Cryptanalysis Strikes Back A Realistic assessment of leakage attacks on Encrypted Search

Abdelkarim Kati†‡

†School of Computer Science,
Mohammed VI Polytechnic University.
‡ Encrypted Systems Lab, Brown University.

January 24, 2023 at Aarhus University.

Some Slides were adapted from A.Trieber RWC'22 Talk.

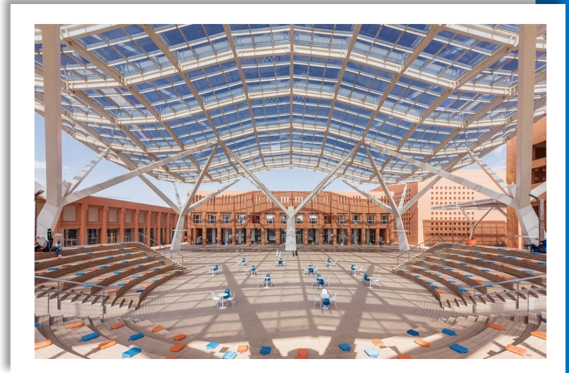

Motivation

Real-World Applications

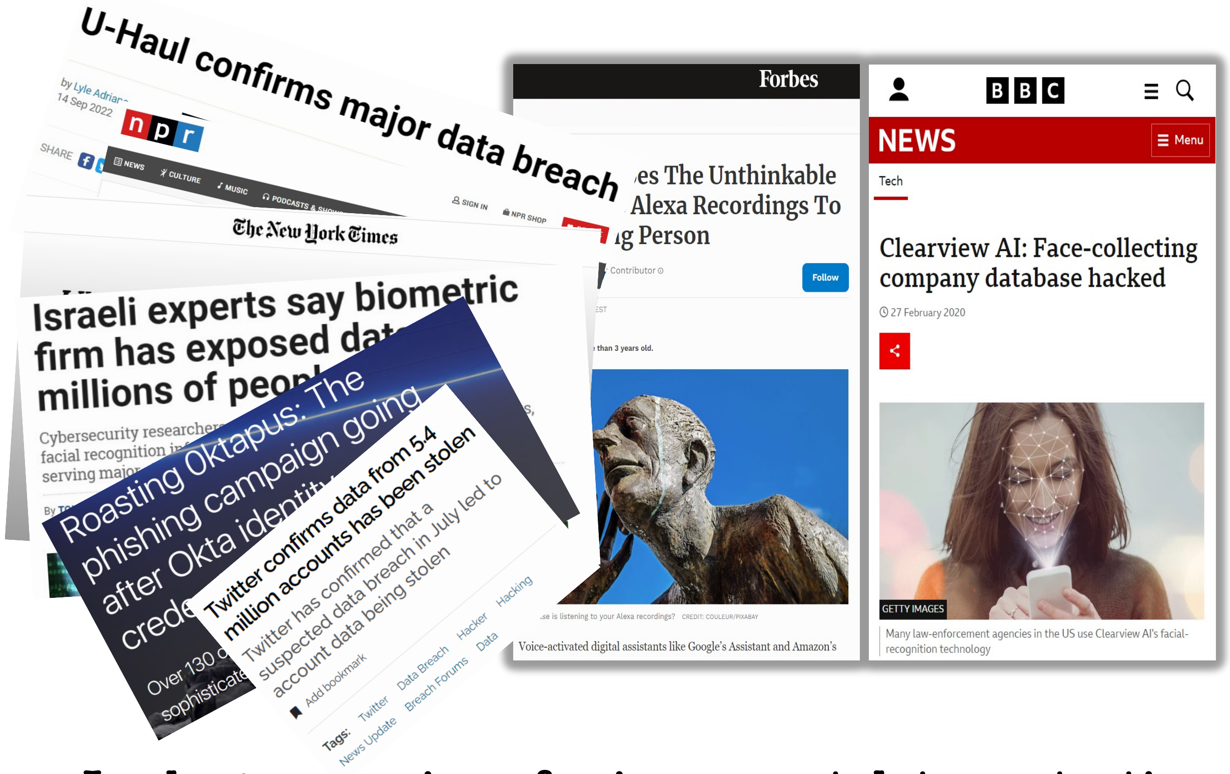

Leakage = erosion of privacy w.r.t data protection

Cryptographic Mechanisms

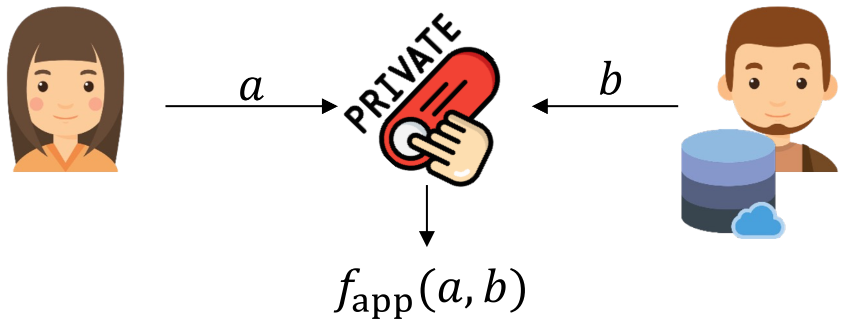

Privacy-Enhancing Technologies (PETs)

A Realistic assessment of leakage attacks on Encrypted Search

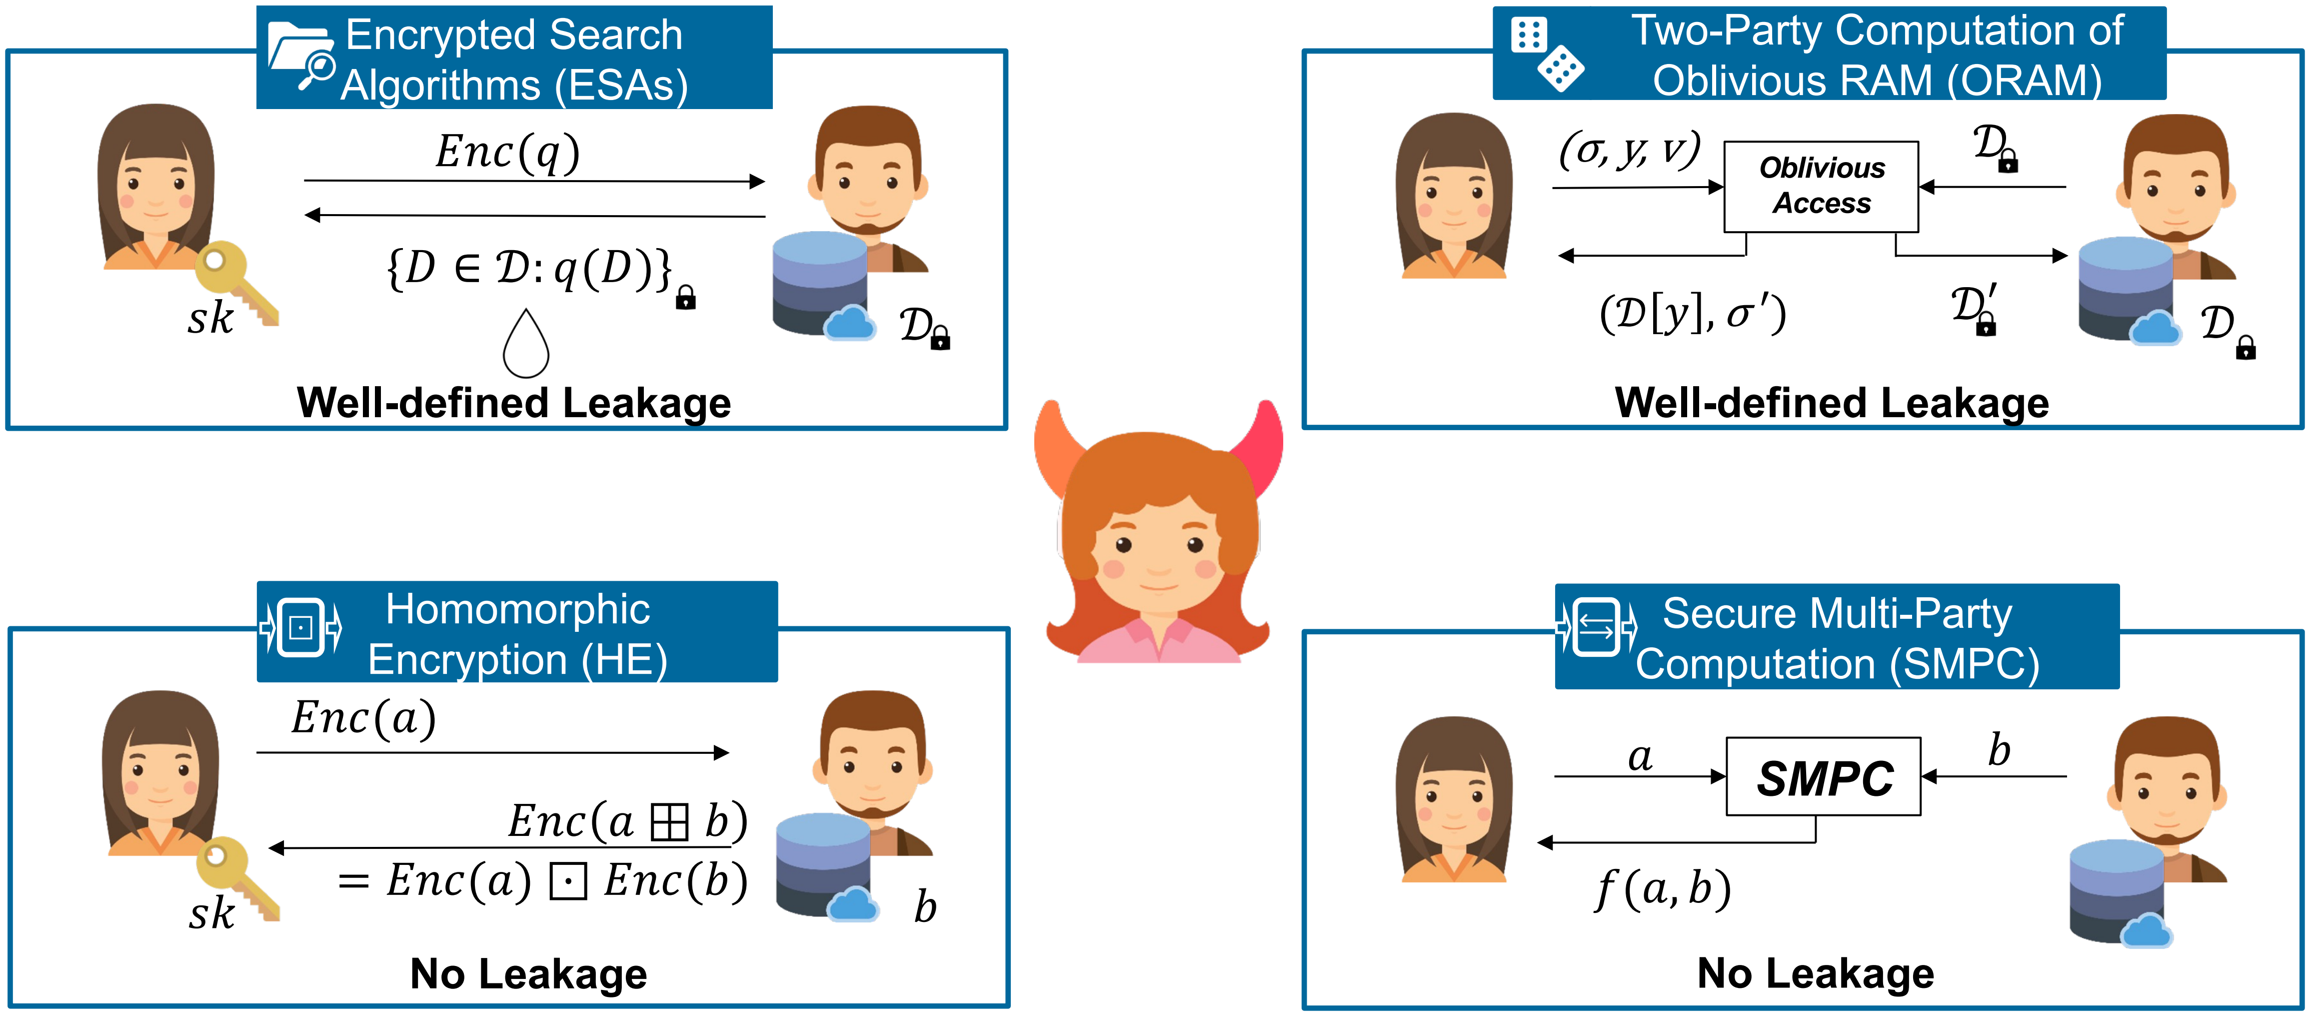

A Realistic assessment of leakage attacks on **Encrypted Search**

Encrypted Search Algorithms (ESAs)

Trusted
Client

Untrusted
Server

D.C

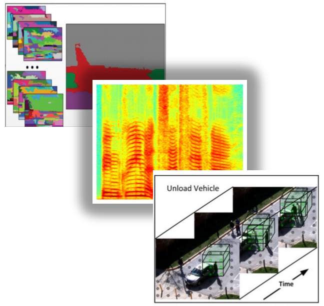

Encrypted Search Algorithms (ESAs)

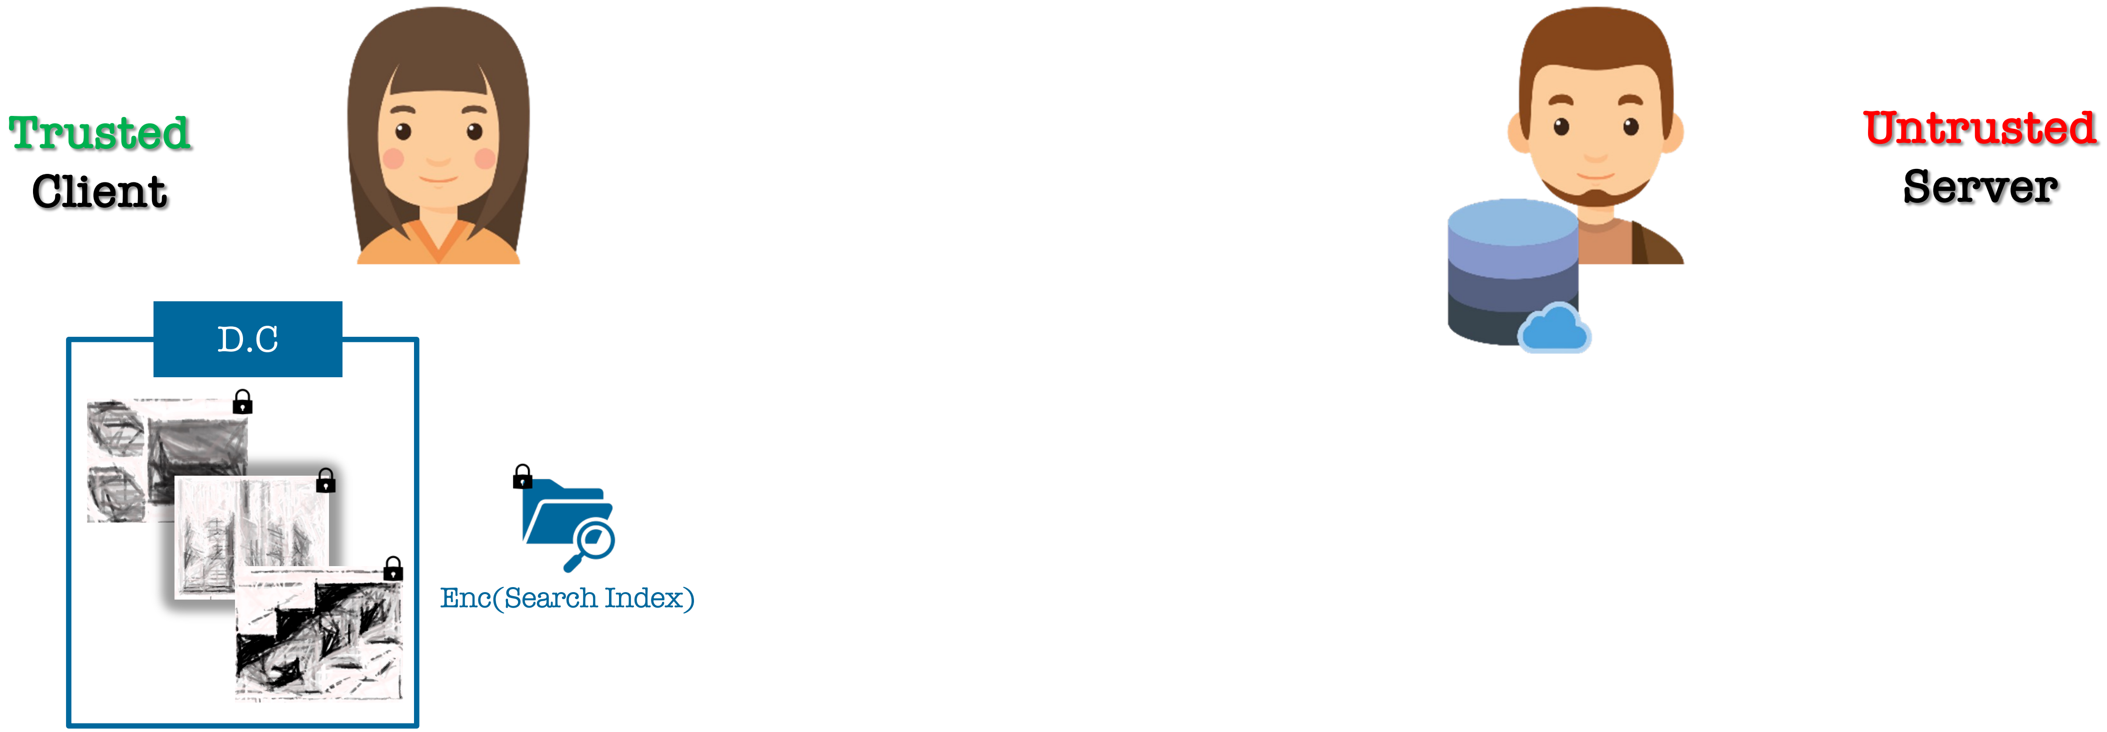

Encrypted Search Algorithms (ESAs)

