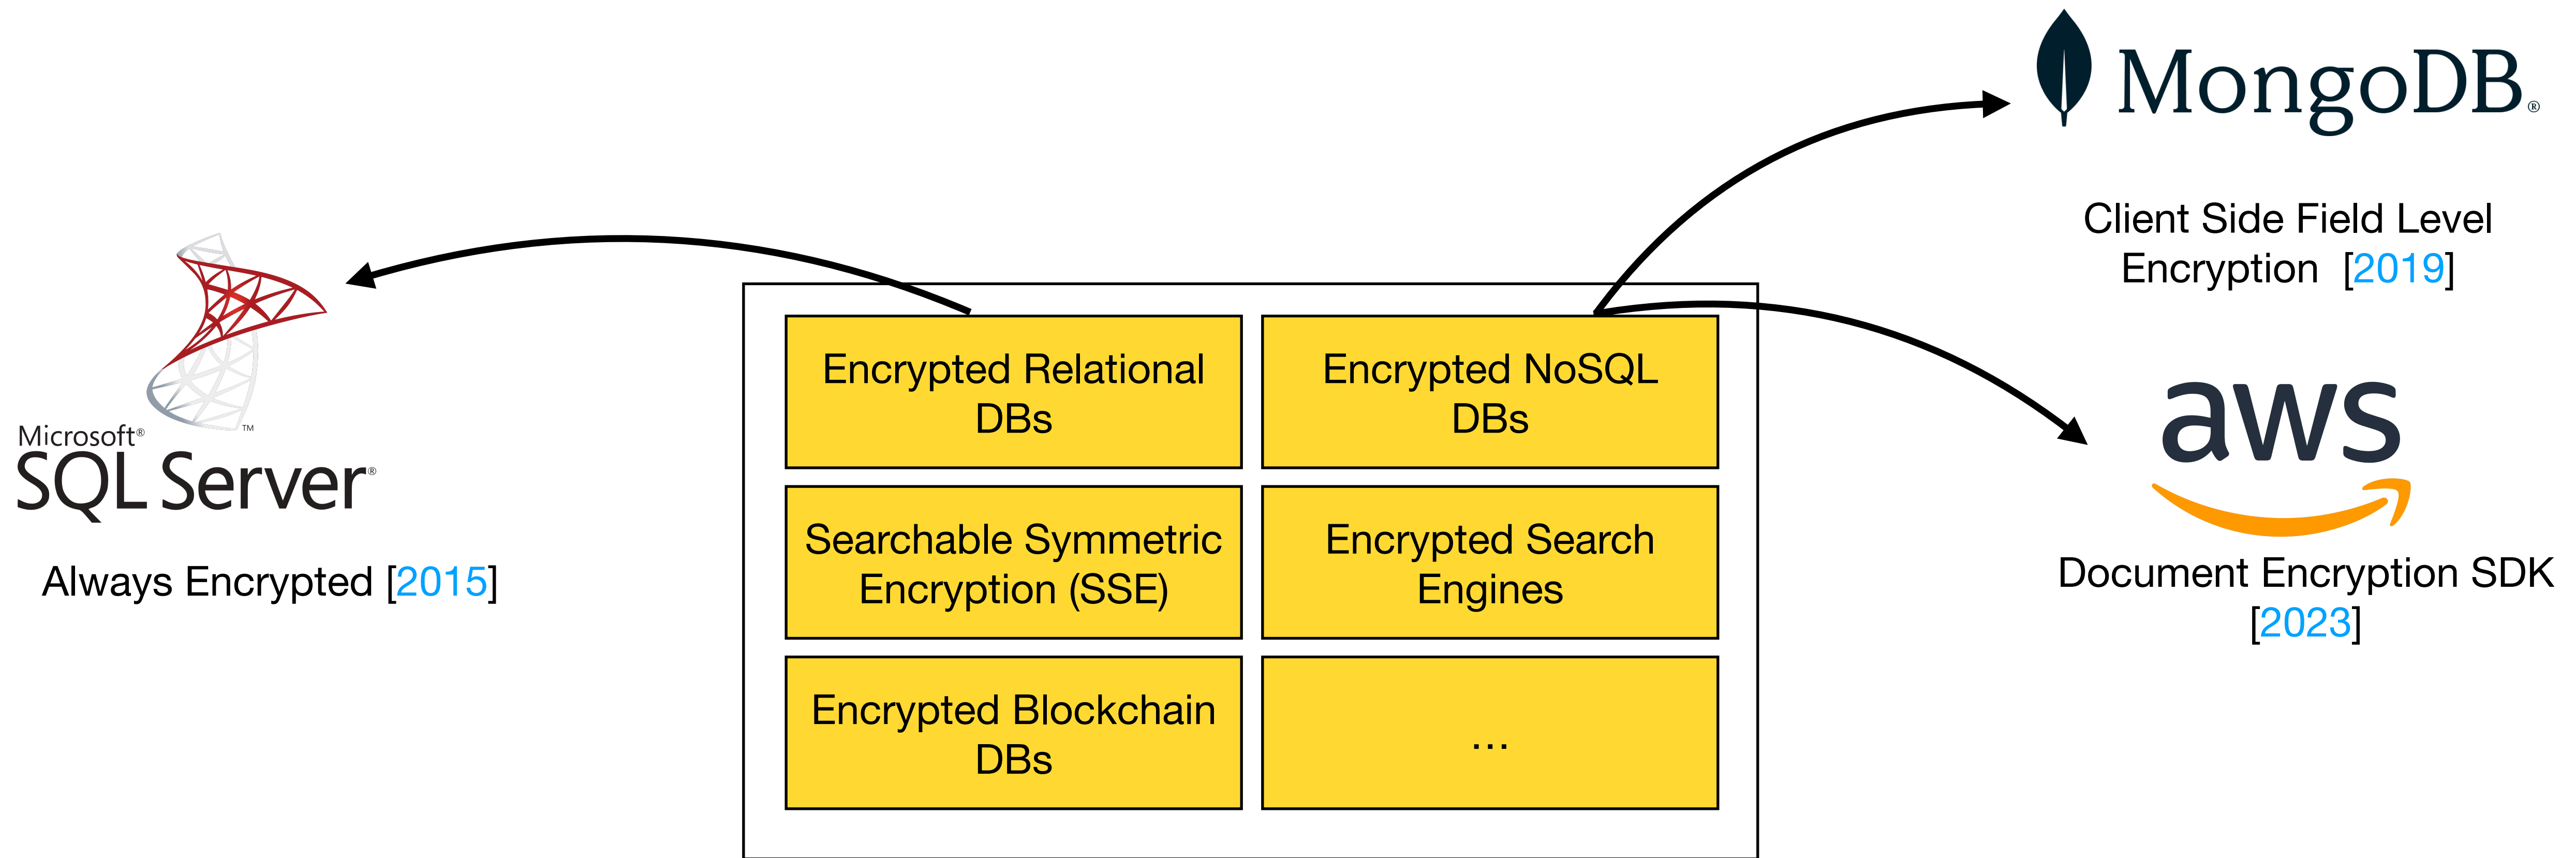
Property-Preserving Encryption (PPE)

Structured Encryption (STE)

Oblivious RAM (ORAM)

Functional Encryption (FE)

Encrypted Search Algorithms



Fully-Homomorphic Encryption (FHE)

