

Probabilistic
data structures

in Adversarial
Settings

Mia Filić ETH Zürich

Probabilistic data structures

in Adversarial Settings

Based on joint work with

Anu Unnikrishnan

Kenny Paterson

Fernando Virdia

ETH Zürich

Universidade
Nova de Lisboa

Sam A. Markelon

Thomas Shrimpton

Jonas Hofmann

University of
Florida

University of
Florida

Ella Kummer
Keran Kocher
Andrea Raguso

Probabilistic Data Structures (PDS)

A way to

compactly represent
(stream of) data

and

provide approximate
answers to queries
about the data

Probabilistic Data Structures (PDS)

A way to

compactly represent
(stream of) data

and

provide approximate
answers to queries
about the data

- Frequency estimation
How many times does x appear in the set?
Count-min sketch, HeavyKeeper

Probabilistic Data Structures (PDS)

A way to

compactly represent
(stream of) data

and

provide approximate
answers to queries
about the data

- Frequency estimation
How many times does x appear in the set?
Count-min sketch, HeavyKeeper
- Membership queries
Is x in the set?
Bloom filter, Cuckoo filter

Probabilistic Data Structures (PDS)

A way to


compactly represent
(stream of) data

and

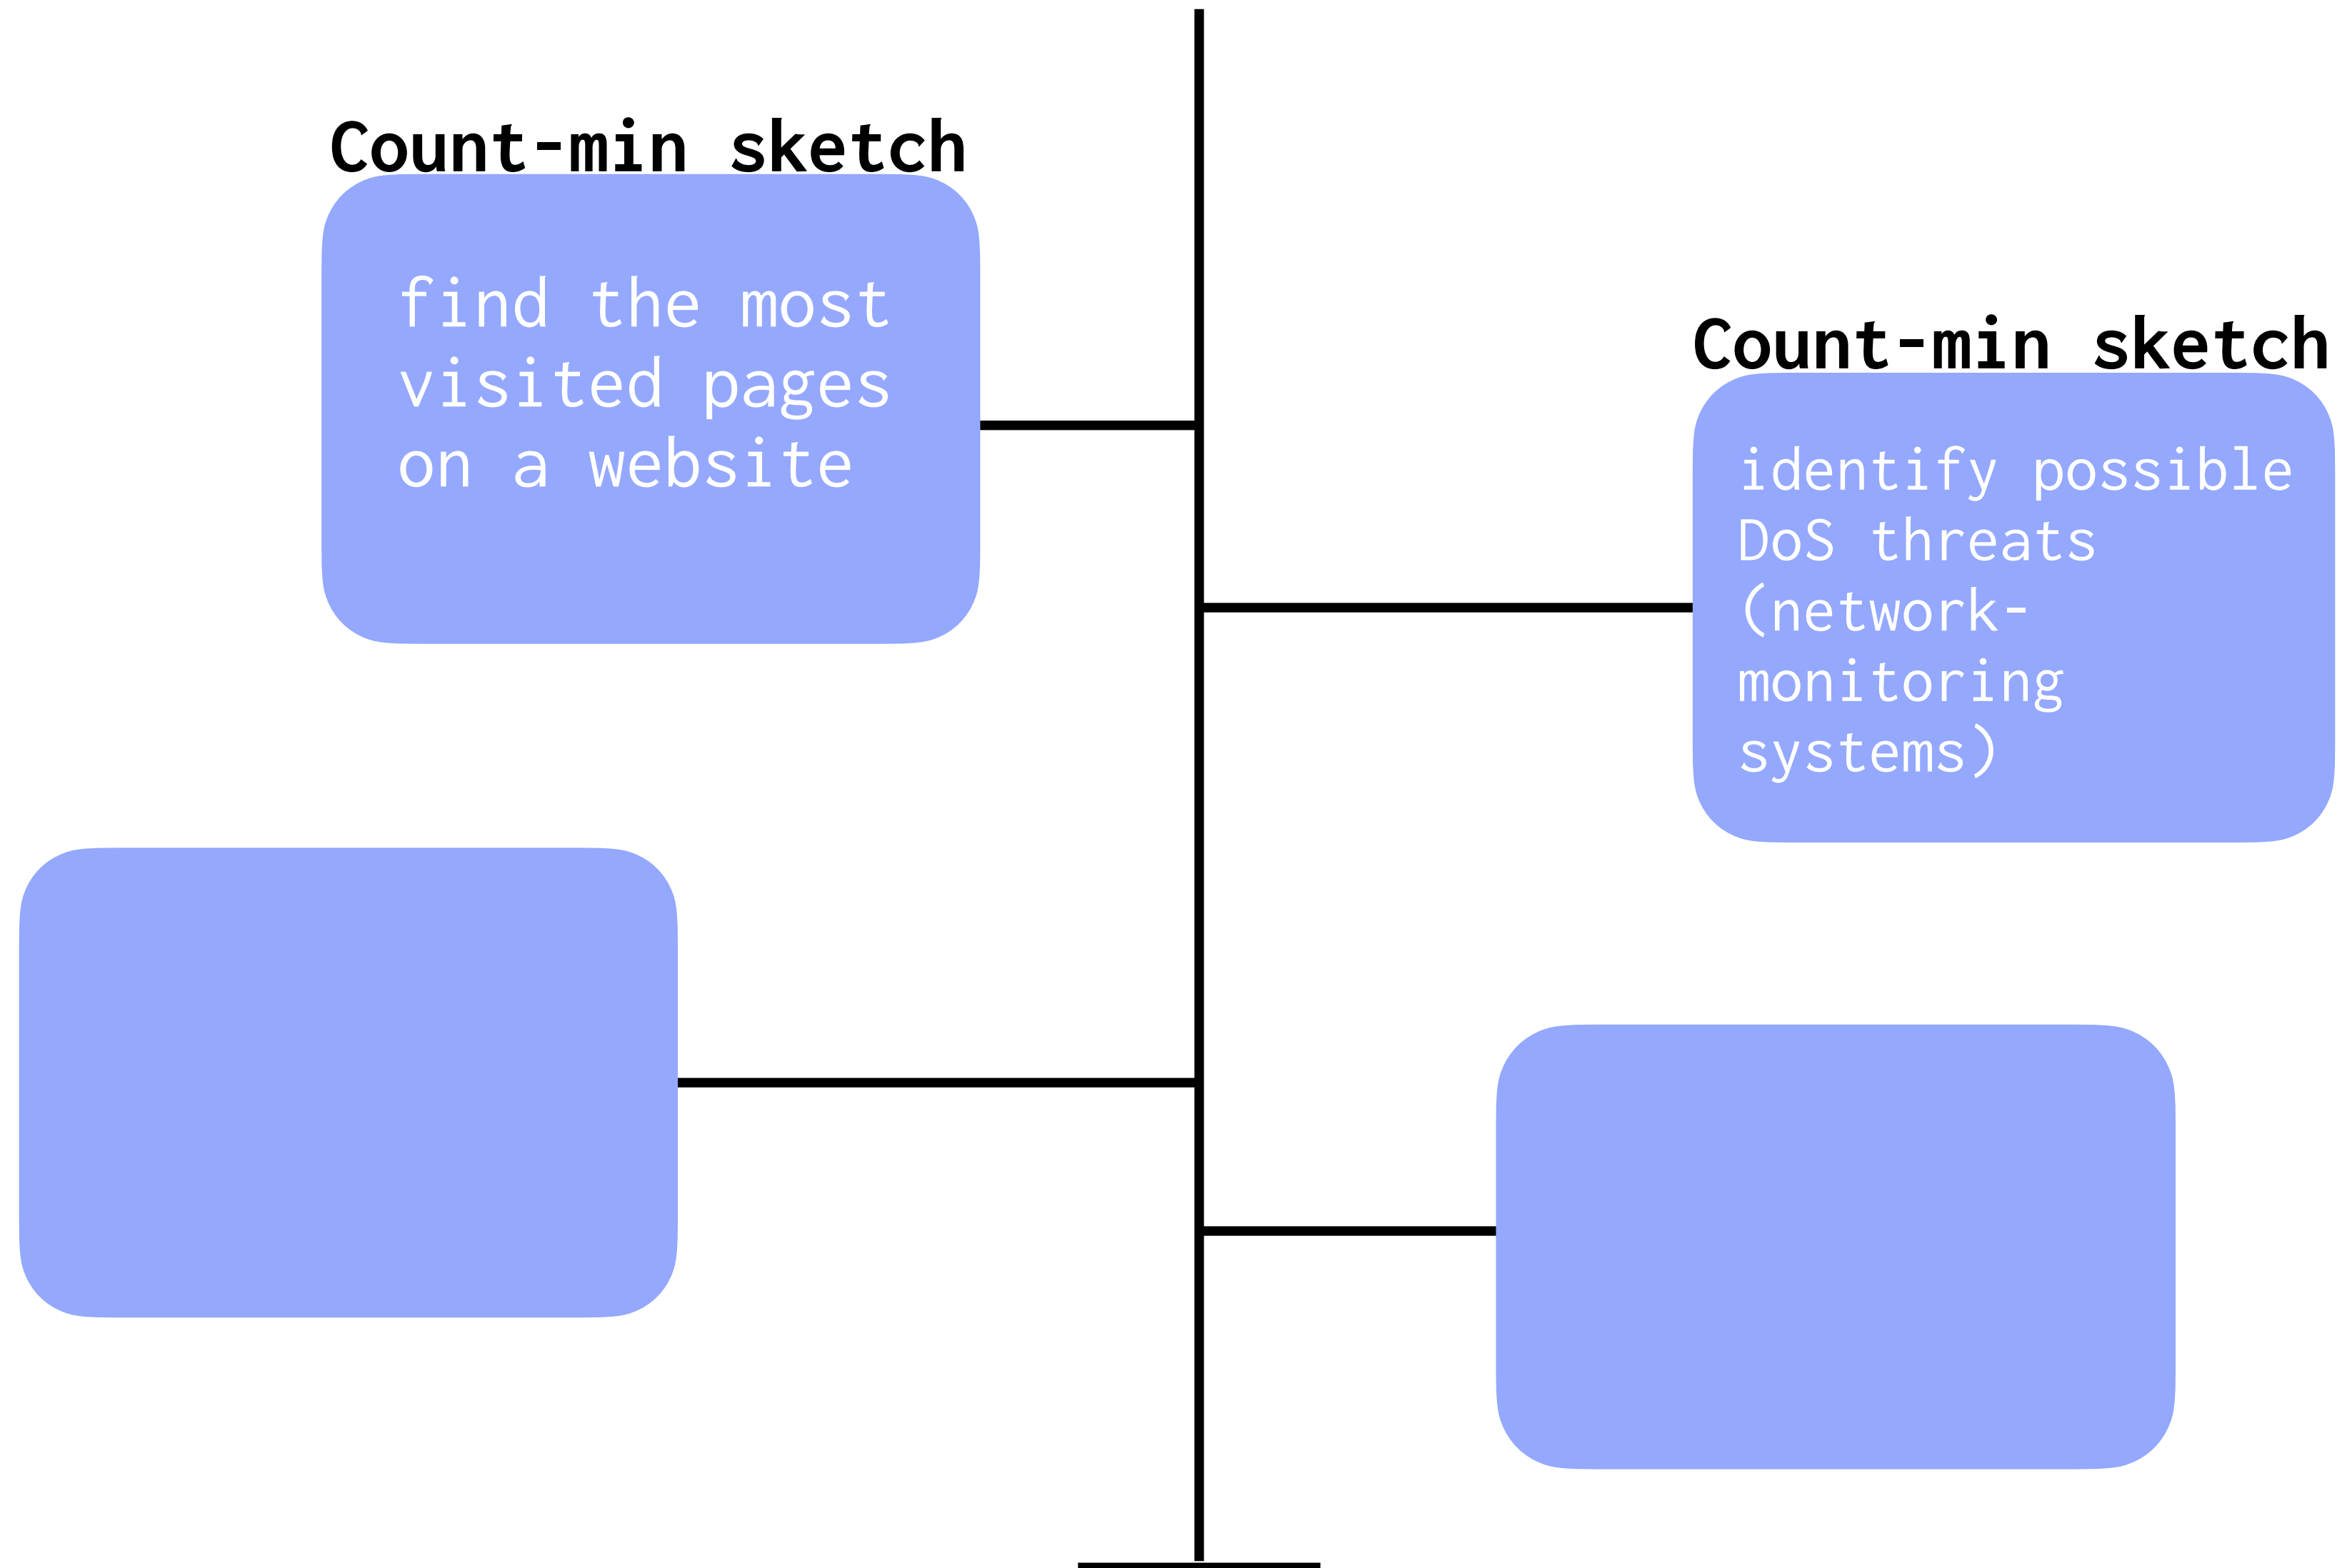
provide approximate
answers to queries
about the data