Encrypted Search Algorithms (ESAs)

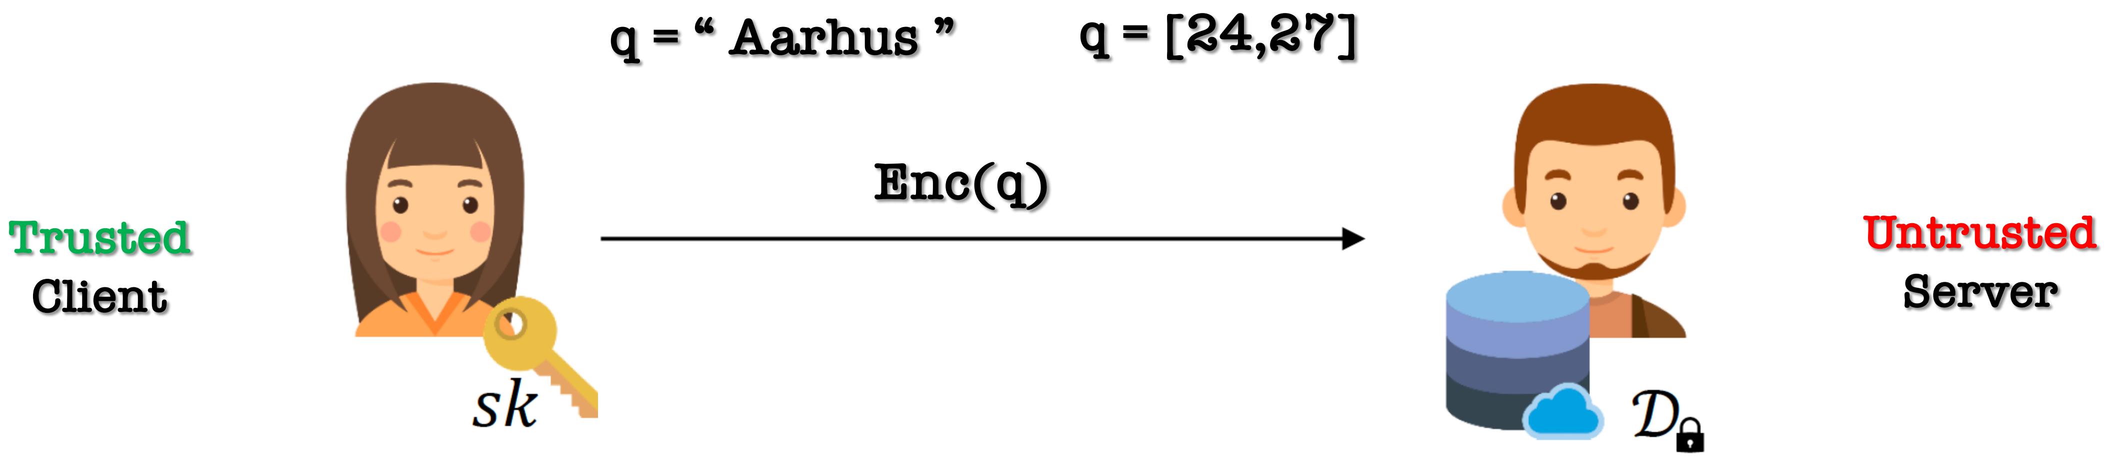

Encrypted Search Algorithms (ESAs)

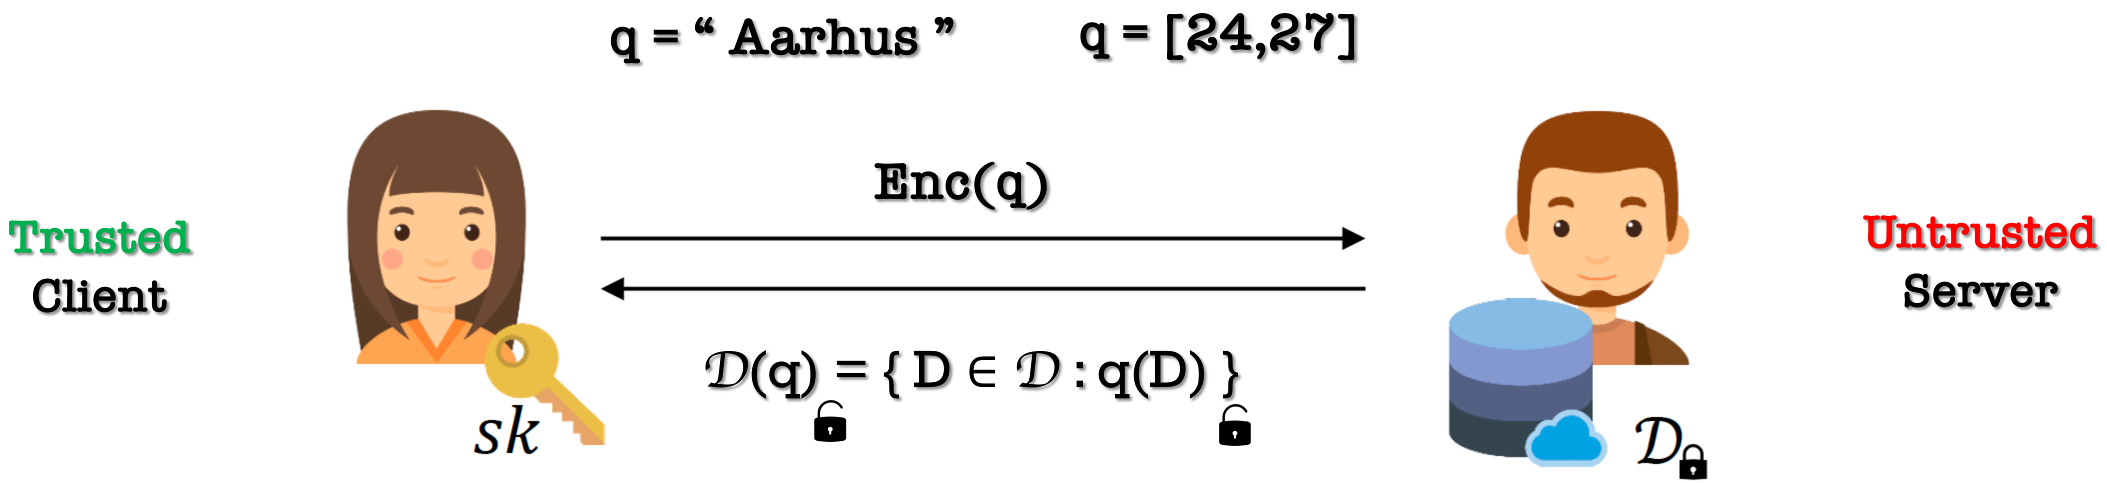

Encrypted Search Algorithms (ESAs)

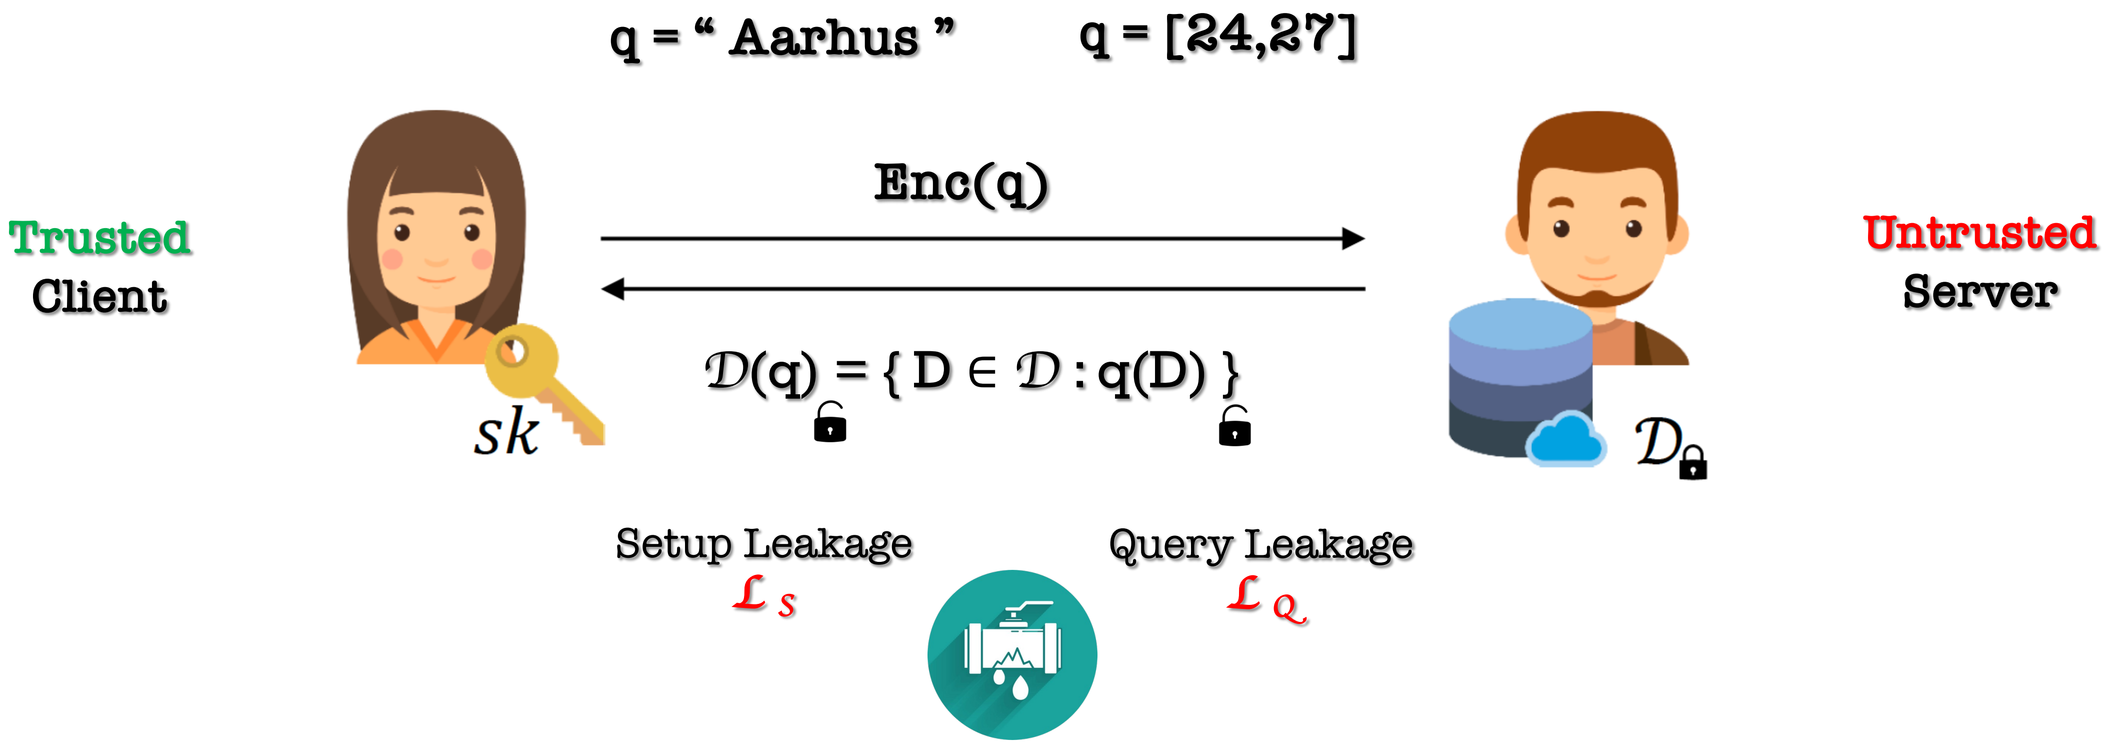

Encrypted Search Algorithms (ESAs)

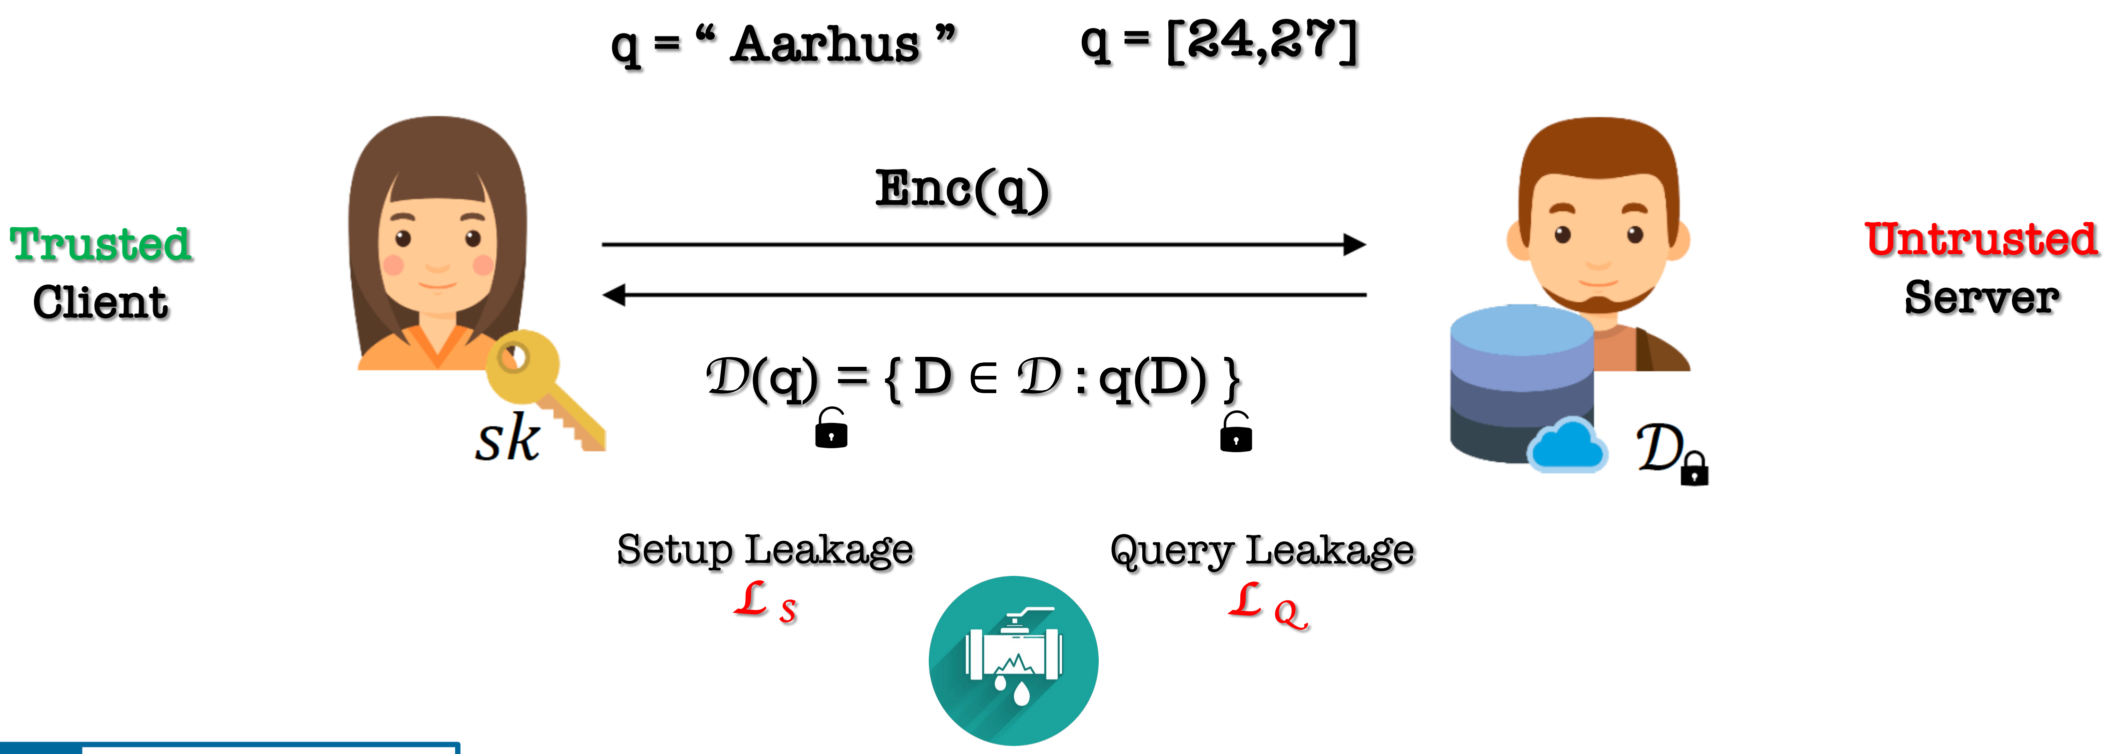

Our work

- Structured Encryption (**STE**)
- Searchable Symmetric Encryption (**SSE**)
- Oblivious RAM (**ORAM**)

