

SQL on Structurally-Encrypted Databases

Seny Kamara

Tarik Moataz



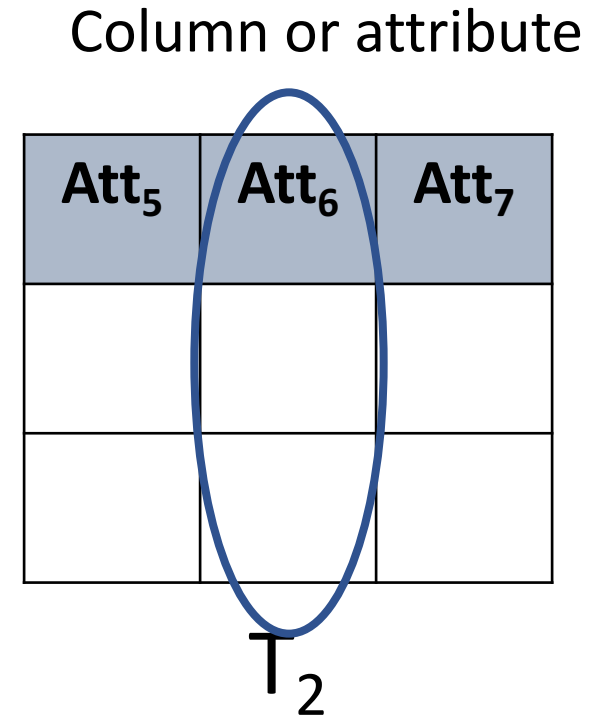
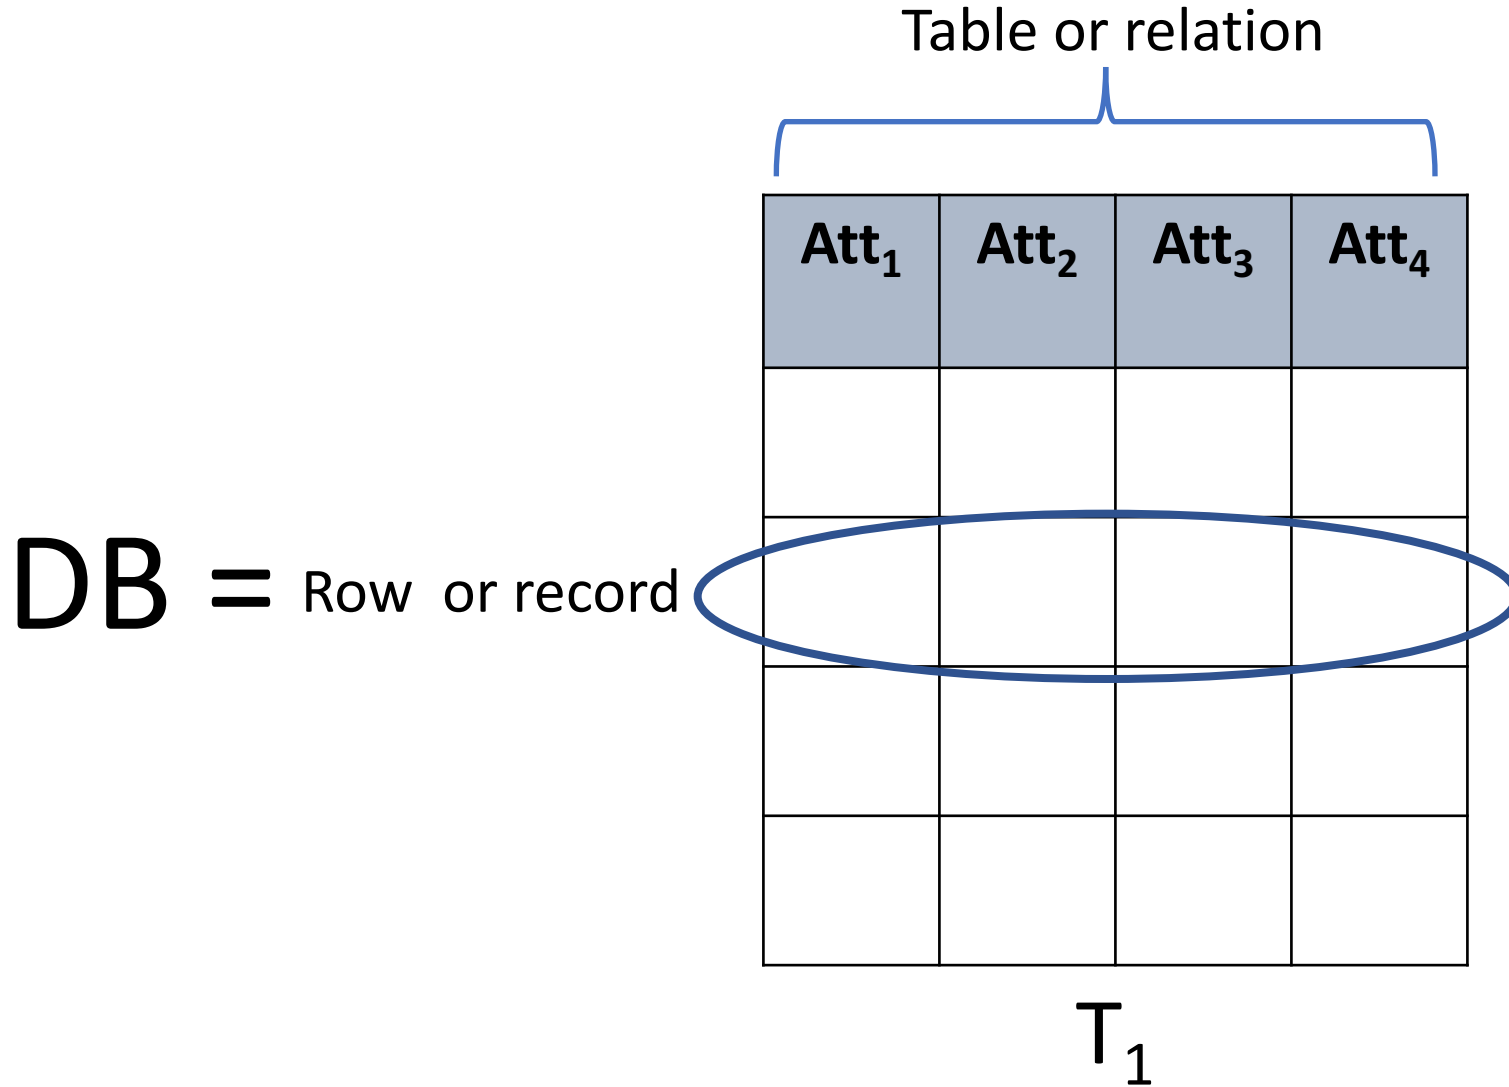
BROWN



ENCRYPTED
SYSTEMS LAB

Q: What is a relational database?

Relational DB



Structured Query Language

- SQL is a language for querying relational DBs

Select attributes From tables Where condition,

- Example:

Select (name, gender, height)

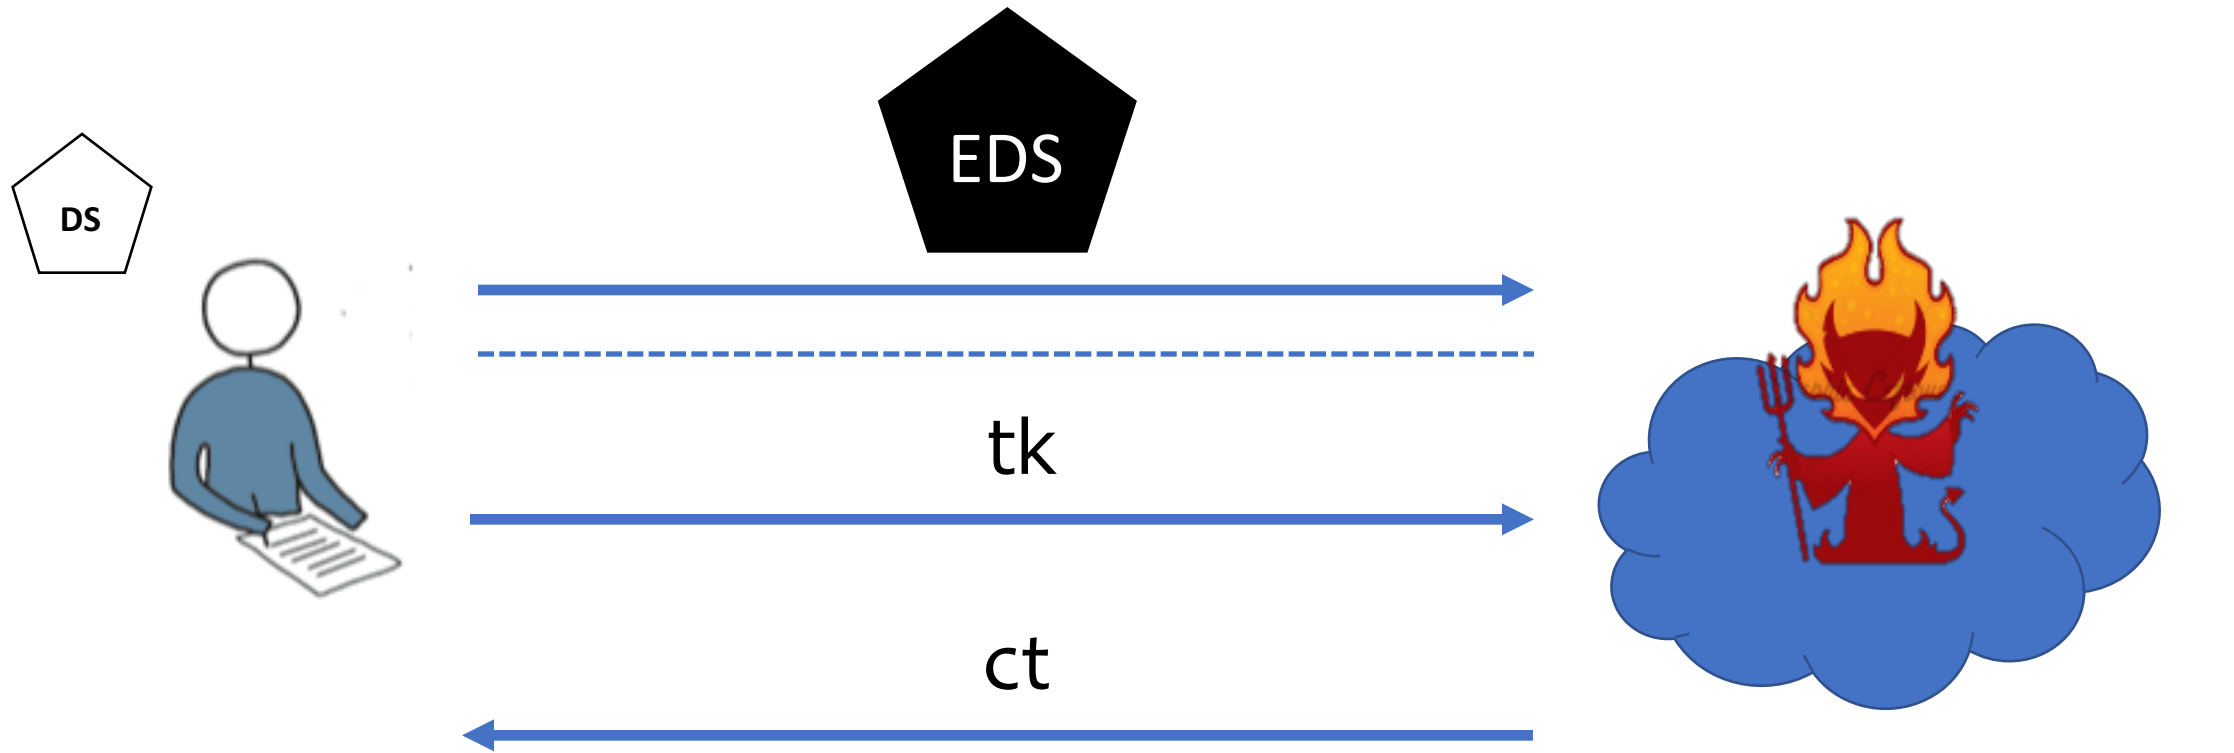
From (T₂, T₈)

Where (age = 36 AND zip = 10040 AND gender = F)

- SQL is the standard way to query a relational DB
 - Standard ANSI/ISO since 1986/1987

Q: What is Structured Encryption (STE)?