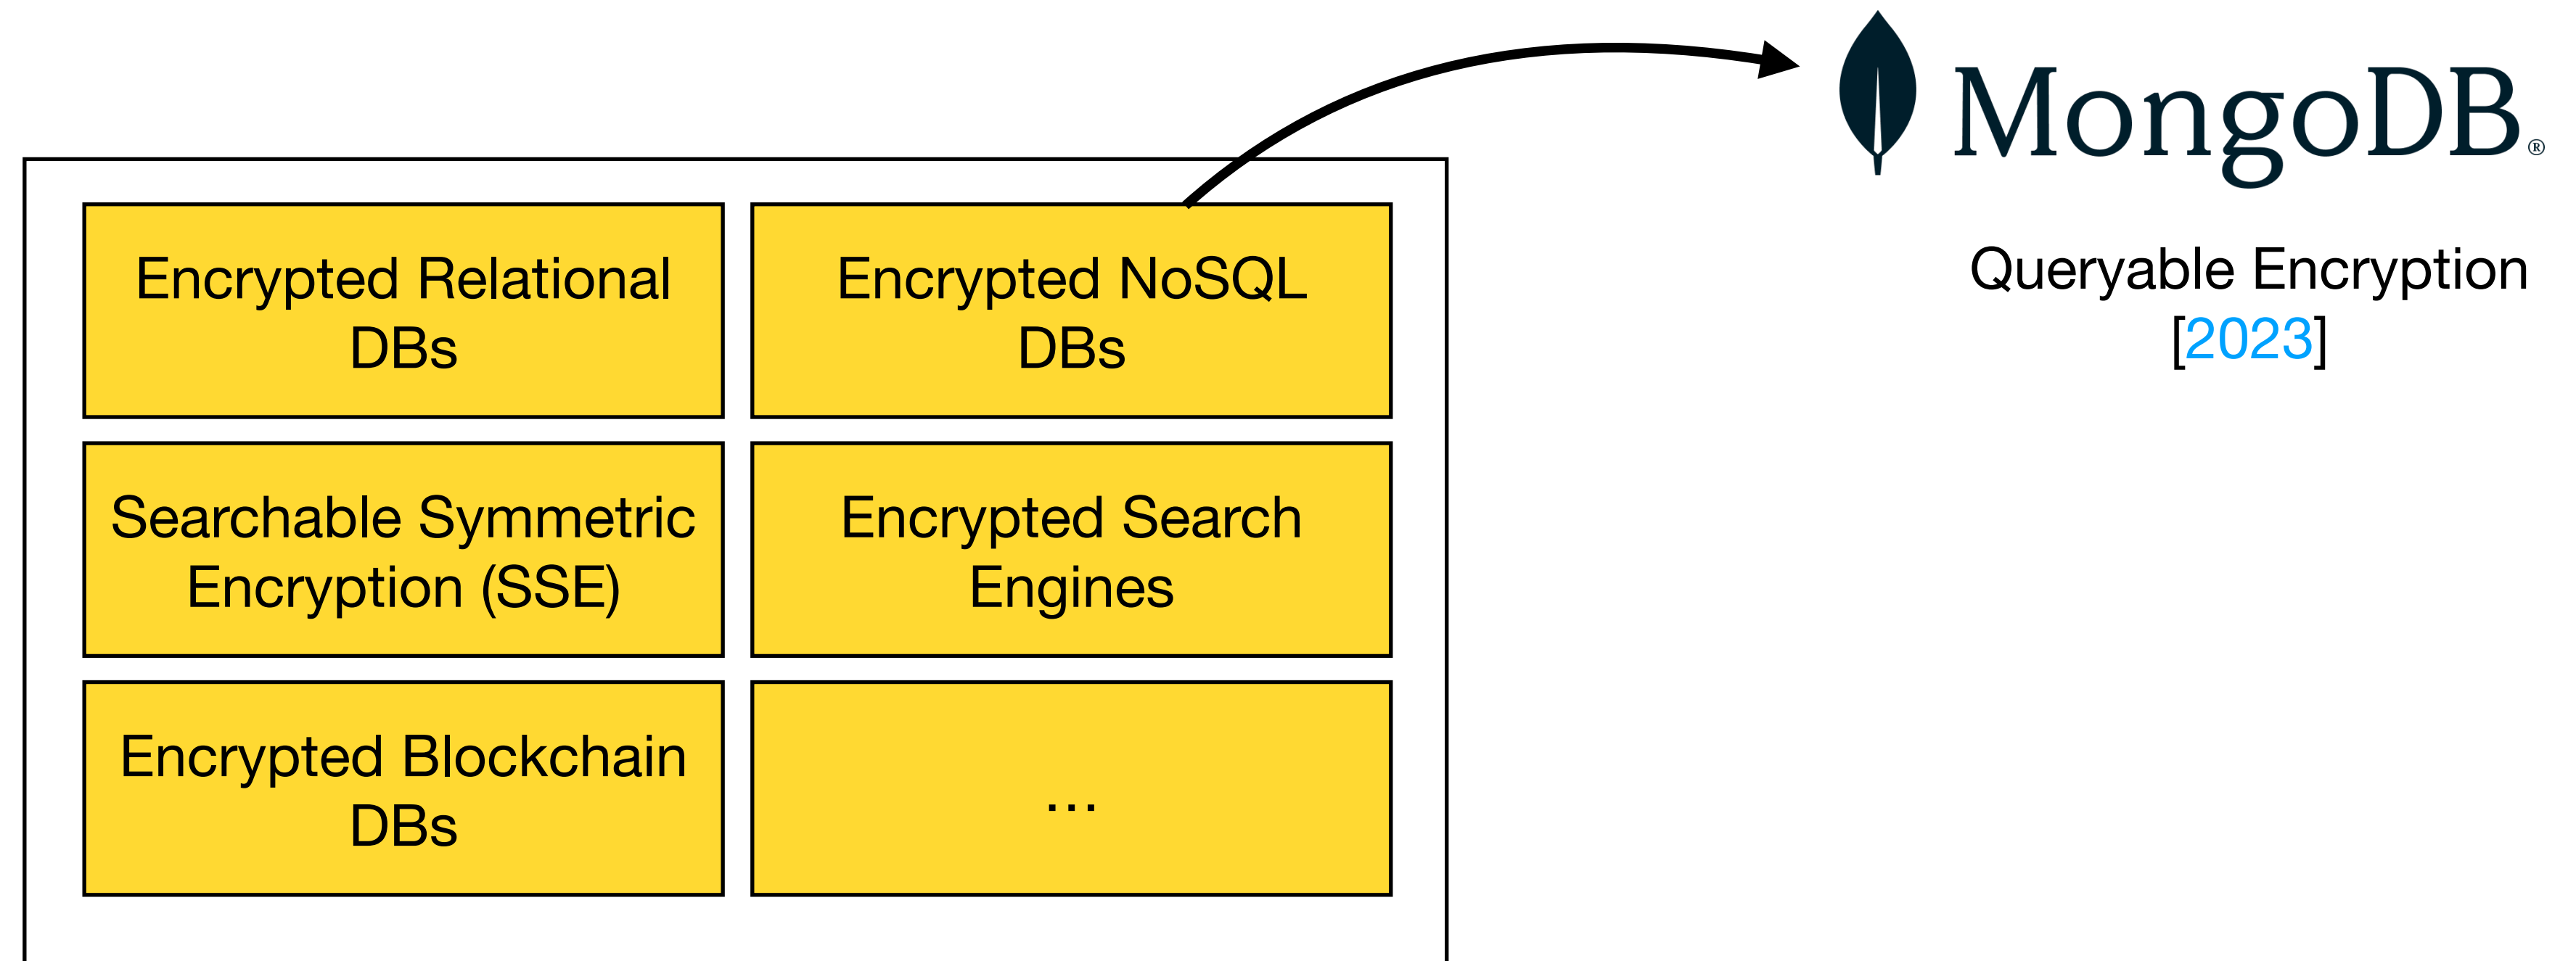
Property-Preserving Encryption (PPE)

Structured Encryption (STE)

Oblivious RAM (ORAM)

Functional Encryption (FE)

