GPU-accelerated PIR with Client-Independent Preprocessing for Large-Scale Applications

Daniel Günther, Maurice Heymann, Benny Pinkas, Thomas Schneider

Contact: guenther@encrypto.cs.tu-darmstadt.de

ENCRYPTO CRYPTOGRAPHY AND PRIVACY ENGINEERING

Bar-Ilan University אוניברסיטת בר אילן

TECHNISCHE UNIVERSITÄT DARMSTADT

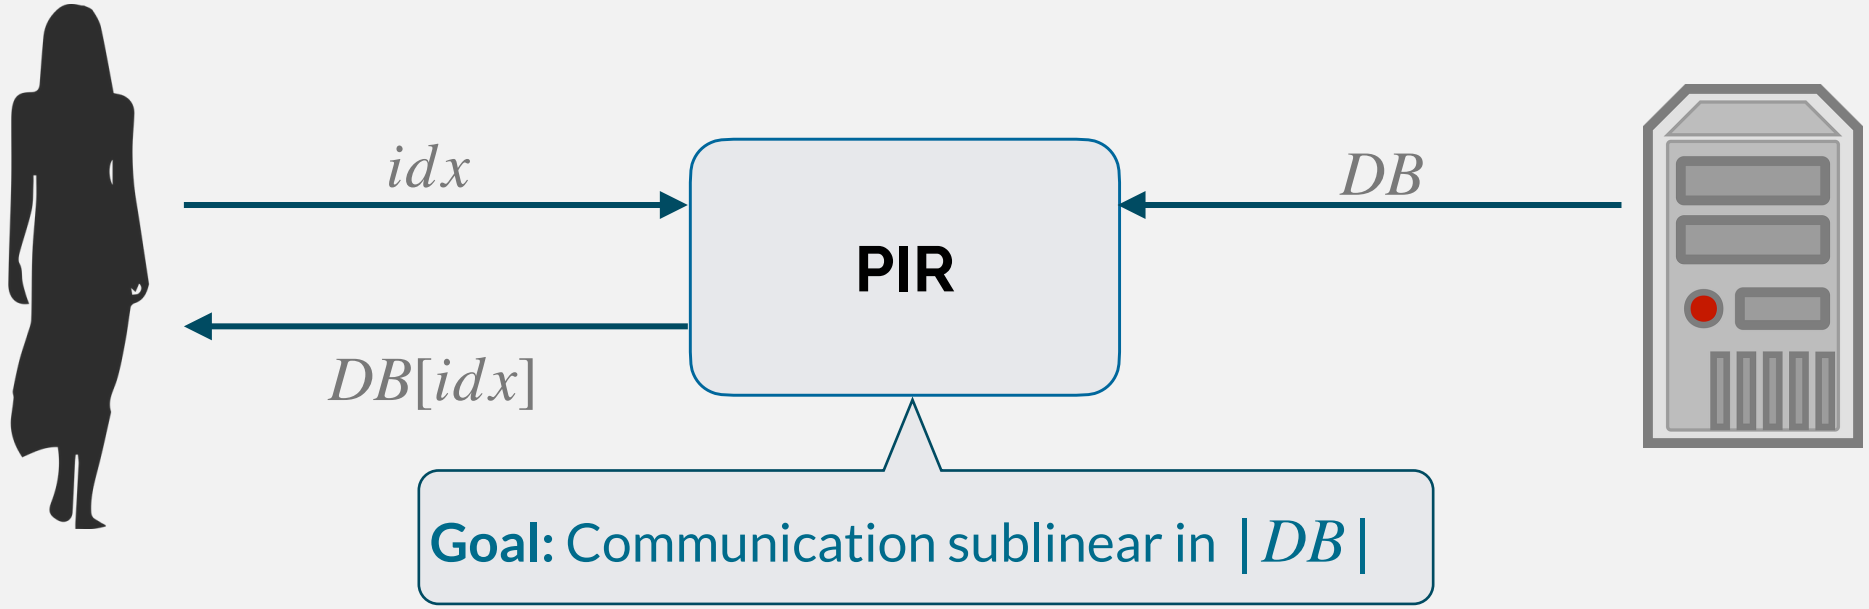

PIR - Many Applications

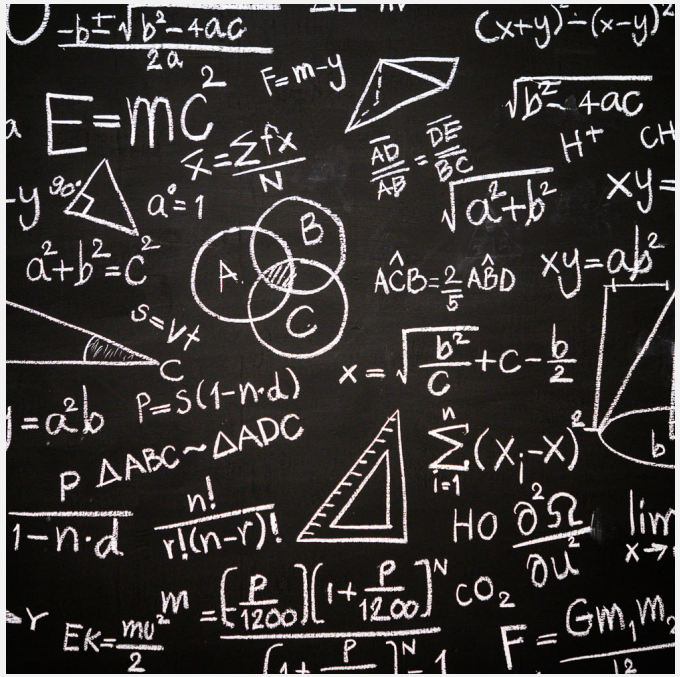

Private Set Membership

Securely check if an entry exists in a public database.

Private Messaging App

Protect messages and metadata.

Leaked Credentials

Securely check if the user's credentials are leaked.

Multi Server PIR [Chor et al., FOCS'95]