- Frequency estimation
How many times does x appear in the set?
Count-min sketch, HeavyKeeper
- Membership queries
Is x in the set?
Bloom filter, Cuckoo filter
- Cardinality estimation
How many distinct elements in the set?
HyperLogLog, KMV estimator

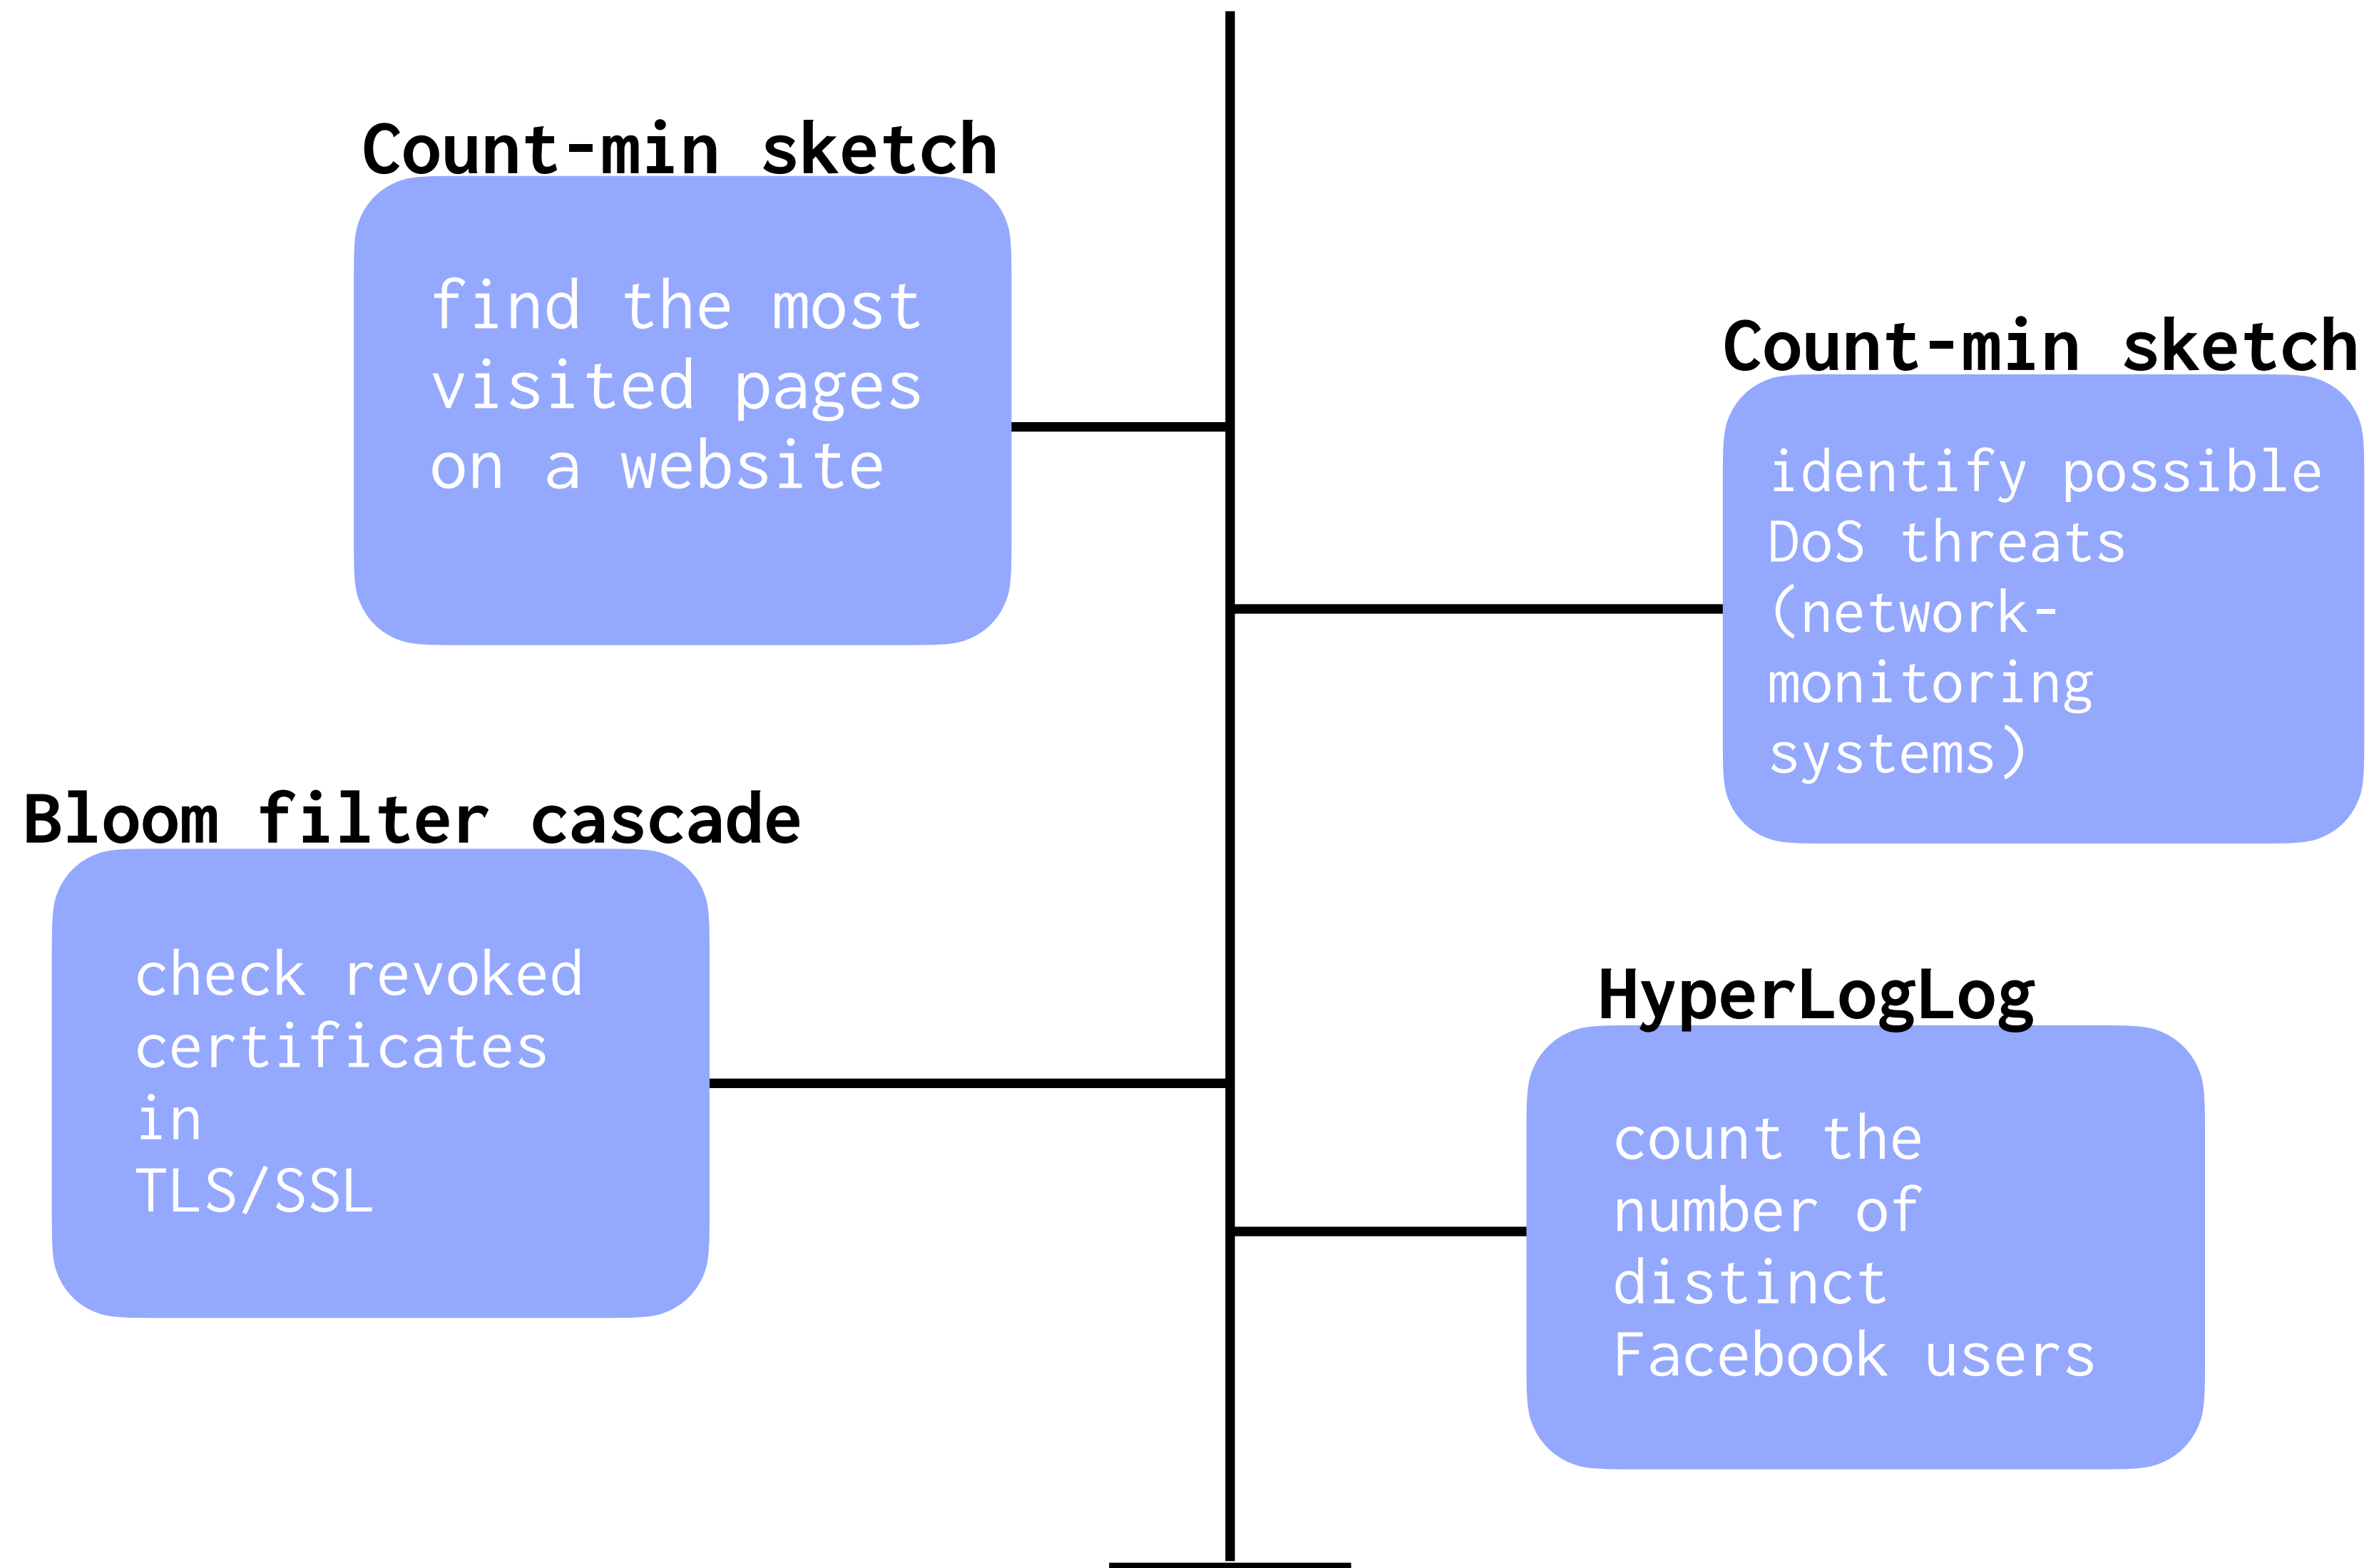
PDS help us



PDS help us



PDS help us



PDS in adversarial settings



PDS in adversarial settings

Adversarial
correctness

- How can an adversary **interfere** with the correct functionality of the PDS?