Encrypted Search Algorithms




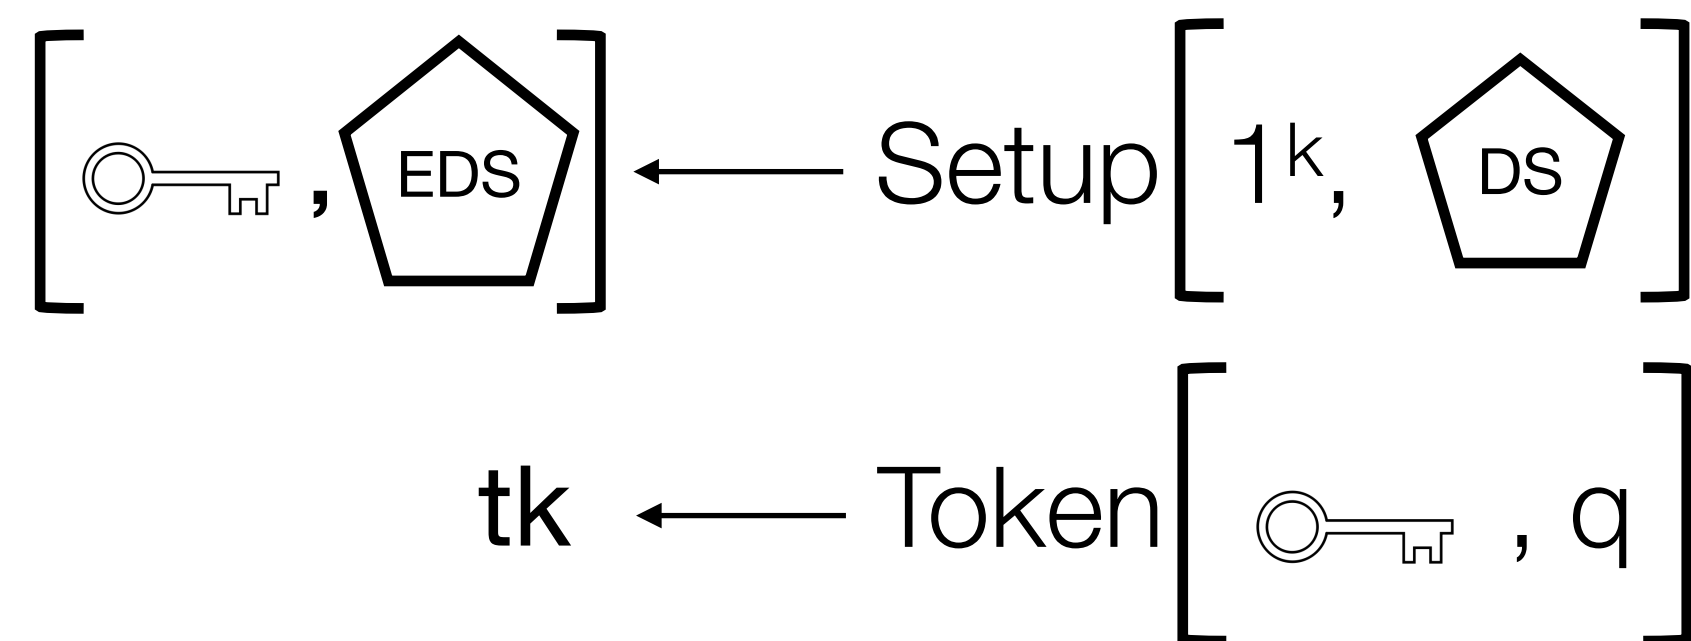
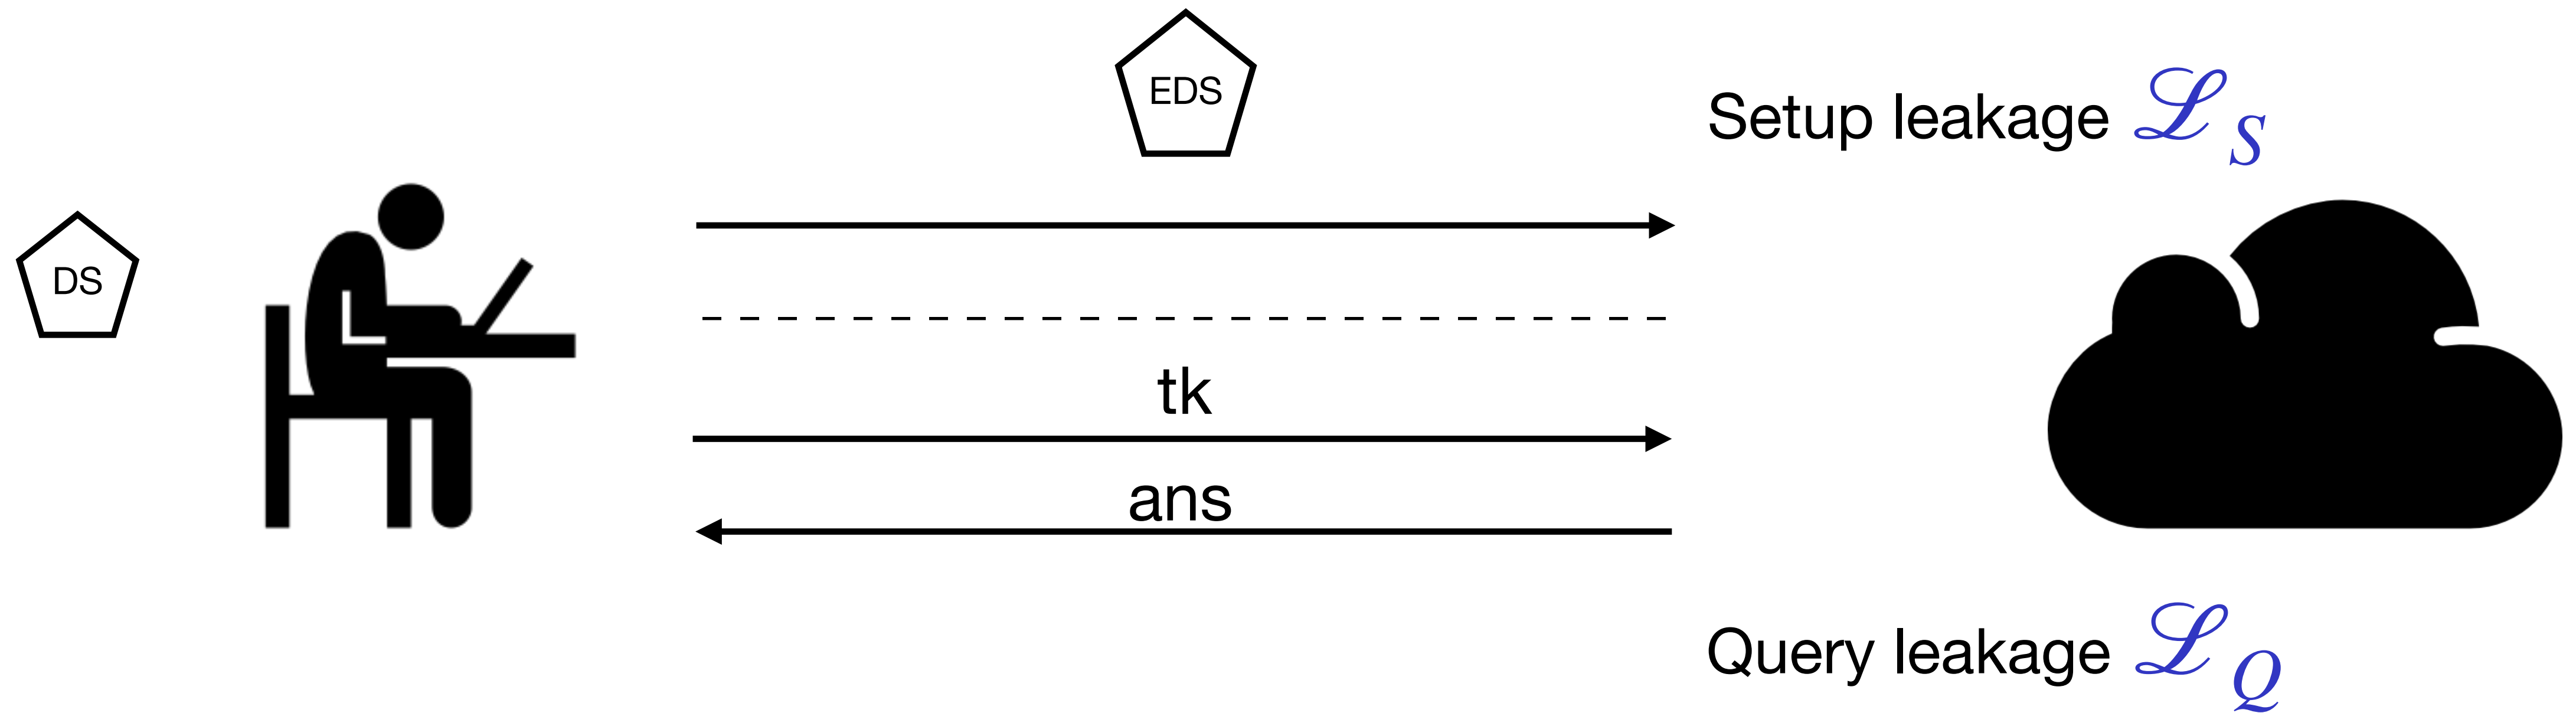
-
- Fully-Homomorphic Encryption (FHE)
 - Property-Preserving Encryption (PPE)
 - Structured Encryption (STE)
 - Oblivious RAM (ORAM)
 - Functional Encryption (FE)

A hand holds a magnifying glass over a landscape. The magnifying glass shows a detailed view of a mountain range with a lake in the foreground. The text 'STE' is overlaid on the magnified scene.

STE

Structured Encryption

Definitions



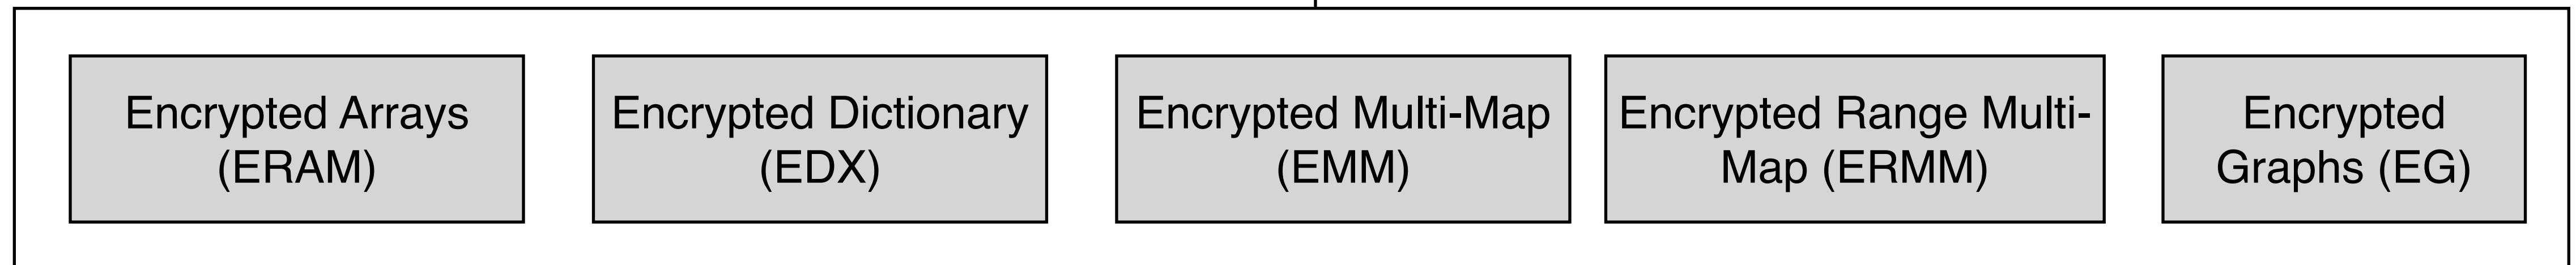

Structured Encryption

Definitions

- An STE scheme is $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_A, \mathcal{L}_D)$ -secure if
 - It reveals no information about the structure beyond \mathcal{L}_S
 - It reveals no information about the structure and queries beyond \mathcal{L}_Q
 - It reveals no information about the structure and adds beyond \mathcal{L}_A
 - It reveals no information about the structure and deletes beyond \mathcal{L}_D