Structured Encryption (STE) [CK10]

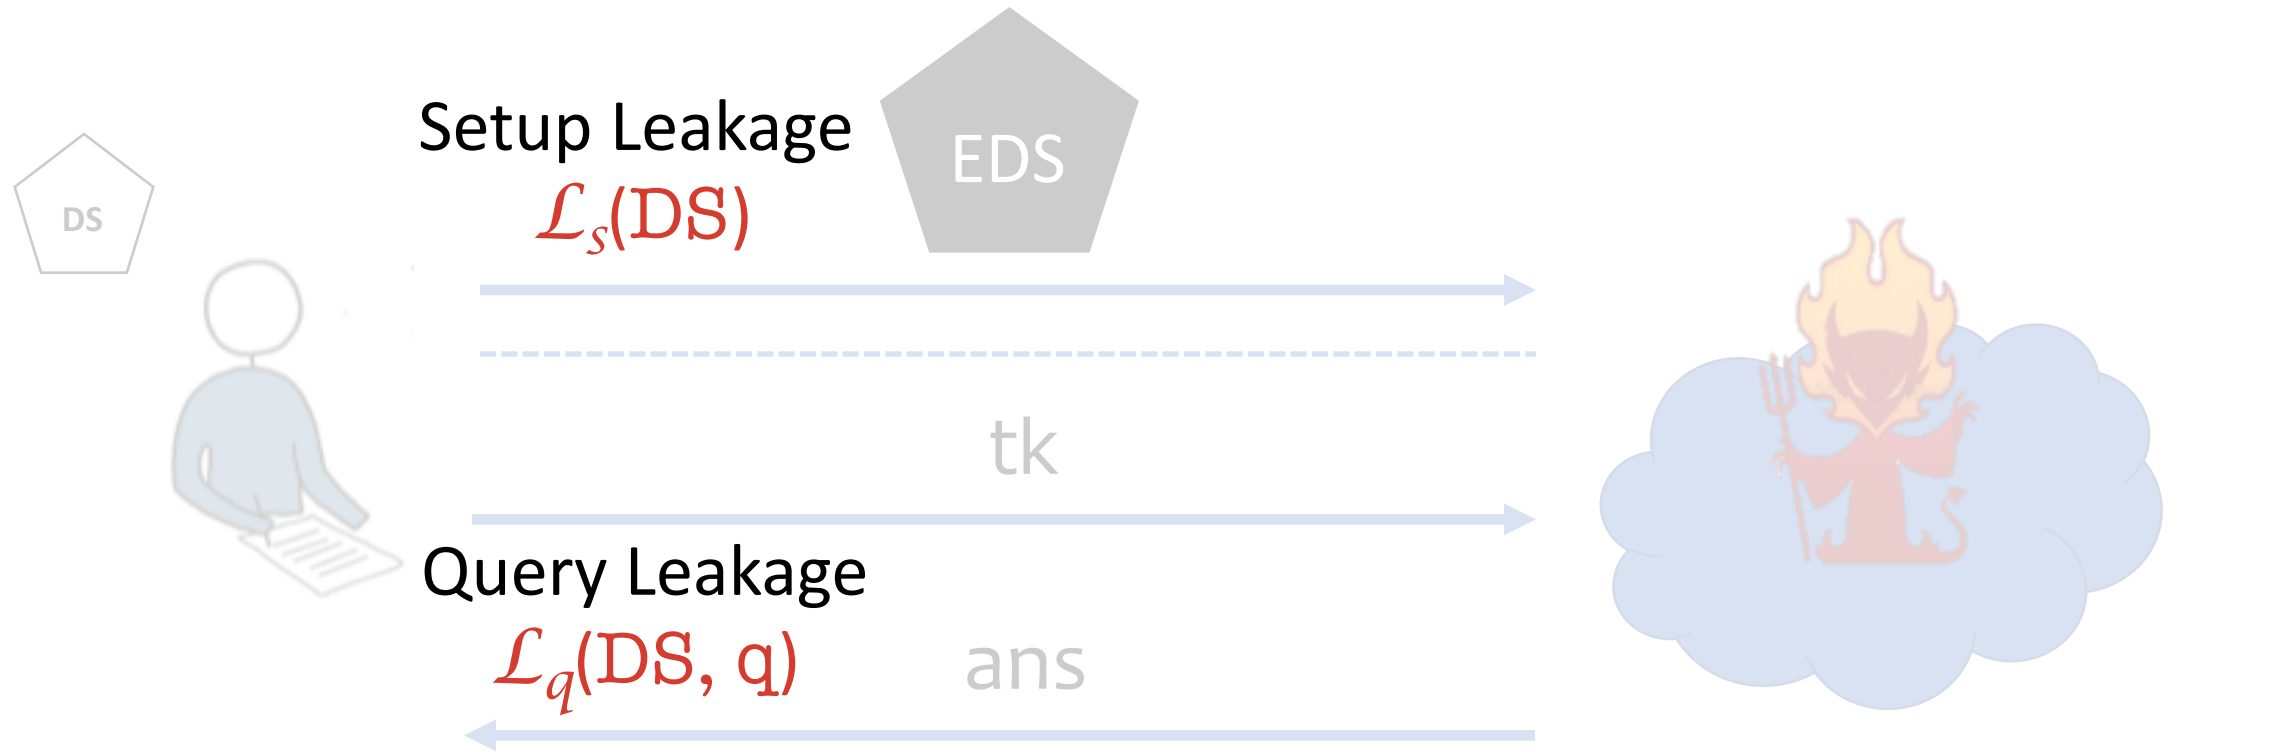


$$\text{Setup}(1^k, \text{DS}) \Rightarrow (\text{K}, \text{EDS})$$

$$\text{Token}(\text{K}, q) \Rightarrow tk$$

$$\text{Query}(\text{EDS}, tk) \Rightarrow ct$$

Structured Encryption (STE) [CK10]



$$\text{Setup}(1^k, \text{DS}) \Rightarrow (\text{K}, \text{EDS})$$

$$\text{Token}(\text{K}, q) \Rightarrow tk$$

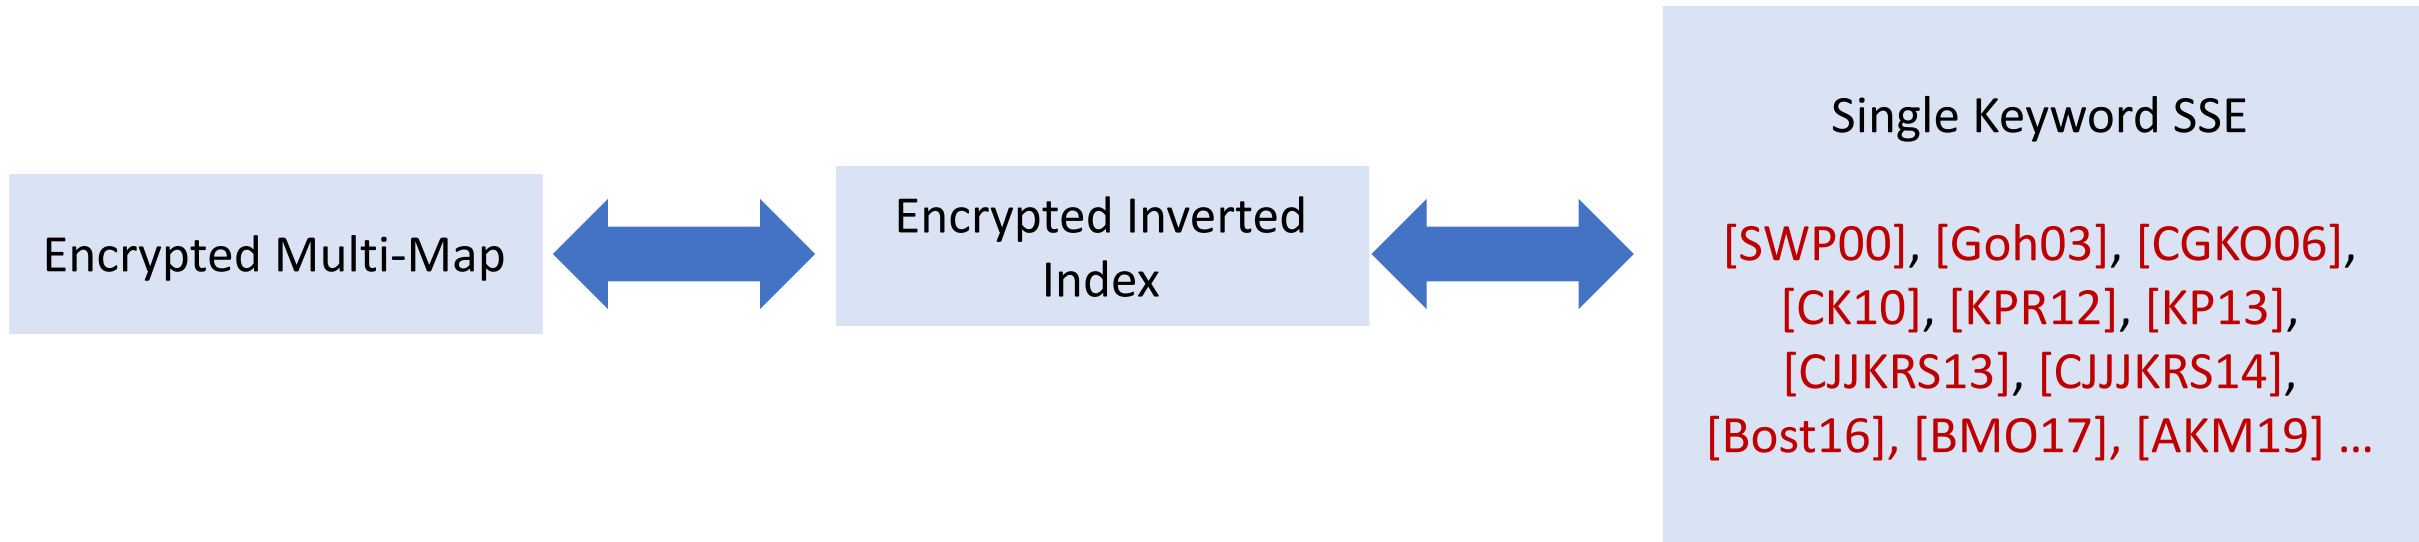
$$\text{Query}(\text{EDS}, tk) \Rightarrow ans$$

Structured Encryption (STE) [CK10]

We say that an STE is $(\mathcal{L}_S, \mathcal{L}_Q)$ -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

Encrypted Multi-Maps [CK10]



Q: How can we encrypt a relational DB?

