

Privacy Enhancing Technologies

5. PIR

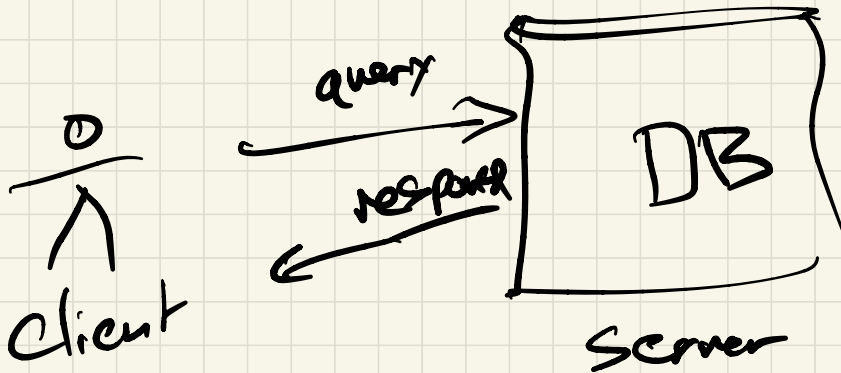
What's a perfect research problem?

1. A beautiful theory ✓
2. It works in practice ~
3. It solves a problem people care about ✓

What's PIR?

"A private Google"

Server learns query



What we want: query a DB
without leaking the query

Sol 1: Client $\xleftrightarrow{\text{MPC}}$ Server

↳ Communication: $O(|DB|)$

trivial Sol 2: Client $\xleftarrow{\text{DB}}$ Server
↳ look up DB

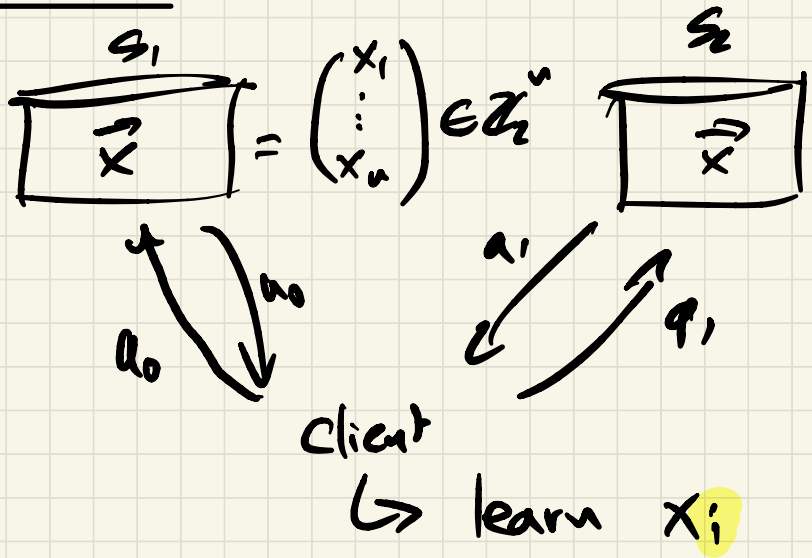
What we really want: sublinear
Search

Client \xrightarrow{q} Server
 $\xleftarrow{\text{res}}$
↳ $O(|DB|)$
↳ DB idx

2 Ways to get PIR:

- public-key crypto
- with non-colluding servers

2-Server PIR



Def: 2-Server PIR

$$q_0, q_1 \leftarrow \text{Query}(j)$$

$$a_i \leftarrow \text{Answer}(\vec{x}, q_i)$$

$$x_j \leftarrow \text{Reconstruct}(a_0, a_1)$$

• Correctness

$$\Pr \left[\text{Reconst}(a_0, a_1) = x_j \mid \right. \\ \left. = 1 \right]$$

$$\begin{aligned} q_0, q_1 &\leftarrow \text{Query}(j) \\ a_0 &\leftarrow \text{Answer}(\vec{x}, q_0) \\ a_1 &\leftarrow \text{Answer}(\vec{x}, q_1) \end{aligned}$$

• Security / privacy:

$$\forall j, j' \in [n], b \in \{0, 1\}$$

view of server b

$$\{ q_b : q_0, q_1 \leftarrow \text{Query}(j) \}$$


$$\{ q_b : q_0, q_1 \leftarrow \text{Query}(j') \}$$

• sublinear comm: $\forall j \in [n], \vec{x} \in \mathbb{Z}_2^n, \forall b$

$$|\text{Query}(j)| + |\text{Answer}(\vec{x}, q_b)| = o(n)$$

Client wants $x_i = \vec{x} \circ e_i$
 \leftarrow with pos
 $[0, \dots, 0, 1, 0, \dots, 0]$

secret-share $e_i = q_0 \oplus q_1$
 \hookrightarrow random in \mathbb{Z}_2^n



$$q_0 \oplus q_1 = \vec{x} (q_0 \oplus q_1) = x_i$$

This isn't sublinear: $q_0, q_1 \in \mathbb{Z}_2^n$

This is non-trivial if each element in the DB has $\gg 1$ bit

Idea: load balancing

Protocol 1 : $\xrightarrow{O(1) \text{ upload}}$
 $\xleftarrow{O(n) \text{ download}}$

Protocol 2 : $\xrightarrow{O(n) \text{ upload}}$
 $\xleftarrow{O(1) \text{ download}}$

Protocol 3 : $\xrightarrow{\Gamma_n \text{ upload}}$
 $\xleftarrow{\Gamma_n \text{ download}}$