Structured Encryption

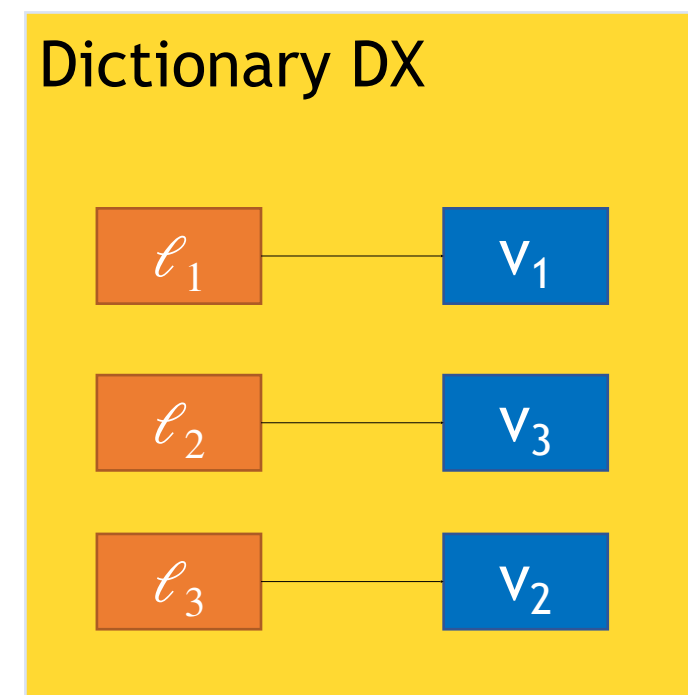
Data Structures



Background

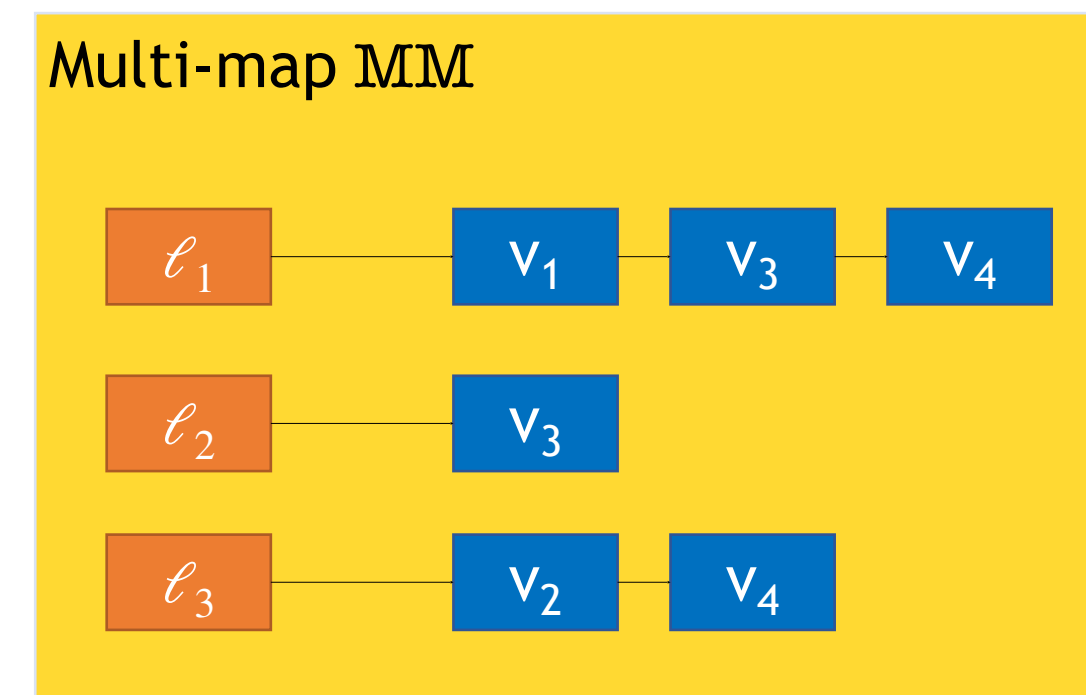
Dictionary and Multi-Map data structures

- DXs map labels to values



- Get: $DX[l_3]$ returns v_2

- MMs map labels to tuples




- Get: $MM[l_3]$ returns (v_2, v_4)

Dynamic EMMs

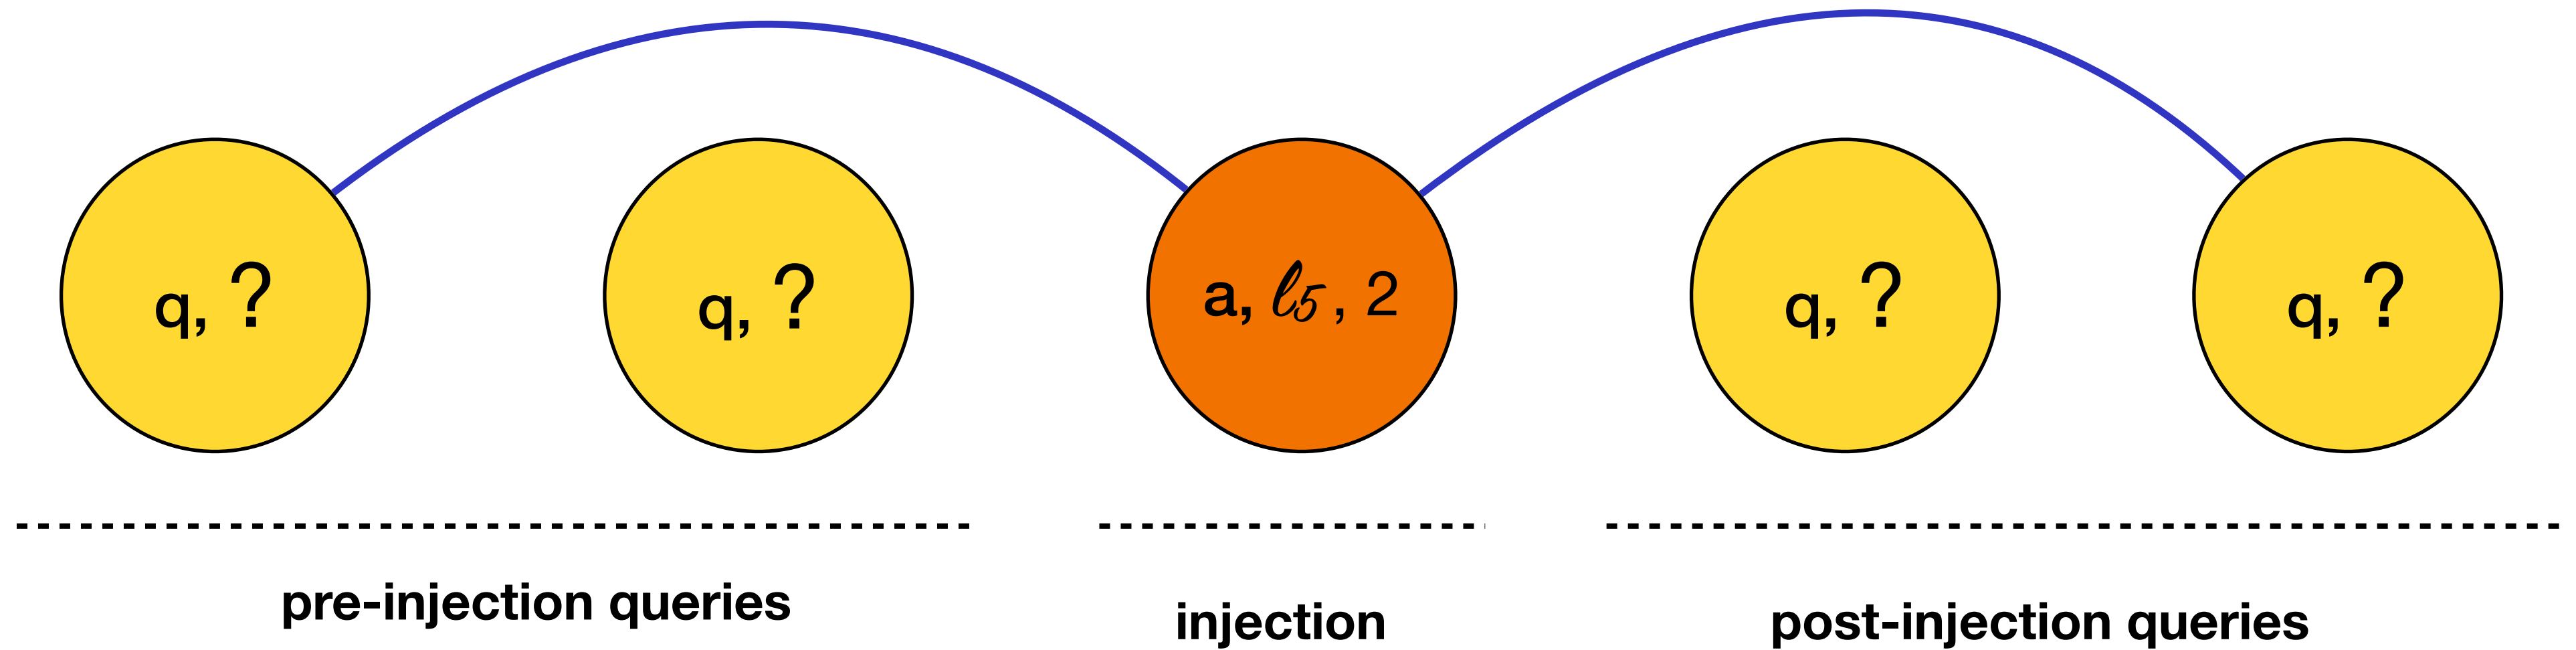
- Dynamic sub-linear EMMs [[KPR12](#), [CKKKRS14](#), ...]
 - Optimal-time queries
 - \mathcal{L}_Q : query operations leak query equality (qeq)
 - If and when two queries are for the same label
 - \mathcal{L}_A : add operations leak add-query equality
 - If and when an add is for the same label as a query