Encrypted Search Algorithms (ESAs)

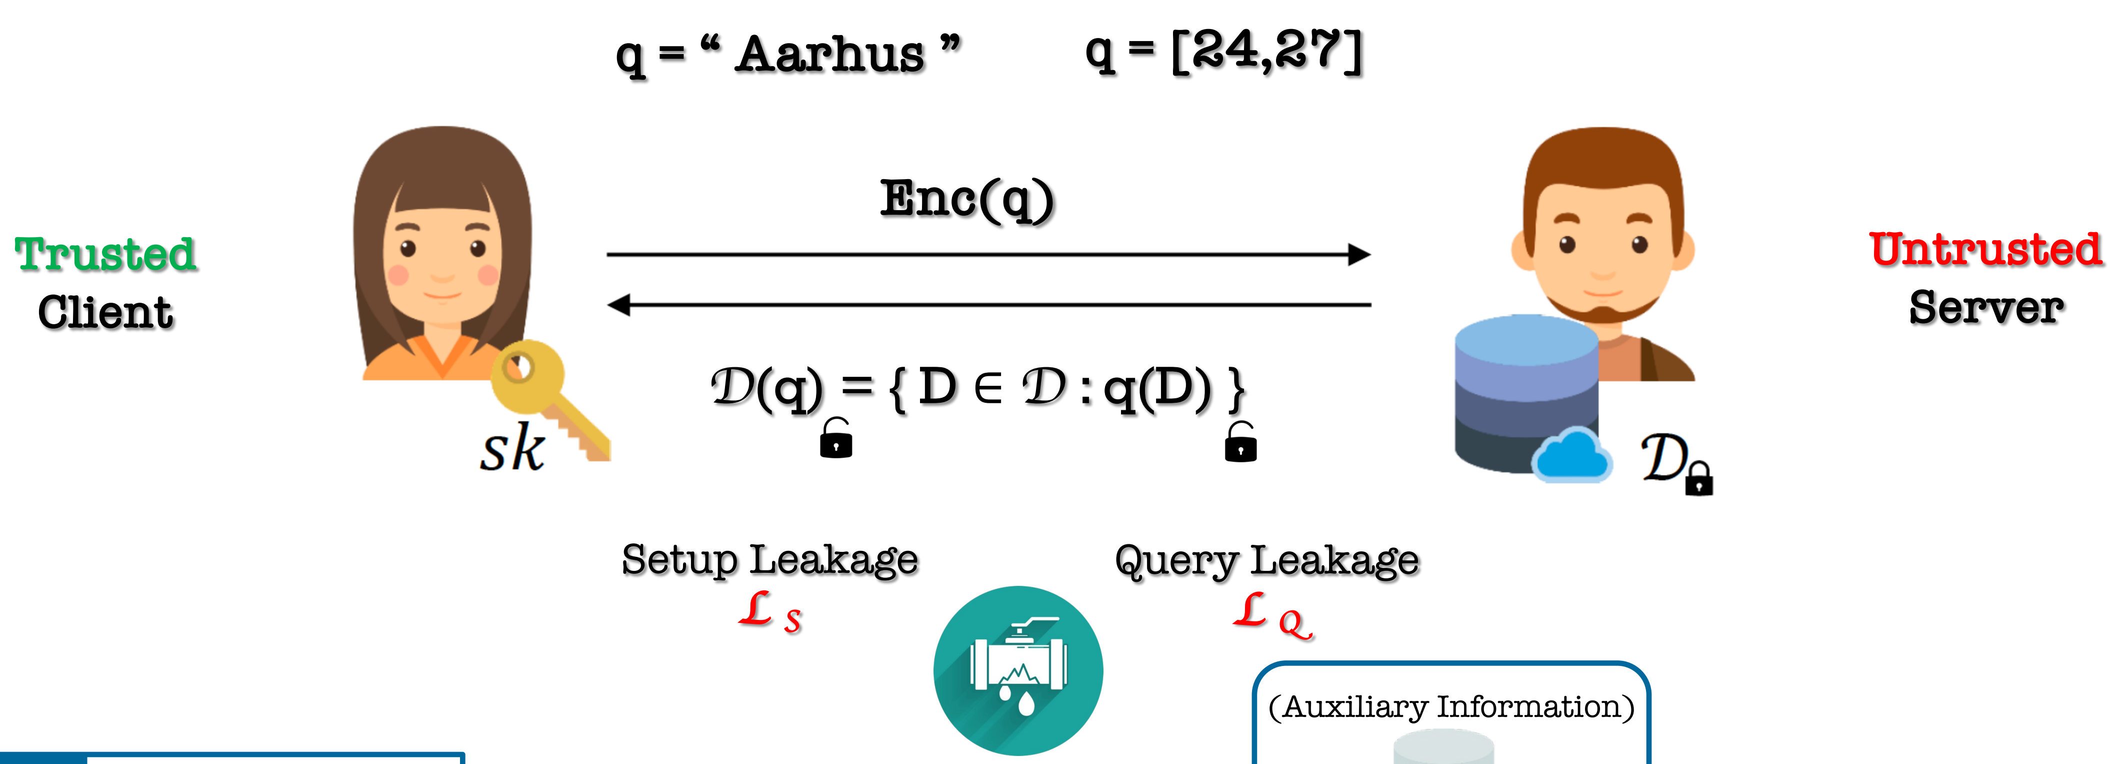

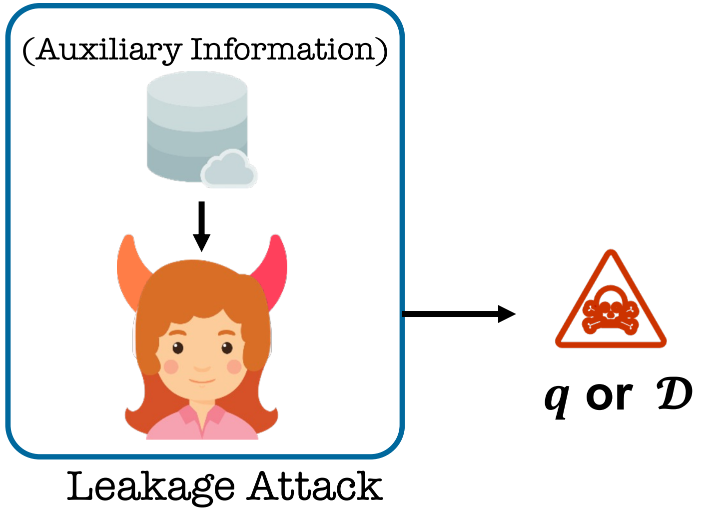

Our work

- Structured Encryption (**STE**)
- Searchable Symmetric Encryption (**SSE**)
- Oblivious RAM (**ORAM**)

Encrypted Search Algorithms (ESAs)

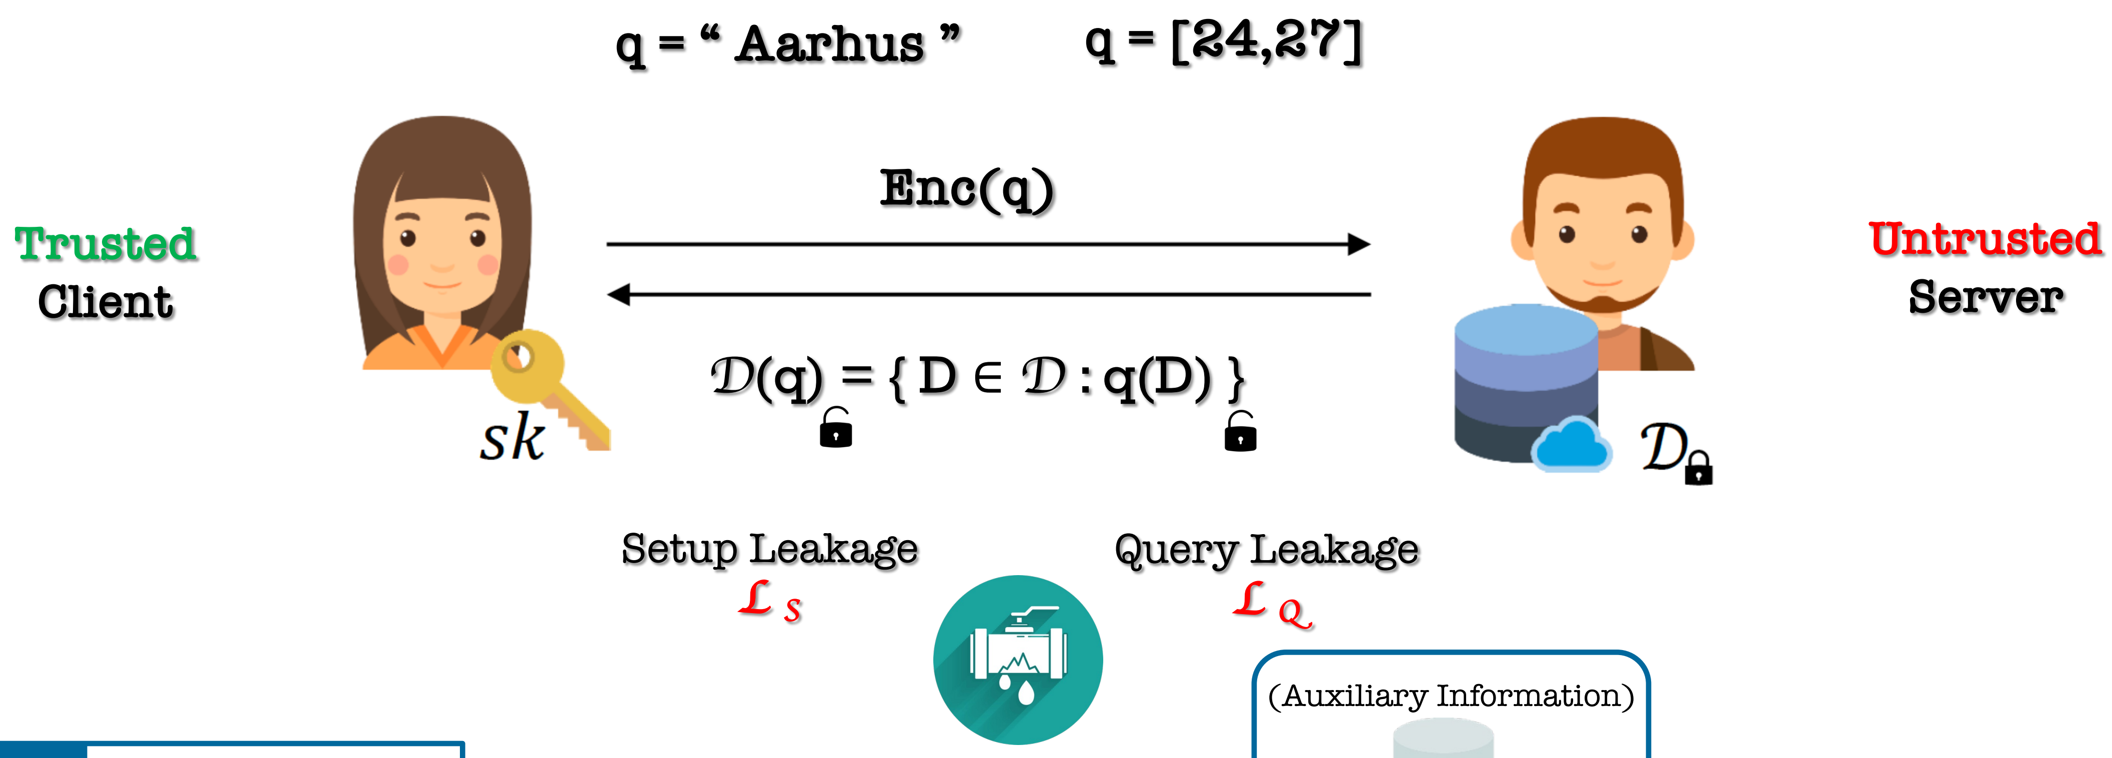

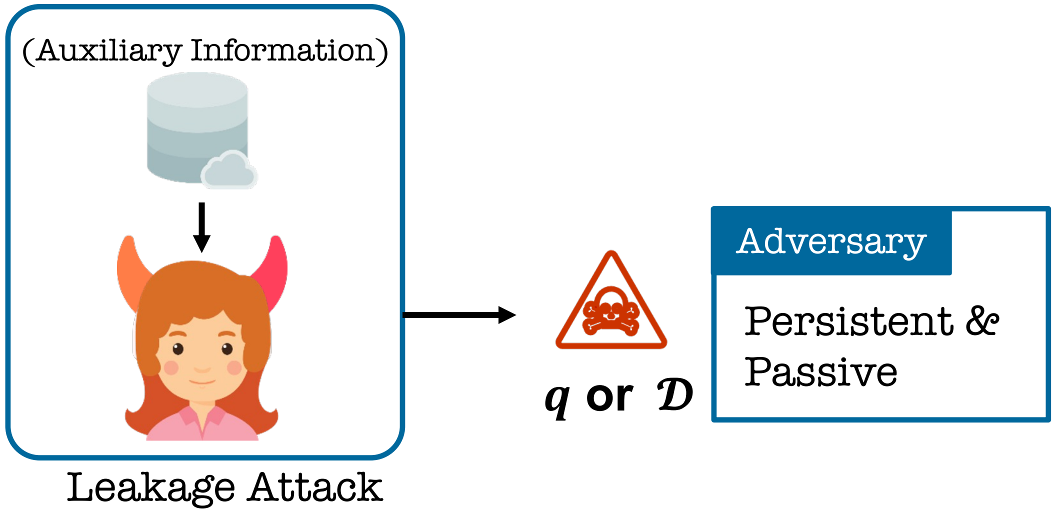

Our work

- Structured Encryption (**STE**)
- Searchable Symmetric Encryption (**SSE**)
- Oblivious RAM (**ORAM**)

Encrypted Search Algorithms (ESAs)