PDS in adversarial settings

Adversarial
correctness

- How can an adversary **interfere** with the correct functionality of the PDS?

Privacy

- What can an adversary **learn** about the elements stored in the PDS?

PDS in adversarial settings

Adversarial
correctness

- How can an adversary **interfere** with the correct functionality of the PDS?

Privacy

- What can an adversary **learn** about the elements stored in the PDS?

Secure PDS

- How can we **provably protect** PDS in adversarial settings?

Our work

- Approximate Membership Query PDS (w/o and w/ deletions)

Adversarial correctness

Privacy

Provable security[!]

Our work

- Approximate Membership Query PDS (w/o and w/ deletions)

Adversarial correctness **Privacy** **Provable security**!

- Compact Frequency Estimation (CFE) PDS

Adversarial correctness **Exploration of a
more robust CFE PDS**
**Attacks against
CMS and HeavyKeeper**

Our work

- Approximate Membership Query PDS (w/o and w/ deletions)

Adversarial correctness **Privacy** **Provable security**!

- Compact Frequency Estimation (CFE) PDS


Adversarial correctness

- Practical implementation

Adversarial correctness

Compact Frequency Estimation (CFE) PDS

Compact Frequency Estimation (CFE) PDS




CFE PDS

Stream

n,z,r,p,t,w,l,l,n,s,k,g,o,i,w,...

Compact Frequency Estimation (CFE) PDS



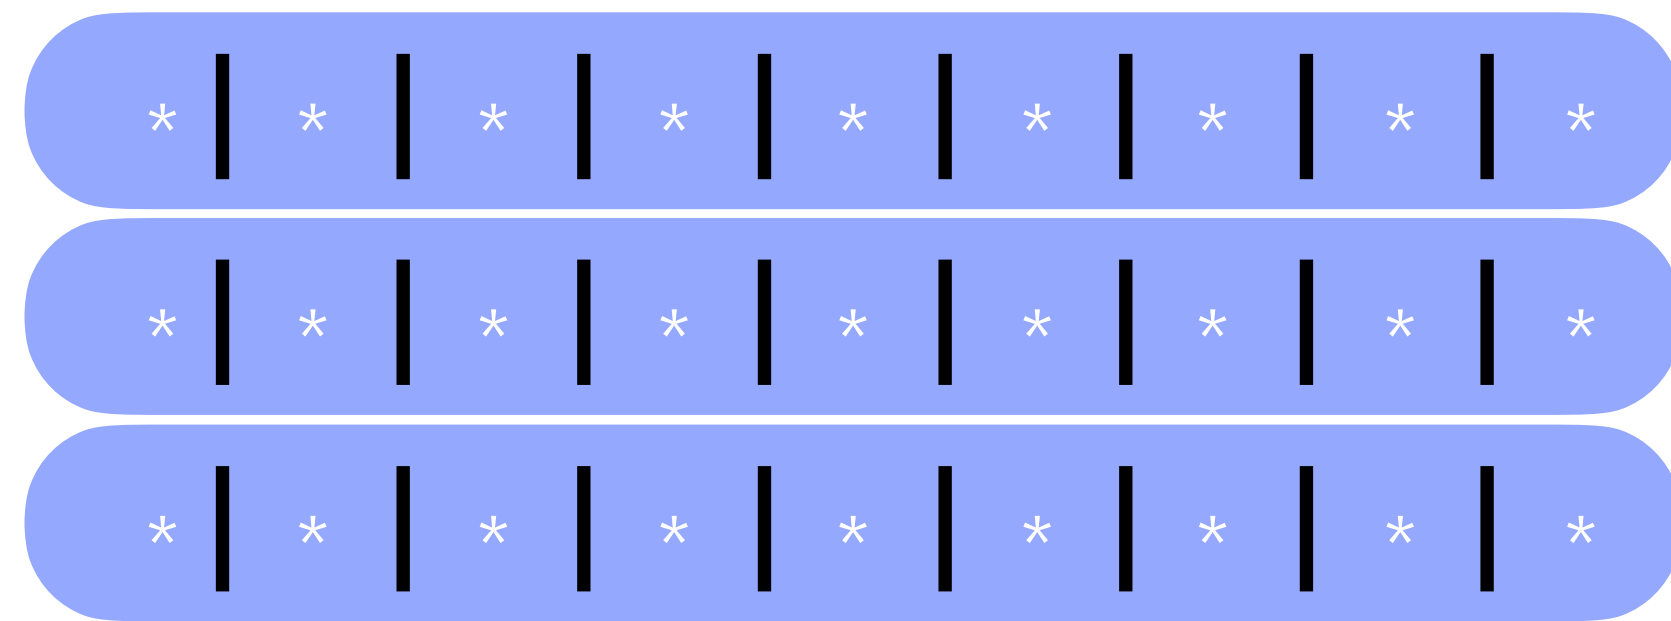
CFE PDS



Stream

n,z,r,p,t,w,l,l,n,s,k,g,o,i,w,...

Compact Frequency Estimation (CFE) PDS




CFE PDS

Stream

n,z,r,p,t,w,l,l,n,s,k,g,o,i,w,...



Can CFE PDS misbehave?




CFE PDS

Stream

n,z,r,p,t,w,l,l,n,s,k,g,o,i,w,...



Can CFE PDS misbehave?



CFE PDS




query(x)

Stream

n,z,r,p,t,w,l,l,n,s,k

Can CFE PDS misbehave?




CFE PDS

query(x)




ans



Stream

n,z,r,p,t,w,l,l,n,s,k

Can CFE PDS misbehave?