Efficiency

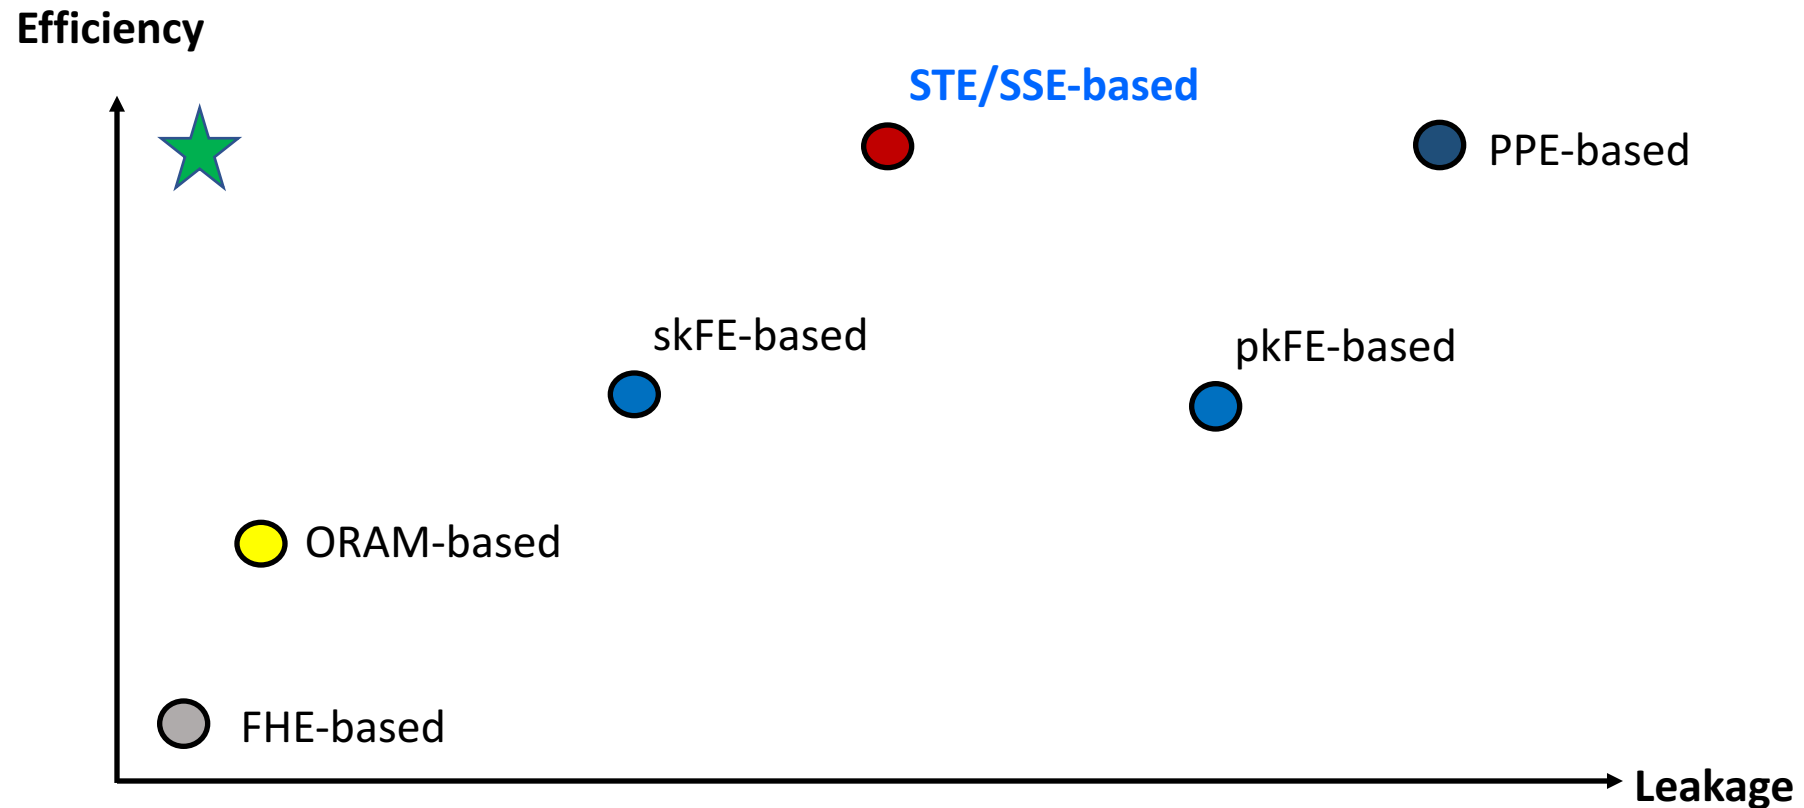


Functionality

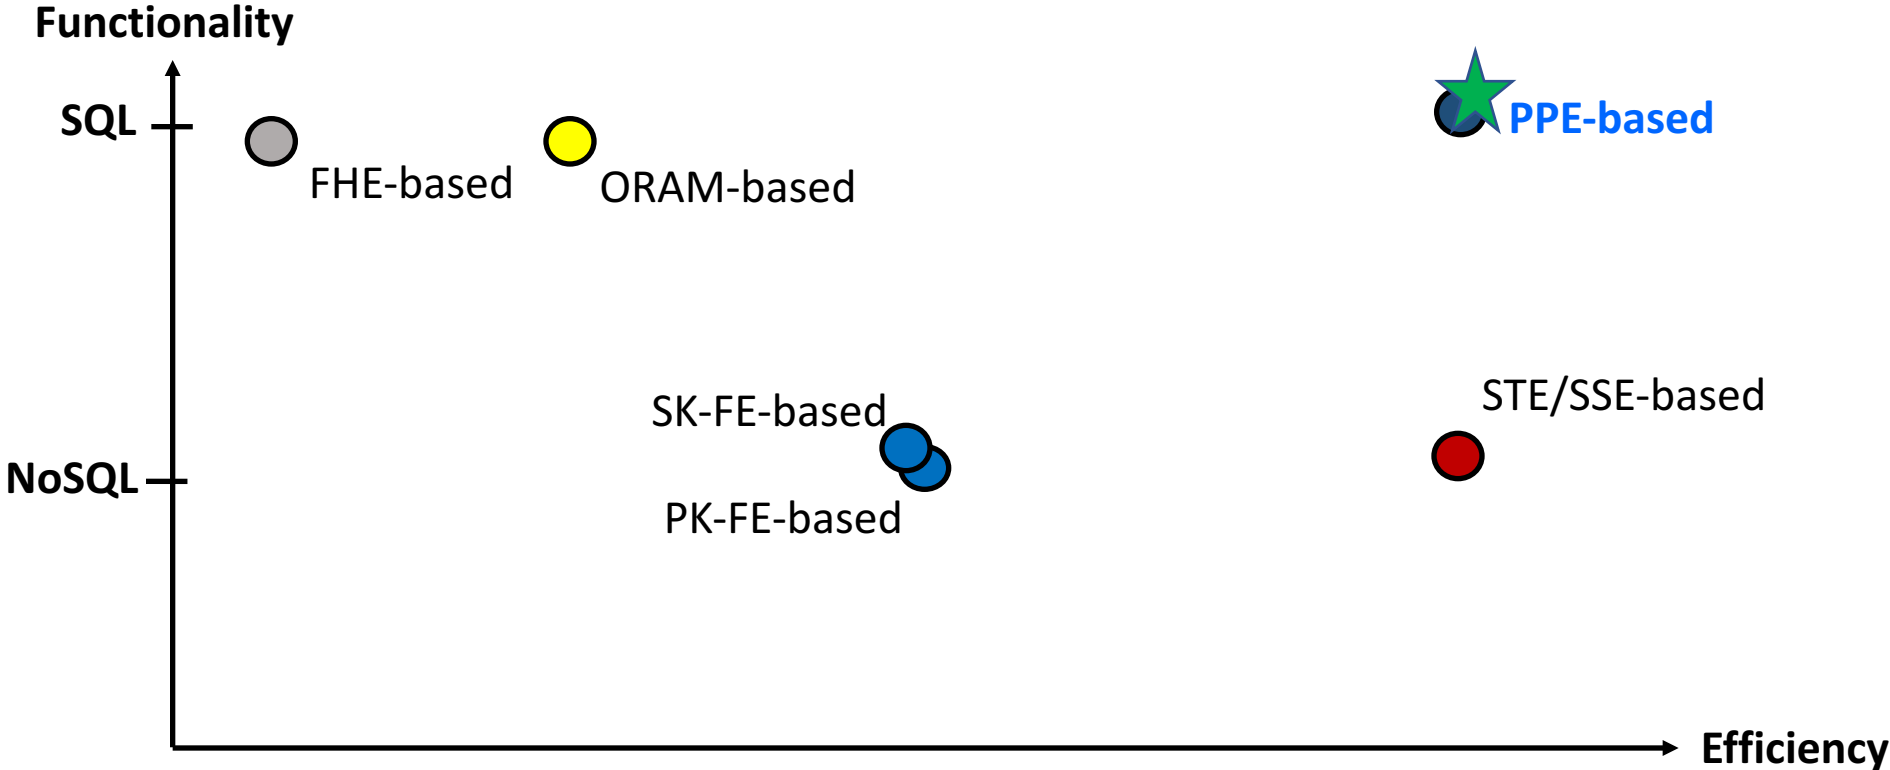
Leakage



Tradeoffs: Efficiency vs. Security



Tradeoffs: Functionality vs. Efficiency



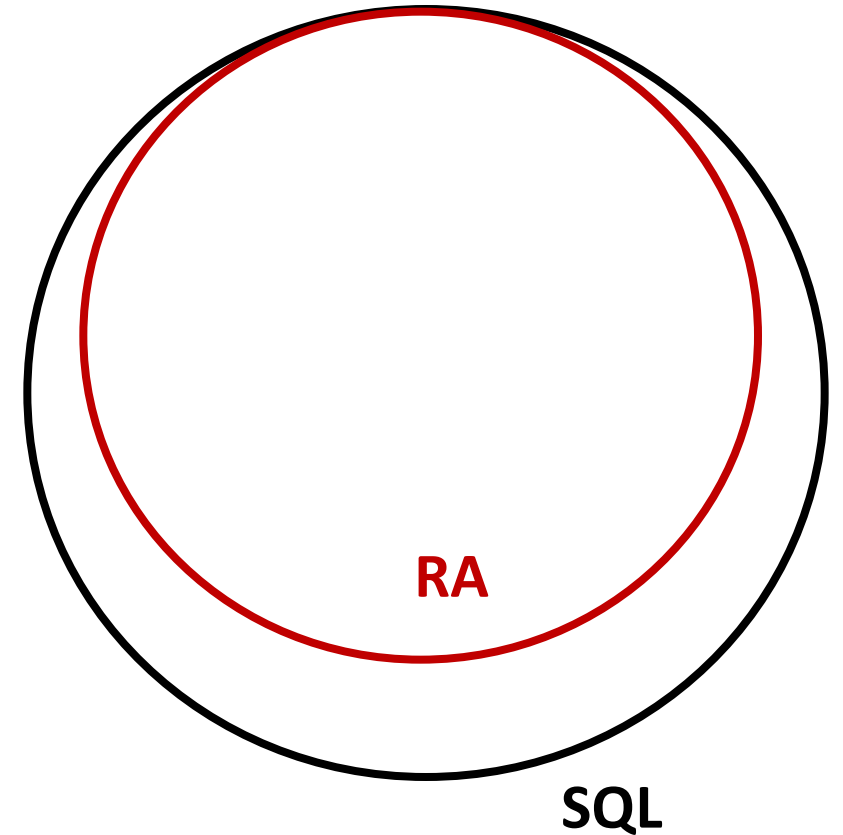
Q: Can we design an **STE**-based **Relational** EDB?

Challenges

- No PPE so no plug-and-play solutions
- SQL is a **declarative** language
 - Where do we even start?
- SQL is **complex**
 - Combination of many basic query types
 - Most STE schemes handle a single type queries
- SQL is “**constructive**”
 - STE has been optimized for “lookup-type” queries

Ch. #1: Declarative => Procedural

- Relational algebra [Codd70]
 - Set of **operations** on relations/tables
 - Union
 - Difference
 - Selection
 - Projection
 - Cross product
 - Join (many kinds)
 - ...



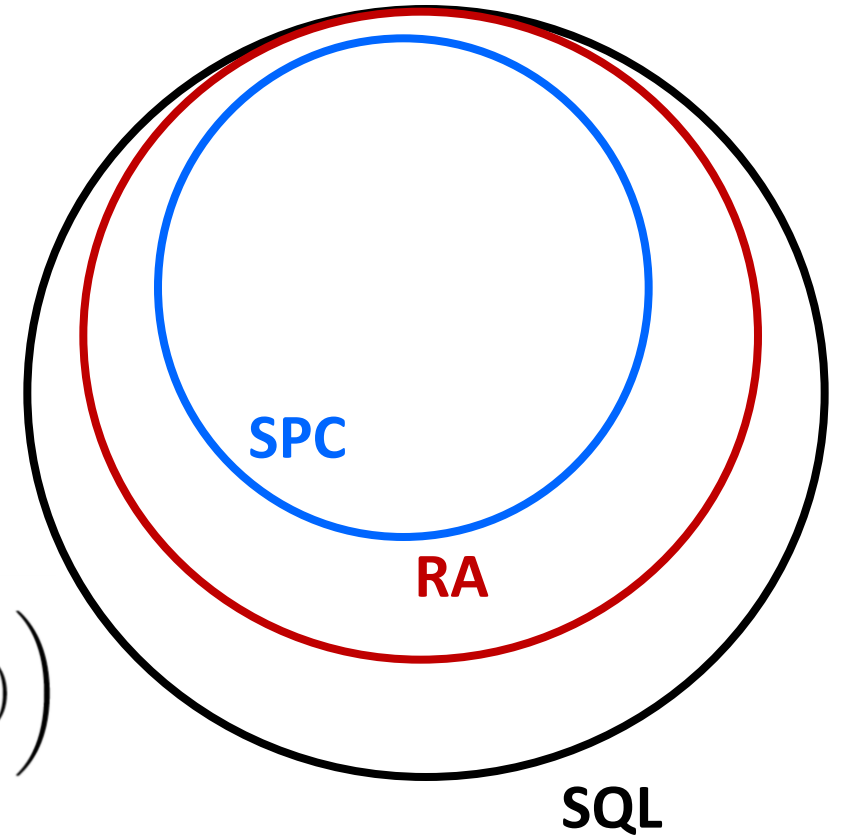
Ch. #2: Complex => Simple

- SPC algebra [Chandra-Merlin77]
 - Selection, Projection, Cross product
 - Equivalent to *Conjunctive* SQL queries