Leakage Graphs



- \mathcal{L}_A : add operations leak add-query equality
- \mathcal{L}_Q : query operations leak query equality

Injection Attacks [ZKP16]




- \mathcal{L}_A : add operations leak add-query equality

Forward-Privacy [SPS14]

Definitions

- Introduced in [SPS14] but not formalized
 - “adds & deletes not correlated with previous queries”
- Formalized in [Bost16] roughly as “adds & deletes leak add & delete size”
- Formalized in [KM18] as “adds & deletes leak nothing”
- [ZKP16] observed that
 - Forward-privacy protects pre-injection queries
 - But not post-injection queries

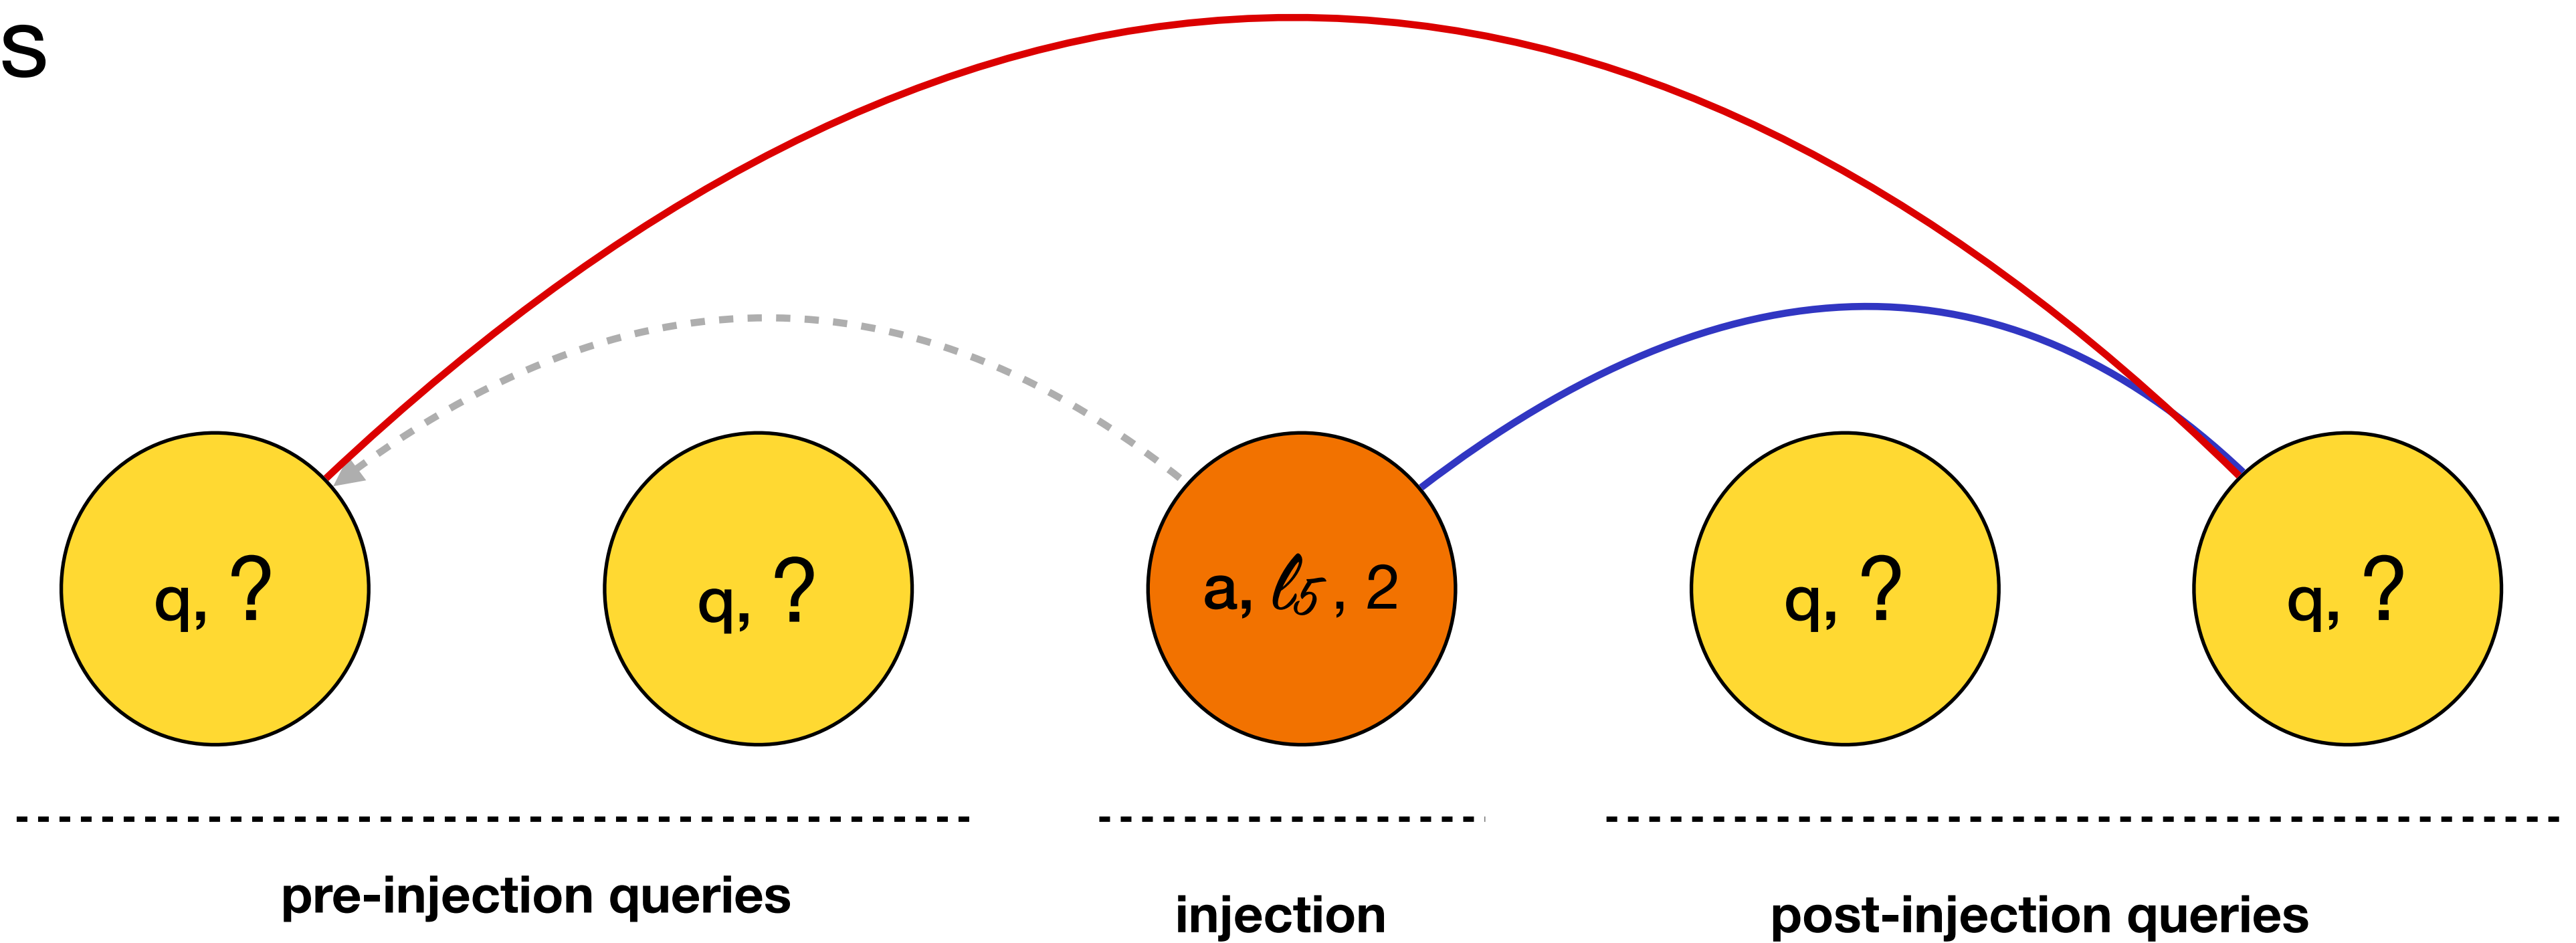
Injection Attacks [ZKP16]



- Forward-privacy
- \mathcal{L}_A : add operations do not leak add-to-query equality but still leak query-to-add equality

Forward Privacy

Limitations



- But what if scheme *also* leaks query equality?
 - \mathcal{L}_A : add operations leak query-to-add equality
 - \mathcal{L}_Q : query operations leak query equality (qeq)

What's Going on?

- Injection attacks are only a subset of a larger class of attacks
- Forward-privacy is only one security notion among a class of related notions

Correlation Attacks

- Learn labels of a subset of *operations*
 - using injections, inference attacks or known-data attacks etc...
- Use leakage to correlate known *operations* with unknown *operations*
- Operations here can be any of queries, adds, deletes, ...