How? view DB as matrix $X \in \mathbb{Z}_2^{\Gamma_n \times \Gamma_n}$

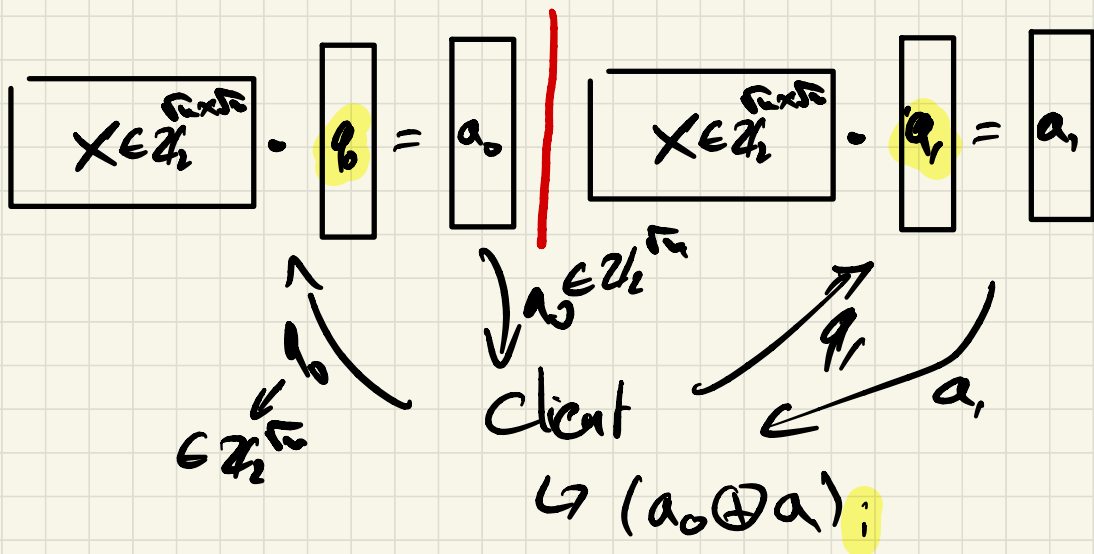
$$X = \begin{bmatrix} x_0 & \dots & x_{\Gamma_n} \\ x_{\Gamma_n+1} & \dots & \dots \\ \vdots & & \vdots \\ & & x_{n-1} \end{bmatrix}$$

- 1) do PIR to learn one of Γ_n cols of X
- 2) do "trivial PIR" on that column

Query (i, j) : $e_j = q_0 \oplus q_1$


Answer (x, a_b) : $a_b \leftarrow X \cdot q_b$

Recons (a_0, a_1) : $(a_0 \oplus a_1)$



total comm: $4 \cdot \sqrt{n}$ bits

Single Server PIR



$$q, sk \leftarrow \text{Query}(i)$$

$$a \leftarrow \text{Answer}(\vec{x}, q)$$

$$x_i \leftarrow \text{Recover}(a, sk)$$

Our strawman scheme

$$\hookrightarrow x_i = \vec{x} \cdot e_i$$

$$\hookrightarrow (q_0 \oplus q_1)$$

Idea: use encryption instead of secret sharing

We need linearly homomorphic encryption

$$E(k, m_1) + E(k, m_2) = E(k, m_1 + m_2)$$

back to strawman scheme:

$$\vec{q} = E(k, \vec{e}_j) = [E(k, 0), \dots, E(k, 1), \dots, E(k, 0)]$$

(u.x) kids

$$\begin{aligned} \vec{x} \cdot \vec{q} &= x_1 \cdot E(k, 0) + \dots + x_j \cdot E(k, 1) + \dots \\ &\quad E(k, 0) \cdot x_n \\ &= E(k, x_j) \end{aligned}$$


Load balancing again

$\leftarrow \{e_{ij}\}$
 $\leftarrow \frac{0}{1}$

$$q = E(k, \vec{e}_j)$$

$\leftarrow \{E(k,0), \dots, E(k,1), \dots, E(k,\sqrt{n})\}$

$\leftarrow O(\lambda \sqrt{n})$ bits



$$\left[X \right] \in \mathbb{Z}_2^{\sqrt{n} \times \sqrt{n}}$$

$$\text{Dec}(k, a); \\ = x_{ij}$$

$$a = Xq = X \cdot E(k, \vec{e}_j) \\ = E(k, X \vec{e}_j) \\ = E(k, x_{ij})$$

Comm: $|q| + |a| = 2\sqrt{n}$ ciphertexts
 $= O(\lambda \sqrt{n})$ bits! ▽

More than 1 recursion is possible

view X as d -dimensional cube
of side $d\sqrt{u}$

limit: $O(\underline{\text{poly}}(\log u))$ comm!

One issue: ciphertext blow up

$$a_0 = X \cdot E(k, \vec{e}_j) = E(k, x_j)$$

$$a_1 = a_0 \cdot E(k, \vec{e}_i) = E(k, a_0 \cdot \vec{e}_i)$$

$$= \underline{E(k, E(k, x_{ij}))}$$

Why does this matter?

Suppose your encryption scheme
encrypts t bits into $O(\lambda + 2t)$ bits

after $\log(n)$ recursions, you get
a ciphertext of $O(n)$ bits

We need encryption schemes
with low (constant) expansion

Q: Does this recursive PIR idea
work for 2-server PIR?

Modern PIR

- 2 server: - Info theoretic: $n^{o(1)}$
- Computational: $O(\log u)$
- Single server: $O(\text{poly}(\log u))$


Computational complexity

All our schemes have server complexity $O(u)$

Maybe interesting: if server doesn't "touch" x_i , then the query can't be i

Solas: (1) Batching: q queries
in $O((1 + \log q)u)$
work

2) pre-processing



LMW23 :