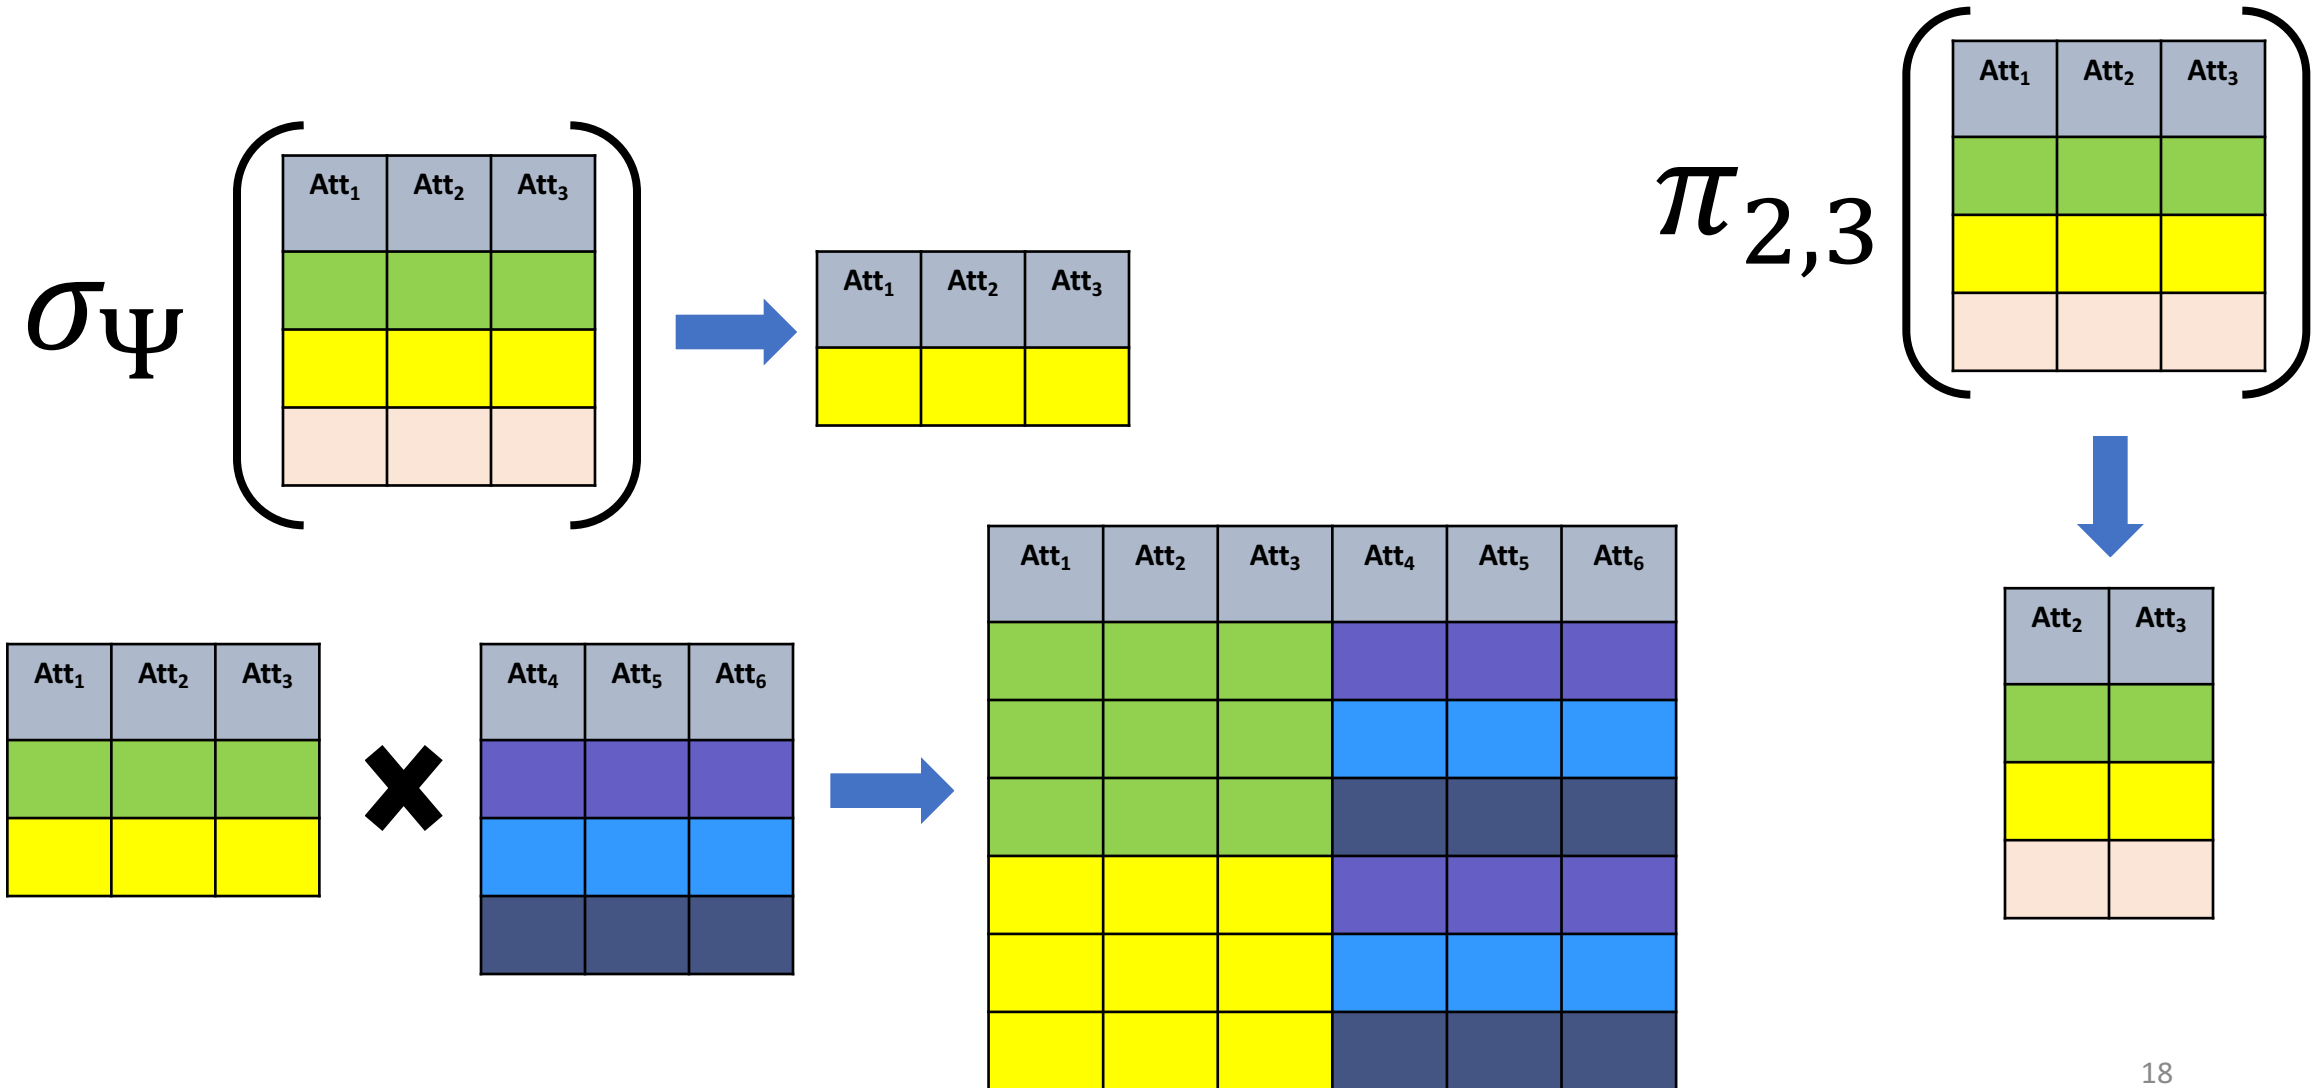
Select *attributes* From *tables* Where $(a_1 = X_1 \wedge \dots \wedge a_\ell = X_\ell)$,

- Any SPC query can be written in a Normal Form:

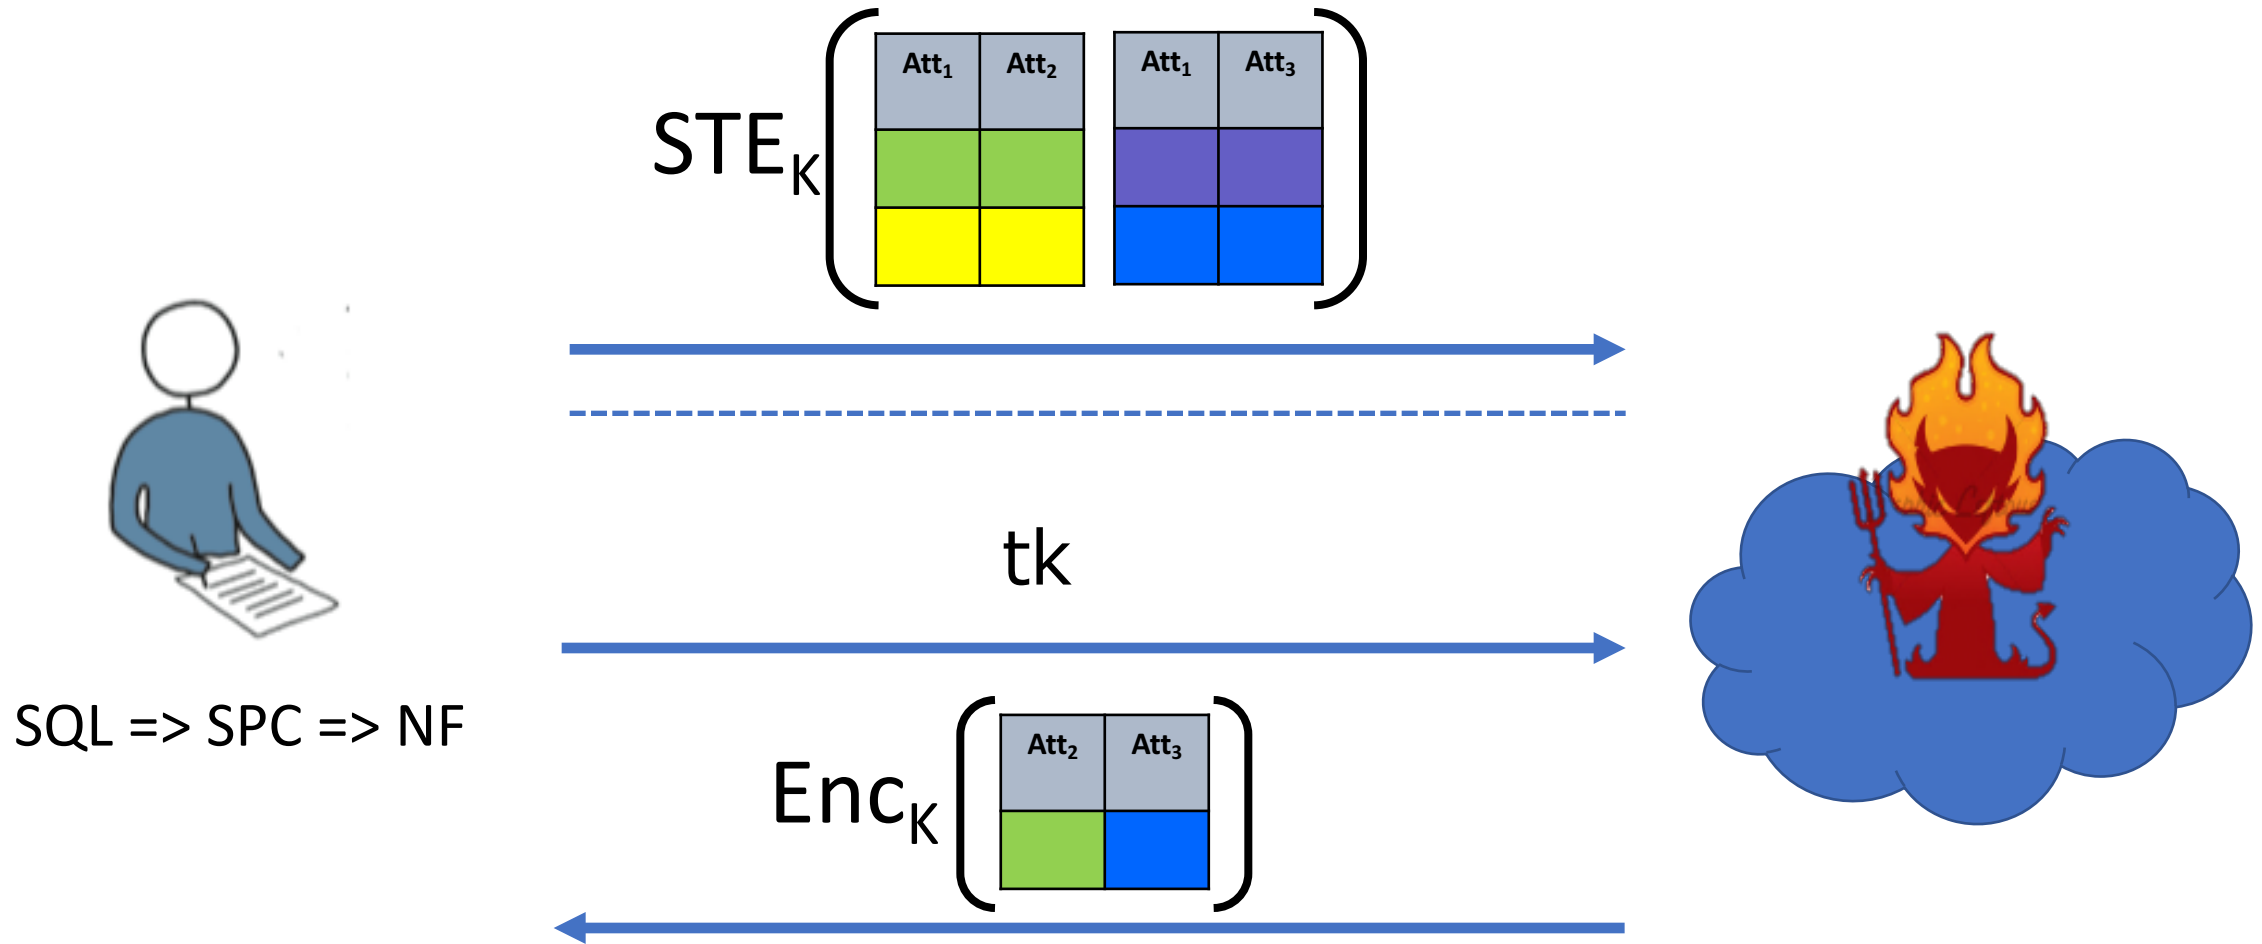
$$\pi_{a_1, \dots, a_h} \left([a_1] \times \dots \times [a_f] \times \sigma_{\Psi} (T_{i_1} \times \dots \times T_{i_t}) \right)$$




Select, Project, Cross Product



Our Goal




Our Results



- **SPX**: Encrypted Relational Database

- **First** STE scheme for relational DBs
- Handles **non-trivial subset** of SQL
- **Sub-linear** search and storage complexity (optimal under certain conditions)
- from any single-keyword SSE




- **SPX+**: dynamic SPX

- **Only** row **addition** and **deletion**
- from any dynamic single-keyword SSE
- **Sub-linear** search and storage complexity (optimal under certain conditions)



- **FP-SPX+**: forward-private dynamic SPX


- **poly-logarithmic** overhead for updates





A: Naïve STE-based Relational EDB

Naïve SPC Algorithm


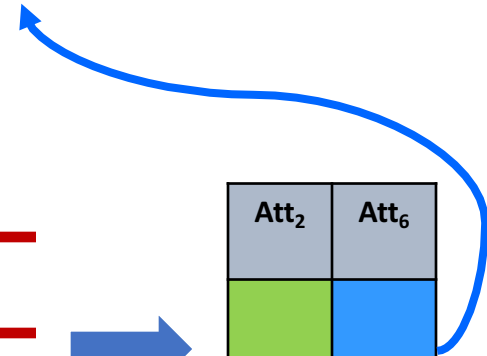
$$\pi_{a_1, \dots, a_h} \left([a_1] \times \dots \times [a_f] \times \sigma_{\Psi}(\mathbf{T}_{i_1} \times \dots \times \mathbf{T}_{i_t}) \right)$$