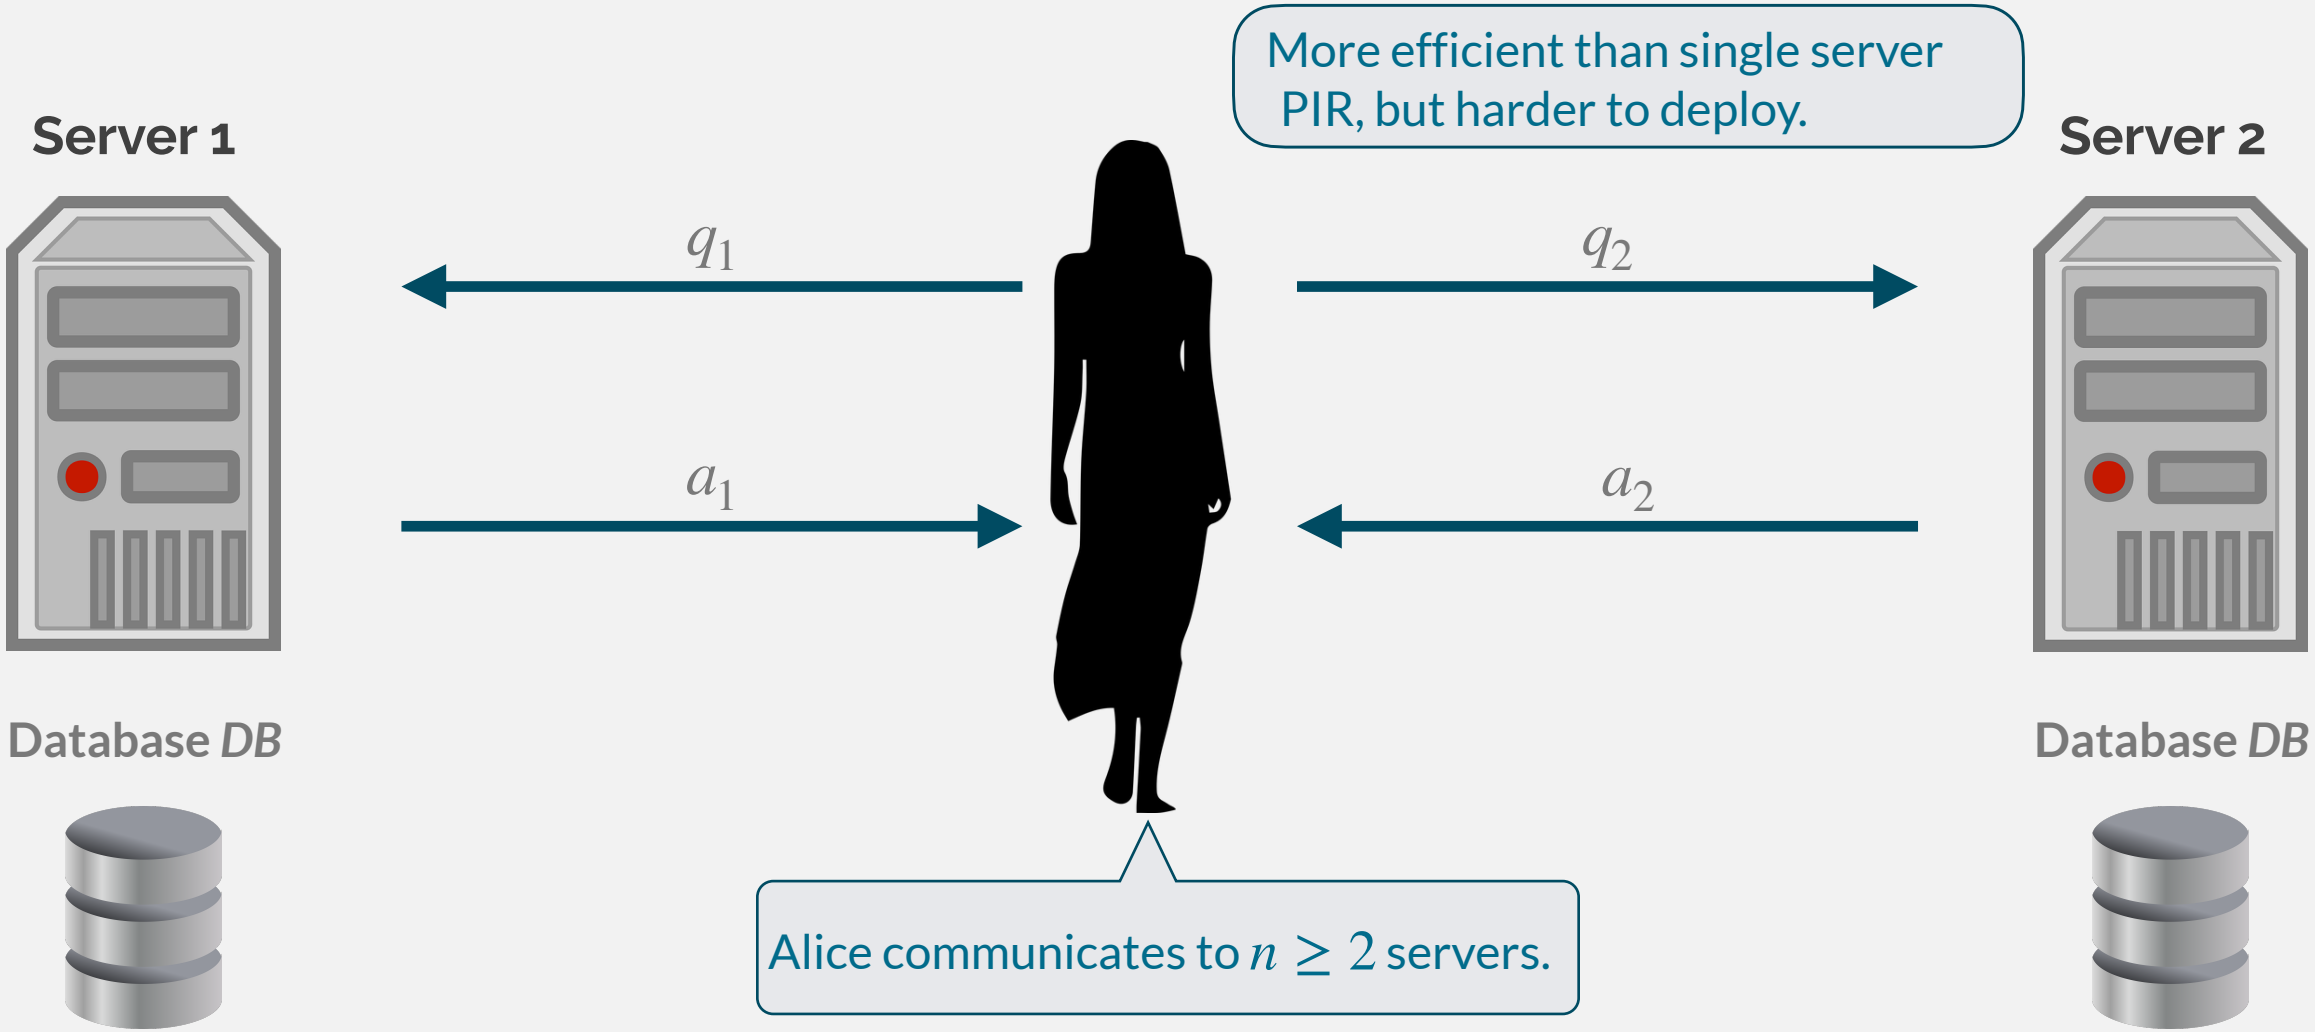

Classical PIR Model

input: index idx

input: database DB

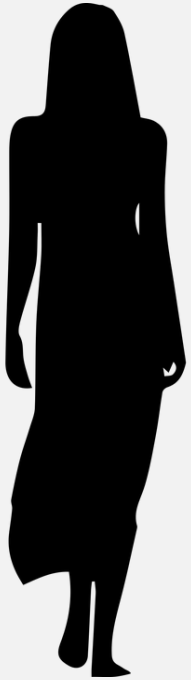

$q_1, \ldots, q_n$ xor up to B-bit zero bit string with single 1

$(q_1, \ldots, q_n) \leftarrow \text{Request}(idx)$

q_i

$a_i \leftarrow \text{Response}(q_i, DB)$

a_i

$d \leftarrow \text{Combine}(a_1, \ldots, a_n)$

Server i

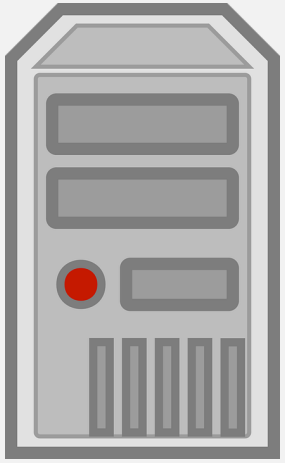

Database DB

DB divided into B blocks, usually $a_i = q_i \cdot DB$

Needs processing of whole DB in the online phase!

output:
$d = DB[idx]$

$a_1, \ldots, a_n$ xor up to $DB[idx]$

Our Contributions

Client-Independent Preprocessing PIR (CIP)

Preprocess at least half of the database before even knowing the client.