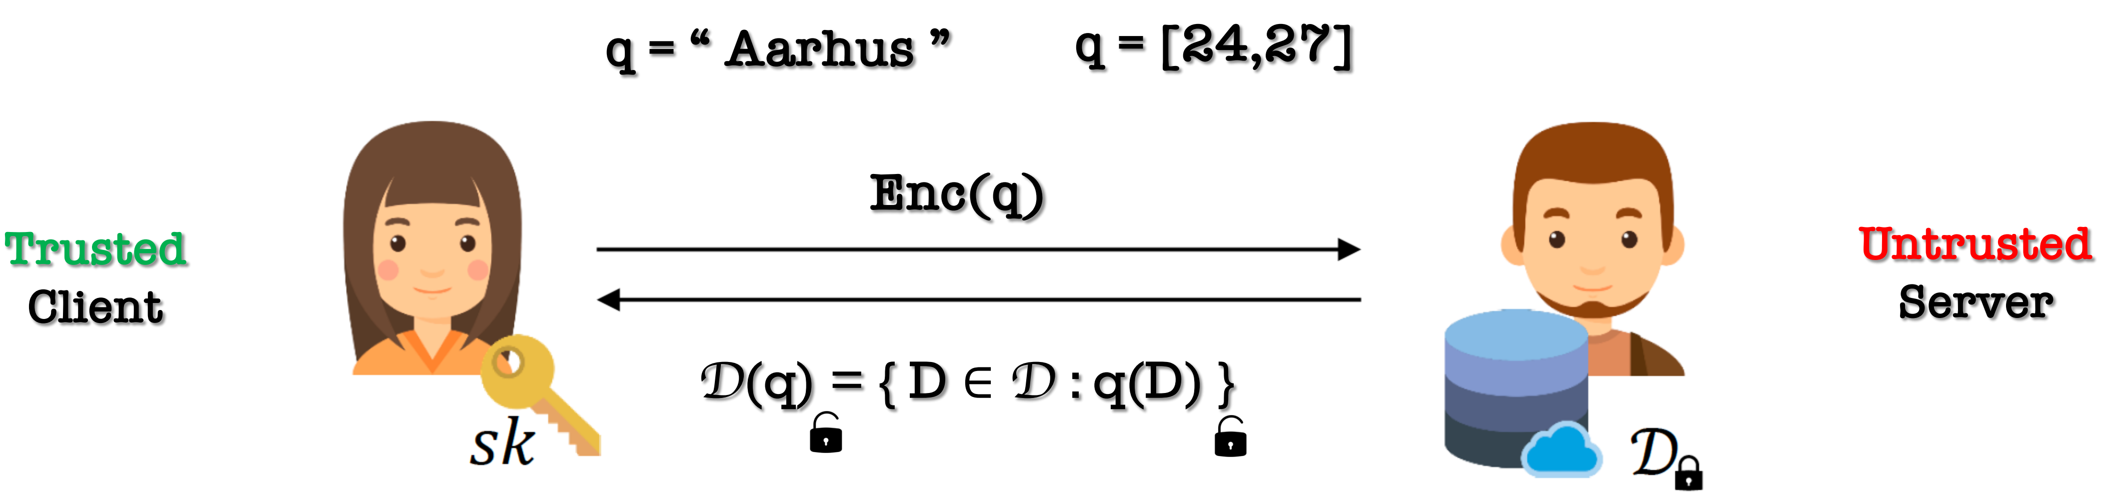

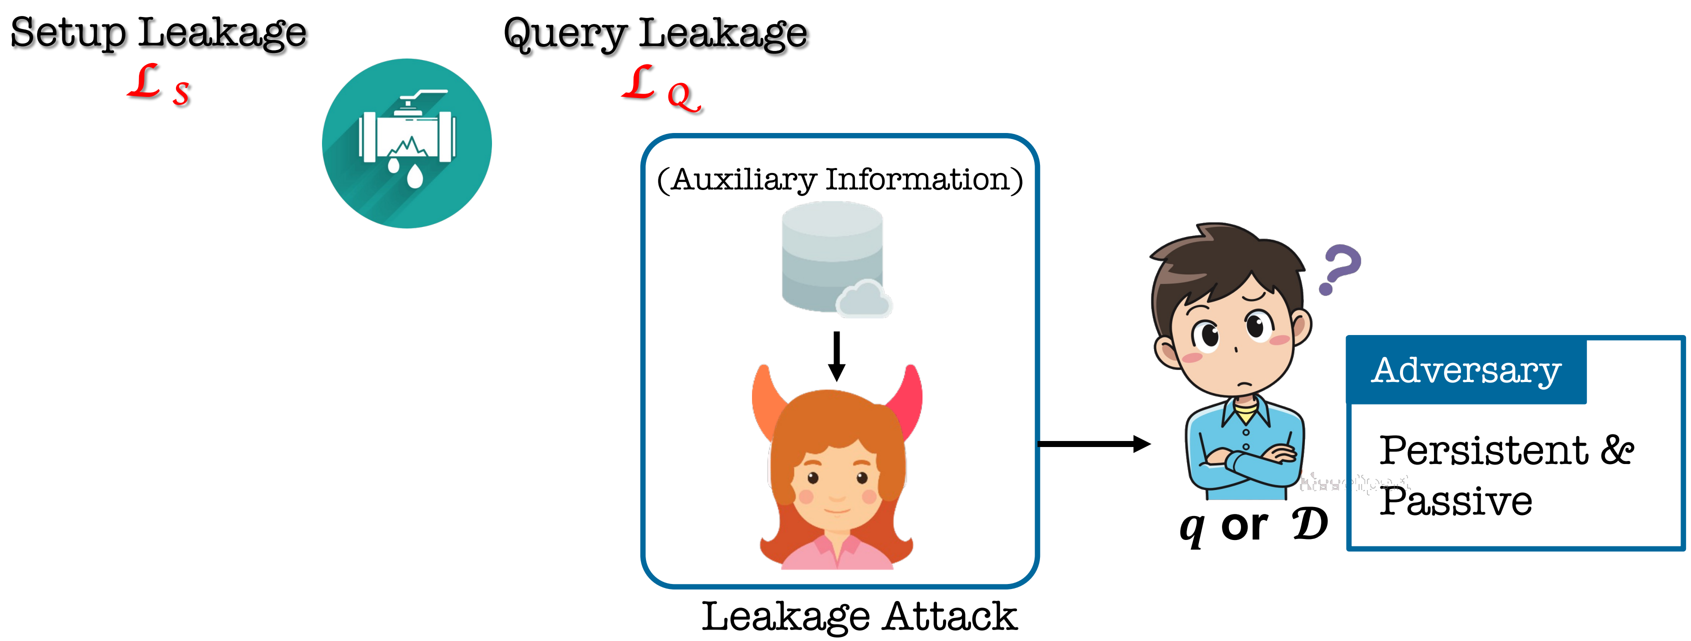

Our work

- Structured Encryption (**STE**)
- Searchable Symmetric Encryption (**SSE**)
- Oblivious RAM (**ORAM**)

A Realistic assessment of **Leakage Attacks** on Encrypted Search

How do we model Leakage ?

- The "Baseline" leakage profile for response-revealing EMMs
  - ✓ $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U)$ = (dsize, (qeq, rid), usize)
- The "Baseline" leakage profile for response-hiding EMMs
  - ✓ $(\mathcal{L}_S, \mathcal{L}_Q, \mathcal{L}_U)$ = (dsize, qeq, usize)
- Several new constructions have better leakage profiles
  - ✓ AZL and FZL [Kamara-Moataz-Ohirimenko'18]
  - ✓ VHL and AVHL [Kamara-Moataz'19]

Leakage	Information
Response Length	$\lvert D(q) \rvert$
Query Equality	$q_i = q_j$
Co-Occurrence	$\lvert D(q_i) \cap D(q_j) \rvert$
Response Identity	$\{i : D_i \in q(D)\}$
Response Volumes	$\{\lvert D_i \rvert_b : D_i \in q(D)\}$

(Simplified)

Leakage Attacks Types

Keyword (point) queries
[IKK12,CGPR15,BKM20,RPH21]

Keyword	Document IDs
'Aarhus'	2,5,11,13,20,31
'systems'	3,5,10,11,13,25
'lab'	5,11,21,27

$q = w$
$\mathcal{D}(q) = \{D \in \mathcal{D}: q \in D)\}$
Recover q

q = 'Aarhus'

Known-data: Adversary knows subset of $\mathcal{D}$

Range queries

[KKNO16,LMP18,GLMP18,
GLMP19,GJW19,KPT20,KPT21]

ID	Age
1	65
2	7
3	27

$q = (a, b)$
$\mathcal{D}(q) = \{r \in \mathcal{D}: a \leq r \leq b\}$
Recover $\mathcal{D}$

q = (18,39)

No auxiliary knowledge

Leakage Attacks against ESAs

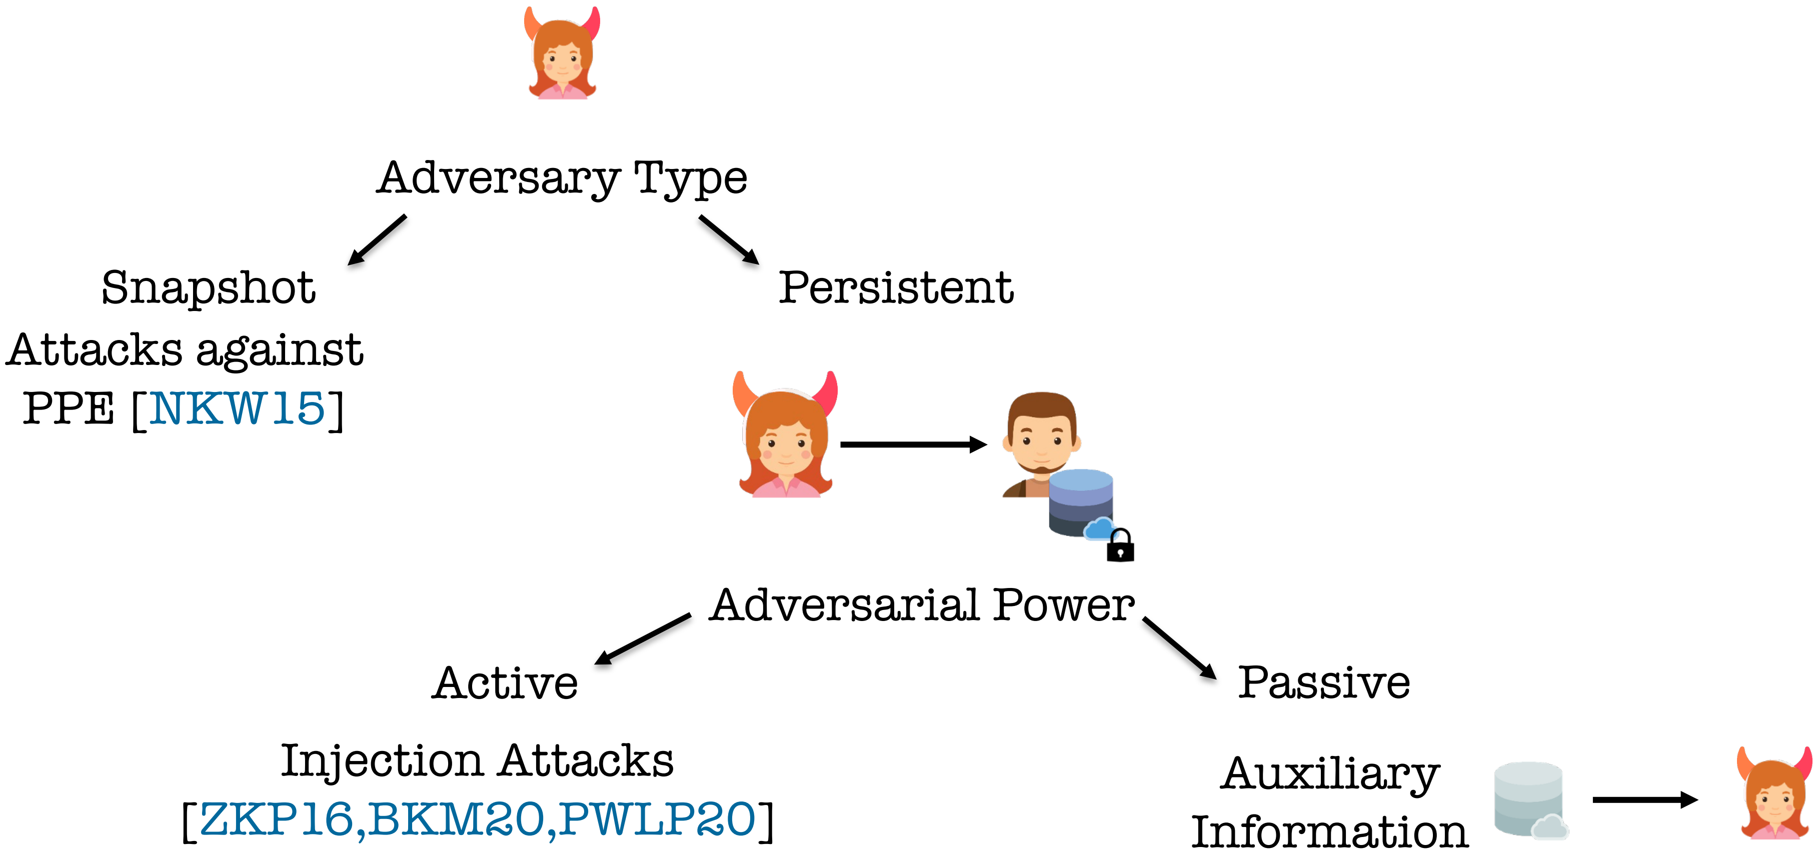

Leakage Attacks against ESAs

Persistent Passive

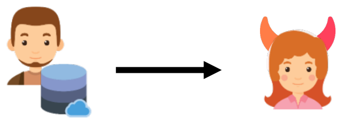

Auxiliary Information

Sampled-data or sampled-query

Keyword & Range attacks
[LZWT14,LMP18,GLMP18,GJW19,OKa21,DHP21,GPP21,OKb21]

This talk

Known-data

Keyword attacks
[IKK12*,CGPR15,BKM20, RPH21]

$q = w$
$\mathcal{D}(q) = \{D \in \mathcal{D}: q(D)\}$
Recover q

None

Range attacks
[KKNO16,LMP18,GLMP18, GLMP19,GJW19,KPT20, KPT21]

[ ]

$q = (a, b)$
$\mathcal{D}(q) = \{r \in \mathcal{D}: a \leq r \leq b\}$
Recover $\mathcal{D}$

ESAs Techniques Overview

Technique	Leakage	Query Time
Fully Homomorphic Encryption (FHE)	• None	Linear
Oblivious RAM (ORAM)	• Response Length + Volume	Sublinear
Structured Encryption (STE)	• Query Equality • Response Identities + Volumes	Optimal
Property-Preserving Encryption (PPE)	• Ciphertext Equality • Ciphertext Order	Optimal

FHE: Considered secure but inefficient

ORAM and STE — Our work: Considered efficient and ???

PPE: Considered efficient but insecure [NKW15]

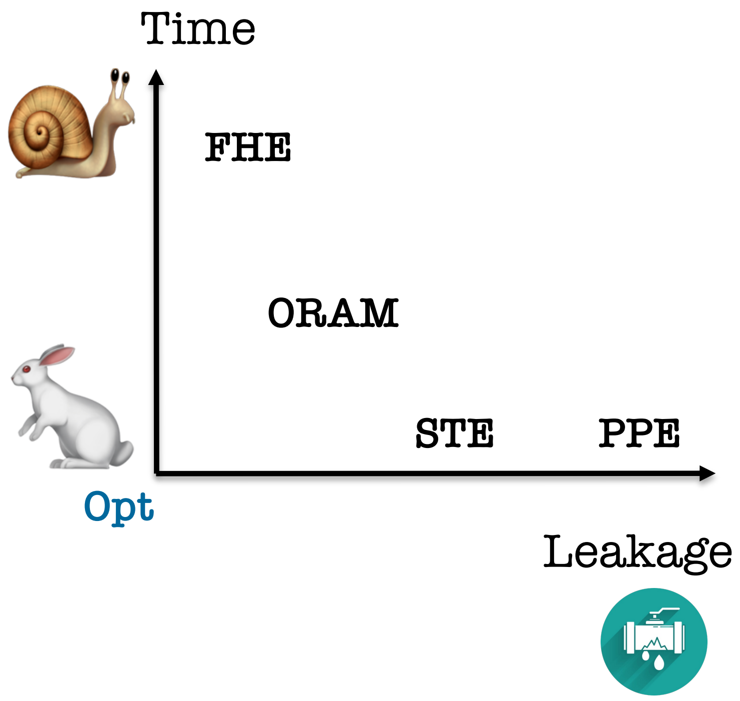

Uncertainty Of Security

Constructions

"Benign leakage"

"Common leakage"

"Standard leakage"

"Accepted leakage"

"[Attacks] assume extremely strong adversarial models"

"Leakages [...] are not exploitable via leakage-abuse attacks in practice"

Attacks & Countermeasures

"Severe threat"

"Devastating results"

"[ESAs] are extremely vulnerable to [attacks]"

"[ESA] schemes should no longer be used without countermeasures"

"Our assumptions on background information are weak"

"With some prior knowledge [...] an honest-but-curious server can recover the underlying keywords"

Uncertainty Of Security

Constructions

"Benign leakage"

"Standard leakage"

"Con"

"Acc"

"[Attacks] assume extremely stron models"

"Leakages [...] are not exploitable abuse attacks in practi"

Attacks & Countermeasures

"hreat"

"Devastating results"

"[ESAs] are extremely vulnerable to [attacks]"

"[ESA] schemes should no longer be used without countermeasures"

"ur assumptions on background information are weak"

"some prior knowledge [...] an honest-but- server can recover the underlying keywords"

A **Realistic Assessment** of Leakage Attacks on Encrypted Search

Previous Evaluations

Usual evaluations for **Keyword attacks:**

1. Enron (& Apache) email data collection
2. Restrict data to 500-3000 keywords
3. Draw 150 queries **from data collection**
→ ??? From which part of the distribution ?
4. Evaluate on **partial knowledge**