×



$$Q = (m^{h'} \cdot \prod_{i=1}^{\ell} \#R_i) \times h$$



$$m^t \times \left(\sum_{i=1}^t \|\mathbf{T}_i\| \right)$$



Sub-Linear SPC Algorithm

- Ideally linear in output size:

$$Q = (m^{h'} \cdot \prod_{i=1}^{\ell} \#R_i) \times h$$

Att ₂	Att ₆
Green	Blue
Green	Dark Blue
Yellow	Dark Blue

- Less than cross product size:

$$m^t \times \left(\sum_{i=1}^t \|\mathbf{T}_i\| \right)$$

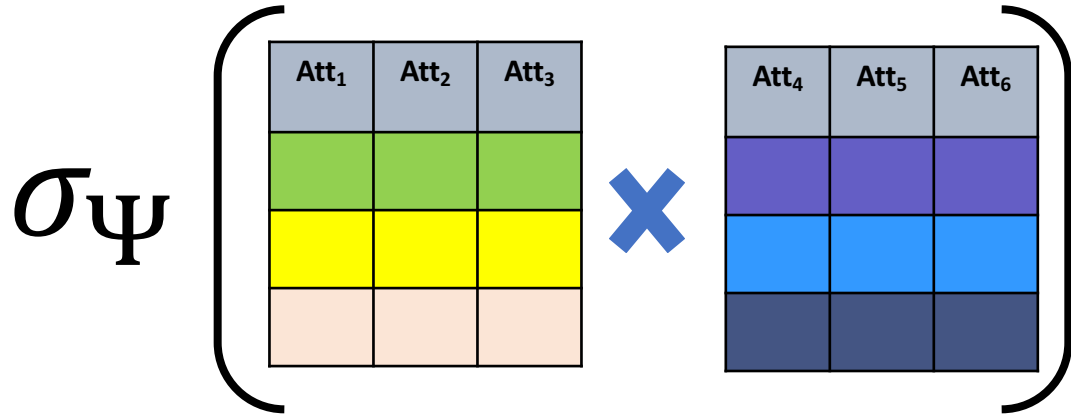
Att ₁	Att ₂	Att ₃	Att ₄	Att ₅	Att ₆
Green	Green	Green	Purple	Purple	Purple
Green	Green	Green	Blue	Blue	Blue
Green	Green	Green	Dark Blue	Dark Blue	Dark Blue
Yellow	Yellow	Yellow	Purple	Purple	Purple
Yellow	Yellow	Yellow	Blue	Blue	Blue
Yellow	Yellow	Yellow	Dark Blue	Dark Blue	Dark Blue
Orange	Orange	Orange	Purple	Purple	Purple
Orange	Orange	Orange	Blue	Blue	Blue
Orange	Orange	Orange	Dark Blue	Dark Blue	Dark Blue

Q: Can we achieve sub-linear **STE**-based EDB?

SPX Overview

- **Step 1.** Heuristic normal form (HNF) instead of the standard normal form
 - Avoid naïve Cartesian product by a “push select through product” method
- **Step 2.** New (plaintext) data structure that supports HNF
 - Different representations of the database to handle different SPC operators
- **Step 3.** Encrypted structure that supports HNF queries
 - Chaining technique with a better control of leakage
 - From any single-keyword SSE schemes

Step 1: Heuristic Normal Form (1)




More complicated

- Correlated/non-correlated
- Different types of select

Push Select through Product

=

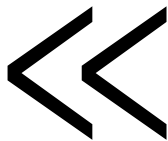
$$\Psi = \Psi_1 \wedge \Psi_2$$



Step 1: Heuristic Normal Form (2)

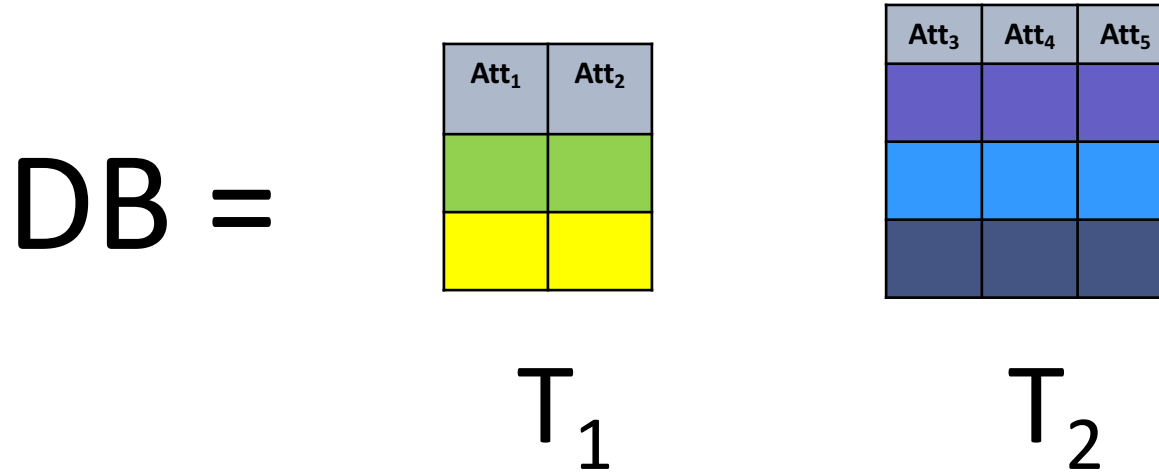
$$\sigma_{\Psi_1} \begin{pmatrix} \text{Att}_1 & \text{Att}_2 & \text{Att}_3 \\ \text{Green} & \text{Green} & \text{Green} \\ \text{Yellow} & \text{Yellow} & \text{Yellow} \\ \text{Orange} & \text{Orange} & \text{Orange} \end{pmatrix} \times \sigma_{\Psi_2} \begin{pmatrix} \text{Att}_4 & \text{Att}_5 & \text{Att}_6 \\ \text{Purple} & \text{Purple} & \text{Purple} \\ \text{Blue} & \text{Blue} & \text{Blue} \\ \text{Dark Blue} & \text{Dark Blue} & \text{Dark Blue} \end{pmatrix} =$$

$$\begin{matrix} \text{Att}_1 & \text{Att}_2 & \text{Att}_3 \\ \text{Green} & \text{Green} & \text{Green} \\ \text{Orange} & \text{Orange} & \text{Orange} \end{matrix} \times \begin{matrix} \text{Att}_4 & \text{Att}_5 & \text{Att}_6 \\ \text{Dark Blue} & \text{Dark Blue} & \text{Dark Blue} \end{matrix} = \begin{matrix} \text{Att}_1 & \text{Att}_2 & \text{Att}_3 & \text{Att}_4 & \text{Att}_5 & \text{Att}_6 \\ \text{Green} & \text{Green} & \text{Green} & \text{Dark Blue} & \text{Dark Blue} & \text{Dark Blue} \\ \text{Orange} & \text{Orange} & \text{Orange} & \text{Dark Blue} & \text{Dark Blue} & \text{Dark Blue} \end{matrix}$$

Size Overhead


Att ₁	Att ₂	Att ₃	Att ₄	Att ₅	Att ₆
Green	Green	Green	Purple	Purple	Purple
Green	Green	Green	Blue	Blue	Blue
Green	Green	Green	Dark Blue	Dark Blue	Dark Blue
Yellow	Yellow	Yellow	Purple	Purple	Purple
Yellow	Yellow	Yellow	Blue	Blue	Blue
Yellow	Yellow	Yellow	Dark Blue	Dark Blue	Dark Blue
Orange	Orange	Orange	Purple	Purple	Purple
Orange	Orange	Orange	Blue	Blue	Blue
Orange	Orange	Orange	Dark Blue	Dark Blue	Dark Blue

Step 2: Database representations



Row
representation

Column
representation

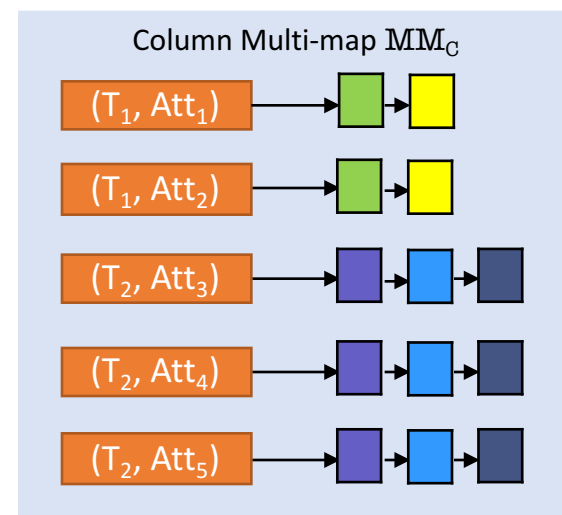
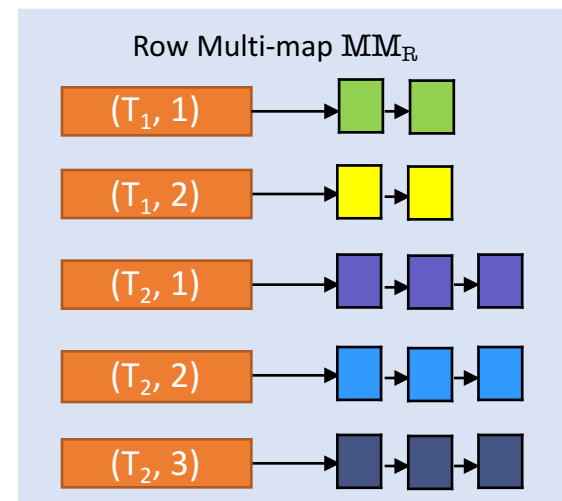
Value
representation

Cross-value
representation

Step 2: Row / Column representation

Att ₁	Att ₂
Green	Green
Yellow	Yellow

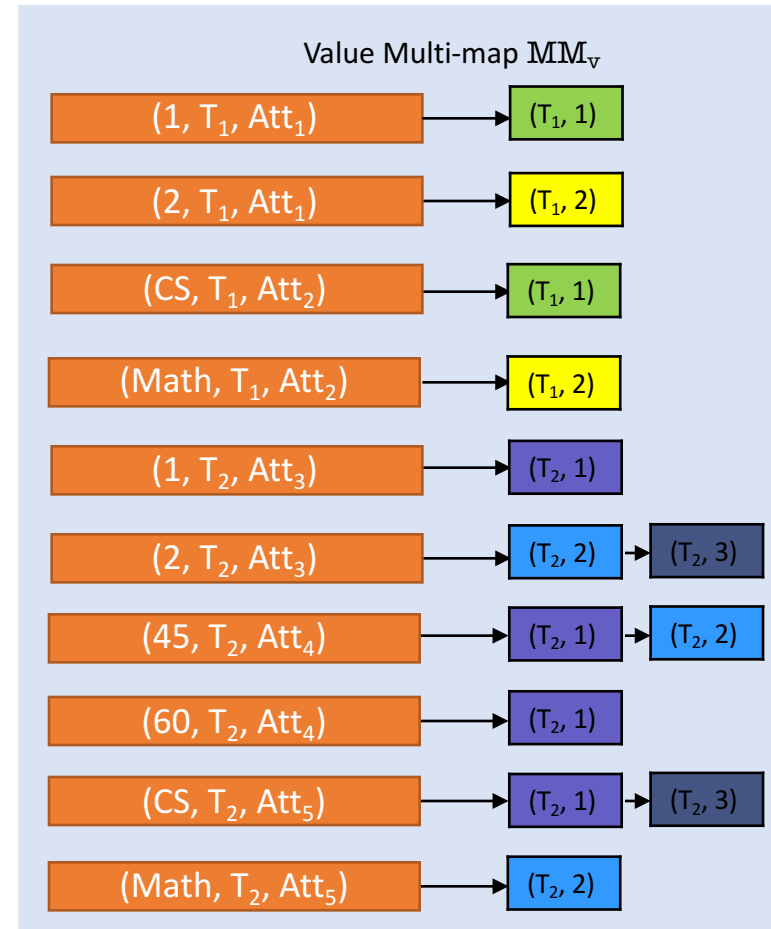
Att ₃	Att ₄	Att ₅
Purple	Purple	Purple
Blue	Blue	Blue
Dark Blue	Dark Blue	Dark Blue



Step 2: Value representation

Att ₁	Att ₂
1	CS
2	Math

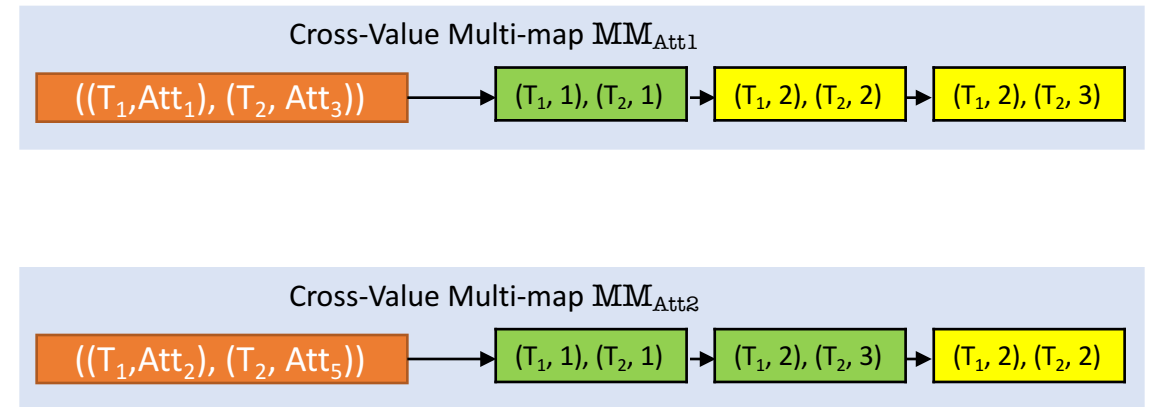
Att ₃	Att ₄	Att ₅
1	45	CS
2	45	Math
2	60	CS