RAID-PIR [Demmler et al., CCSW'14] - Database

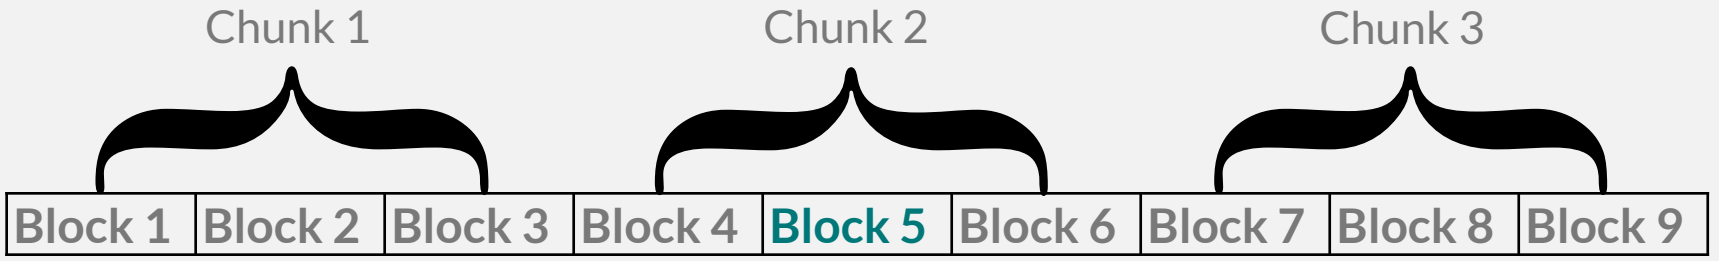

q_1	0	1	1	0	1	1	1	0	0
q_2	0	1	0	0	1	1	0	1	1
q_3	0	0	1	0	1	0	1	1	1
$\oplus Q$	0	0	0	0	1	0	0	0	0