CFE PDS

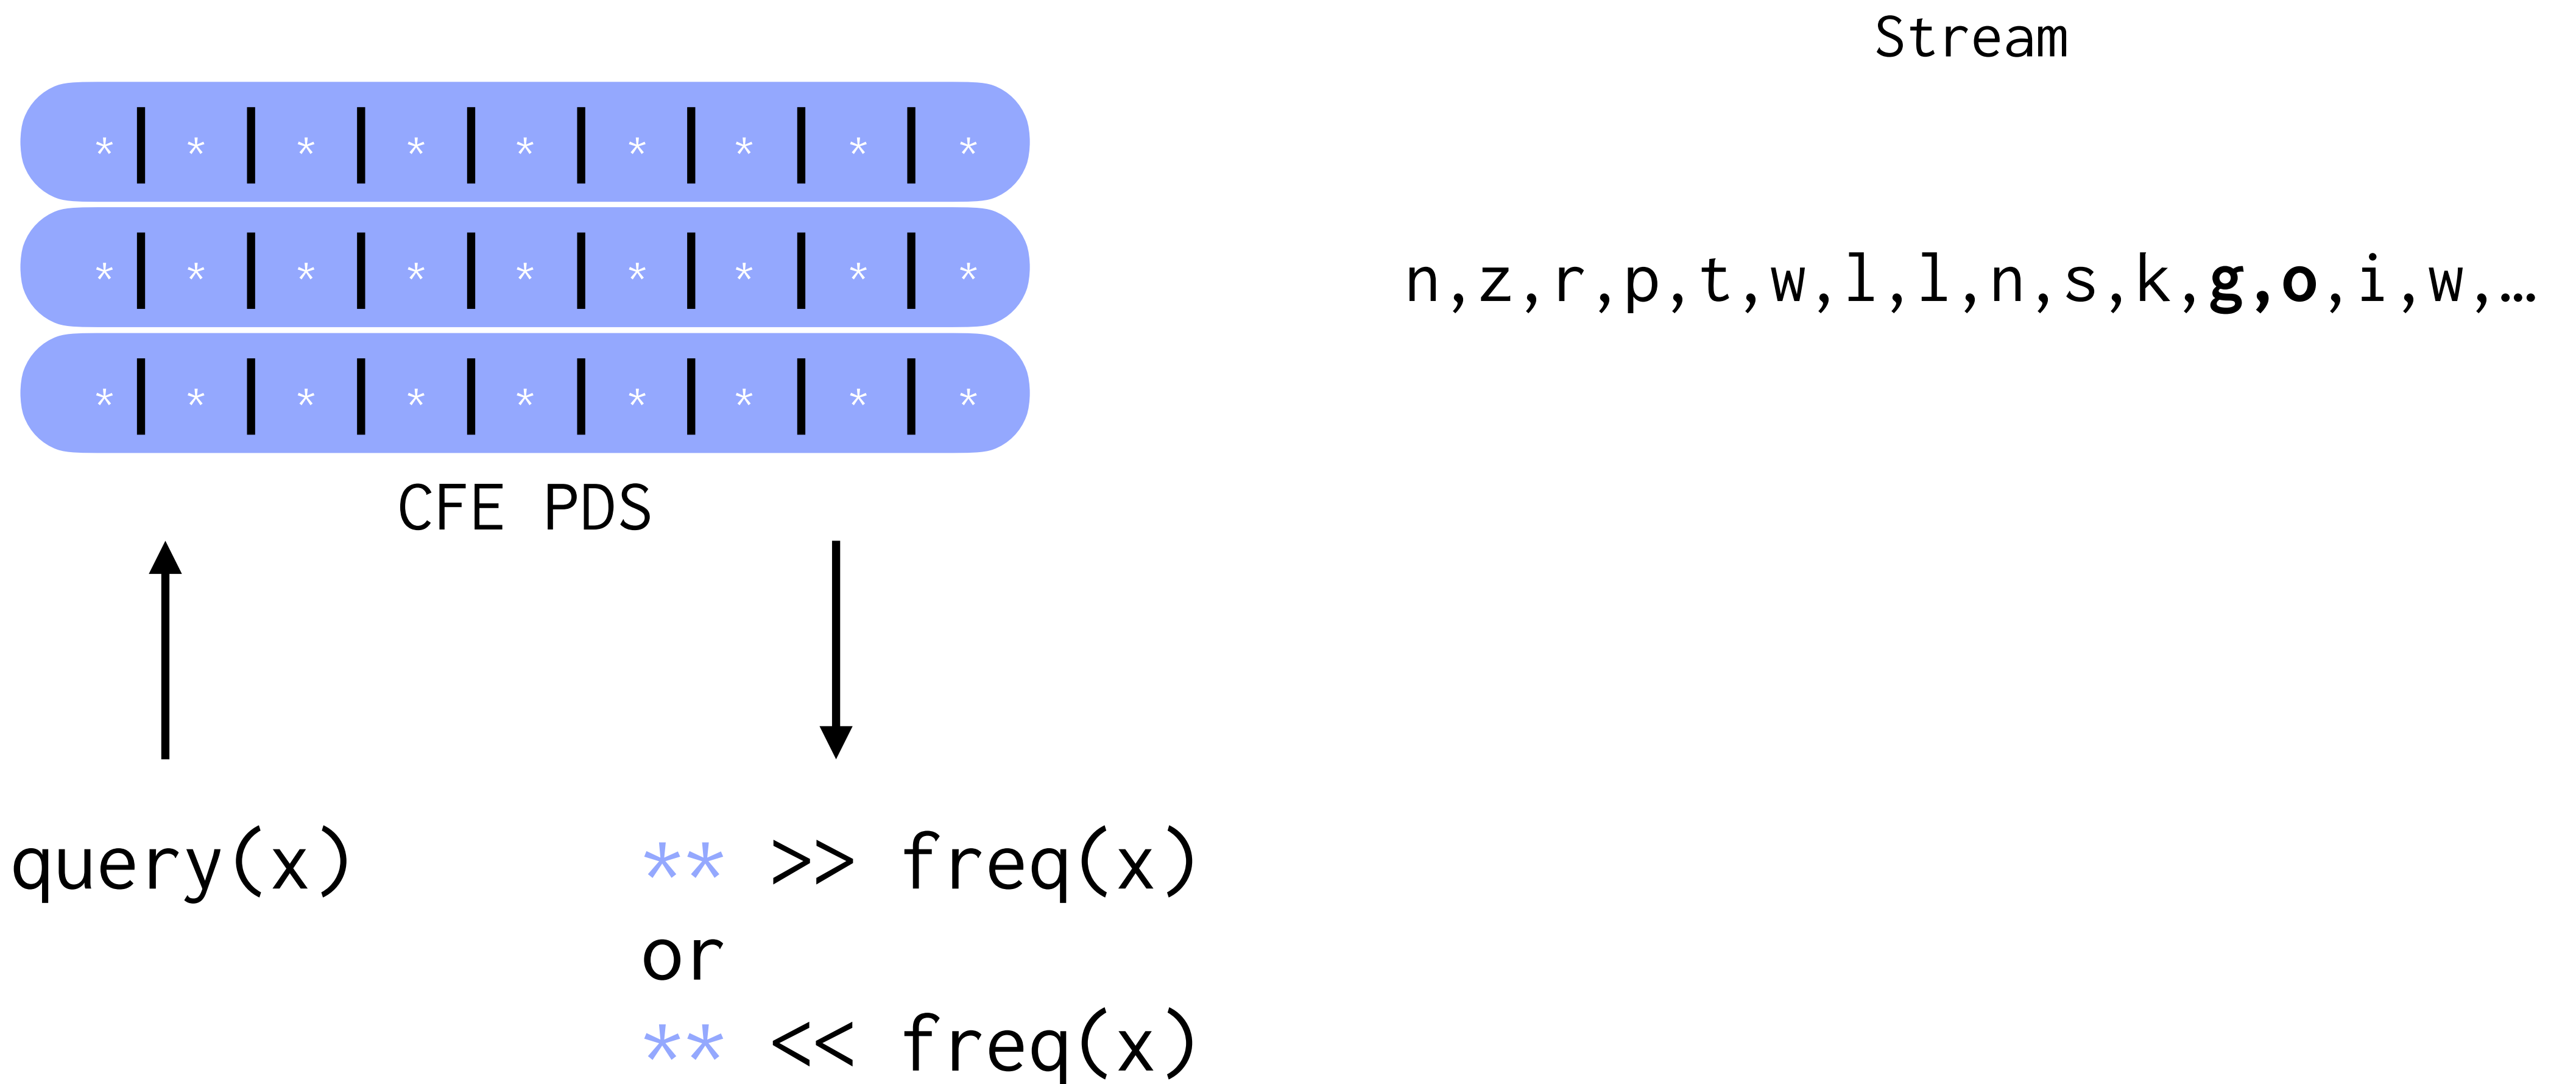
Stream

`n,z,r,p,t,w,l,l,n,s,k,f(ans),g(ans)...`

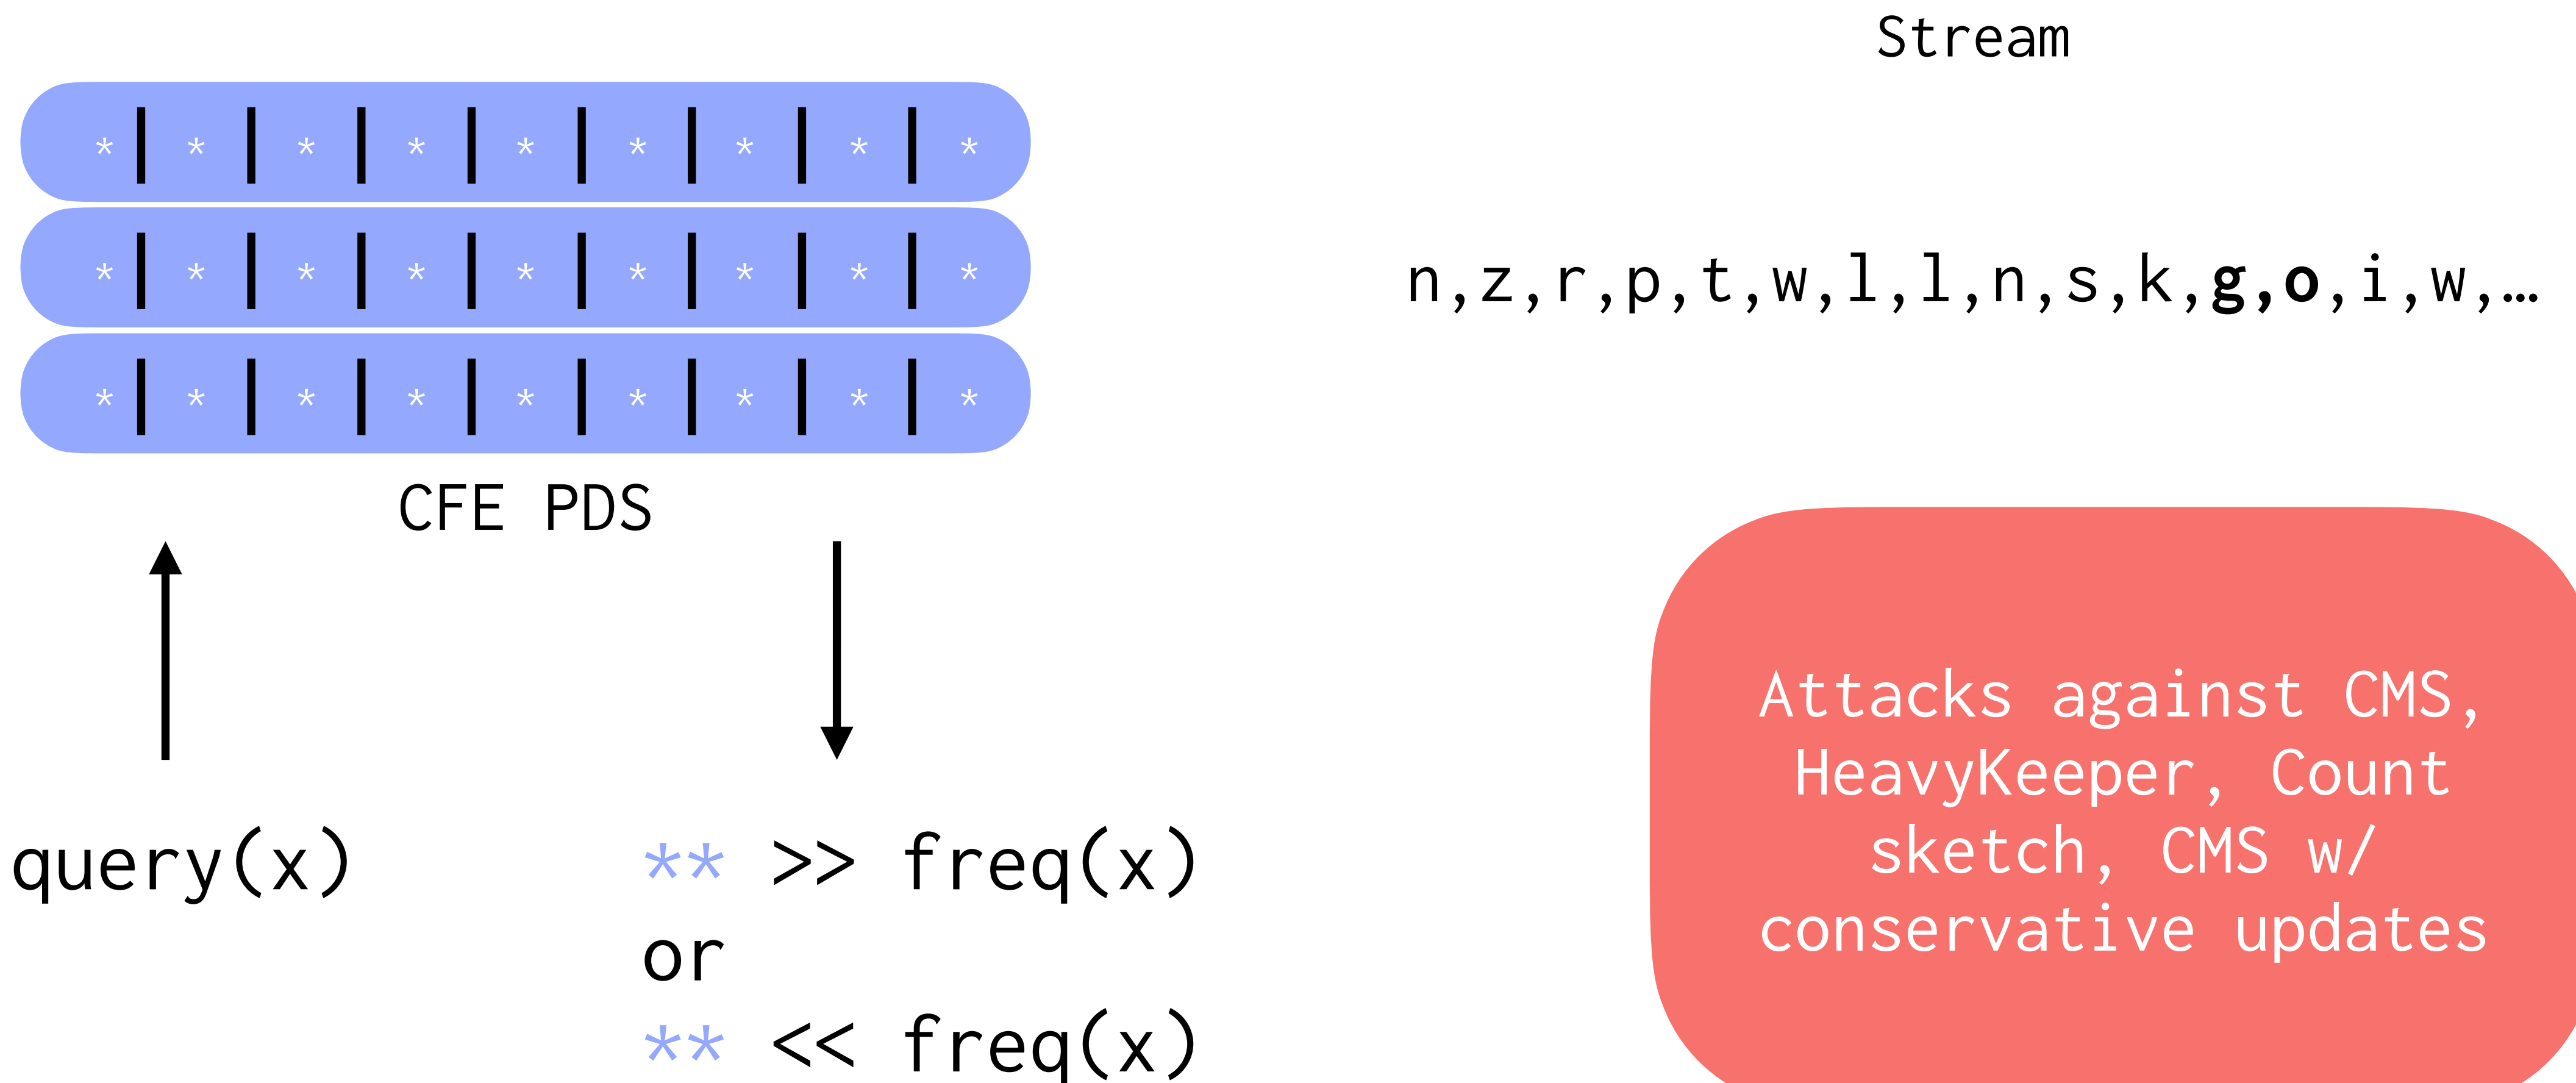
`ans`



Can CFE PDS misbehave?




Can CFE PDS misbehave?




Attacks against CMS, HeavyKeeper, Count sketch, CMS w/ conservative updates

CMS: how does it work?


5 ← h1(x)



9 ← h2(x)




3 ← h3(x)




insert(x)

CMS: how does it work?


5 ← h1(x)



9 ← h2(x)




3 ← h3(x)




insert(x)

CMS: how does it work?


$5 \leftarrow h_1(x)$



insert(x)




$9 \leftarrow h_2(x)$




$3 \leftarrow h_3(x)$

CMS: how does it work?


5 ← h1(x)



9 ← h2(x)




3 ← h3(x)




query(x)

CMS: how does it work?


5 ← h1(x)



9 ← h2(x)




3 ← h3(x)




query(x)

CMS: how does it work?


5 ← h1(x)



9 ← h2(x)




3 ← h3(x)




CMS(x) = 39

CMS


h1(.)



h2(.)



h3(.)