- Special case of injection attacks
 - execute *add* operations with known label
 - use *add-query* leakage to correlate adds to unknown *queries*

Correlation Security

- A leakage profile $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_A, \mathcal{L}_D)$ is correlation-secure
 - “if it doesn’t reveal correlations between certain types of operations”
 - “if the leakage graph has no paths between certain types of operations”
 - Formalized using a (complex) parametrizable game-based definition

Injection Security

- Special case of correlation security
- A leakage profile $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_A, \mathcal{L}_D)$ is injection-secure
 - “if there are no correlations between adds and pre- & post-injection queries”
 - “if no paths exist between adds and pre- and post-injection queries”
 - Formalized by parameterizing the definition of correlation security

Other Approaches

- ORAM-based
 - Store multi-map in an ORAM
 - but in a volume-hiding manner (see paper for more)
- SWiSSSE [[GPPW20](#)]
 - argument of security against injection attacks

Contributions

Theorem 1

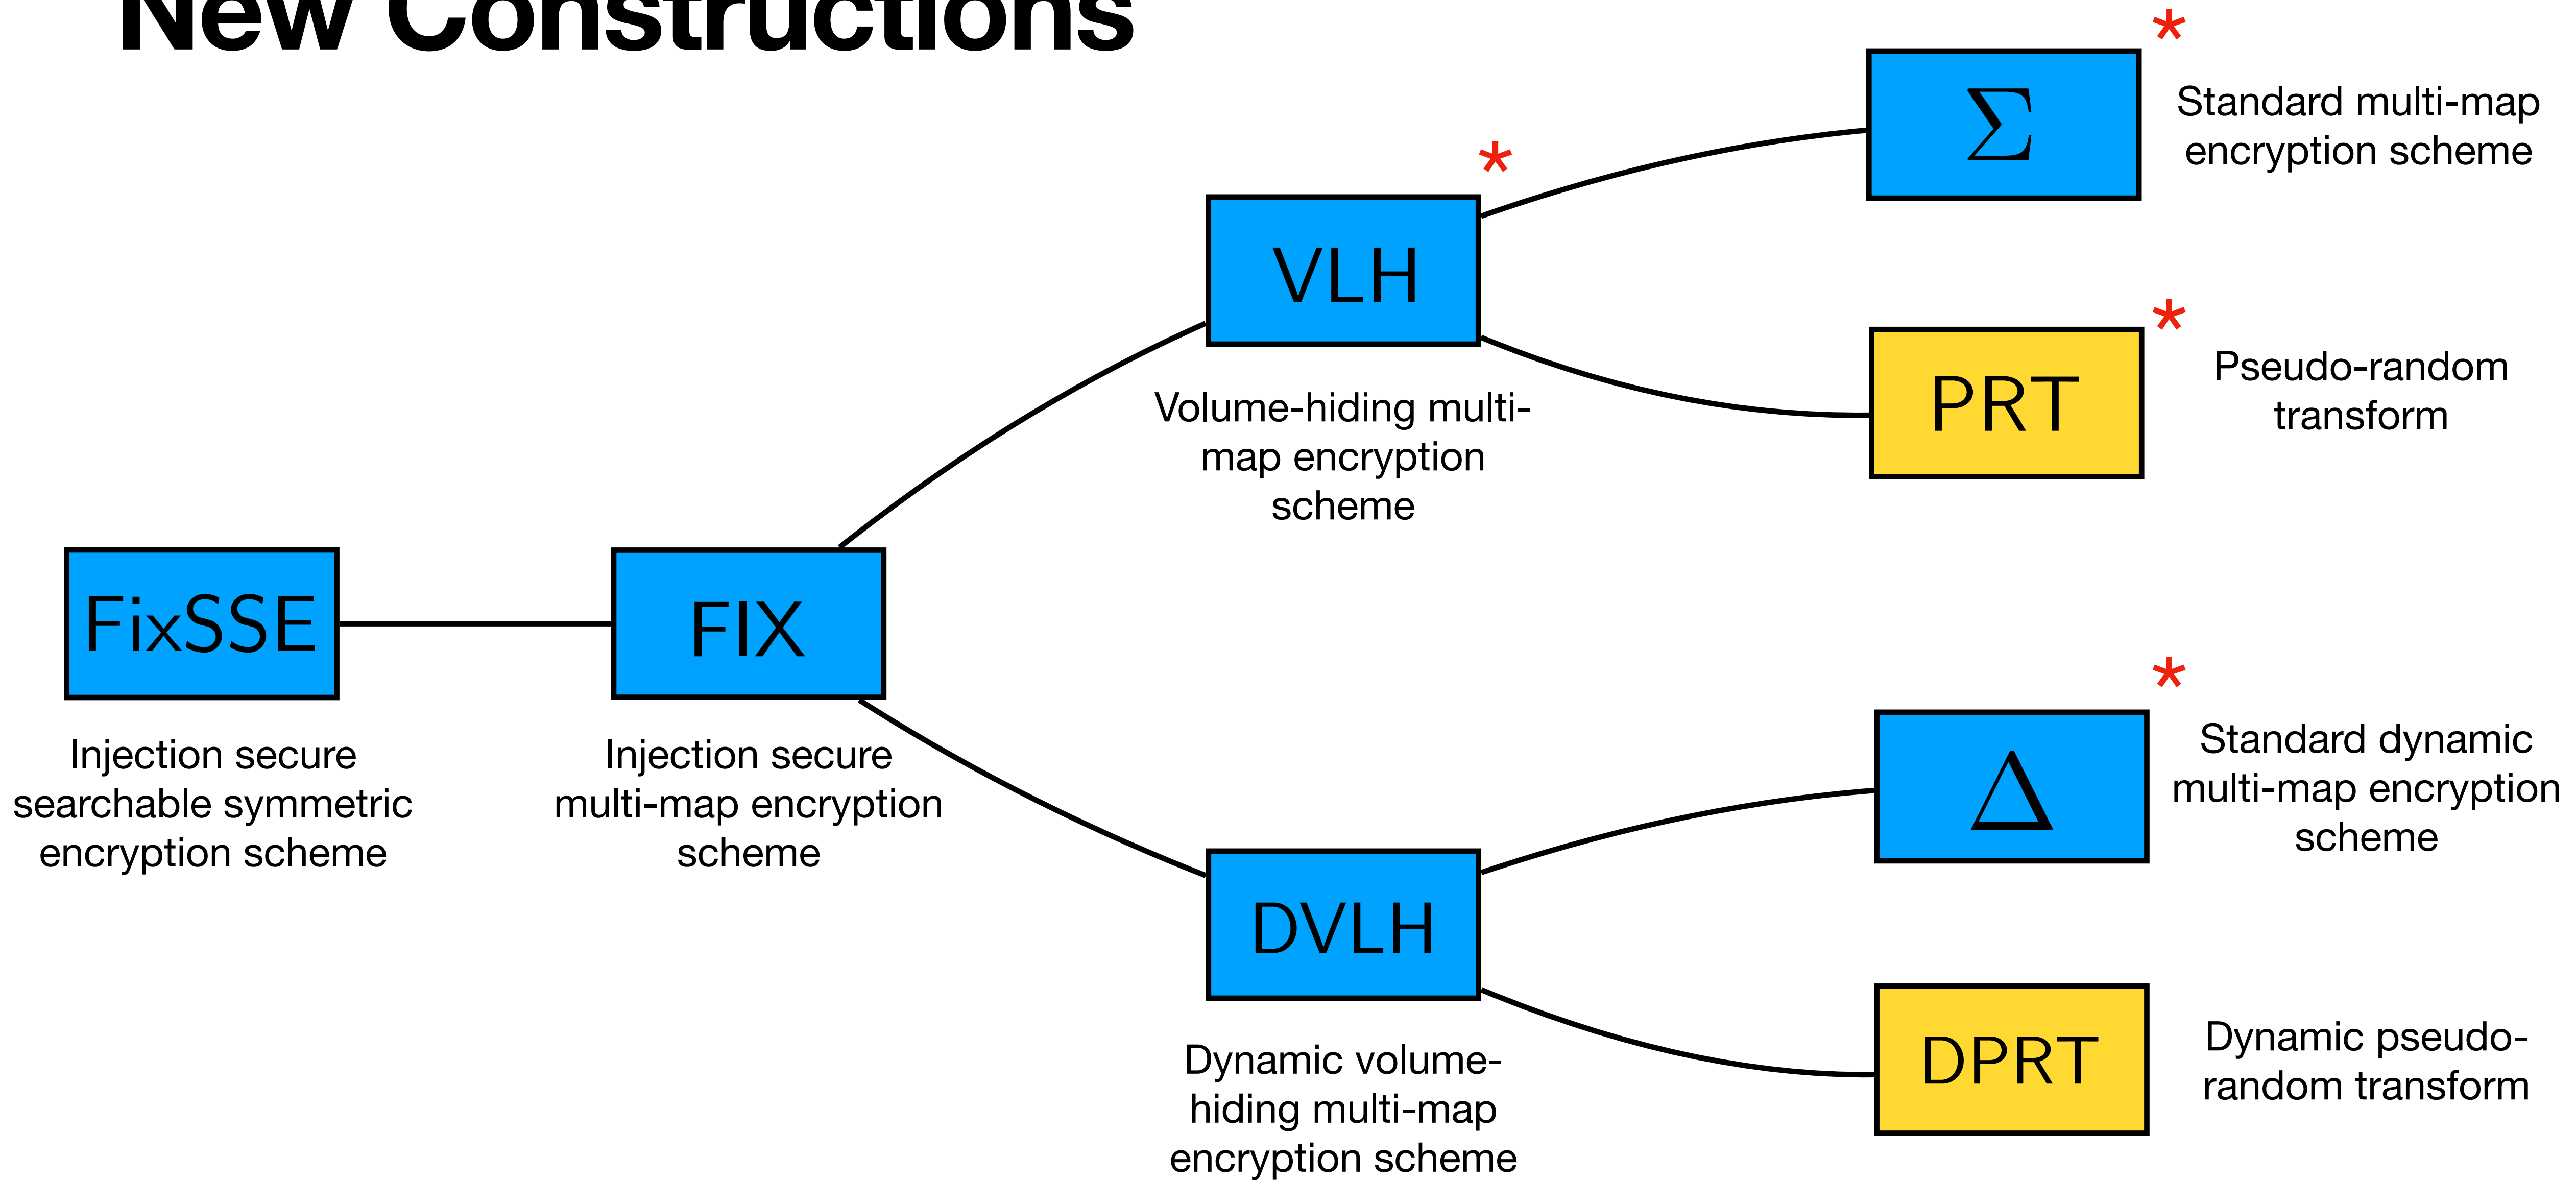
Theorem 1. If a leakage profile is such that $\mathcal{L}_Q = \text{qeq}$ and $\mathcal{L}_A = \mathcal{L}_D = \perp$ then it is injection-secure.