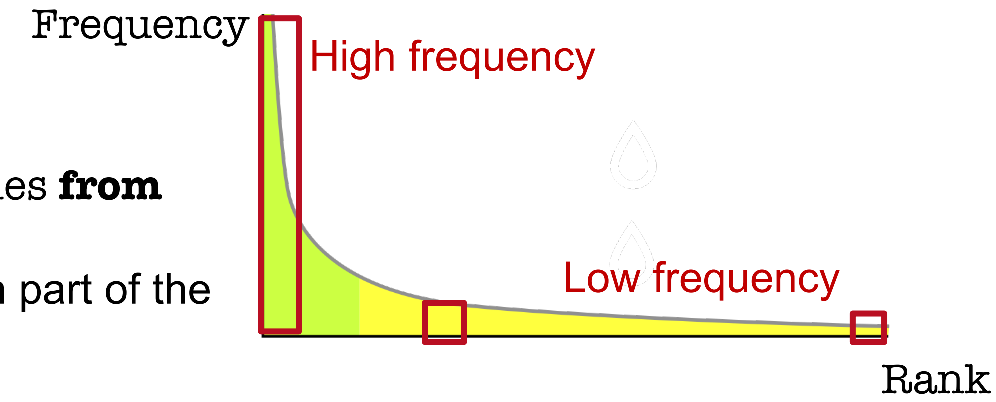

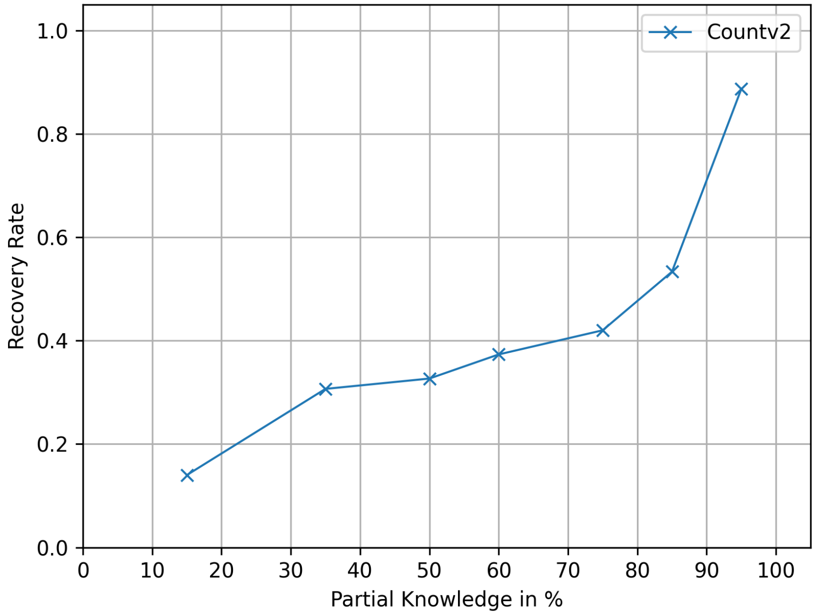

High frequency

or

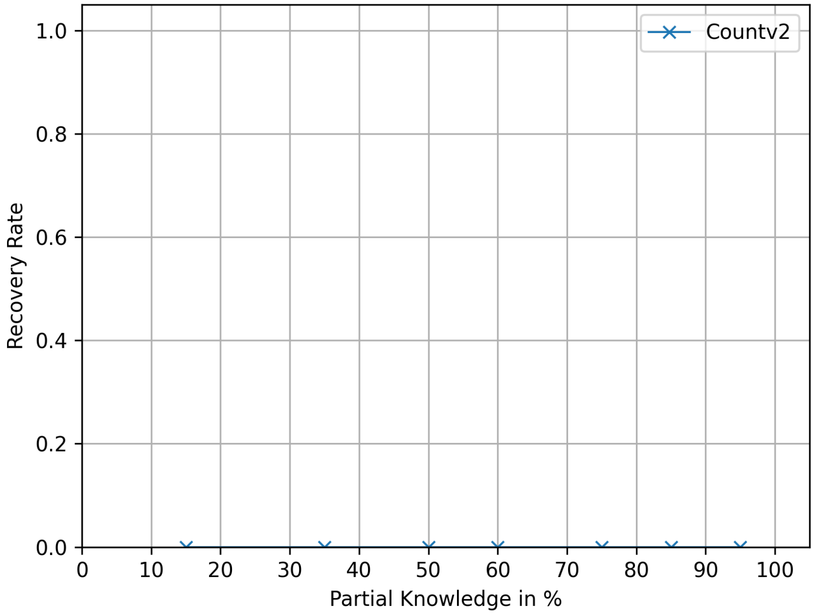

Low frequency

Previous Evaluations

Usual evaluations for **Range attacks:**

1. Subset of HCUP Data collection

2. Pick Artificial query distribution (Uniform/Zipf/...)

3. Evaluate for different amounts of queries

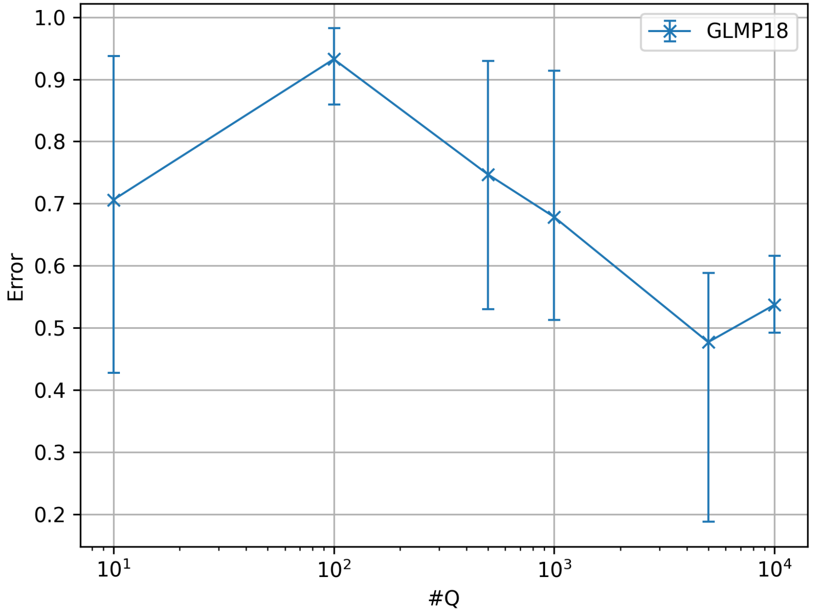

or

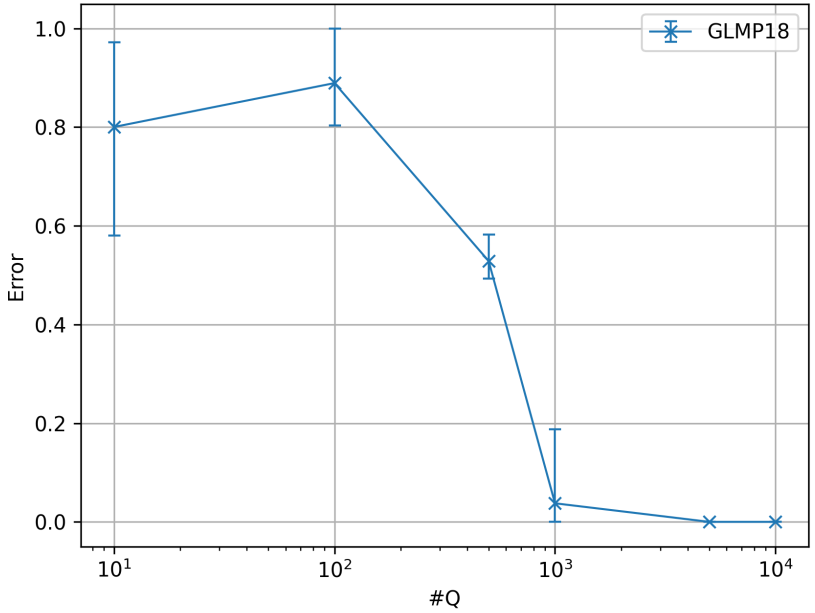

Limitations & Contributions

Limitations

- Systematization Lacking
- Single Use Case
- Few Comparisons
- Closed-Source Code

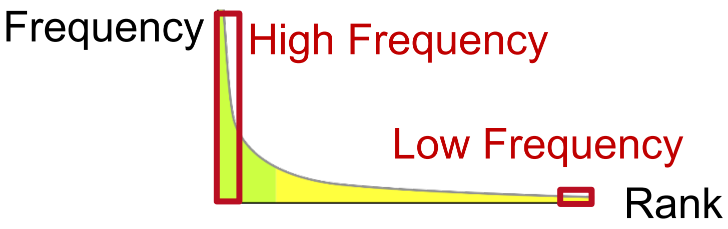

Artificial Queries

Our Contributions

- Survey of ESA Cryptanalysis
- New Attacks New Data
- Systematic Re-Evaluation
- Open-Source Framework

```
User,Query
216,'Aarhus'
216,'University'
106,'Visit'
216,'Cryptanalysis'
```

First Real-World Query Logs

LEAKER Framework

https://encrypto.de/code/LEAKER

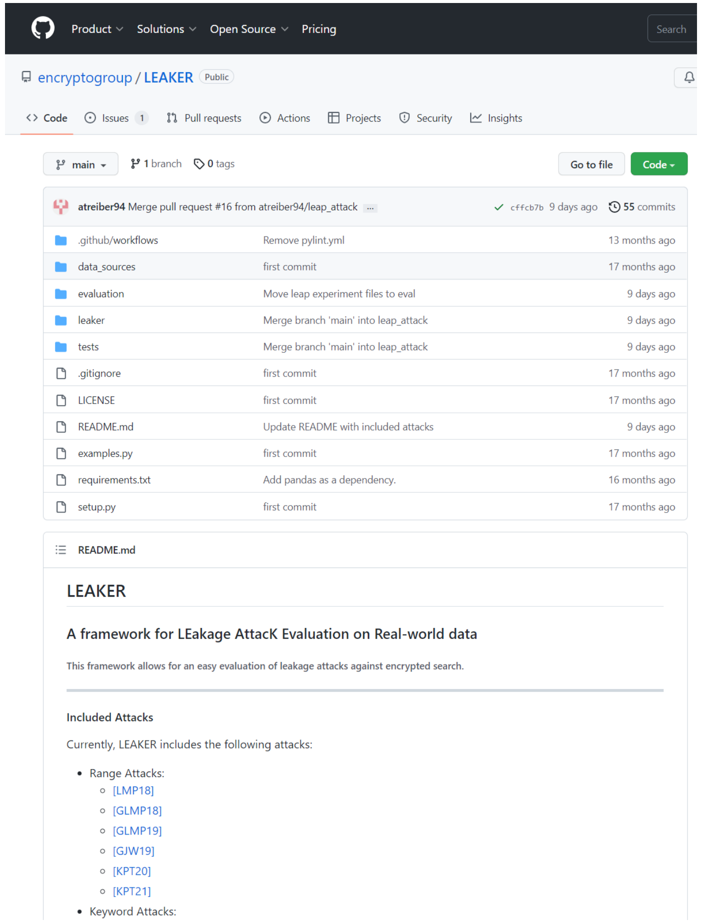

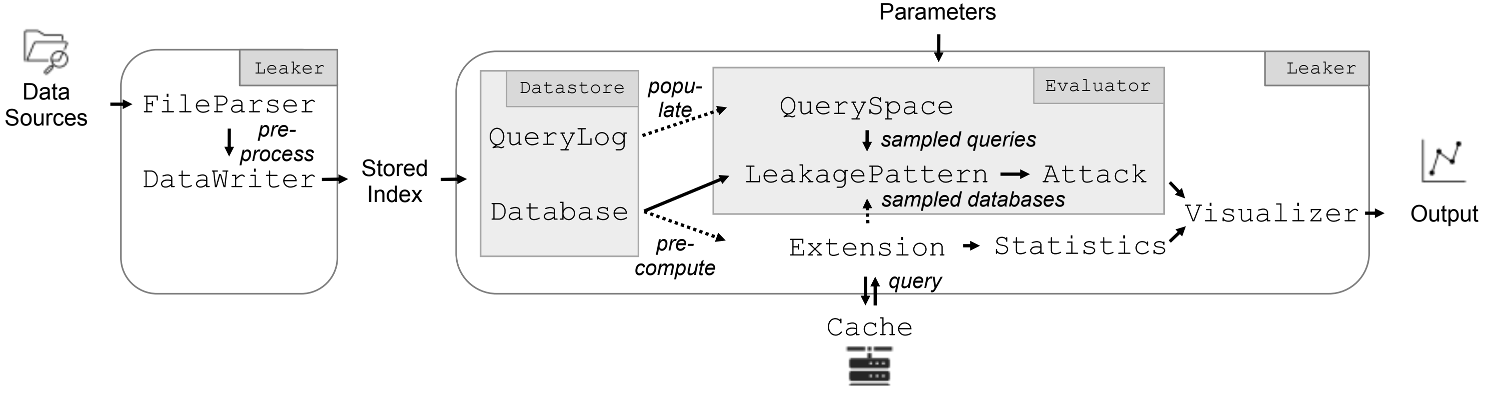

- Re-Implementation of major attacks in open-source Framework

  On Release: [IKK12, CGPR15, LMP18,GLMP18, GLMP19, GJW19, BKM20,KPT20,KPT21,RPH21]
  Since then: [KPT19,FMA+20,NHP+21,Sie22]
  In development: [OK21,DHP21,OK22,???]

- Modular design & supports interoperability
- Easy to implement new attacks & Countermeasures
- Easy to pre-process & use new data types.

Data Sources

Keyword Queries

Email/Cloud

My Google activity

7 Query Logs &
7 Data Collections
7 Users
16-100 Queries
200-47k Documents
19k-895k Keywords

Web

AOL

1 Query Log &
1 Data Collection
656k Users
2.9M Queries
151k Documents
268k Keywords

Genetic

tair

1 Query Log &
1 Data Collection
1.3k Users
54k Queries
115k Documents
690k Keywords

Range Queries

Scientific

SDSS

3 Query Logs &
1 Data Collection
3 Users
215-8k Queries
5M Records
Domain N = 10k
Density 96%

Medical

MIMIC

3 Data Collections
2k-8k Records
Domain N = 73 – 10k
Density 3.3%-81%

Salaries

DATA.GOV.UK Beta
Opening up Government

1 Data Collection
536 Records
Domain N = 395
Density 2.3%

Sales

Walmart

1 Data Collection
143 Records
Domain N = 6.3k
Density 2.3%

Insurance

NYC OpenData

1 Data Collection
886 Records
Domain N = 25k
Density 1.2%

Evaluation Summary

None of the attacks worked against low-[frequency] keywords

[BKM20]

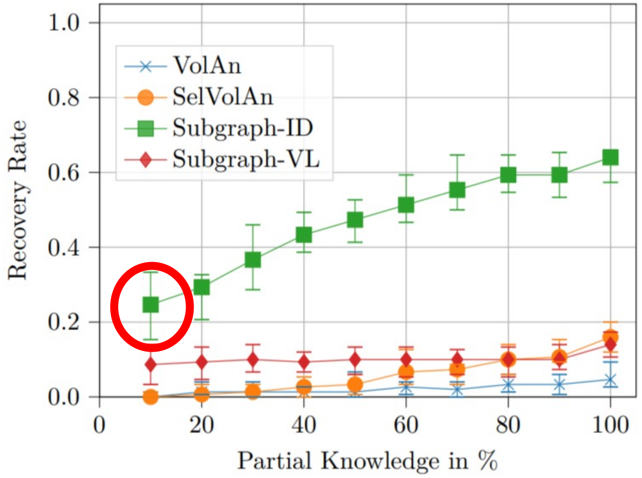

(Lowest) Mean Frequency: 1.54!

(On the TAIR Dataset)

Users are more likely **to search for a specific email**

[RPH21]

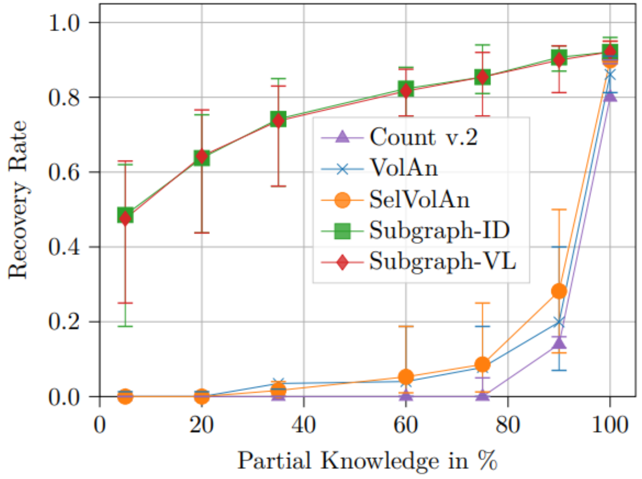

Mean Frequency:

326!

(On the Gmail Datasets)

[BKM20] L. Blackstone, S. Kamara, T. Moataz. Revisiting leakage abuse attacks. NDSS'20
[RPH21] R.G. Roessink, A. Peter, F. Hahn. Experimental review of the IKK query recovery attack: Assumptions, recovery rate and improvements. ACNS'21

Evaluation Summary (Keyword Search)

(subjective)

Attacks	Leakage	Success Cases	Risk
• VolAn [BKM20] • SelVolAn [BKM20]	• Response length • Response volume	• High adversarial knowledge	Low
• [IKK12] • Count V.2 [CGPR15] • DetIKK [RPH21]	• Co-occurrence	• High adversarial knowledge	Low
• SubgraphID [BKM20] • SubgraphVL [BKM20]	• Response identities • Response volumes	• Low adversarial knowledge	High

=> Suppression of identifier and volume leakage of responses necessary!