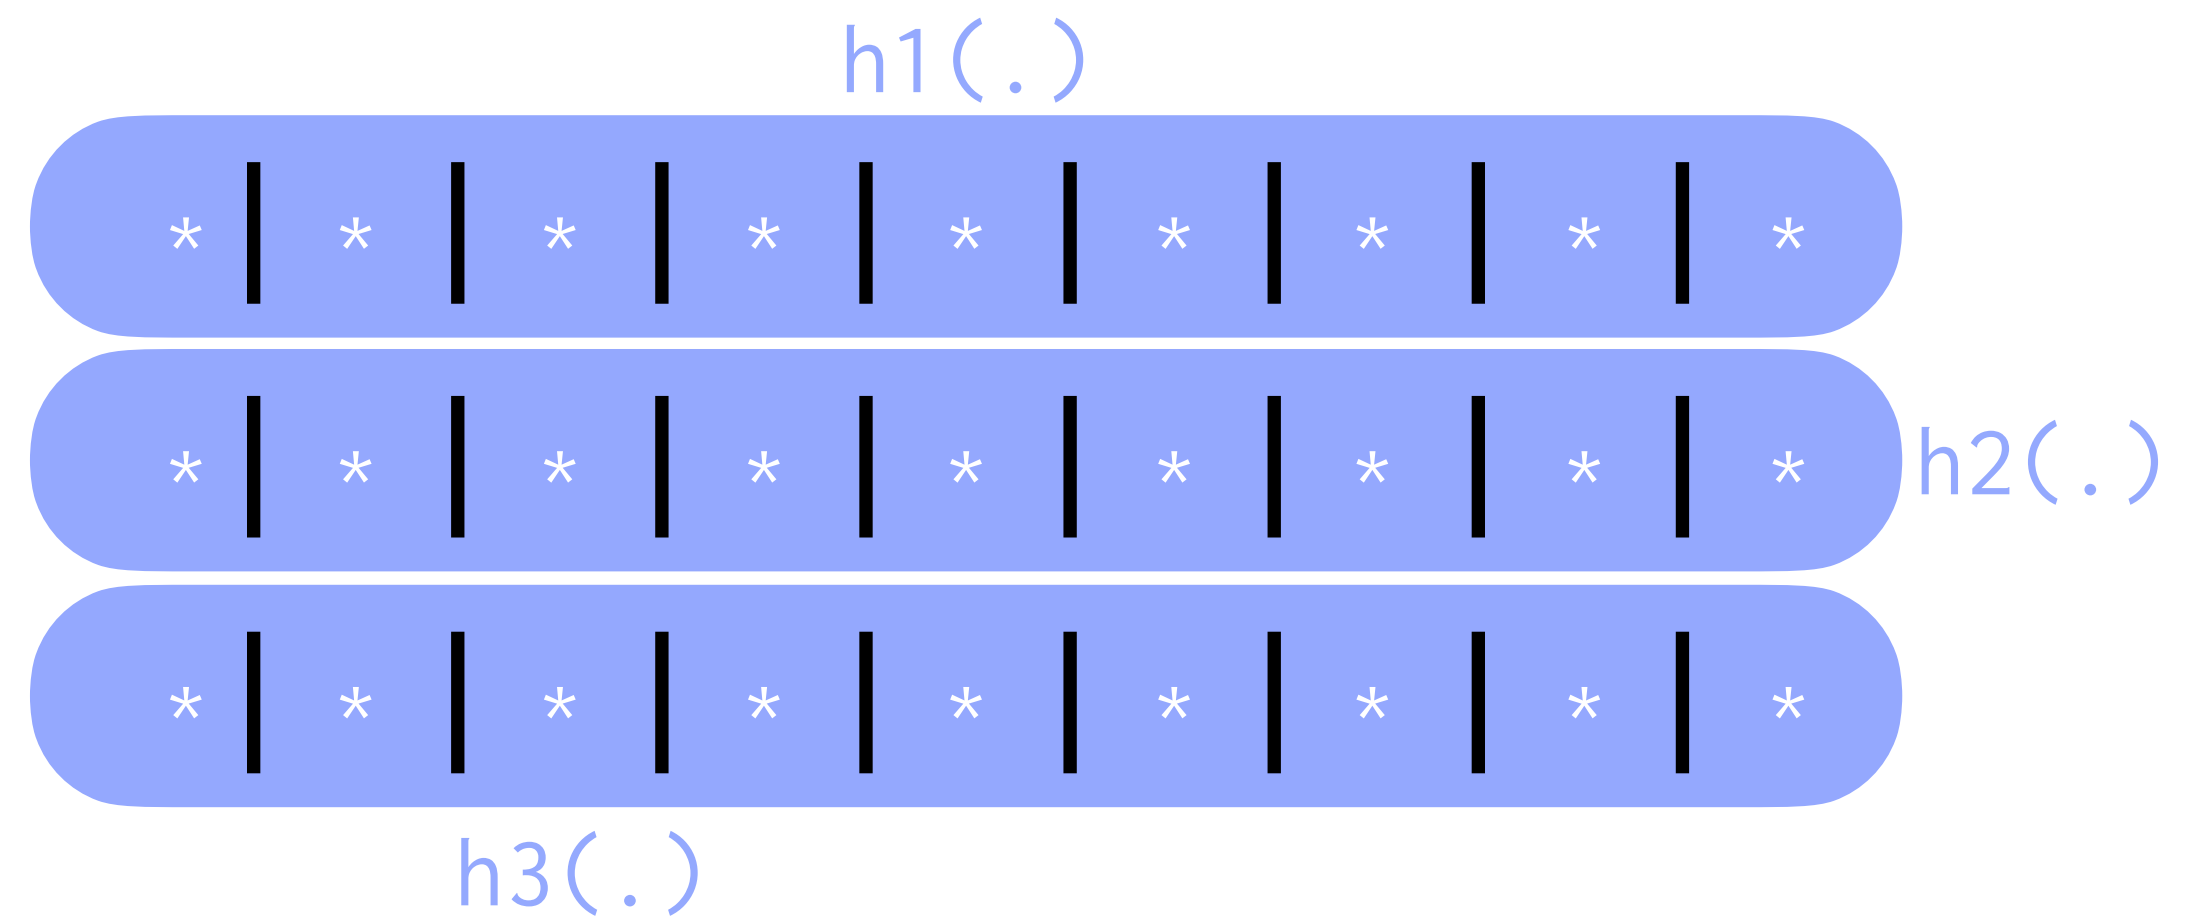


CMS: attack model




CMS[]:
insert
query

Hash




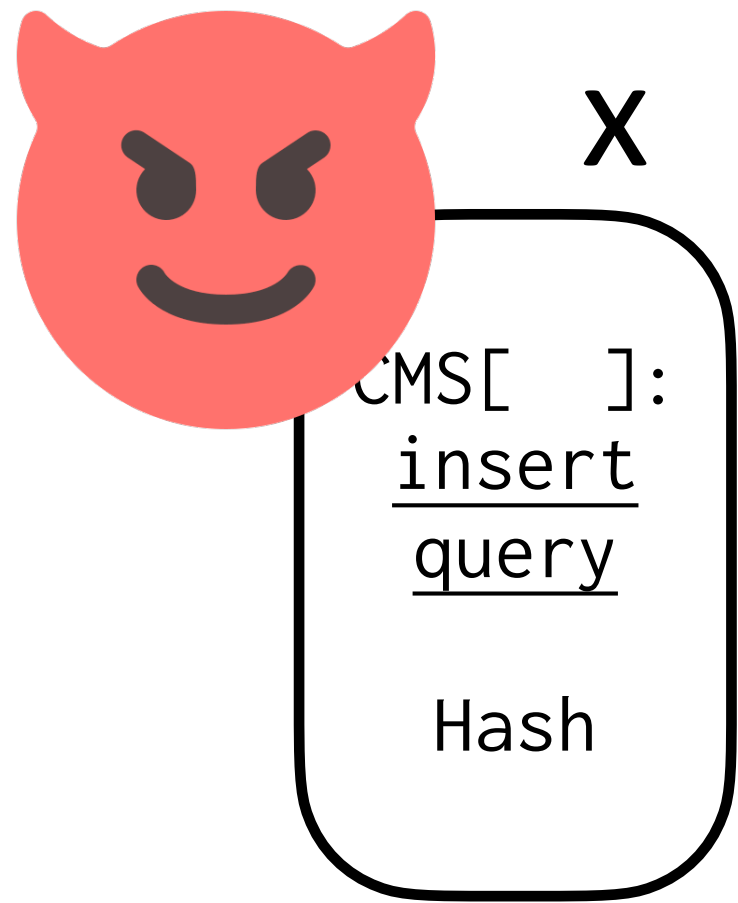
CMS: attack goal



Maximise
CMS error

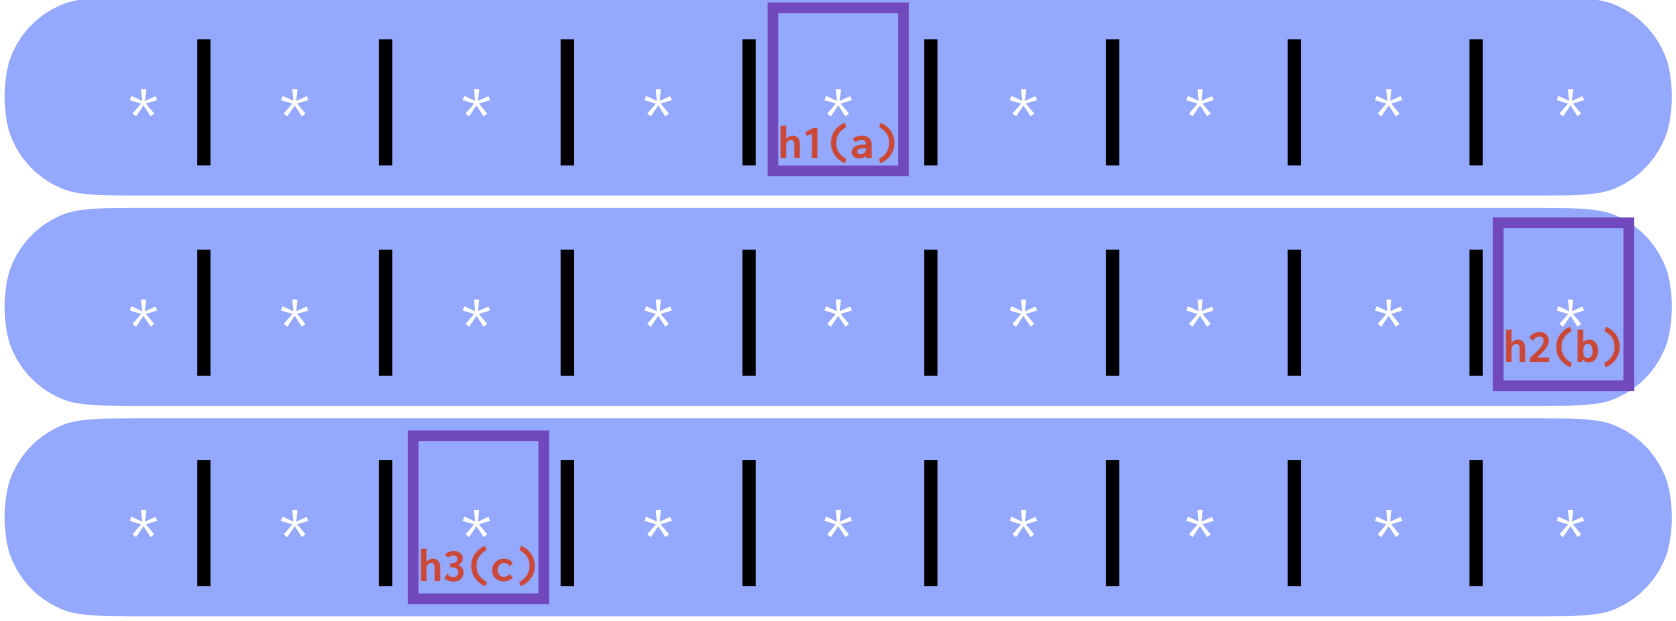
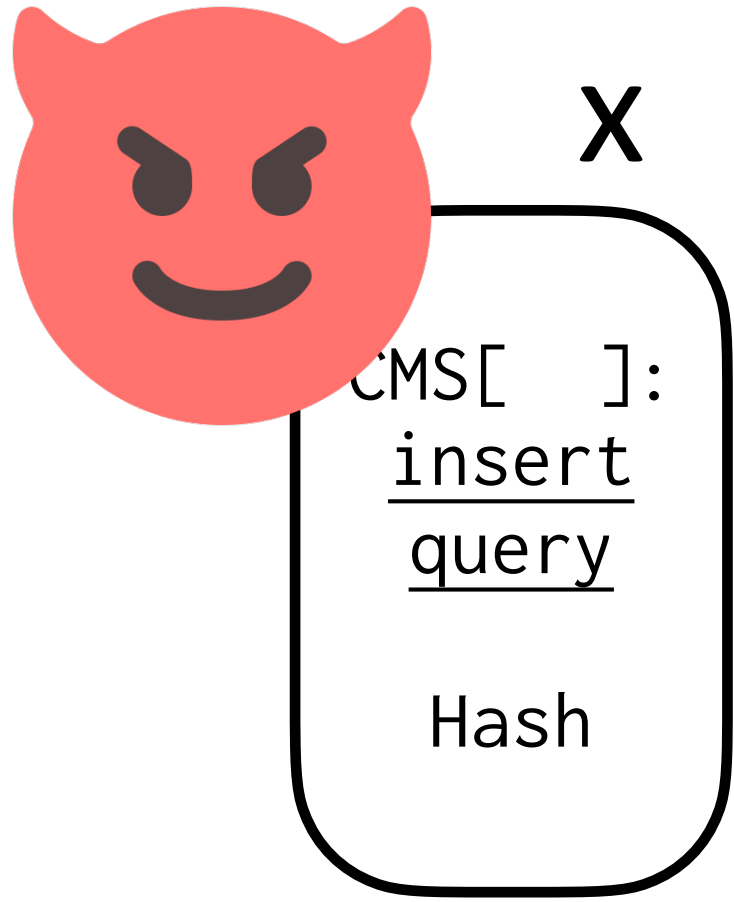
$$\text{query}(x) \gg \text{true_frequency}(x)$$