New Constructions

FIX, FixSSE and DVLH

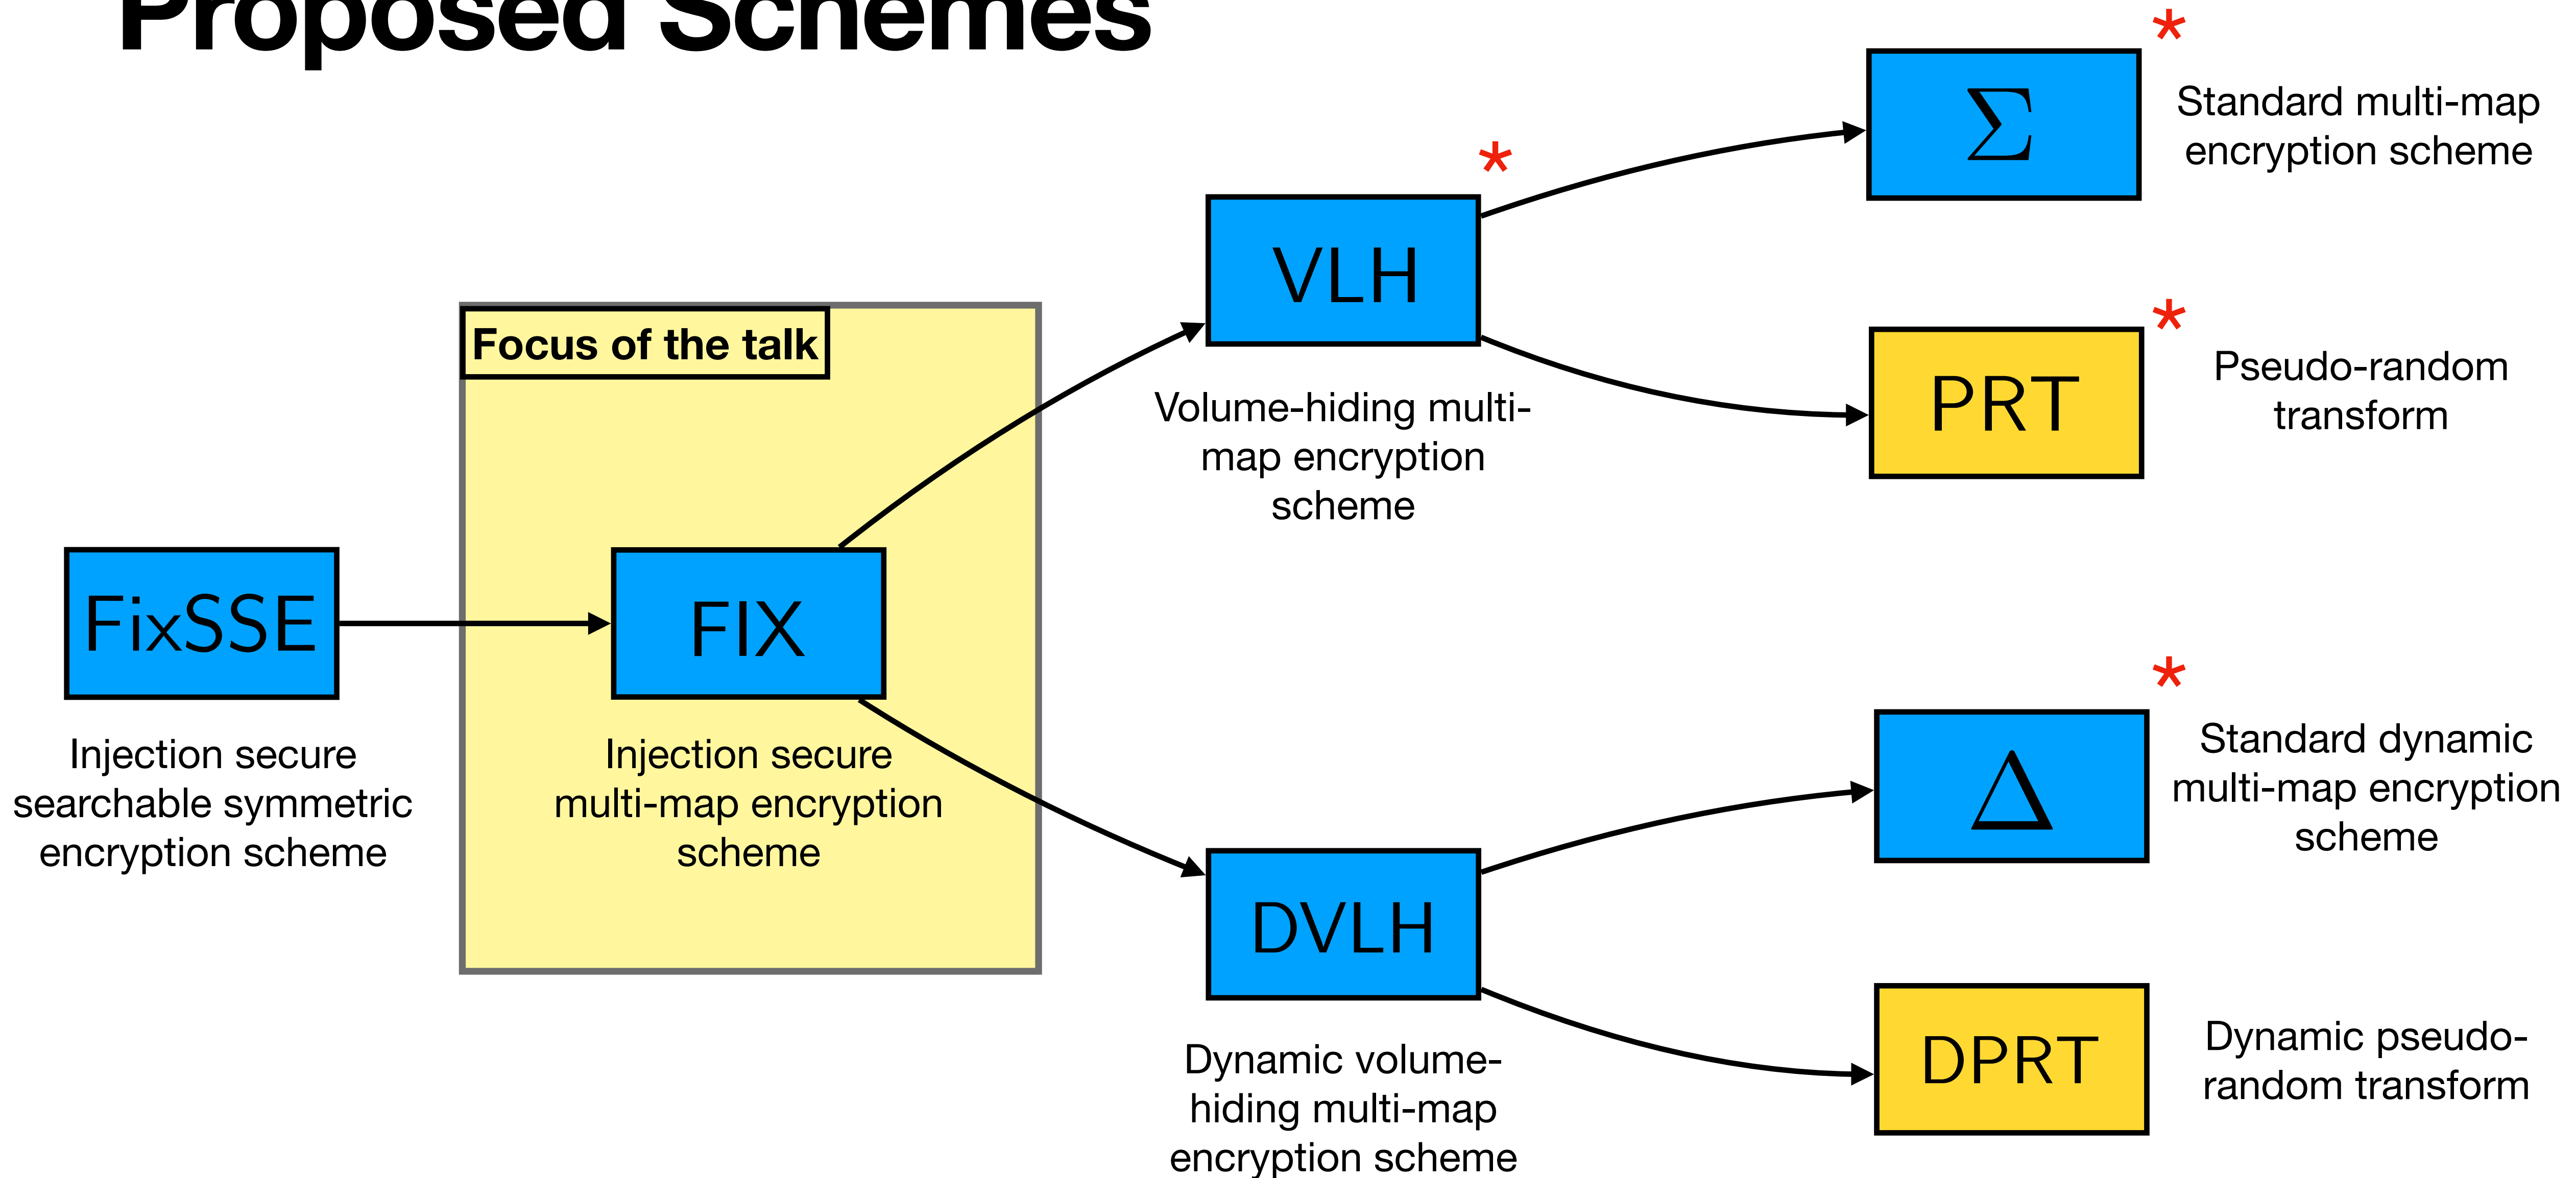
- FIX is the first (dynamic) injection-secure multi-map encryption scheme
- FixSSE is the first (dynamic) injection-secure searchable symmetric encryption scheme
- DVLH is a dynamic volume-hiding multi-map encryption scheme
 - DVLH is based on the dynamic pseudo-random transform (DPRT)
- All schemes above can achieve sub-linear search overhead under some assumptions