Chunk i = flip chunk of server i

RAID-PIR [Demmler et al., CCSW'14] - Seed Expansion

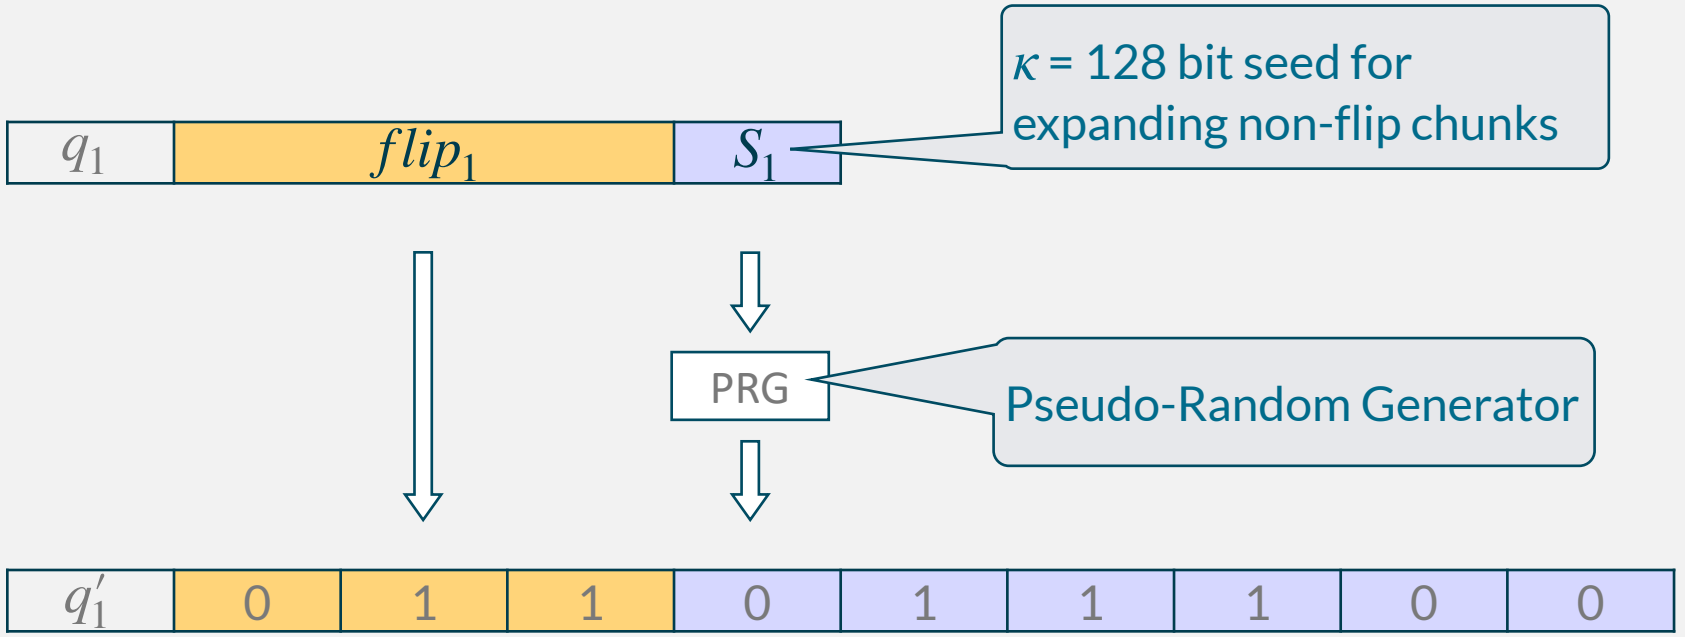

Client chooses seed S_i for each of the n servers.

Our idea: let the servers choose the seeds in advance!

Client-Independent Preprocessing (CIP)

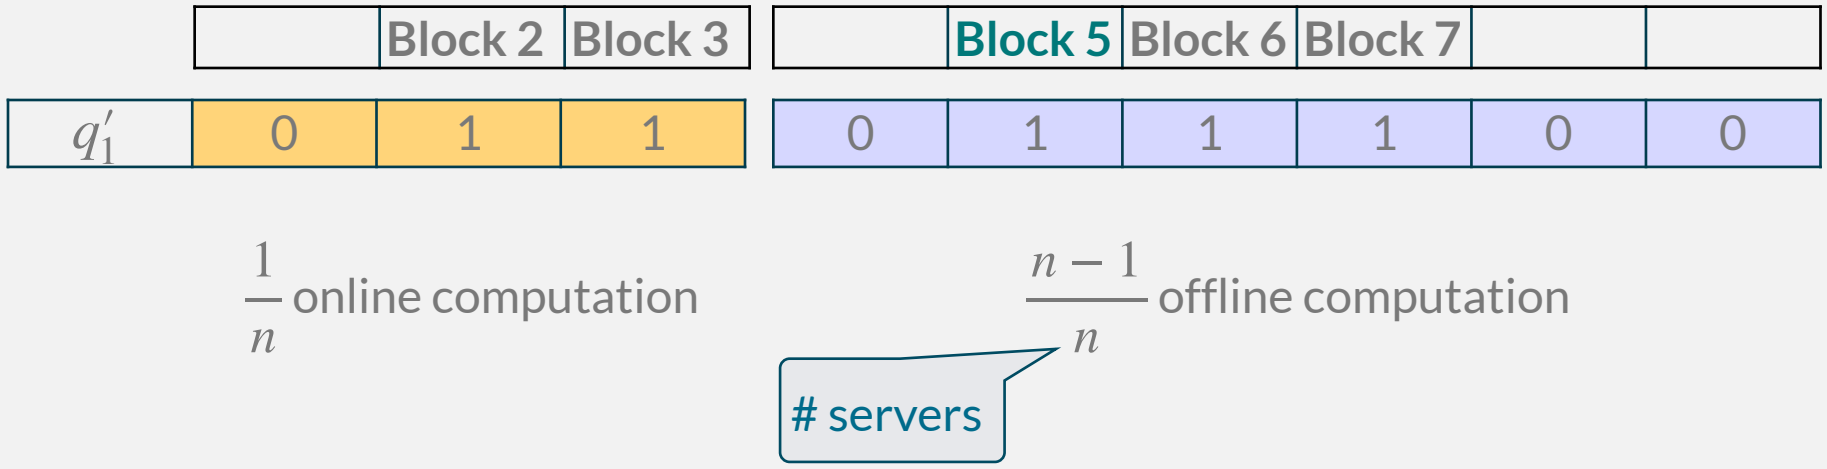

CIP-PIR Model

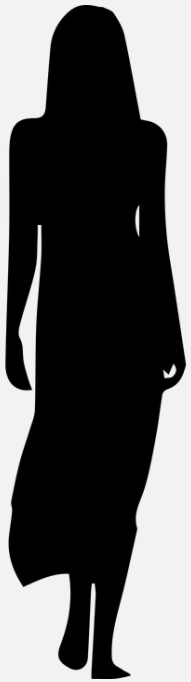

input: index idx

input: database DB, queue Q_i with (S_i, A_i)-pairs

Preprocessing independent of the client.