Step 2: Cross-Value representation

Att ₁	Att ₂
1	CS
2	Math

Att ₃	Att ₄	Att ₅
1	45	CS
2	45	Math
2	60	CS




Step 3: SPX Setup


$$\text{Setup}^{\text{SPX}} \left[\mathbf{1}^k, \begin{array}{|c|c|} \hline \text{Att}_1 & \text{Att}_2 \\ \hline \text{Green} & \text{Green} \\ \hline \text{Yellow} & \text{Yellow} \\ \hline \end{array}, \begin{array}{|c|c|c|} \hline \text{Att}_3 & \text{Att}_4 & \text{Att}_5 \\ \hline \text{Purple} & \text{Purple} & \text{Purple} \\ \hline \text{Blue} & \text{Blue} & \text{Blue} \\ \hline \text{Dark Blue} & \text{Dark Blue} & \text{Dark Blue} \\ \hline \end{array} \right]$$

Step 3: SPX Setup

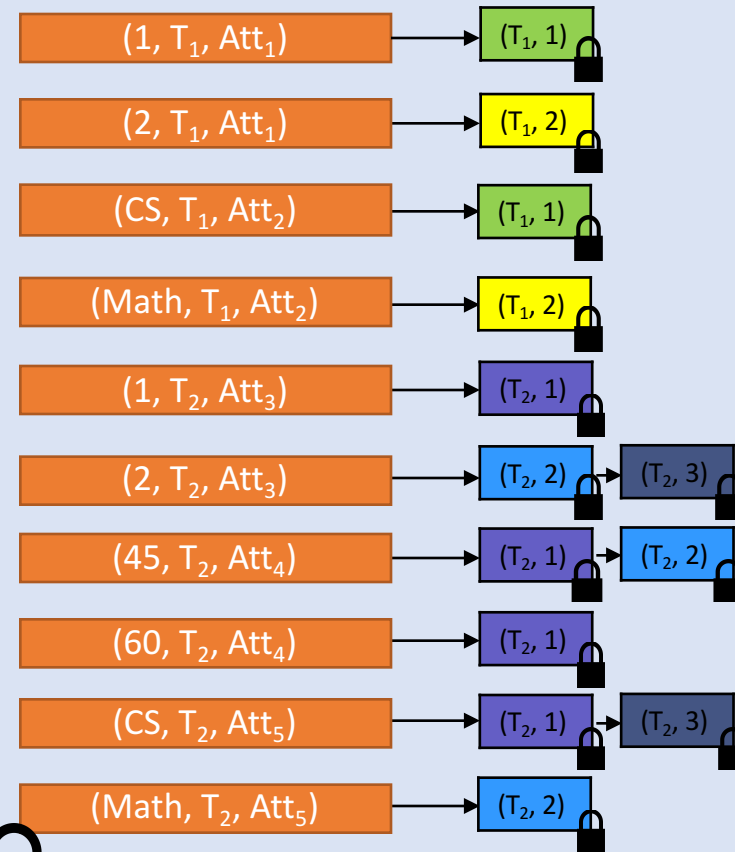
Encrypted Row Multi-map EMM_R




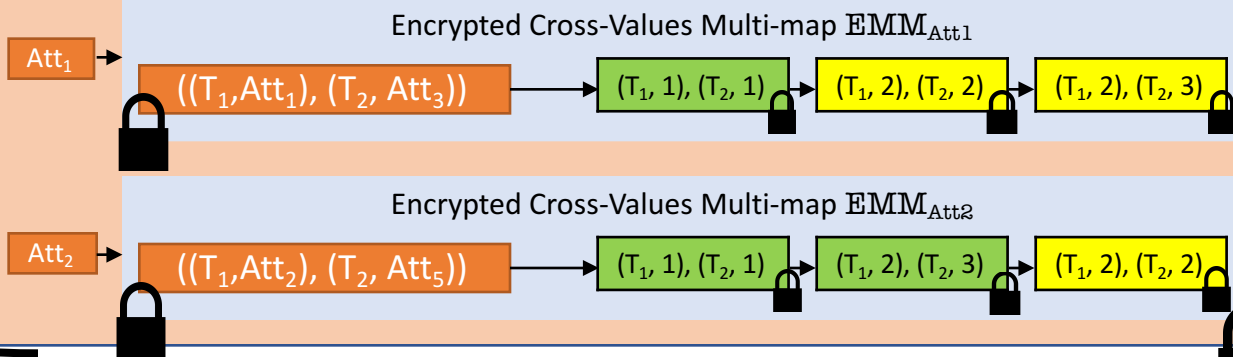
Encrypted Column Multi-map EMM_C



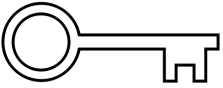
Encrypted Value Multi-map EMM_V



Encrypted dictionary EDX



Step 3: SPX Token (1)

Token^{SPX} [ , **Select Att₃**
From (T₁, T₂)
Where T₁.Att₂ = T₂.Att₅]

Step 3: SPX Token (2)

1. Rewrite SQL query to Normal Form

$$\pi_{att_3} \left(\sigma_{att_2=att_5} (T_1 \times T_2) \right)$$

2. Rewrite Normal Form to Heuristic Normal Form


3. Generate the token

Att₂ 

Dictionary
sub-token

3

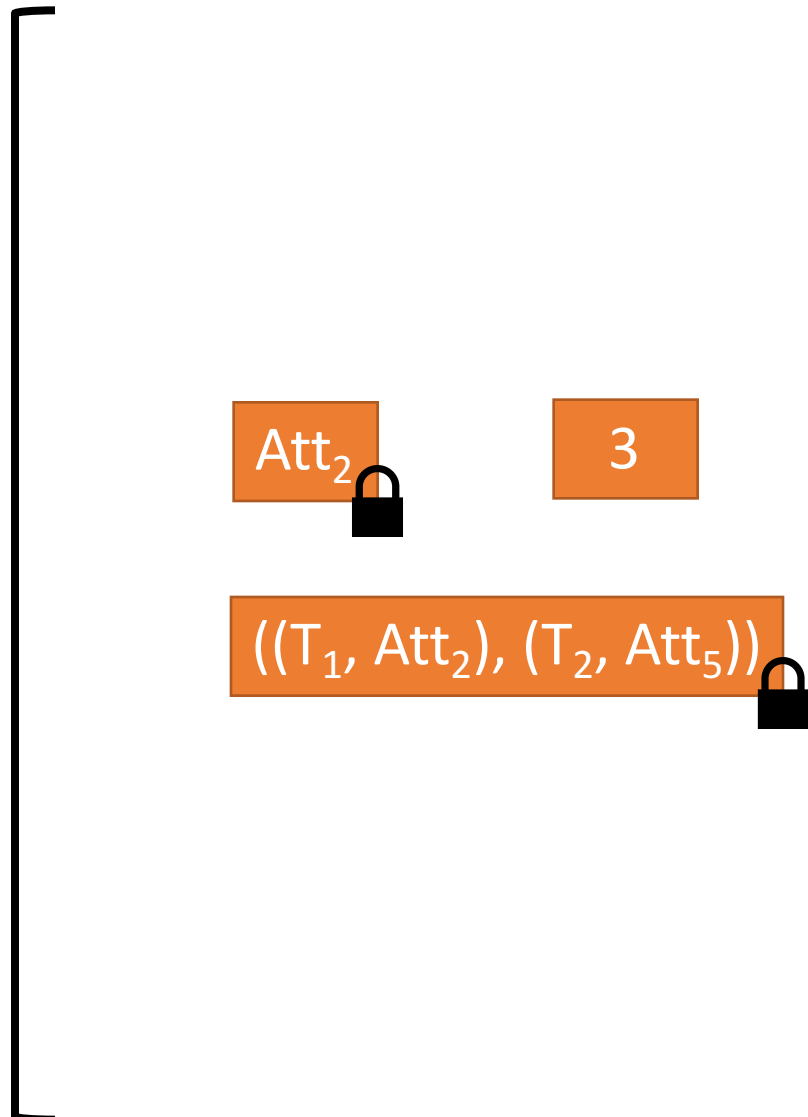
Projection
Sub-token

((T₁, Att₂), (T₂, Att₅)) 

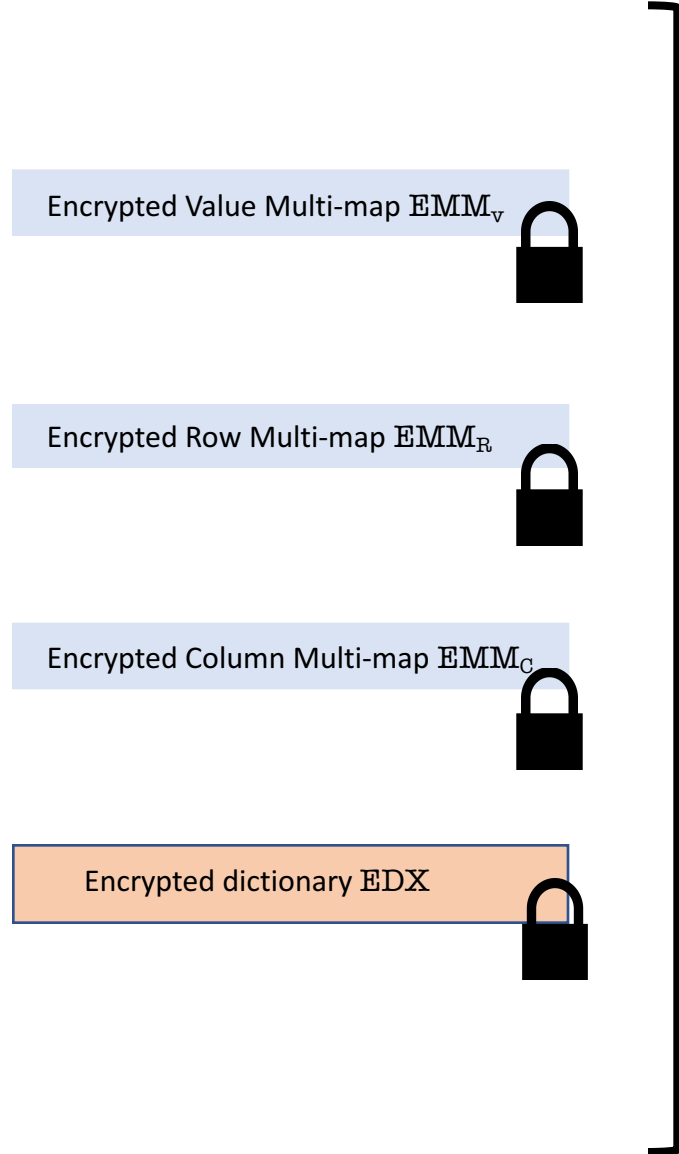
Select
Sub-token

Step 3: SPX Query (1)