Evaluation Summary (Keyword Search)

AOL
single user & low frequency

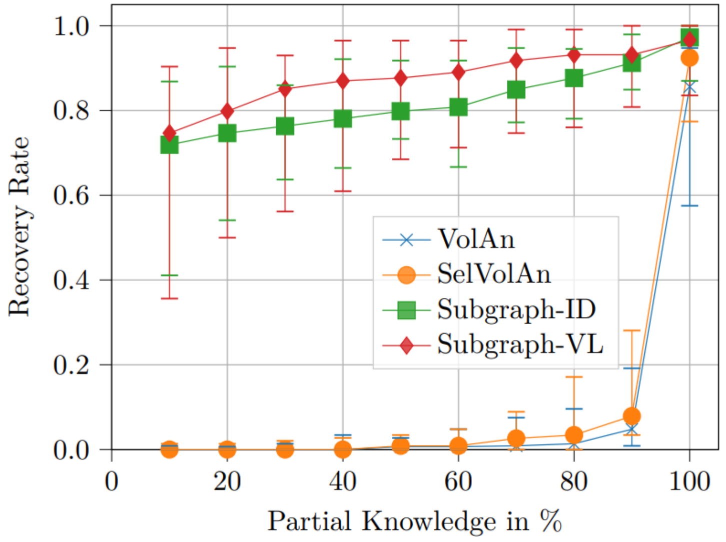

including queries outside of partial knowledge

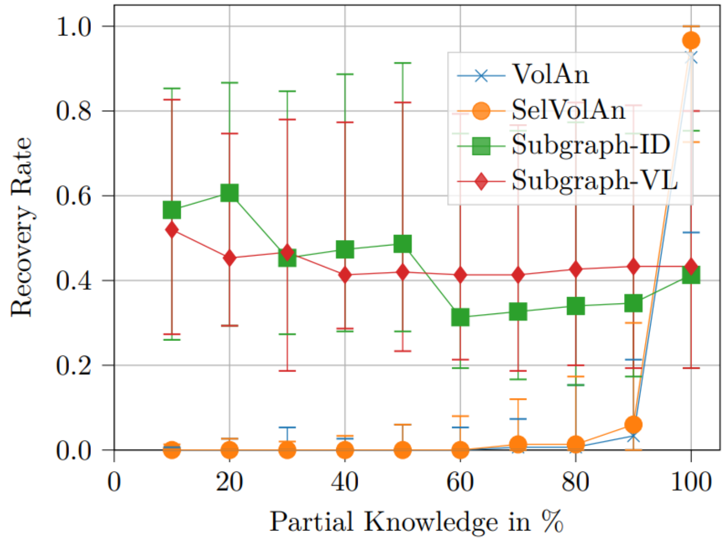

with repeating queries

AOL
single user & high frequency

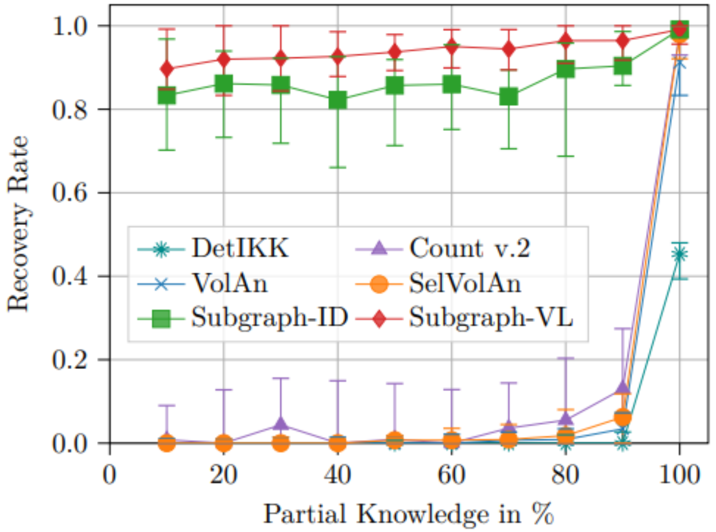

Evaluation Summary (Range Search)

(subjective)

Attacks	Leakage	Success Cases	Risk
• [GLMP18] • [GJW19]	• Response length	• None	Very low
• APA [KPT21]	• Response length • Query equality	• Evenly distributed data	Medium
• [LMP18]	• Response identities	• Dense	Medium
• GenKNNO [GLMP19] • ApprValue [GLMP19] • ARR [KPT20] + ApprOrder [GLMP19]	• Response identities	• Large widths • Skewed values	Medium
• ARR [KPT20]	• Response identities • Order	• Most cases	High

=> Research on order reconstruction + Leakage suppression for range case!

Nuanced highlights given LEAKER

[BKM20] attacks on identifier or volume leakage work much better than previously anticipated

[IKK12,CGPR15] keyword attacks perform much worse than previously anticipated

Range attacks rarely work on our data and success highly depends on query/data distributions

[OK22] attacks recovery rate given a specific leakage profile highly depends on prior assumption over query/data

ESA cryptanalysis is very nuanced

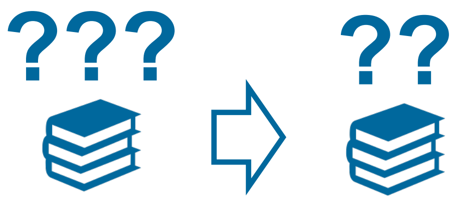

[BKM20] L. Blackstone, S. Kamara, T. Moataz. Revisiting leakage abuse attacks. NDSS'20
[IKK12] M. S. Islam, M. Kuzu, M. Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. NDSS'12
[CGPR15] D. Cash, P. Grubbs, J. Perry, T. Ristenpart. Leakage-abuse attacks against searchable encryption. CCS'15
[OK22] S. Oya and F. Kerschbaum. IHOP: Improved Statistical Query Recovery against Searchable Symmetric Encryption through Quadratic Optimization. USENIX'22

Statistical query recovery attacks

Statistical-based query recovery attacks achieve [lower] accuracy and are [not] considered a serious threat.

[OK22]

Auxiliary Information

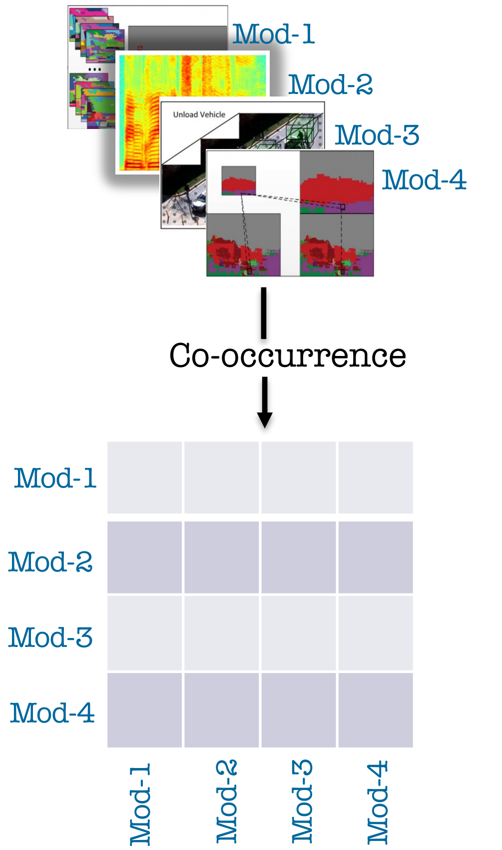

Observations

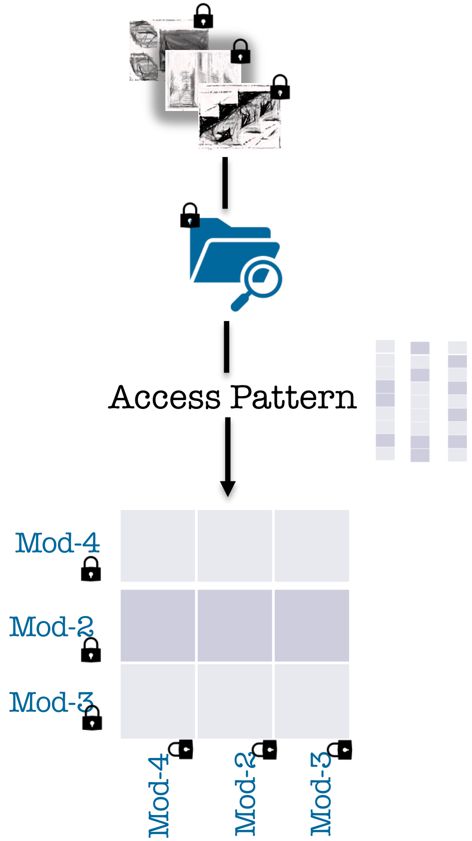

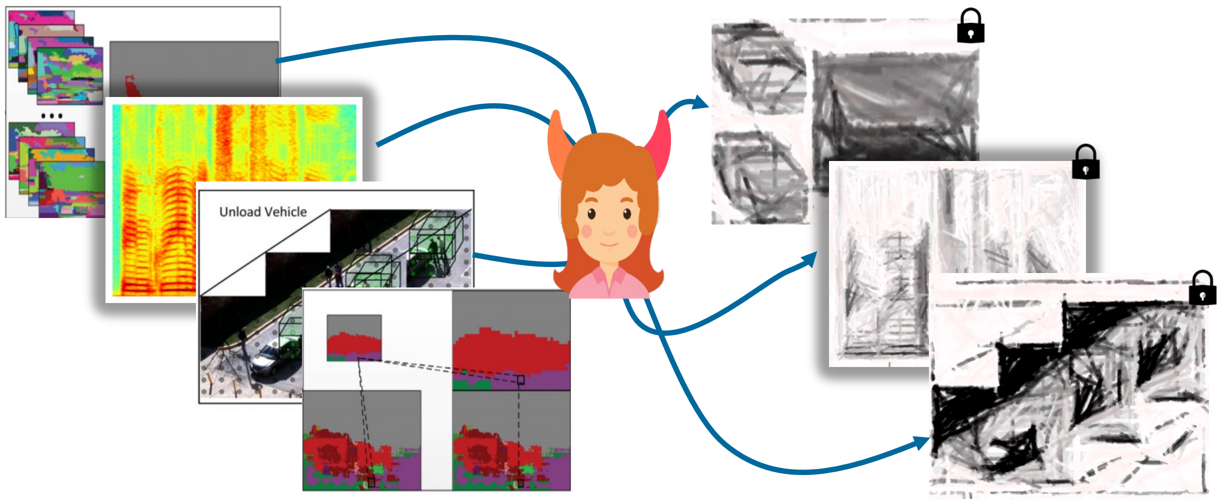

Statistical query recovery attacks

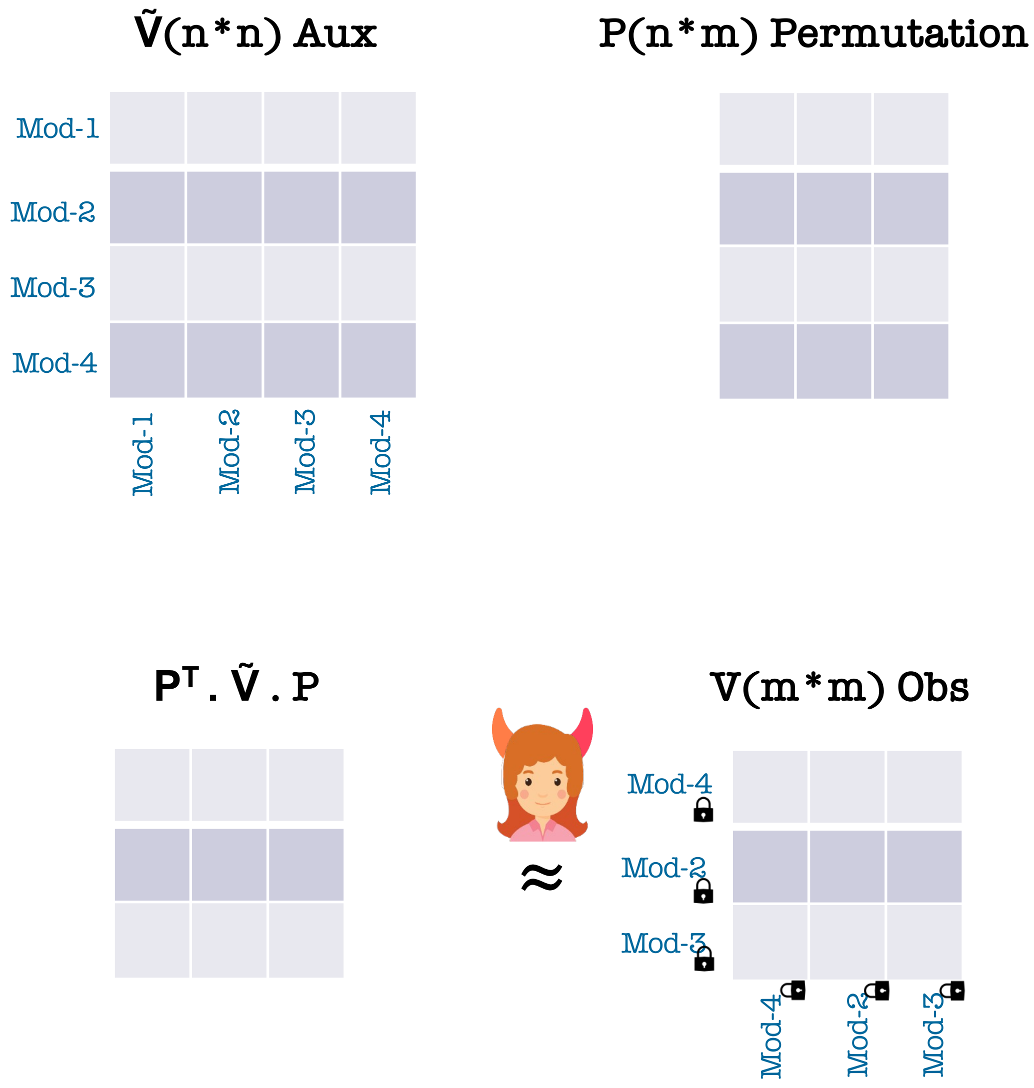

The query recovery is formulated as a quadratic assignment problem and iteratively solved via linear assignments.

[OK22]

Q.A.P

$$\mathbf{P} = \arg\min_{\mathbf{P} \in \mathcal{P}} \sum_{i,i' \in [n]} \sum_{j,j' \in [m]} c_{i,i',j,j'} \cdot \mathbf{P}_{i,j} \cdot \mathbf{P}_{i',j'}$$

Examples:

- IKK : $P = \text{argmin} \, || \tilde{V} - P^T . \tilde{V} . P ||_2$
  --> simulated annealing
- graphM : $P = \text{argmin} \, || \tilde{V} - P^T . \tilde{V} . P ||^2_2 - tr(CP)$
  --> convex-concave rel.

[IKK] Islam et .al. Access pattern disclosure on searchable encryption: ramifications, attacks and mitigation. NDSS12.
[graphM] Pouliot and wright. The shadow nemesis: inference attacks on efficiently deployable, efficiently searchable encryption. CCS16.

Statistical query recovery attacks

$\tilde{V}$(n*n) Aux

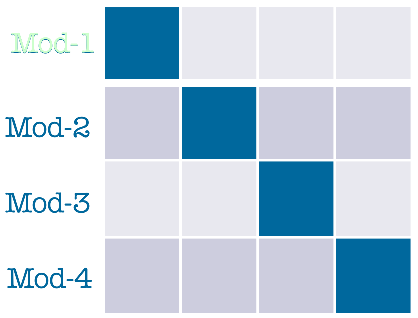

P(n*m) Permutation

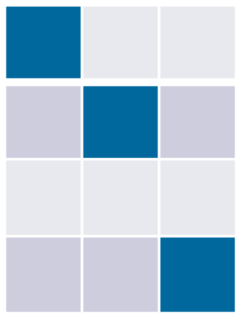

$$\mathbf{P} = \underset{\mathbf{P} \in \mathcal{P}}{\arg\min} \sum_{i \in [n]} \sum_{j \in [m]} c_{i,j} \cdot \mathbf{P}_{i,j} .$$

L.A.P

This very efficient, but a lot of information is wasted because of not using the off-diagonal terms.

0,5	0,51	0,18	0,2	
0,01	0,02	0,31	0,29	0,49
0,01	0	0,33	0,31	0,51
0,3	0,31	0,02	0	0,2

Hungarian algorithm

V(m*m) Obs

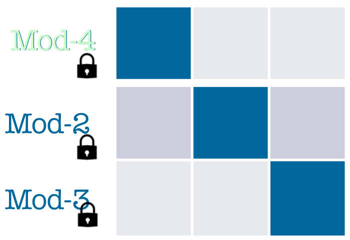

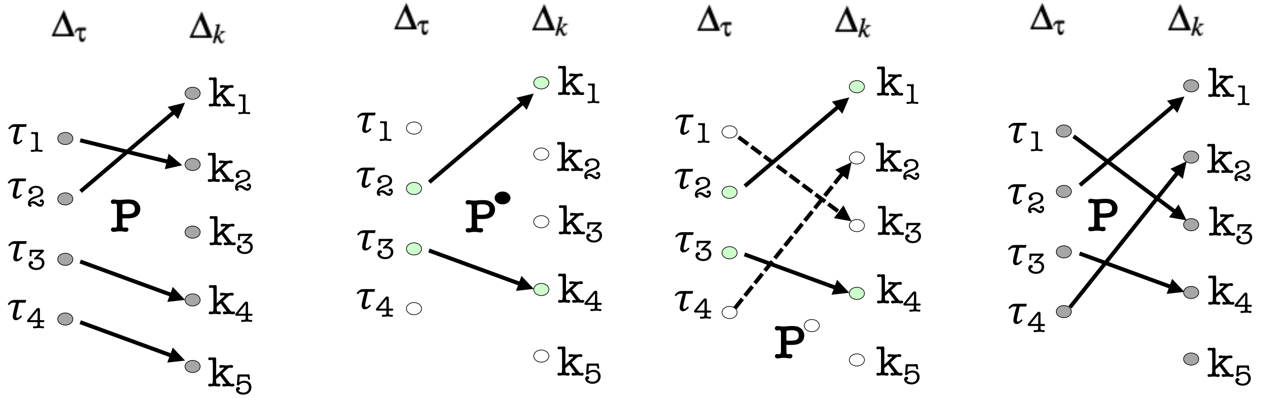

Full Assignment | Select fixed mappings | Compute mapping for free tokens L.A.P/Freeze | After iteration Solve L.A.P

$\Delta_\tau^\circ = \{\tau_1, \tau_4\}$ $\Delta_\tau^\bullet = \{\tau_2, \tau_3\}$ $\Delta_k^\circ = \{k_2, k_3, k_5\}$ $\Delta_k^\bullet = \{k_1, k_4\}$

Statistical query recovery attacks

Adversary can exploit Qeq in the dependent setting where the client's queries are correlated, even when access obfuscation defenses are applied.