CMS: attack




Cover set = {a, b, c}

CMS: attack

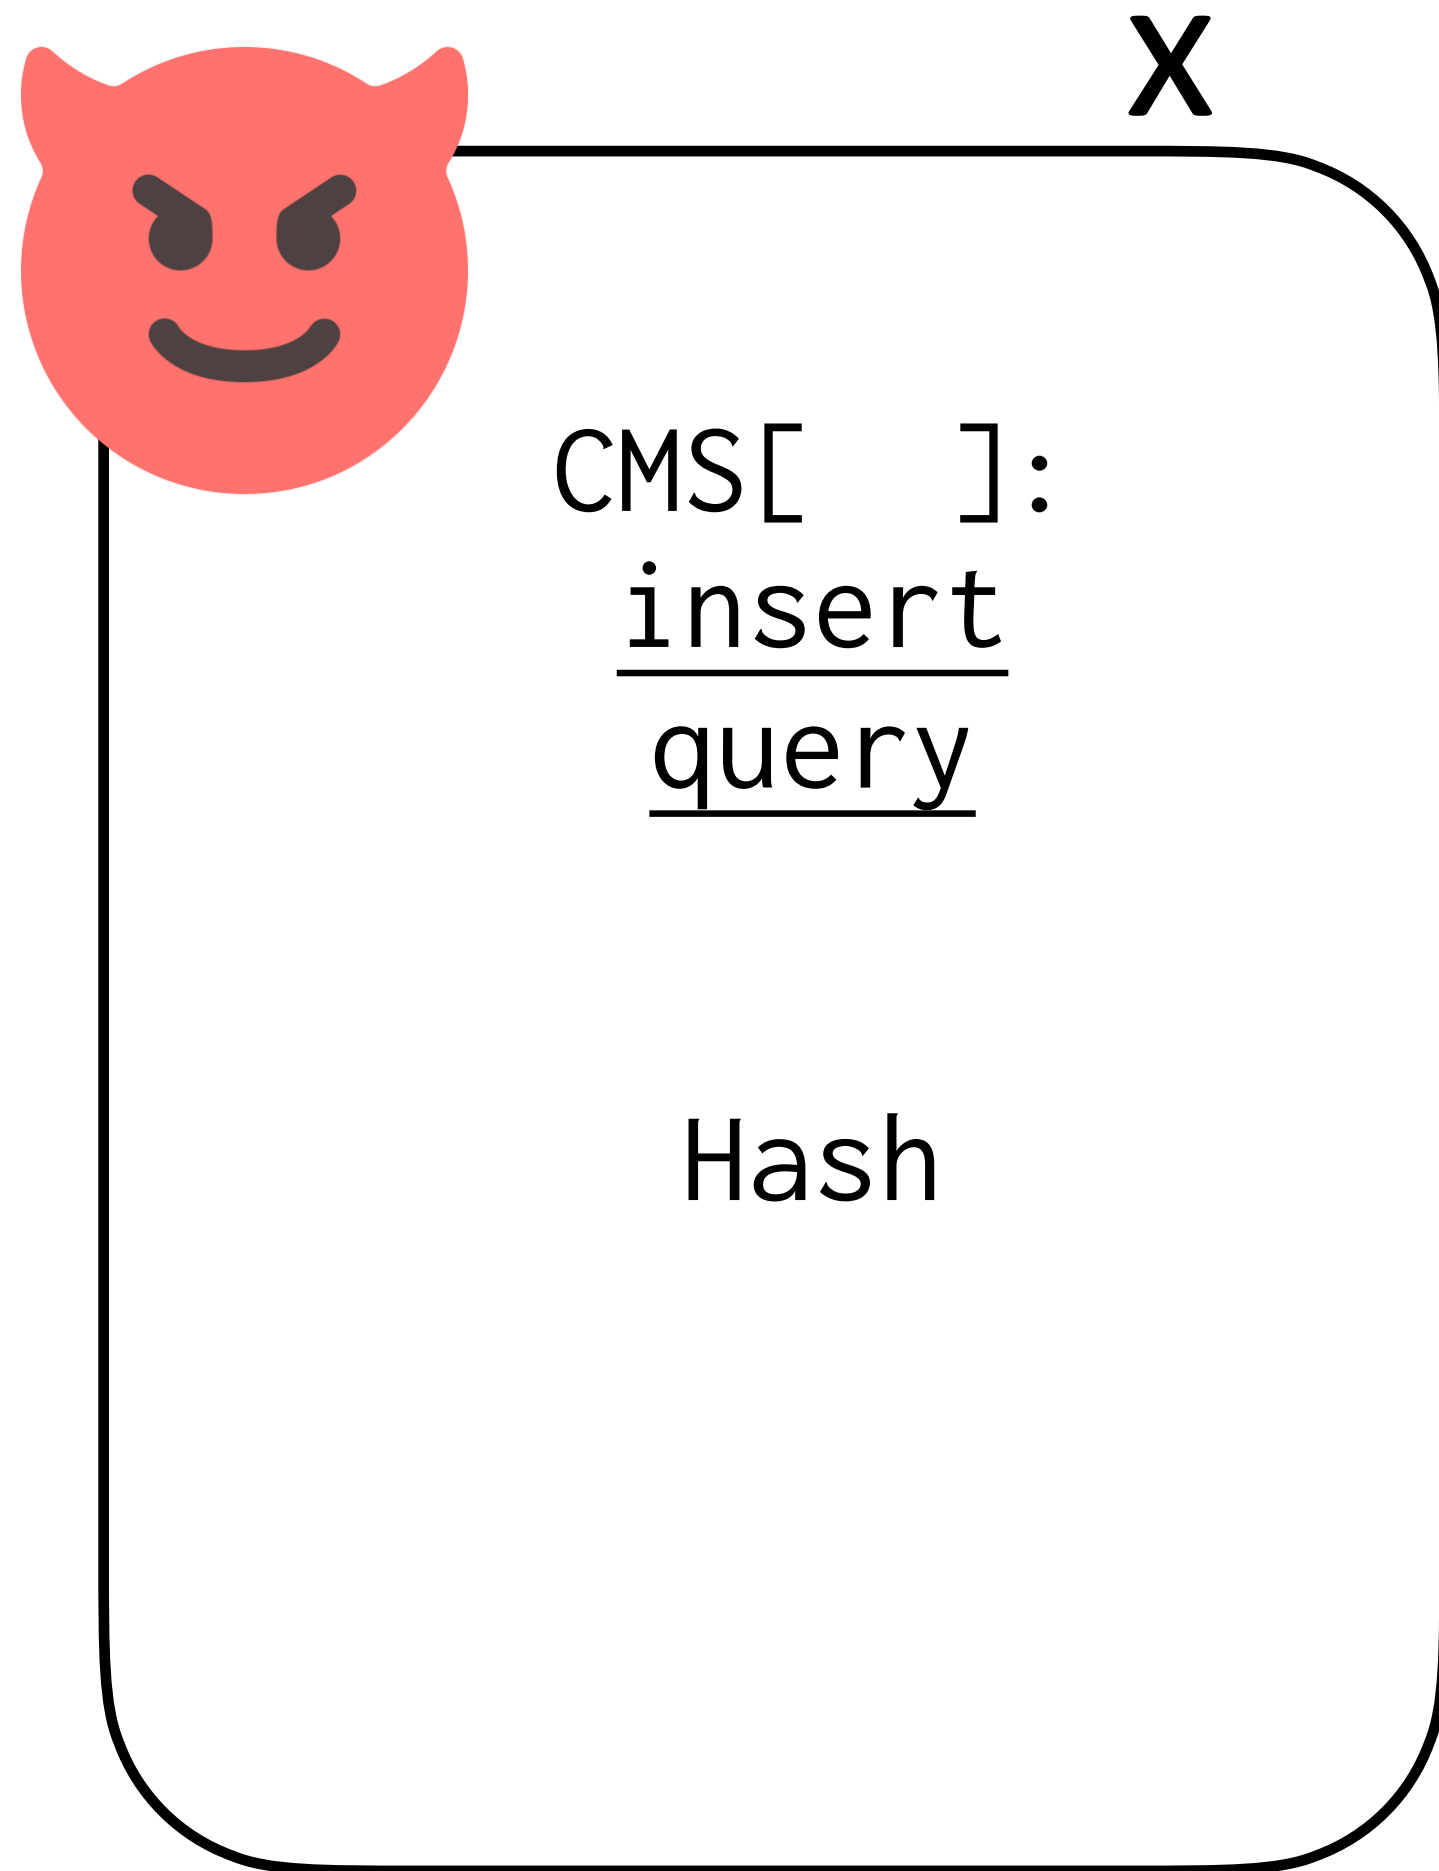


Cover set = {a, b, c}

Cover set



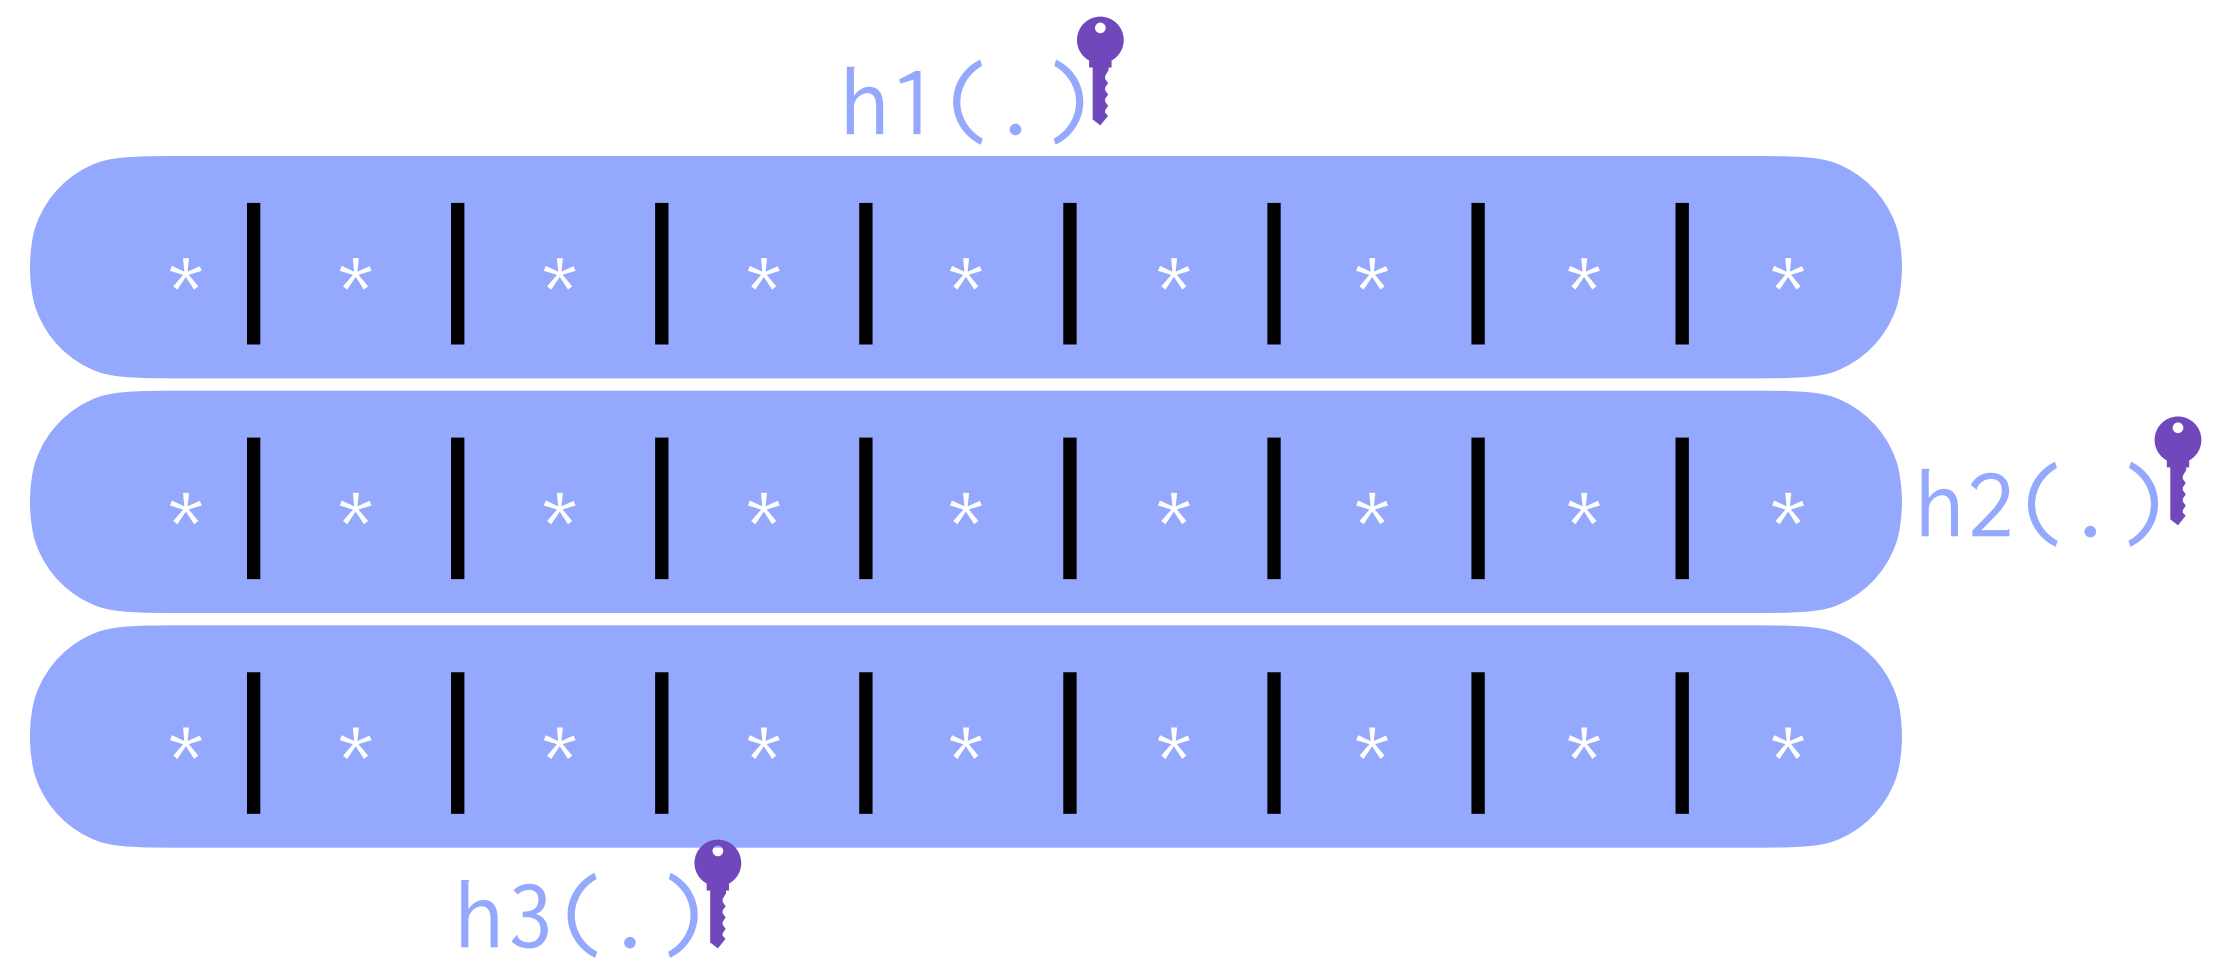

CMS: attack




Err:

insertions/k

CMS: attack model cont.



CMS: attack



Err:

insertions/k - m Hk

Our attacks make


elements **absent** from the stream marked as **heavy**

Our attacks make

elements **absent** from the stream marked as **heavy**
or
high-frequency elements marked as **absent**.


More robust CFE PDS: Overestimator
+ Underestimator

More robust CFE PDS: Overestimator + Underestimator




CMS est & *HeavyKeeper est $\xrightarrow{\text{refine}}$ final est

Count Keeper (CK)



CMS M

&




*HeavyKeeper A



Honest setting
experiments