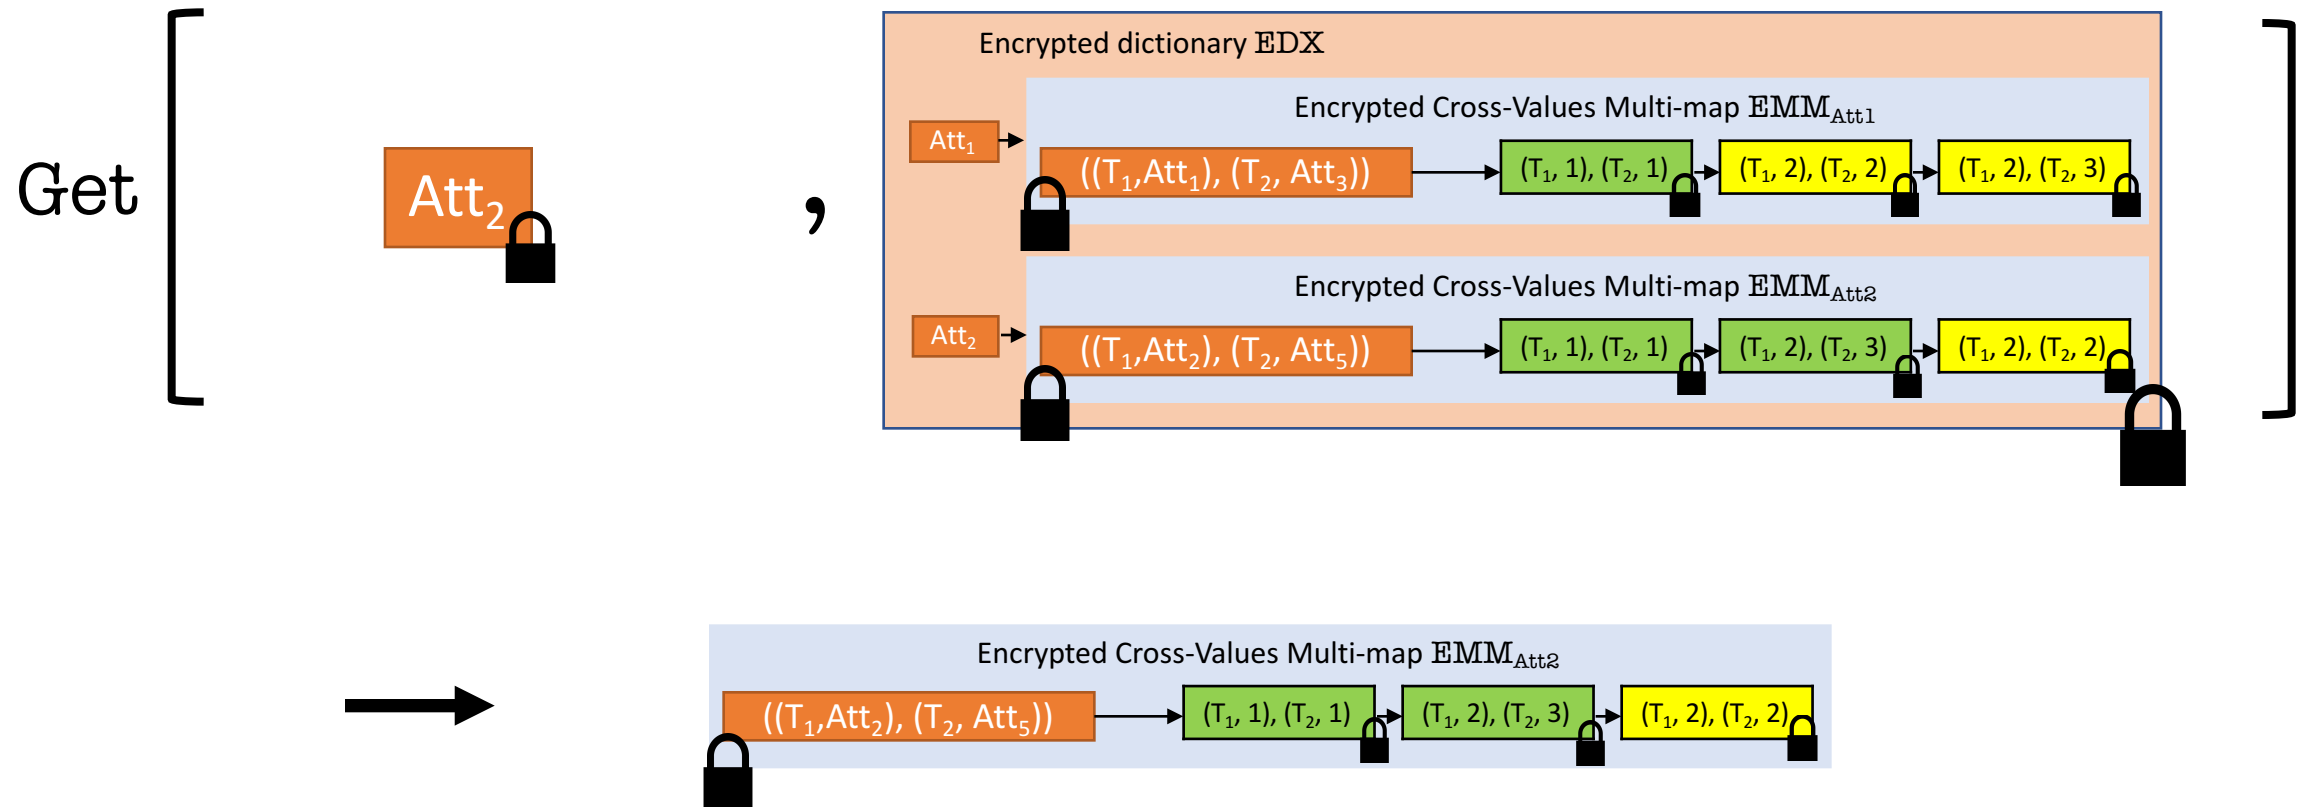
Query^{SPX}



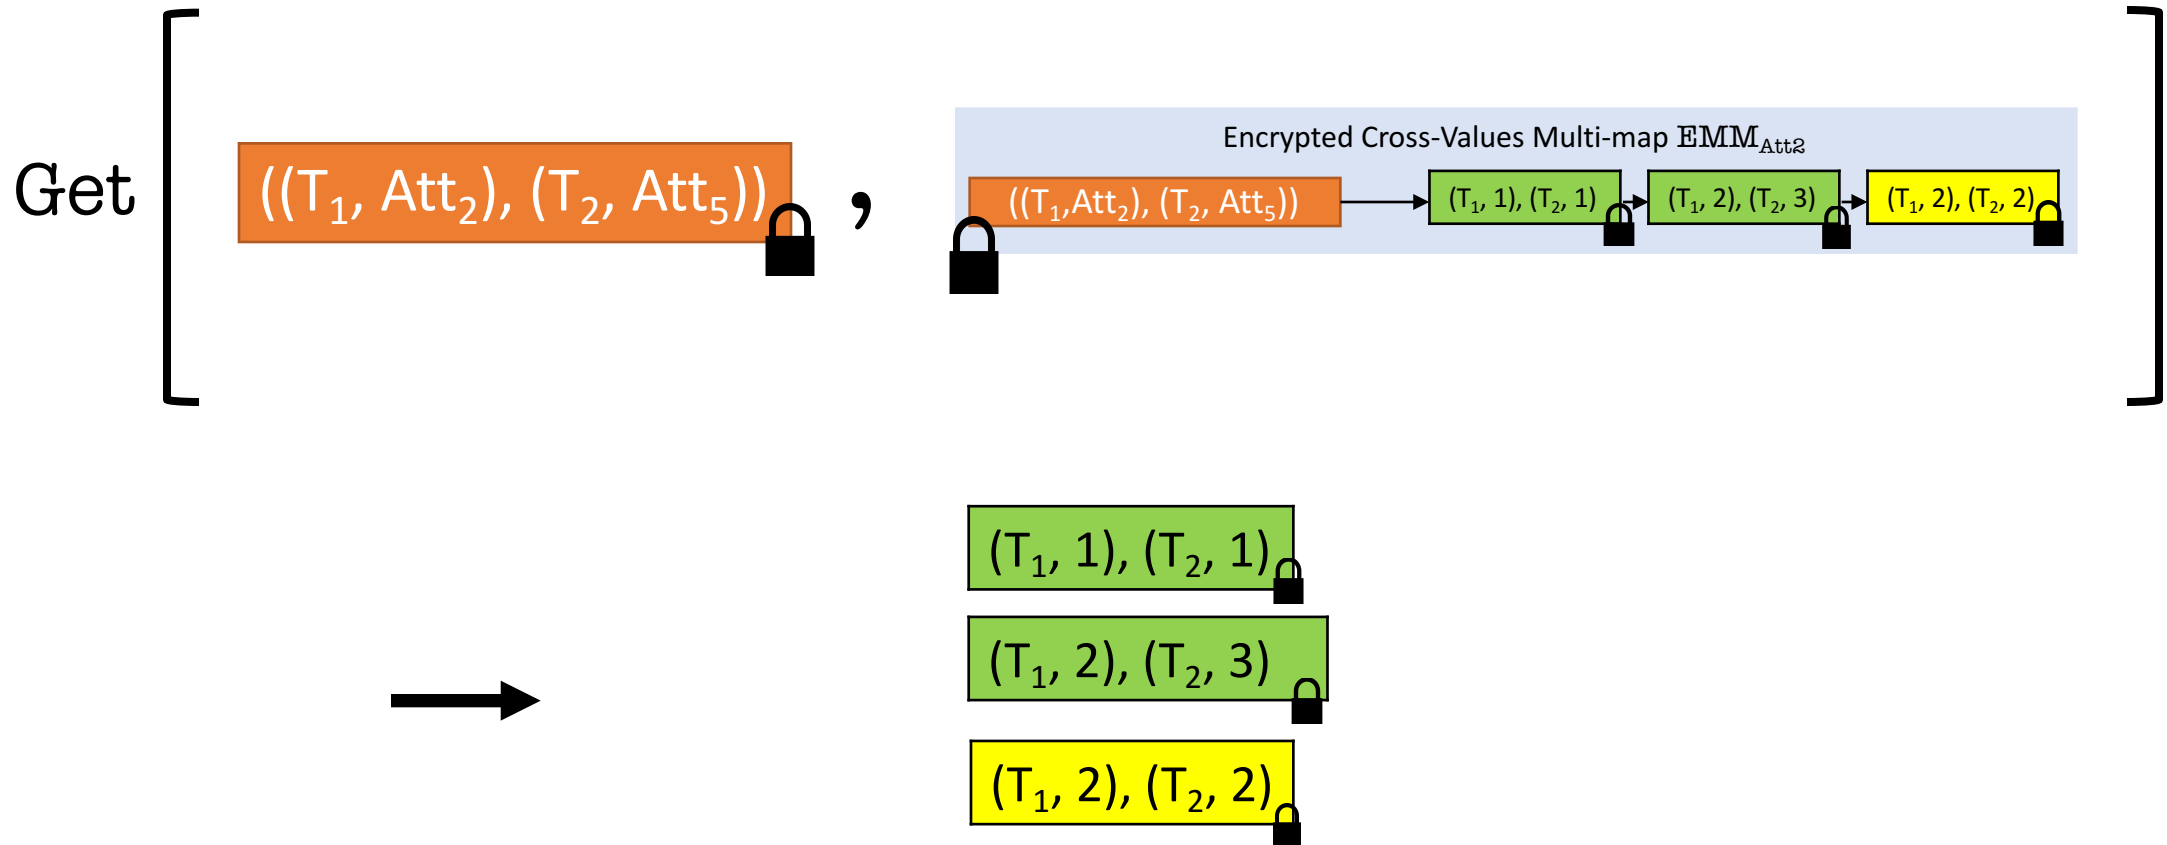
,



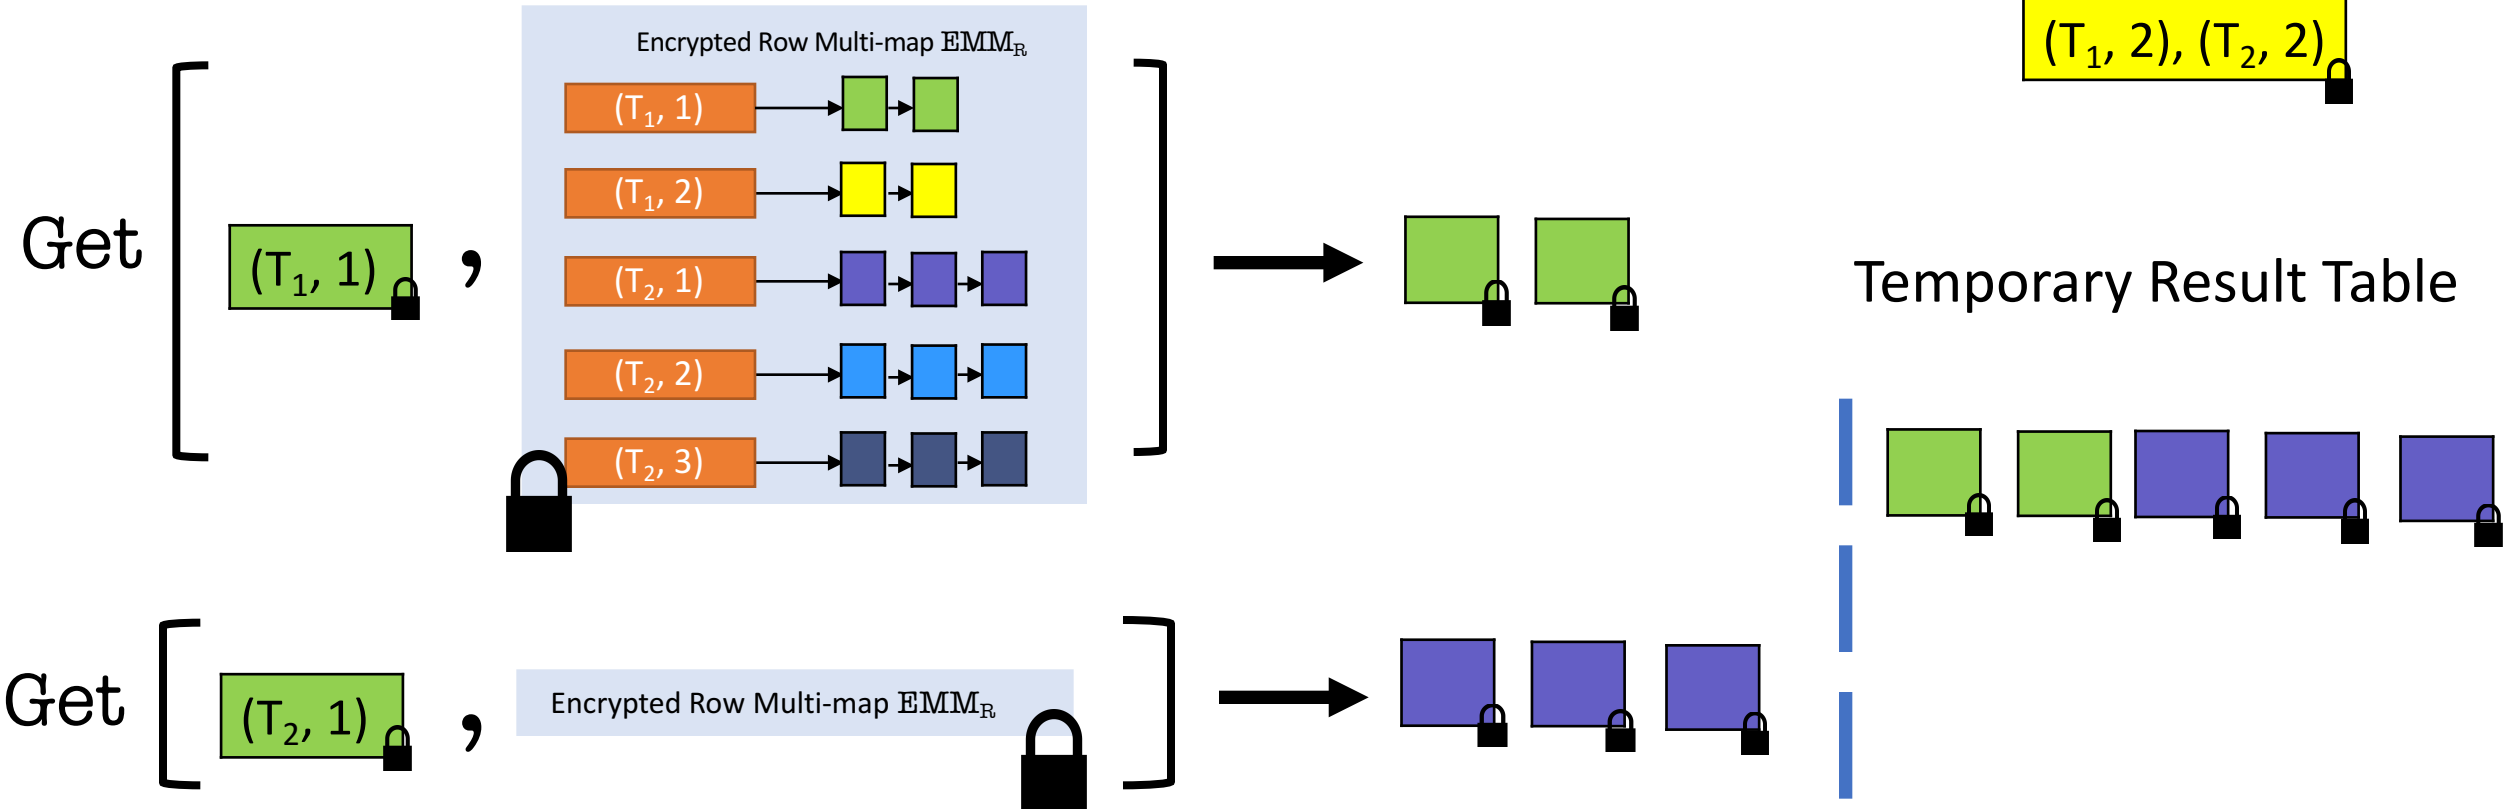
Step 3: SPX Query (2)