Server i

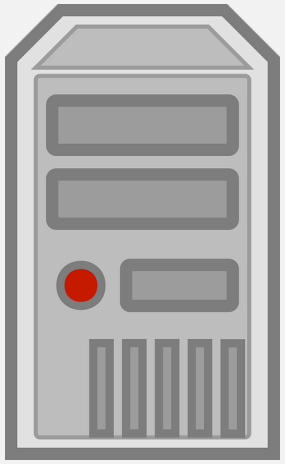

Database DB

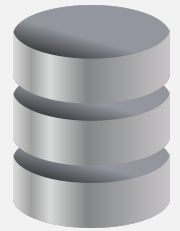

$(S_i, A_i) \leftarrow Q_i.pop()$

S_i

$(q_1, \ldots, q_n) \leftarrow \text{Request}(idx, S_i)$

q_i

$a_i \leftarrow \text{Response}(q_i, DB, A_i)$

a_i

$d \leftarrow \text{Combine}(a_1, \ldots, a_n)$

output:
$d = DB[idx]$

Our Contributions

Client-Independent Preprocessing PIR (CIP)

Preprocess at least half of the database before even knowing the client.

CIP-PIR Protocol

- Allows very efficient batch-preprocessing
- Only 1/n of the database needs to be touched in the online phase