[OK22]

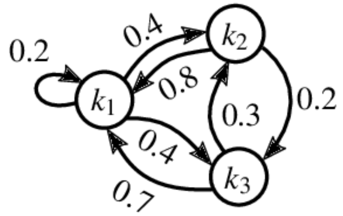

Markov Model

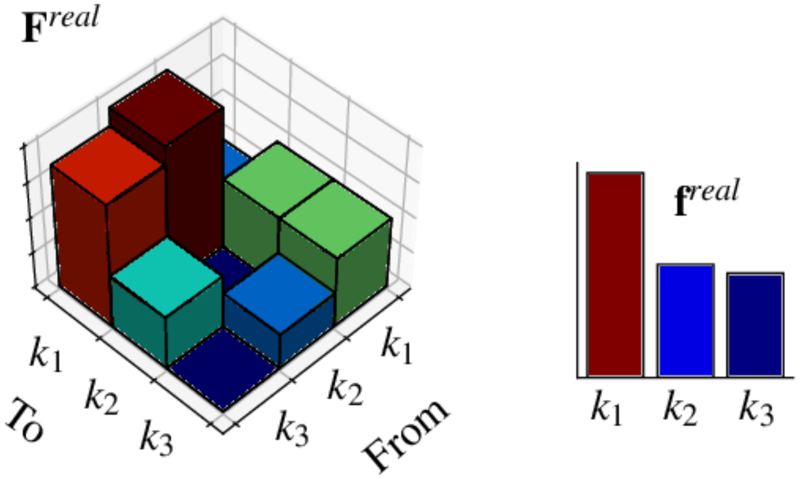

Markov matrix ($\mathbf{F}^{\text{real}}$) and its stationary distribution ($\mathbf{f}^{\text{real}}$) of the queried keywords.

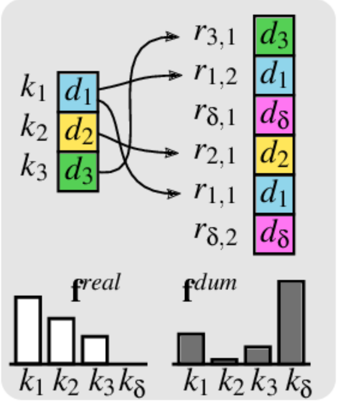

PANCAKE setup.

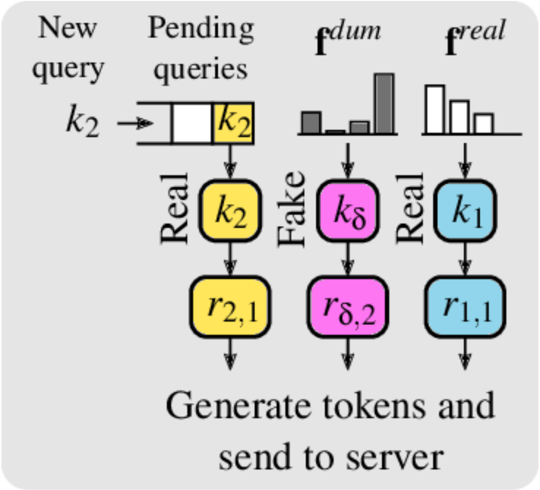

PANCAKE query.

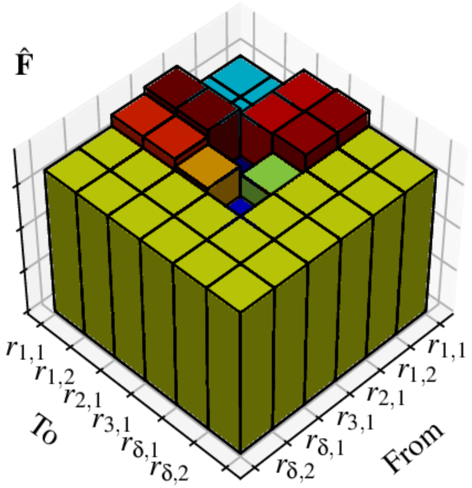

Markov matrix ($\hat{F}$) of the queried replicas by following PANCAKE protocol.

Statistical query recovery attacks

Sorting

Step 1 :
- Initializes an empty mapping

Step 2 :
- Computes the stationary distribution π,

Step 3 :
- Calculate the histogram of the sequence of queries v.
  - ≈ to the average number of visits over the M.C states)

Step 4 :
- Map the closest value in π to vi, for all $i \in [t]$;
  - the average number of visits to the i^{th} state is approximately equal to the i^{th} component of the stationary distribution π.

Step 5:
- output the mapping and the error score
  - Error: the total distance between the avg.# visits and the selected component of the stationary distribution

Decoding

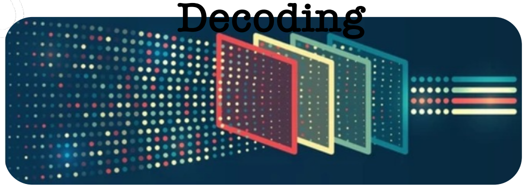

Step 1 :
- Initializes an empty mapping

Step 2 :
- Computes the Observation matrix of HMM O=$(o_{i,j})$ $i \in [m], j \in [\#I]$,
  - The frequency fj, of each unique query $j \in [\#I]$, is first calculated using query equality leakage.
  - Set $o_{i,j}$ to $1-|f_j-\pi_i|$ i.f.f $|f_j-\pi_i|_1 < \epsilon$, error parameter.
  - Normalize O, s.t the sum of each row is equal to 1.

Step 3:
- Compute transition matrix P^A and a uniform initial distribution μ to form HMM parameters $\Theta := (P^A, O, \mu)$.

Step 4 :
- (Mapping α the attacked query sequence to the state identifiers of unique queries via the equality leakage, the likelihood s of this mapping given the observation) ← Viterbi .
  - Generate a sequence of observed states that matches the set of observation states of the created HMM parameters

Step 5 :
- A new map α'translates the states α maps to actual keywords using the adversary's knowledge.
  - error parameter, we set $s'=1-s$ such that the result with the maximum likelihood will correspond to the lowest score.

Evaluation results (R.W Q-log)

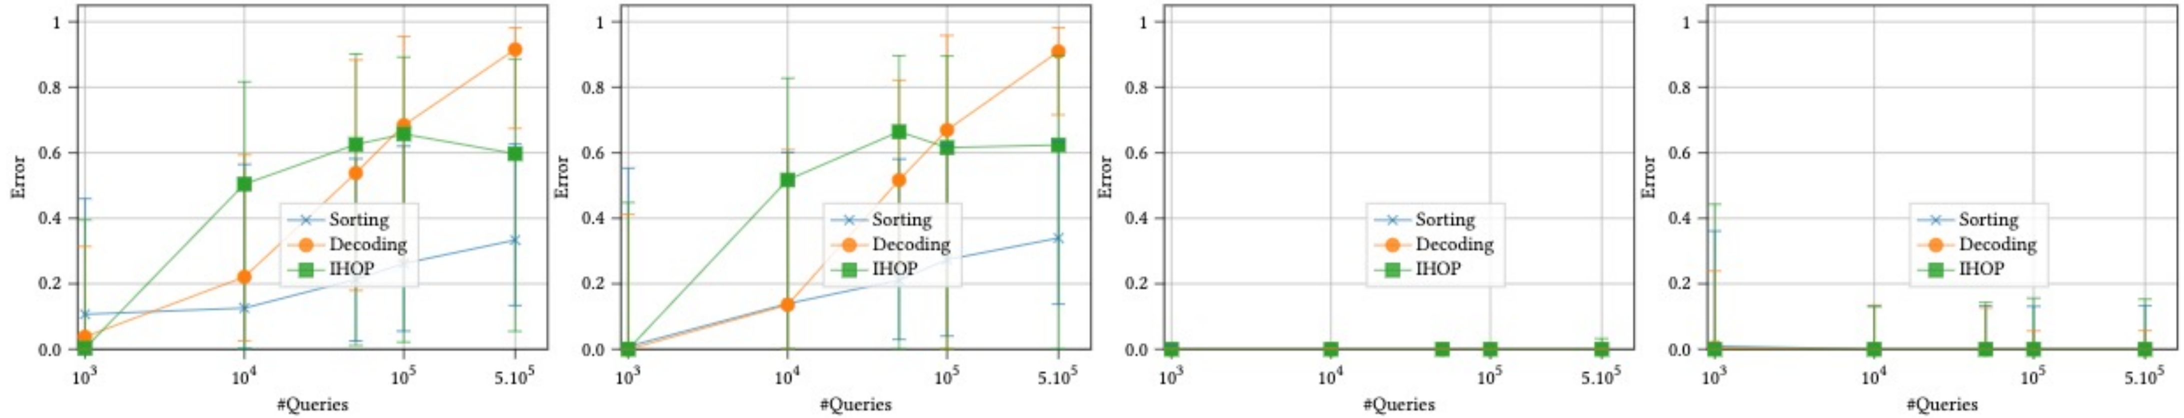

(E) (A) (O) (S)

Evaluation for each of 5 users on AOL

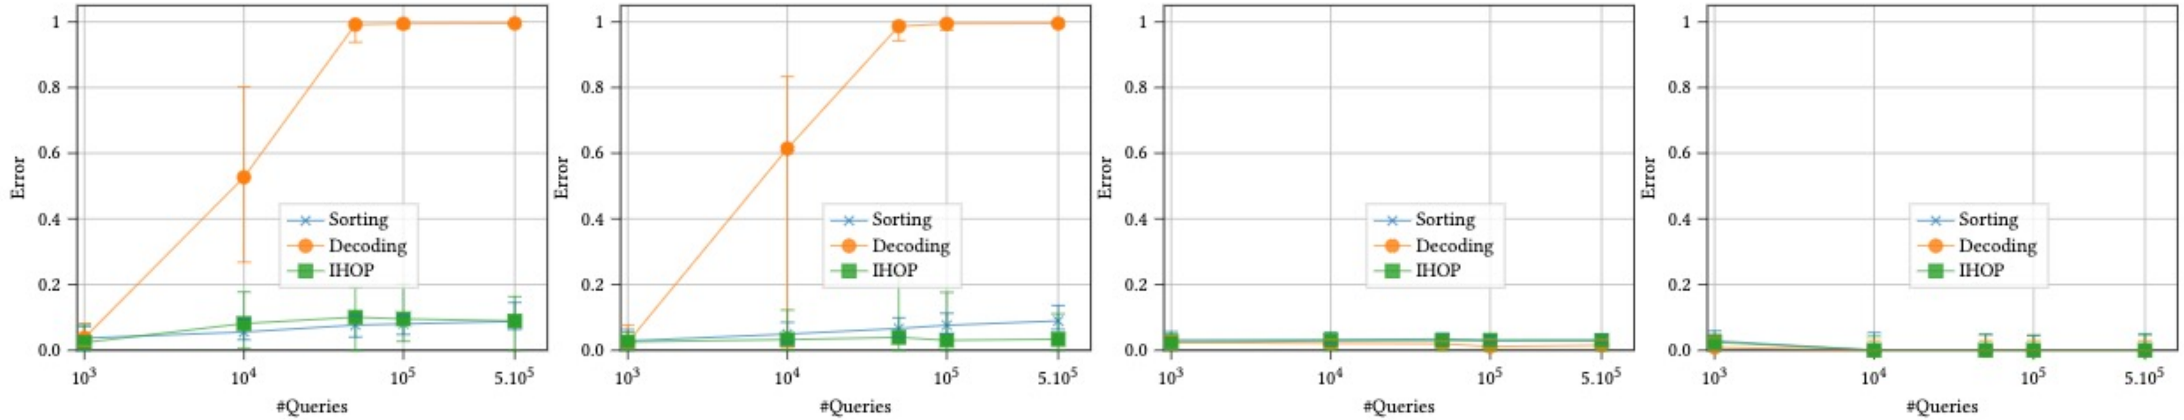

(E) (A) (O) (S)

Evaluation for each of 5 users on TAIR

Evaluation results (Art.Distributions)

Evaluation for Zipf-Zipf Artificial distribution with fixed H-W

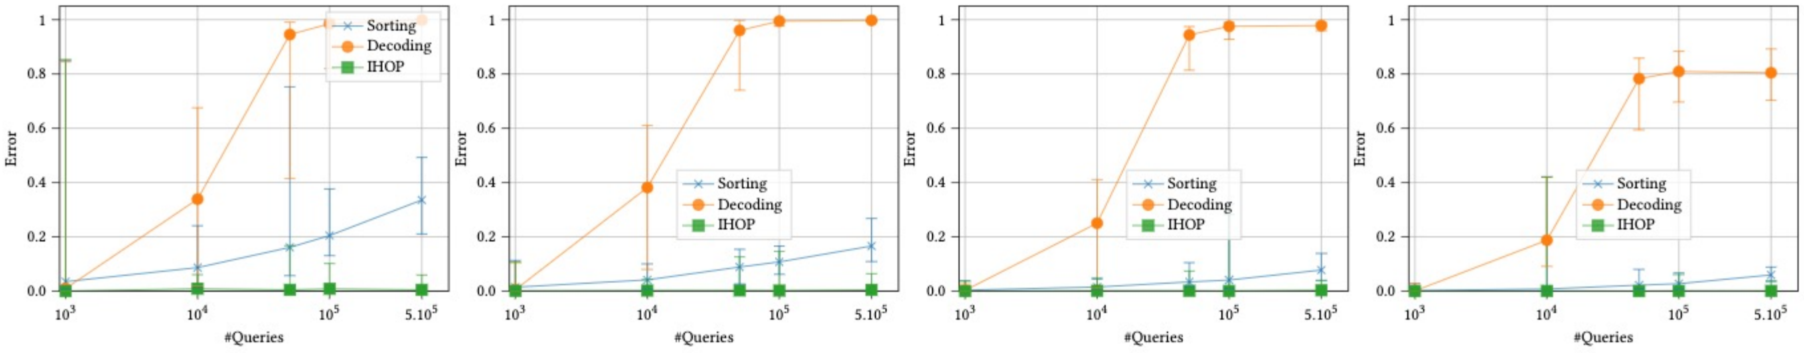

250 States | 500 States | 1 000 States | 1 500 States

Evaluation for Zipf-Zipf Artificial distribution with variable H-W

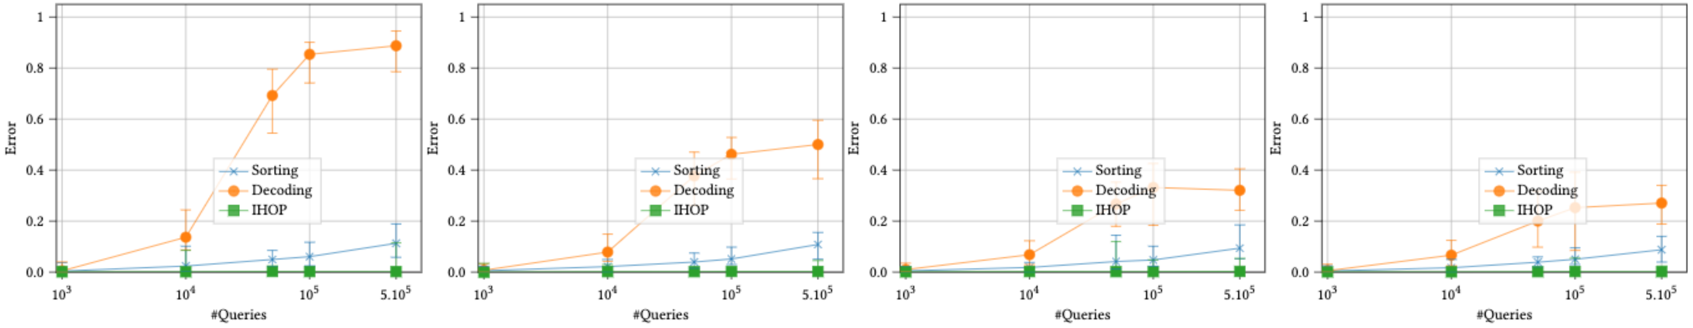

H-W = 5 | H-W = 10 | H-W = 15 | H-W = 20

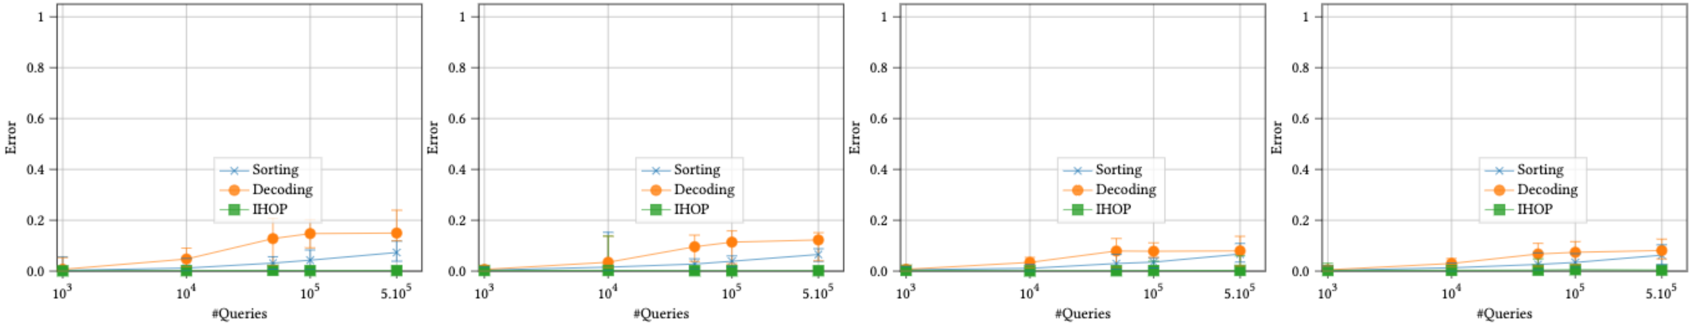

H-W = 50 | H-W = 100 | H-W = 350 | H-W = 450

Evaluation results (Art.Distributions)

Evaluation for *Erdos* Artificial Distribution.