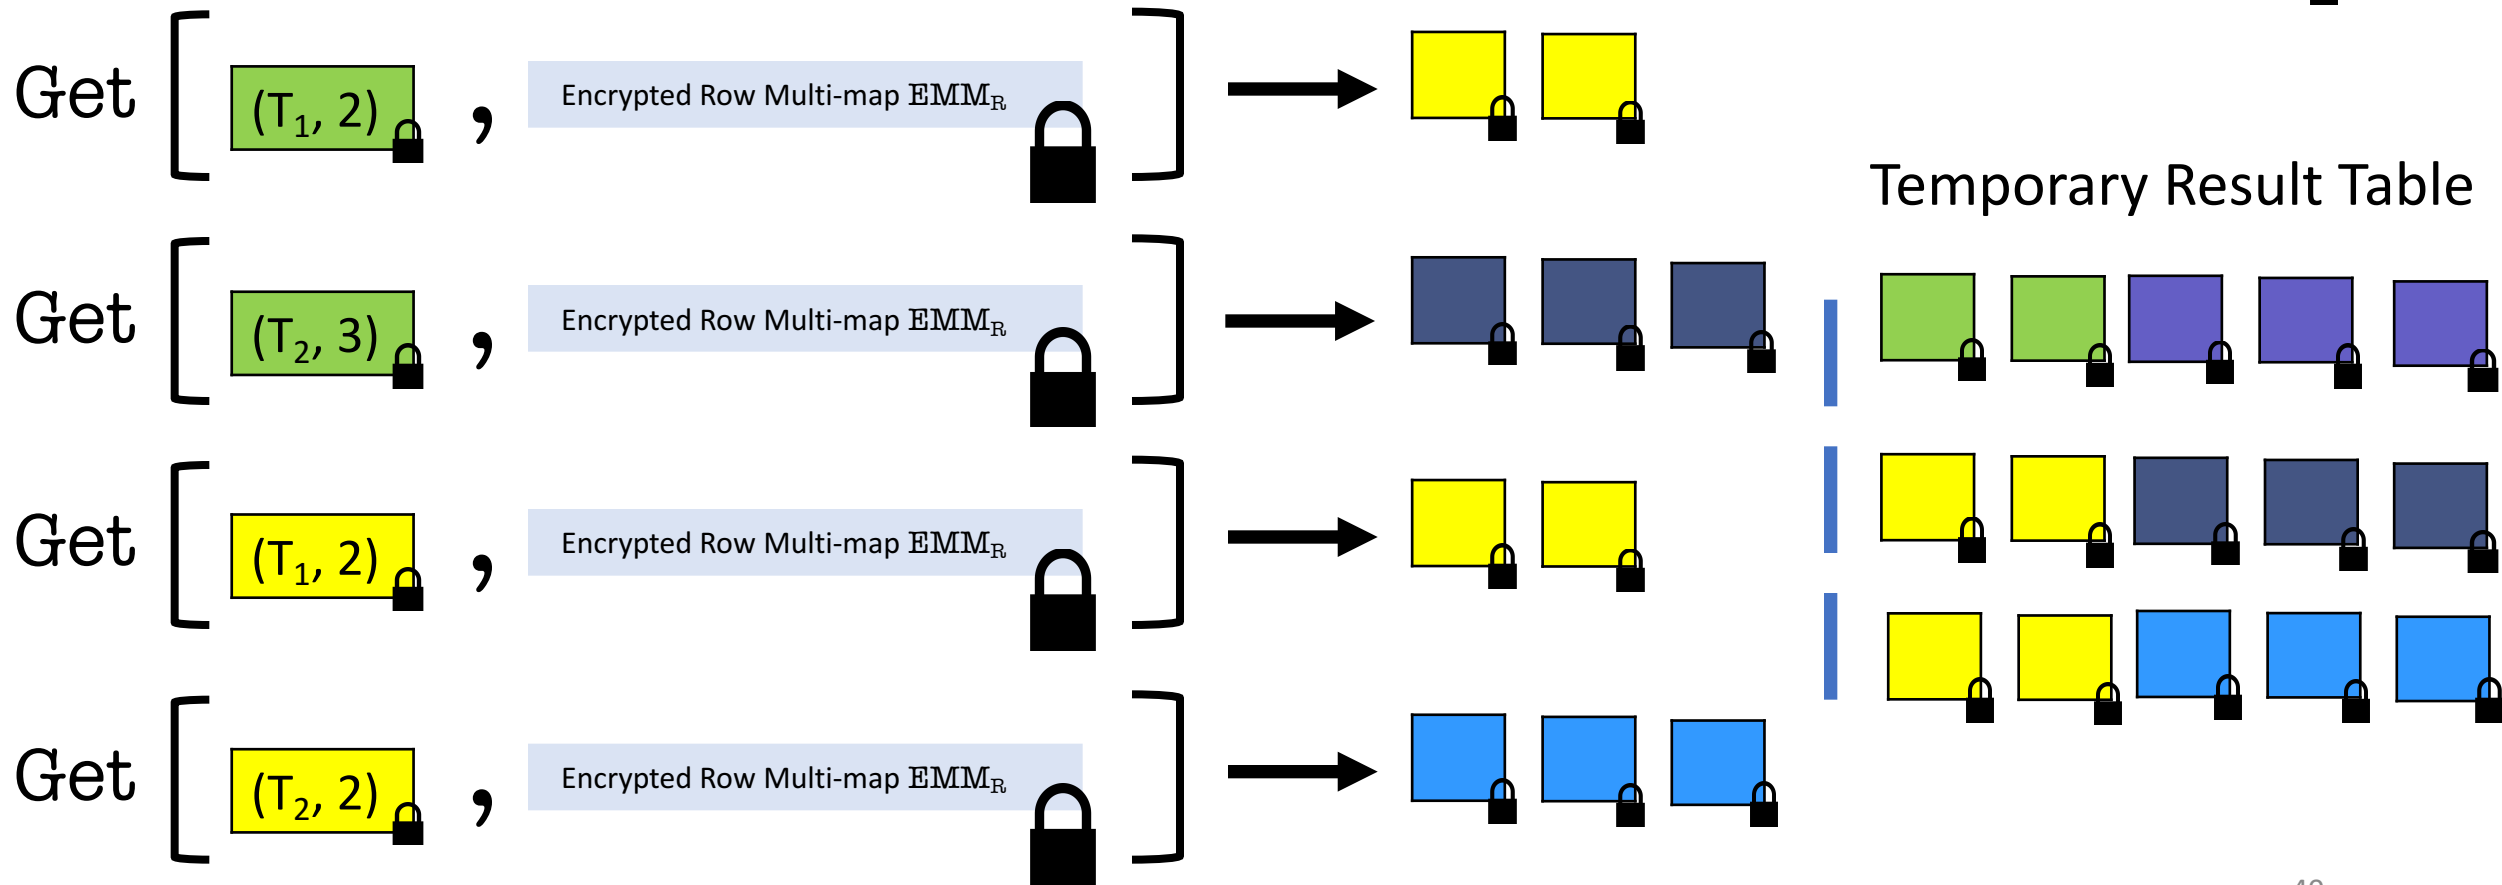
Step 3: SPX Query (3)



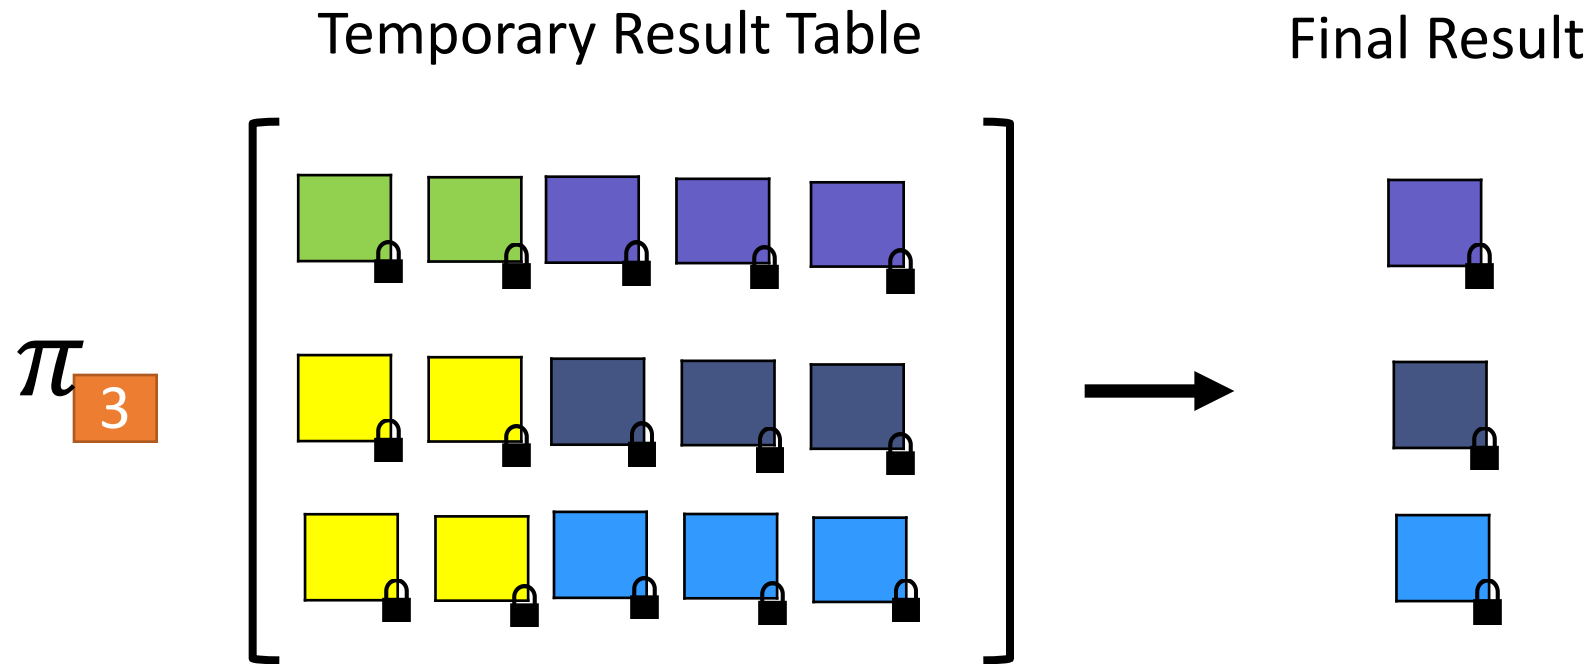
Step 3: SPX Query (4)



Step 3: SPX Query (5)



Step 3: SPX Query (6)



Leakage: SPX-OPT vs. PPE-based

- Query leakage of SPX
 - Cross product pattern
 - Projection pattern
 - Selection pattern



- Query leakage of PPE-based schemes
 - Cross product pattern
 - Projection pattern
 - Selection pattern
 - Frequency pattern
 - Persistent
 - Existing very strong attacks

Modularity: SPX-Obliv vs. SPX-OPT

Query leakage of SPX-Obliv

- When the EMMs are oblivious
[[GO96](#),[SvDS+13](#),[GMP16](#),[KMO18](#)]
- But comes with extra overhead



Query leakage of SPX-OPT

SPX-OPT Asymptotics

- Worst-case query complexity with
 - Assuming optimal MM and DX encryption schemes [CK10,CJJ+14]
 - h projected attributes
 - t tables with size s_i each
 - $\#R$ is the size of the result on a plaintext database

$$O\left(\frac{\#R}{h} \cdot \sum_{i=1}^t s_i\right)$$

- Mild condition

- If $h^{-1} \cdot \sum_{i=1}^t s_i$ is constant in $\#R$, then query time is **optimal**

SPX-OPT Asymptotics

- Communication complexity is **optimal**
- Storage depends on the data distribution

$$O(\#\text{DB} + \sum_{\text{att} \in \mathcal{S}} \#\text{MM}_{\text{att}})$$

- Concretely, the big-O hides a multiplicative factor of 3

Takeaways and Future Work

- First STE-based encrypted relational database
- Sub-linear query time (optimal under certain conditions)
- First (forward-private) dynamic STE-based encrypted relational database
- Better leakage profile than PPE-based
- Future research questions:
 - Extend SPX to handle relational algebra
 - Remove interaction in SPX⁺
 - Extend SPX to handle range sub-predicates
 - Design of new encrypted range solutions is required [[KKNO16](#), [LMP17](#)]

Thank you!

<https://eprint.iacr.org/2016/453>