CIP-PIR - Batch Preprocessing

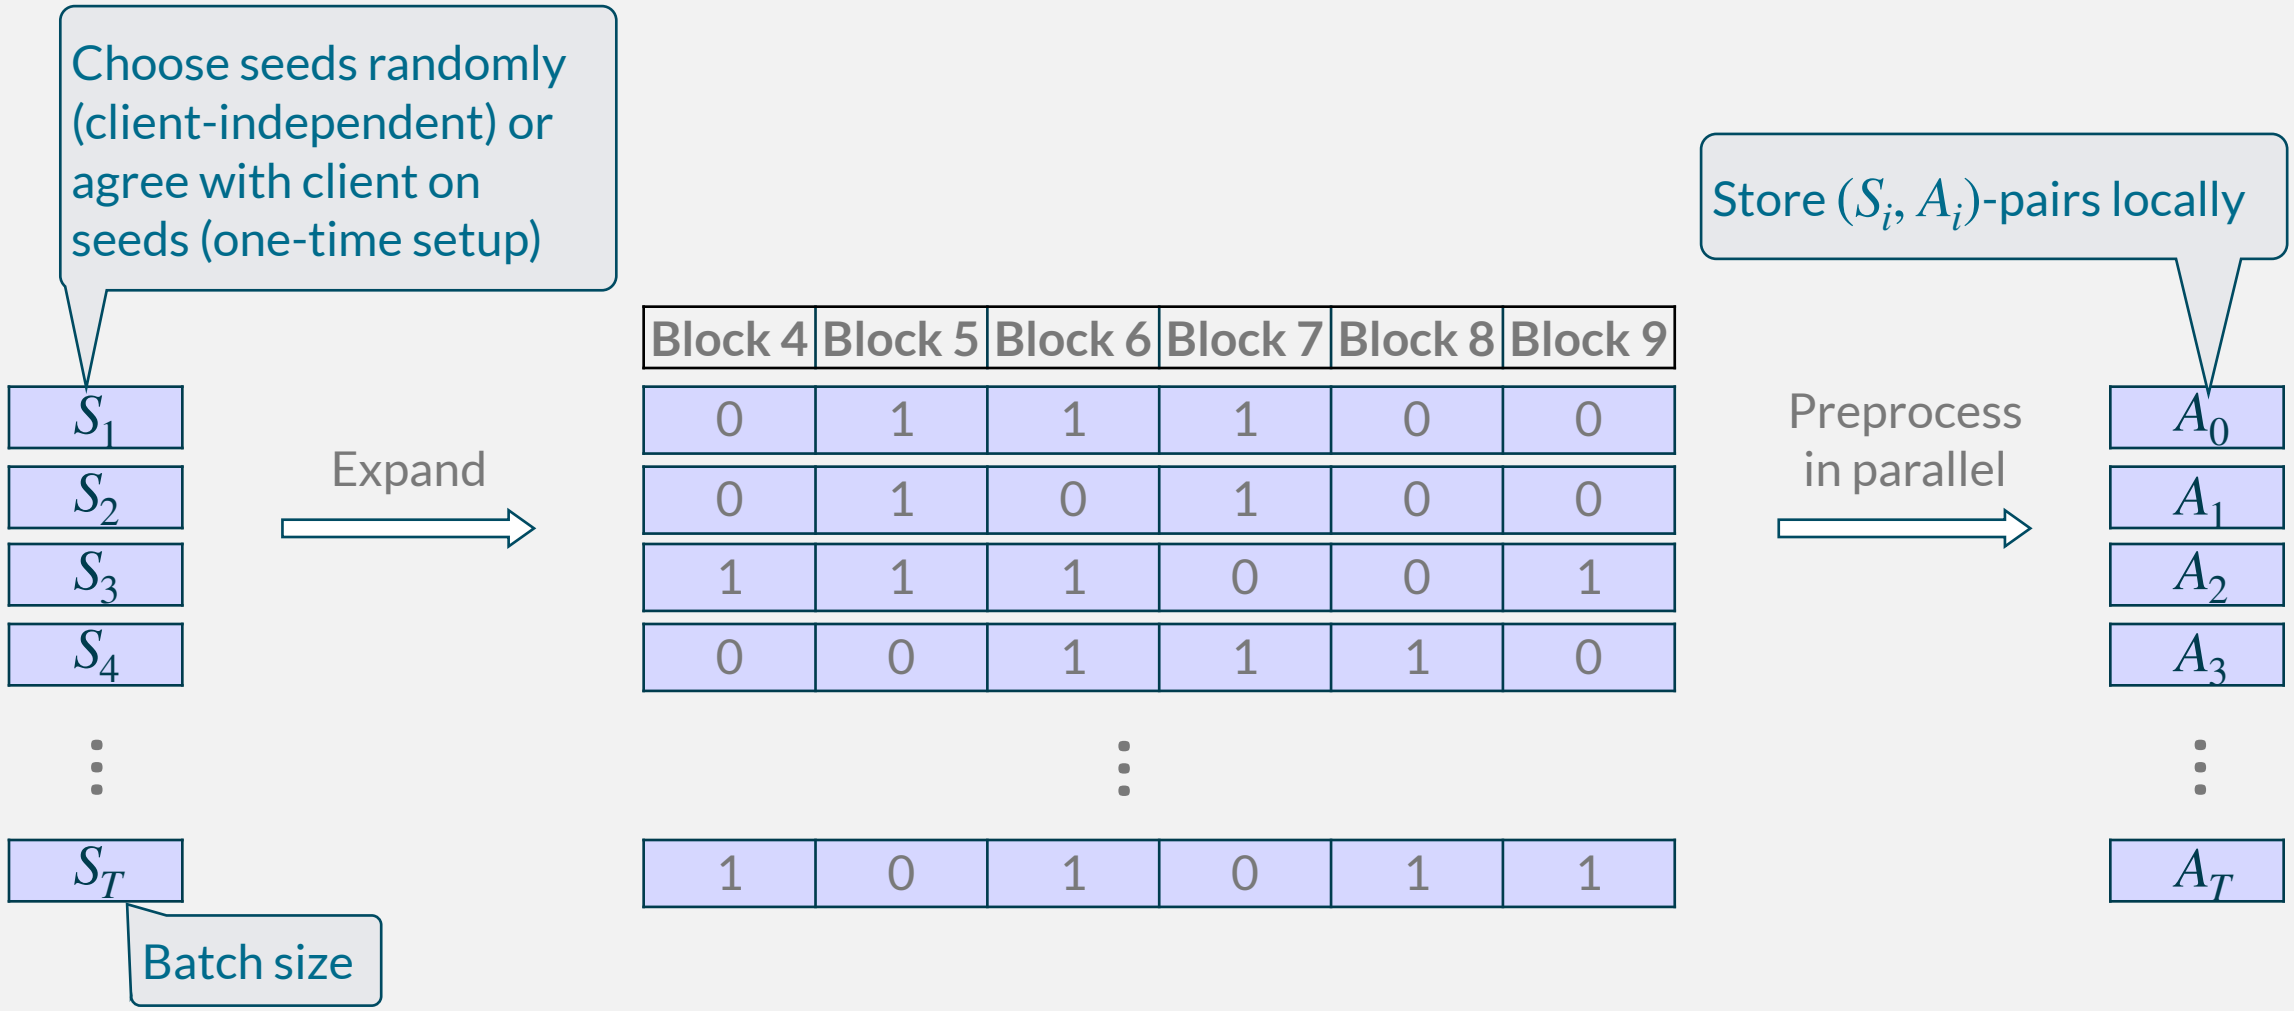

Our Contributions

ENCRYPTO
CRYPTOGRAPHY AND PRIVACY ENGINEERING

1

Client-Independent Preprocessing PIR (CIP)

Preprocess at least half of the database before even knowing the client.

2

CIP-PIR Protocol

- Allows very efficient batch-preprocessing
- Only 1/n of the database needs to be touched in the online phase

3

Database Compression in PIR

1. Shorter hashes
2. Difference sequence

Bar-Ilan University
אוניברסיטת בר אילן

TECHNISCHE UNIVERSITÄT DARMSTADT

Database Compression in PIR

ENCRYPTO
CRYPTOGRAPHY AND
PRIVACY ENGINEERING

Shorter Hashes

entries

Store only $\omega + \log_2 N$ bits of hash values

Statistical security parameter

Difference Sequence

Sorted Blocks

x_1	x_2	x_3	x_4	...	x_N

Improvement by factor 3.3x

Stored Blocks

x_1	$x_2 - x_1$	$x_3 - x_2$	$x_4 - x_3$	...	$x_N - x_{N-1}$

Bar-Ilan University
אוניברסיטת בר אילן

TECHNISCHE UNIVERSITÄT DARMSTADT

Our Contributions

ENCRYPTO
CRYPTOGRAPHY AND
PRIVACY ENGINEERING

1

Client-Independent Preprocessing PIR (CIP)

Preprocess at least half of the database before even knowing the client.