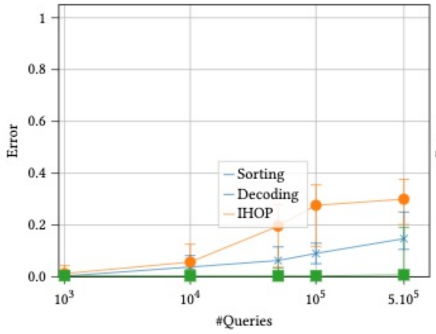

250 States

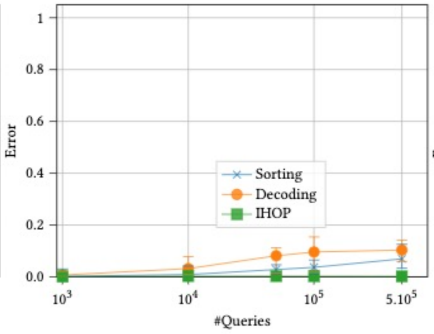

500 States

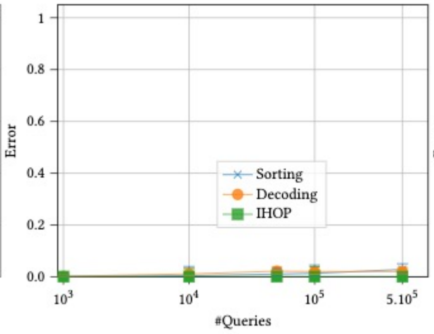

1 000 States

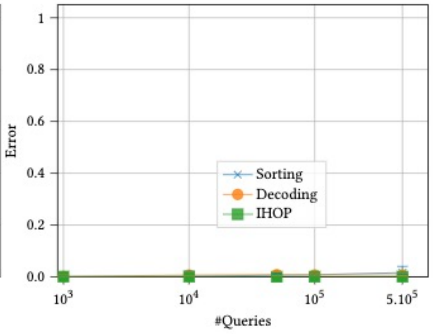

1 500 States

Evaluation for *Uniform* Artificial distribution.

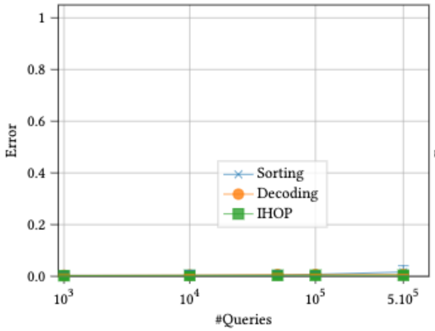

250 States

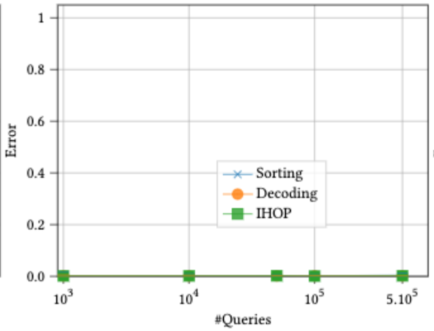

500 States

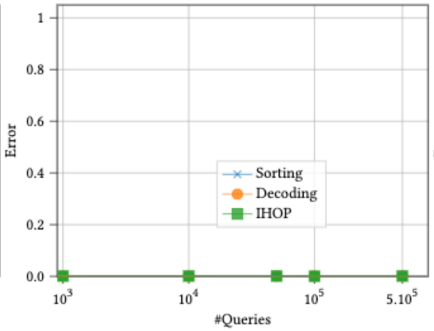

1 000 States

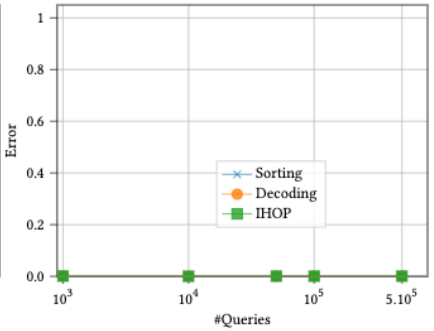

1 500 States

Evaluation for *Zipf* Artificial distribution.

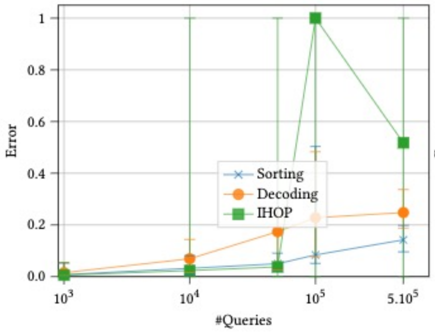

250 States

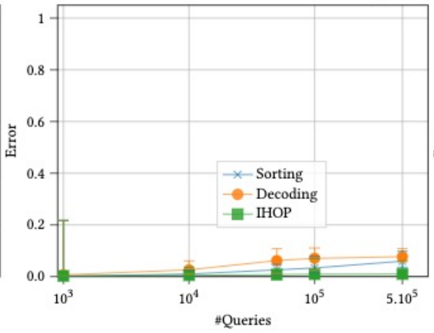

500 States

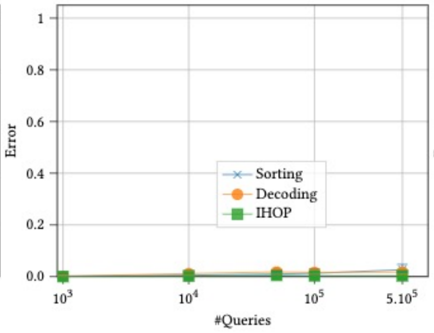

1 000 States

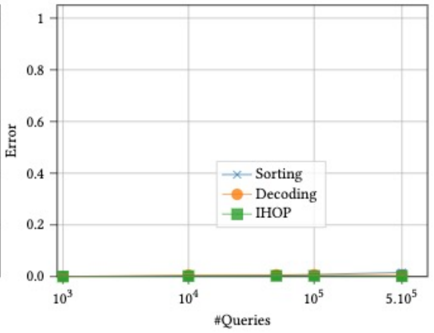

1 500 States

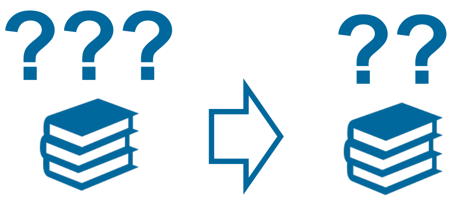

Thank you
for your attention

Cryptanalysis Strikes Back A Realistic assessment of leakage attacks on Encrypted Search

Abdelkarim Kati†‡
together with T. Moataz, S. Kamara and A. Treiber.

†School of Computer Science,
Mohammed VI Polytechnic University.
‡ Encrypted Systems Lab, Brown University.

January 24, 2023 at Aarhus University.

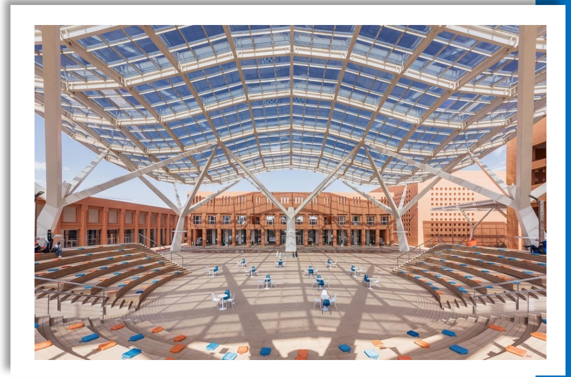

Viterbi Algorithm (Uncovering Problem)

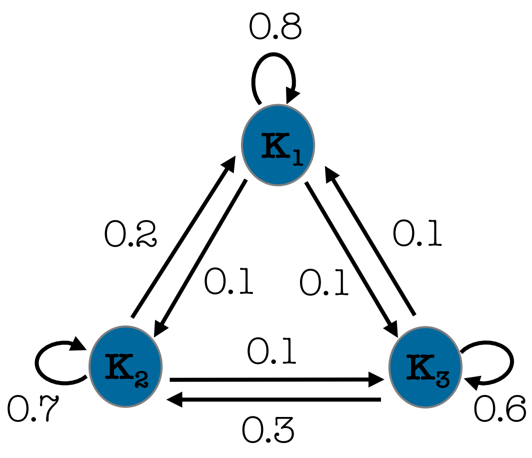

States
(Hidden)

Observations
(Visible)

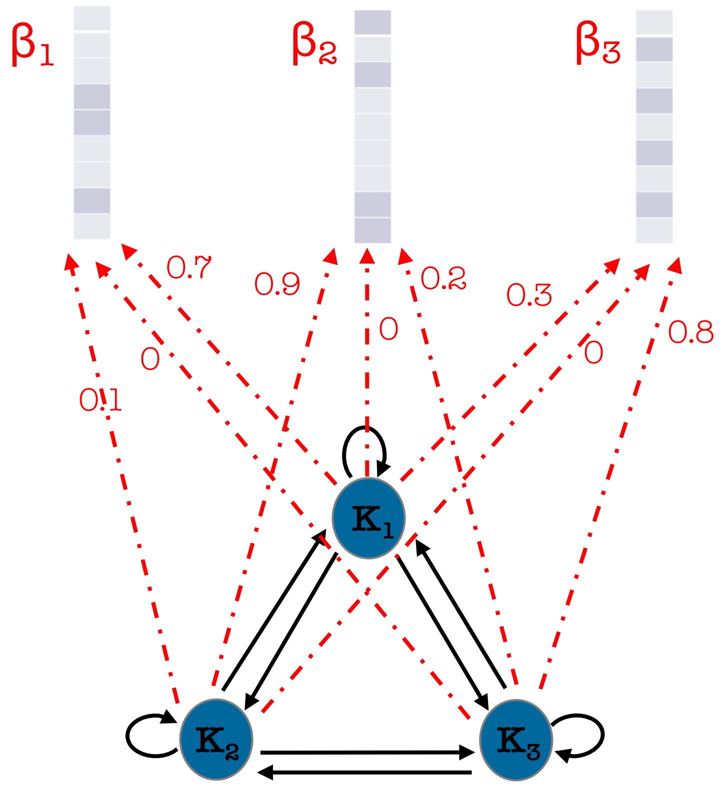

	K_1	K_2	K_3
K_1	0.8	0.1	0.1
K_2	0.2	0.7	0.1
K_3	0.1	0.3	0.6

State transition probabilities

	K_1	K_2	K_3
	0.6	0.2	0.2

Initial state probabilities

	β_1	β_2	β_3
K_1	0.7	0	0.3
K_2	0.1	0.9	0
K_3	0	0.2	0.8

Emission probabilities

* MAPLE: Markov Process Leakage attacks on Encrypted search (under submission)

Viterbi Algorithm (Uncovering Problem)

Input

Viterbi

Output

Observation Sequence

$O = (o_1, o_2, o_3, o_4, o_5, o_6)$

	$o_1=\beta_1$	$o_2=\beta_3$	$o_3=\beta_1$	$o_4=\beta_3$	$o_5=\beta_3$	$o_6=\beta_2$
K_1	0.4200	0.1008	0.0564	0.0135	0.0033	0
K_2	0.200	0	0.0010	0	0	0.0006
K_3	0	0.0336	0	0.0045	0.0022	0.0003

Accumulated probability matrix

	$o_1=\beta_1$	$o_2=\beta_3$	$o_3=\beta_1$	$o_4=\beta_3$	$o_5=\beta_3$
K_1	1	1	1	1	1
K_2	1	1	1	1	3
K_3	1	3	1	3	3

$I_6 = 2$

Backtracking matrix

Observation Sequence

$S^* = (K_1, K_1, K_1, K_3, K_3, K_2)$

Baum-Welch Algorithm (Estimation Problem)

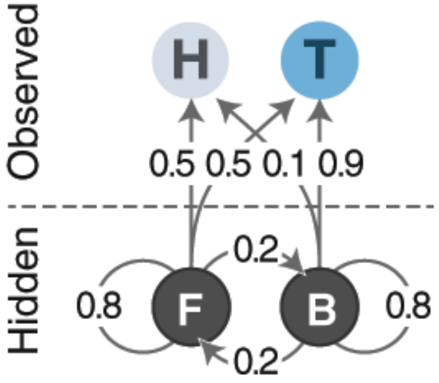

Hidden Markov Model
of an unstable coin

Ground truth | Initial estimates

$$E = \begin{array}{c} \\ F \\ B \end{array}\begin{array}{c}\begin{array}{cc} H & T \end{array}\\ \begin{bmatrix} 0.5 & 0.5 \\ 0.1 & 0.9 \end{bmatrix}\end{array} \qquad \hat{E}_0 = \begin{array}{c} \\ F \\ B \end{array}\begin{array}{c}\begin{array}{cc} H & T \end{array}\\ \begin{bmatrix} 0.5 & 0.5 \\ 0.3 & 0.7 \end{bmatrix}\end{array}$$

$$T = \begin{array}{c} \\ F \\ B \end{array}\begin{array}{c}\begin{array}{cc} F & B \end{array}\\ \begin{bmatrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{bmatrix}\end{array} \qquad \hat{T}_0 = \begin{array}{c} \\ F \\ B \end{array}\begin{array}{c}\begin{array}{cc} F & B \end{array}\\ \begin{bmatrix} 0.6 & 0.4 \\ 0.4 & 0.6 \end{bmatrix}\end{array}$$

HMM true parameters
And initial estimations

Forward probability

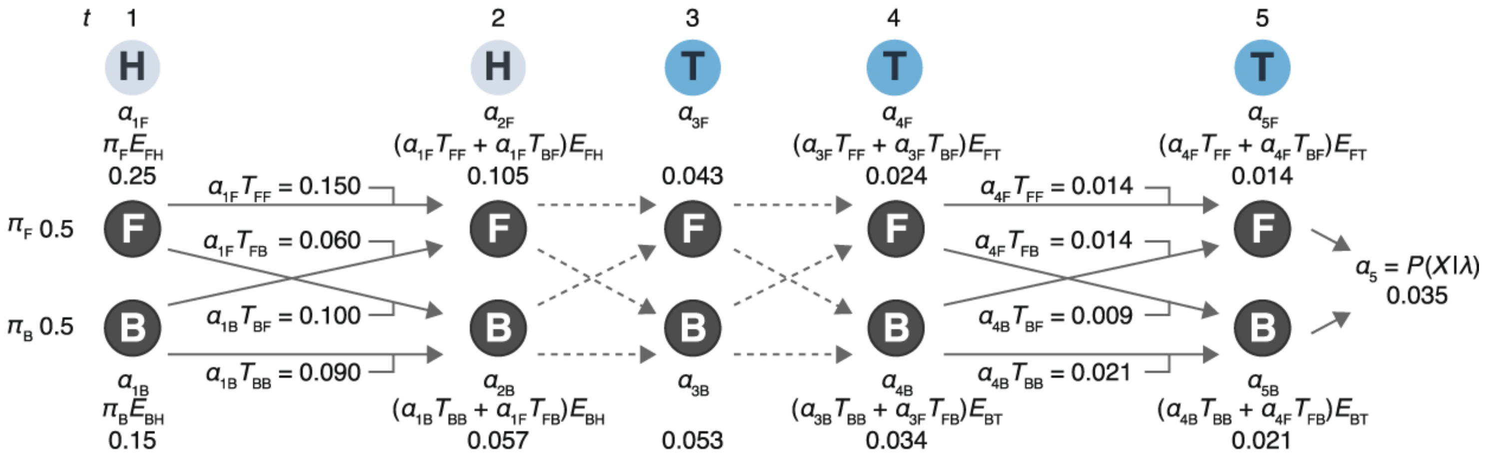

Backward probability