New Constructions



* existing schemes from the literature

Proposed Schemes



* existing schemes from the literature

Design Techniques

Randomized but Fixed Schedule

- Adds, deletes etc are stored in local stash...
 - ...and pushed to server-side encrypted structures following a **randomly** sampled but **fixed** schedule
 - Scheduled defined by a random permutation π
- Every label ℓ has a scheduled push time/epoch defined by $\pi(\ell)$
 - Add, erase or edit tokens for ℓ are stored in stash until ℓ 's scheduled push epoch
 - Once we reach the last epoch, we go back to the first one
- The client-side state can **grow** as a function of the distribution of adds, deletes etc.

FIX

Building Blocks

- A static volume-hiding multi-map encryption scheme
 - Can be instantiated using VLH [[KM19](#)], AVLH [[KM19](#)] or dprfMM [[PPYY19](#)]
 - In this talk we use VLH as the underlying instantiation.
 - Hides volume (response length) and reveals query equality
- A dynamic volume-hiding multi-map encryption scheme
 - Can be instantiated using the DVLH ([this work](#)) or 2ch [[APPYY23](#)]
 - In this talk, we use DVLH as the underlying instantiation.
 - Hides volume and reveals the query equality

FIX

Setup

Setup $\left[1^k, \text{MM} \right]$:

- 1 Sample a random permutation π
- 2 Create two empty multi-map MM_n and MM_{st} and a counter cnt
- 3 Compute
$$\left[K_o, \text{st}_o, \text{EMM}_o \right] \leftarrow \text{VLH.Setup} \left[1^k, \text{MM} \right]$$
$$\left[K_n, \text{st}_n, \text{EMM}_n \right] \leftarrow \text{DVLH.Setup} \left[1^k, \text{MM}_n \right]$$
- 4 The client receives $K = (K_o, K_n)$ and $\text{st} = \left[\text{st}_o, \text{st}_n, \text{MM}_{st}, \text{cnt}, \pi \right]$
The server receives $\text{EMM} = \left[\text{EMM}_o, \text{EMM}_n \right]$

FIX

Append

Append $\left[K, st, \ell, \bar{v}; EMM \right]:$

1 For all v in \bar{v} , add $v || \text{add}$ to $MM_{st}[\ell]$

2 Compute

$\left[st'; EMM' \right] \leftarrow \text{Push} \left[K, st, \mu; EMM \right]$

3 Client outputs the updated state st'

Server outputs the updated multi-map EMM'

FIX

Erase

Erase $\left[K, st, \ell, \bar{v}; EMM \right]:$

- 1 For all v in \bar{v} ,
 - if $v || \text{add} \in MM_{st}[\ell]$,
 - Remove $v || \text{add}$ from $MM_{st}[\ell]$
 - Else
 - Add $v || \text{del}$ to $MM_{st}[\ell]$
- 2 Compute $\left[st'; EMM' \right] \leftarrow \text{Push} \left[K, st, \mu; EMM \right]$
- 3 Client outputs the updated state st'
Server outputs the updated multi-map EMM'

FIX

Get

- 1 Initialize an empty set Result
- 2 The client and server execute

$$\begin{aligned} \left[R_o; \perp \right] &\leftarrow \text{VLH.Get} \left[K_o, \text{st}_o, \ell; \text{EMM}_o \right] \\ \left[R_n; \perp \right] &\leftarrow \text{DVLH.Get} \left[K_n, \text{st}_n, \ell; \text{EMM}_n \right] \end{aligned}$$

$$\text{Get} \left[K, \text{st}, \ell; \text{EMM} \right]:$$

- 3 The client computes the following sets

$$R_{\text{st}}^+ = \left[v : v \parallel \text{add} \in \text{MM}_{\text{st}}[\ell] \right] \quad R_{\text{st}}^- = \left[v : v \parallel \text{del} \in \text{MM}_{\text{st}}[\ell] \right] \quad R_0^- = \left[v : v \parallel \text{del} \parallel \text{old} \in \text{MM}_{\text{st}}[\ell] \right]$$

- 4 The client outputs

$$\text{Result} = \left[R_o \cup R_n \cup R_{\text{st}}^+ \right] \setminus \left[R_{\text{st}}^- \cup R_0^- \right]$$

FIX

Push (Part 1)

Push $\left[K, \text{st}, \mu ; \text{EMM} \right] :$

1 For all i in $\{1, \dots, \mu\}$,

- Compute $j = \text{cnt} + i \pmod{\#\mathbb{L}_{\text{MM}}}$
- Compute

$$\begin{aligned} \left[R_o ; \perp \right] &\leftarrow \text{VLH.Get} \left[K_o, \text{st}_o, \ell_{\pi(j)} ; \text{EMM}_o \right] \\ \left[R_n ; \perp \right] &\leftarrow \text{DVLH.Get} \left[K_n, \text{st}_n, \ell_{\pi(j)} ; \text{EMM}_n \right] \end{aligned}$$

- For all $\mathbf{v} \mid \mid \text{flag} \in \text{MM}_{\text{st}}[l]$,
- If $\text{flag} = \text{add}$, $R_{\text{st}}^+ := R_{\text{st}}^+ \cup \mathbf{v}$
- If $\text{flag} = \text{del}$,
 - If $\mathbf{v} \in R_o$, set $R_{\text{st}}^+ := R_{\text{st}}^+ \cup \mathbf{v} \mid \mid \text{del} \mid \mid \text{old}$
 - Else, set $R_{\text{st}}^- := R_{\text{st}}^- \cup \mathbf{v}$

FIX

Push (Part 2)

Push $\left[K, \text{st}, \mu ; \text{EMM} \right] :$

2 Remove the tuple of $\ell_{\pi(j)}$ from MM_{st}

3 Set $\bar{v}_n := \left[R_n \cup R_{\text{st}}^+ \right] \setminus R_{\text{st}}^-$

4 Compute

$$\begin{array}{l} \left[\perp ; \text{EMM}' \right] \leftarrow \text{DVLH.Delete} \left[K_o, \text{st}_o, \ell_{\pi(j)} ; \text{EMM}_o \right] \\ \left[\perp ; \text{EMM}' \right] \leftarrow \text{DVLH.Insert} \left[K_n, \text{st}_n, \ell_{\pi(j)}, \bar{v}_n ; \text{EMM}_n \right] \end{array}$$

4 Increment the counter cnt by μ

5 Client outputs the updated state st'

Server outputs the updated multi-map EMM'

FIX

Leakage

- Setup leakage
 - The size of the input multi-map
 - The size of the label space
- Get leakage
 - The query equality pattern
- Inserts and Erases leakage
 - None

FIX

Efficiency when $\mu = O(\log N)$

- Get complexity $O(\log N)$
- Insert/Eraser complexity $O(\log^3 N)$
- Storage complexity $O(\log N \cdot \#\mathbb{L}_{MM})$
- Interactive
 - 2 rounds
- Stateful
- Lossy

FIX

Client Stash Analysis

Theorem 2. If the update labels are sampled uniformly at random and if their update lengths are sampled uniformly at random from U , then the expected stash size of FIX is at most $H_m \cdot s_{\max}/2\mu$ where

$$U = \left\{ \frac{s_0}{1^s \cdot H_{m,s}}, \frac{s_0}{2^s \cdot H_{m,s}}, \dots, \frac{s_0}{m^s \cdot H_{m,s}} \right\}$$

Theorem 3. If the update labels sampled uniformly at random and if their tuple lengths are sampled from a Zipf $\mathcal{Z}_{m,1}$ distribution, then, with probability at least $1 - \epsilon$ the stash size is at most

$$\frac{H_m \cdot s_{\max}}{2\mu} + s_{\max} \cdot \sqrt{\frac{(m - 2\mu) \cdot \ln(1/\epsilon)}{2\mu}}$$

Thank you

<https://eprint.iacr.org/2023/533>