Count Keeper (CK)




Err:

CK < 1/2 CMS
CK << 1/2 HeavyKeeper

Attack
experiments


Count Keeper (CK)



+ error related properties (see CCS23 paper) :)




Count Keeper (CK)



CK can detect suspicious estimates




Count Keeper (CK)




CK can detect suspicious estimates 

Open problems & Future work




&



Overestimator ?


Underestimator ?

Open problems & Future work



PDS A ?

&



PDS B ?

Our work / Open problems

- Approximate Membership Query PDS (w/o and w/ deletions)

Adversarial correctness Privacy Provable security[!]

- Compact Frequency Estimation (CFE) PDS

Adversarial correctness Privacy Provable security[!]


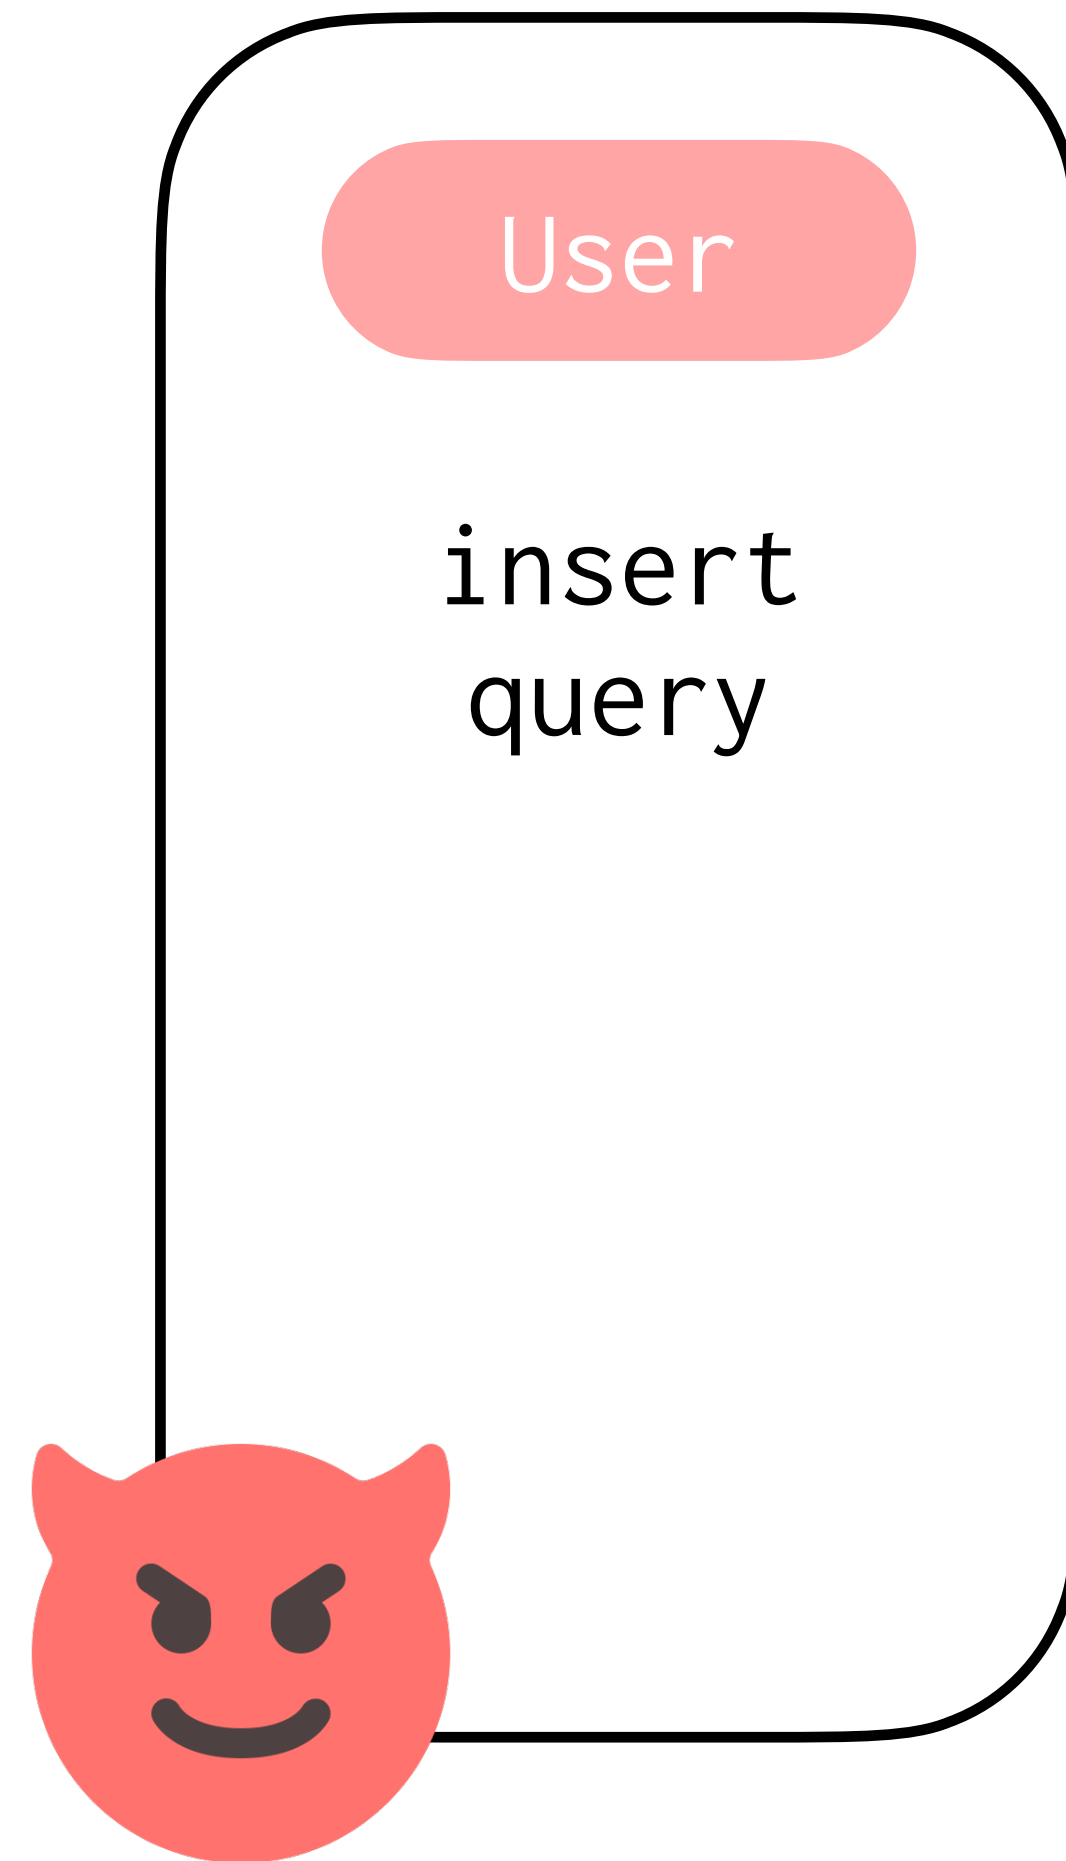
- Other PDS

Adversarial correctness Privacy Provable security[!]