2

CIP-PIR Protocol

- Allows very efficient batch-preprocessing
- Only 1/n of the database needs to be touched in the online phase

3

Database Compression in PIR

1. Shorter hashes
2. Difference sequence

4

CIP-PIR Implementation

- CPU implementation parallelized via OpenMP
- GPU implementation

Bar-Ilan University
אוניברסיטת בר אילן

TECHNISCHE UNIVERSITÄT DARMSTADT

Benchmarks

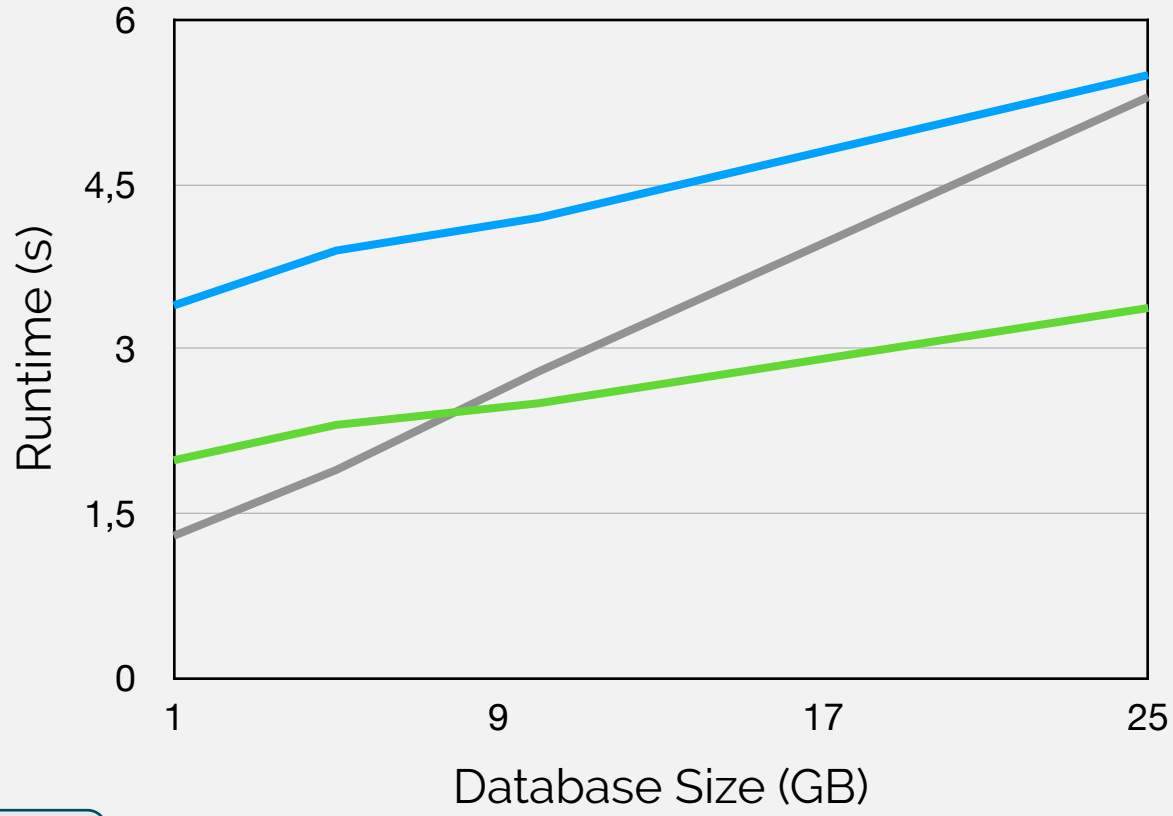

Communication Complexity:
$O(\sqrt{|DB|})$

- CIP-PIR (amortised total cost)
- CIP-PIR (online cost)
- FSS-PIR [Boyle et a., EUROCRYPT'15]

Communication Complexity:
$O(\log(|DB|))$

Our Contributions

ENCRYPTO
CRYPTOGRAPHY AND
PRIVACY ENGINEERING

1

Client-Independent Preprocessing PIR (CIP)

Preprocess at least half of the database before even knowing the client.

2

CIP-PIR Protocol

- Allows very efficient batch-preprocessing
- Only 1/n of the database needs to be touched in the online phase

3

Database Compression in PIR

1. Shorter hashes
2. Difference sequence

4

CIP-PIR Implementation

- CPU implementation parallelized via OpenMP
- GPU implementation

Bar-Ilan University
אוניברסיטת בר אילן

TECHNISCHE UNIVERSITÄT DARMSTADT

Code: https://encrypto.de/code/cip-pir

Thank you

Contact:

Daniel Günther
https://encrypto.de/dguenther

ENCRYPTO
CRYPTOGRAPHY AND
PRIVACY ENGINEERING