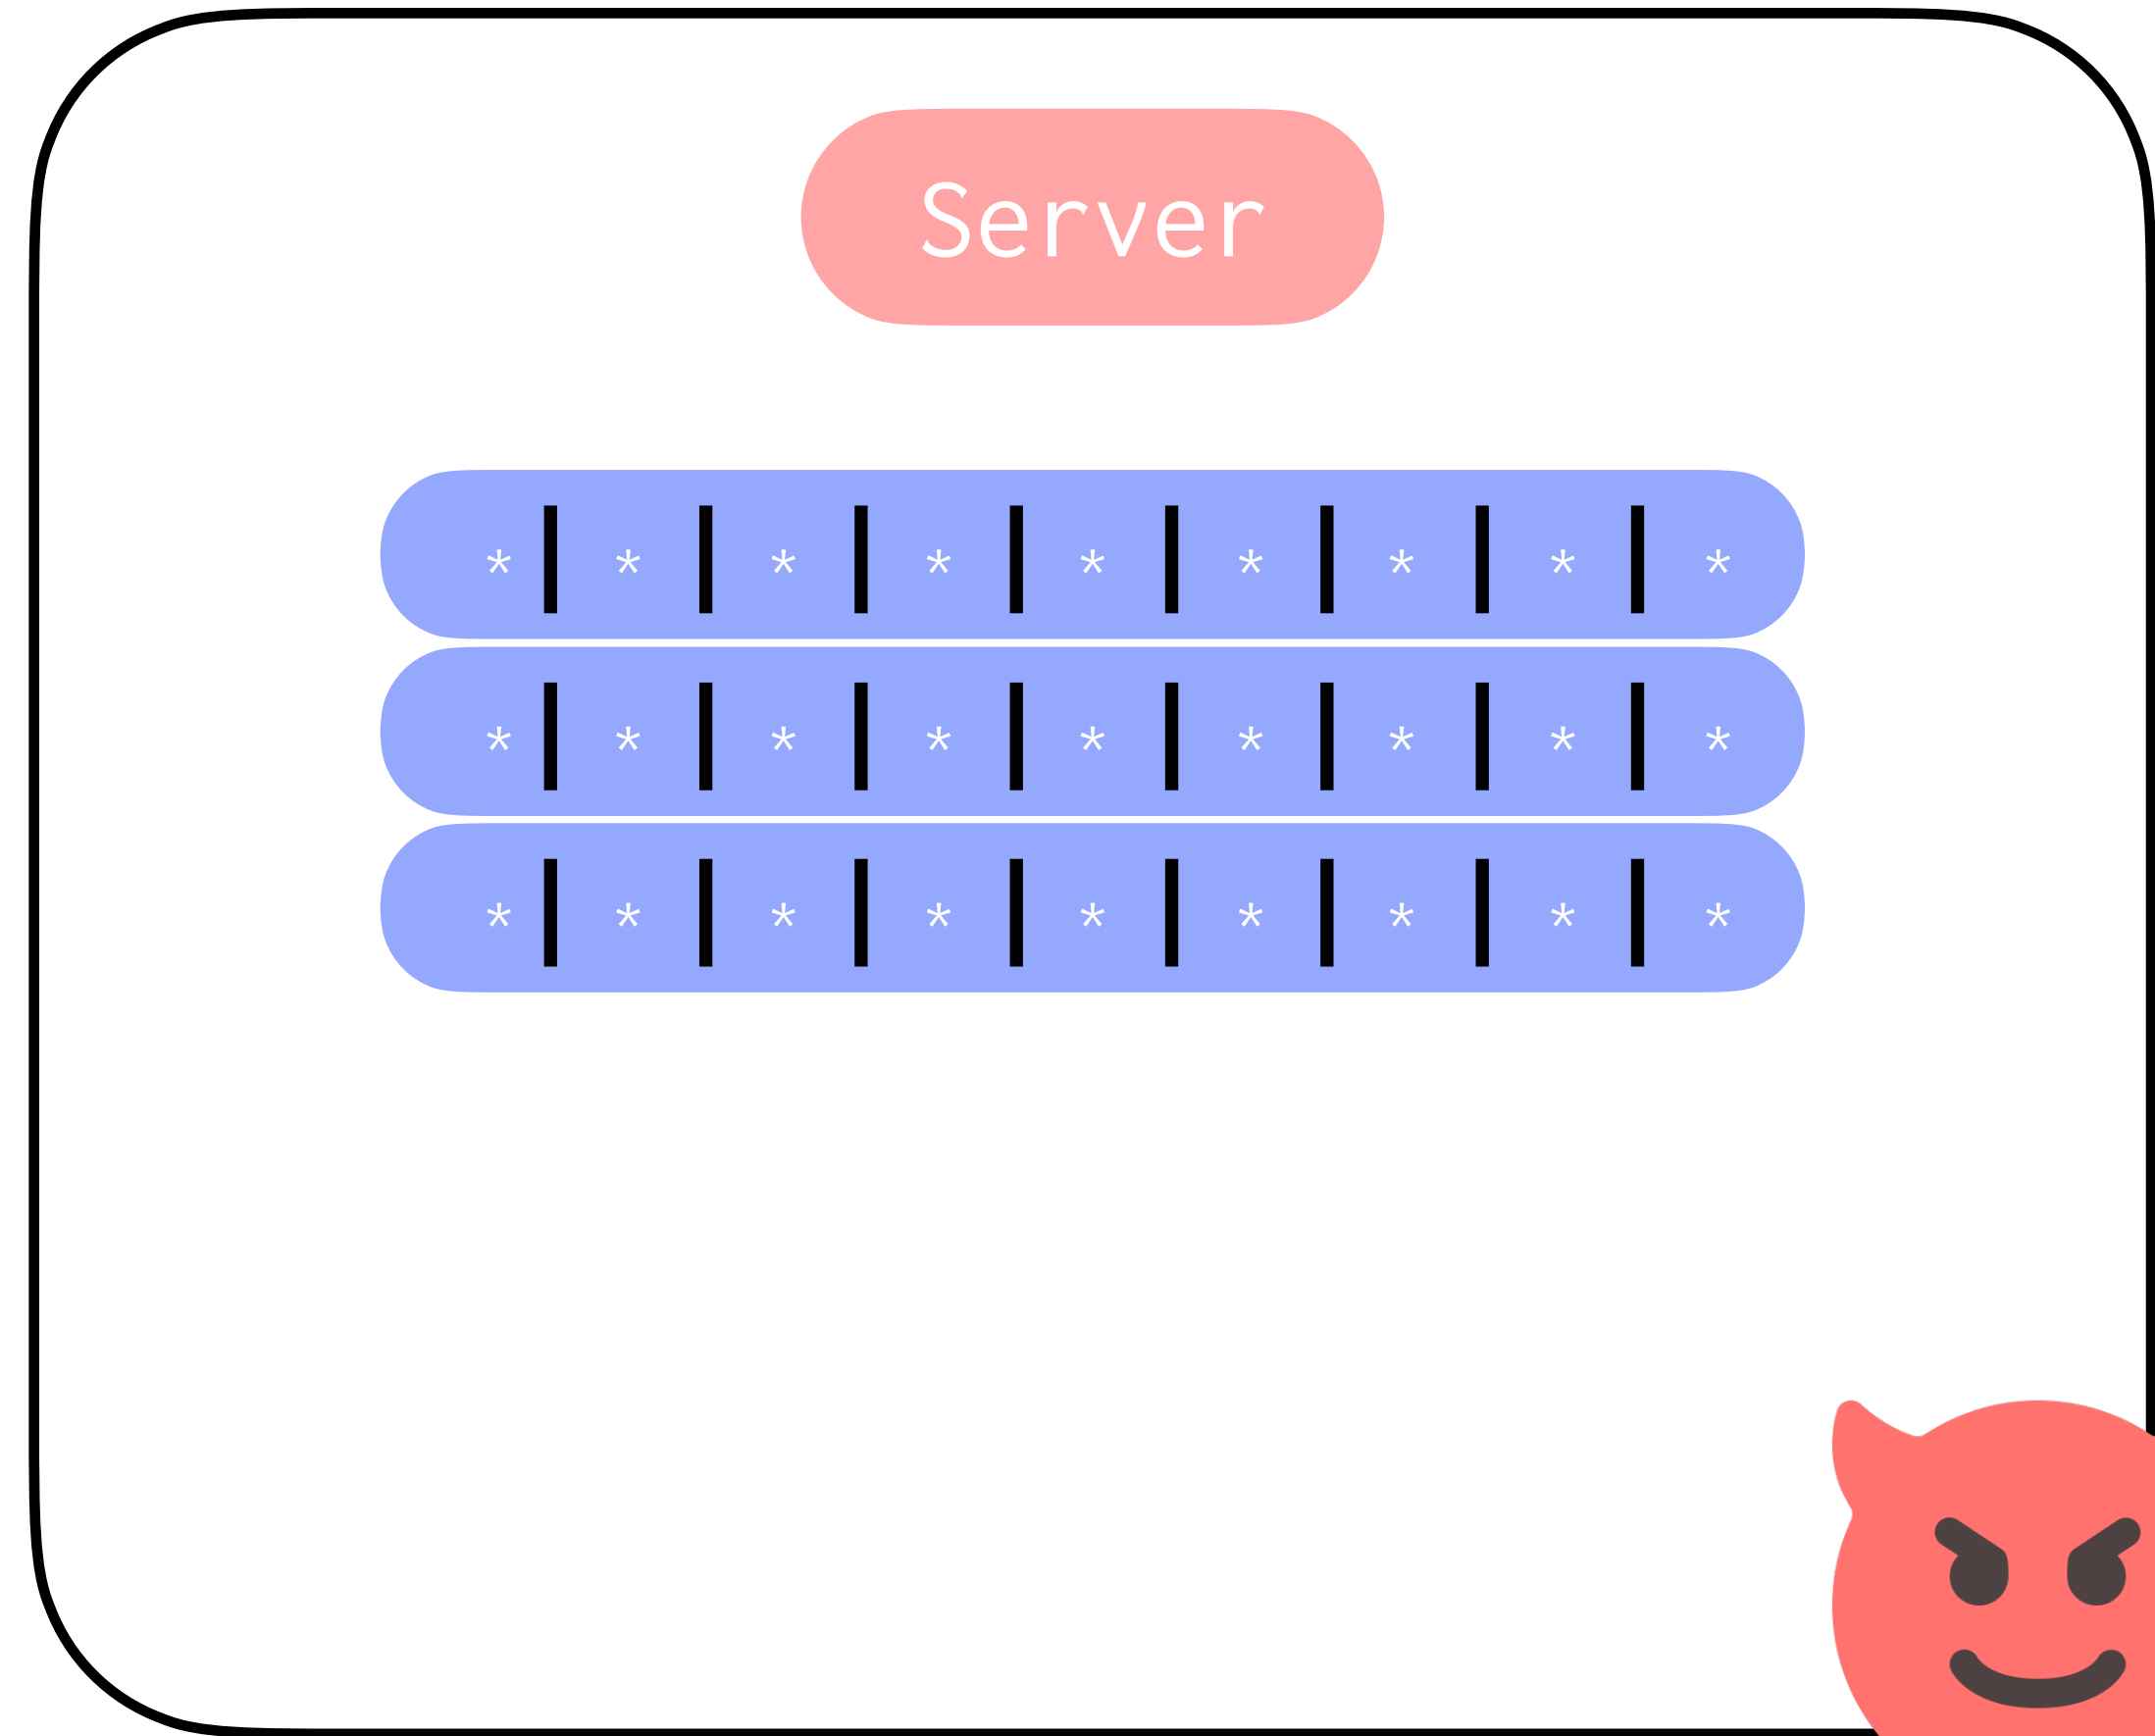

- Practical implementation

Adversarial correctness

Future work



Future work



What if the server is malicious
and the user is honest?

Thank you!

Thank you!

Approximate Membership Query PDS (CCS22)



Compact Frequency Estimation (CFE) PDS (CCS23)

