

Lecture 5: Offline/Online PIR

MIT - 6.893
Fall 2020
Henry Corrigan-Gibbs

Plan

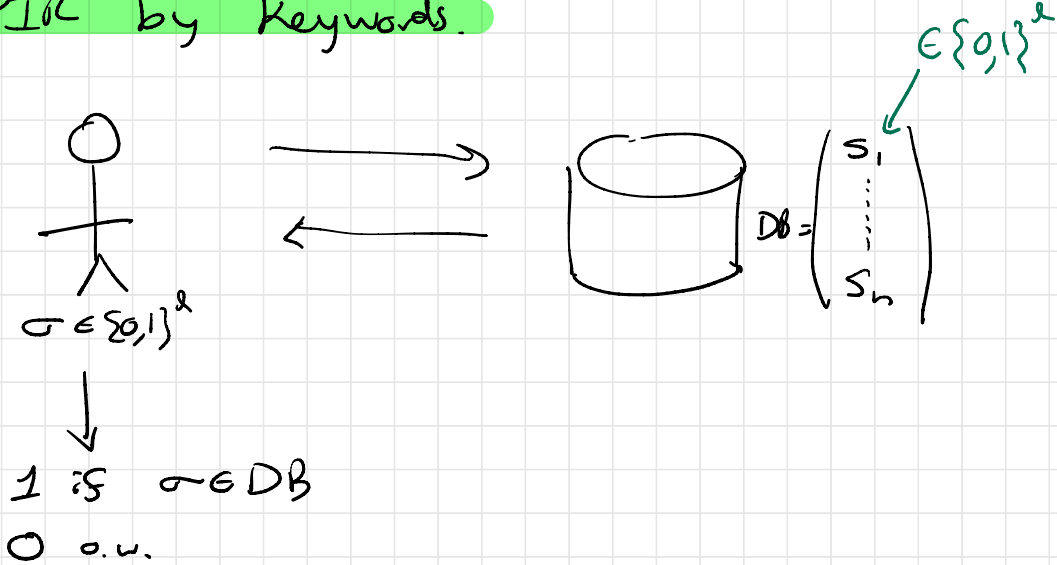
- Recap: Batch PIR (+ PIR by keywords)
- Computation in PIR
- Stretch break
- Offline/Online PIR

Logistics

- Last chance for HW1: today 5pm
- HW2 Posted
↳ Start early!
I am not joking!
Due 10/2 5pm via Gradescope
- Breakout sessions:
will aim for 1/week
- Guest lecture on W... please look at reading & bring one Q.
- Computer issues...

Recap:

PIR by keywords.




Bottom Line: Cost of PIR by keywords
 \approx
Cost of normal PIR

- Approach: Hash twice...
- Once into buckets (few collisions)
 - Once within buckets (no collisions.)

Recap: Batch PIR

Idea: Answer q queries at server-side cost of answering 1 query. to save DB

Observation:



1 query to DB
of size n

q queries to DB of
size n/q .

Strategy: If client wants to make q queries, partition DB into q chunks at random.

As long as client's desired bits fall in diff chunks \Rightarrow can recover.

\Rightarrow A little more work gets correctness
w) \ll but negl prob.

Server-Side Computation in PIR

BIM'04: If PIR servers store DB in original form, servers must probe every DB bit in responding to client's query.

Intuition: Servers don't touch bit i^*
 \Rightarrow Client is probably not reading bit i^* .

Takes some work to make precise.

How to get around?

1. Batching: Amortize cost of linear scan over many queries.
2. Preprocessing: Server does linear scan in a preprocessing phase.

\hookrightarrow Per client (today) — Offline/online
 \hookrightarrow Per database (HW) — "PIR w/ Preprocessing"

\Rightarrow Active area of research.
Lots to do still...



Offline-Online PIR (with D. Kogan)

↳ Will discuss two-server setting... also makes sense in single-server setting

Idea: Push heavy work to an **offline phase**... takes place before client even knows which DB element it wants

→ Push heavy linear scan to a more convenient time (out of critical path)

→ Servers can still store DB in unmodified form

↳ Other approaches that you'll see on HW blow up server storage
 n -bit DB \rightarrow n^3 -bit preprocessed data structure

Nice open Qs here on whether it's possible to avoid these blowups

Notation: $\tilde{O}(f(n)) = O(f(n) \cdot \text{poly}(\log(n)))$

The background of the page is decorated with several yellow, hand-drawn squiggly lines of varying lengths and curves, scattered around the text.

Stretch

Break

Will present a two-server scheme
with computational security (PRG)

comm $\delta(\sqrt{n})$
online time $\delta(\sqrt{n})$

We show that for PIR schemes in
which servers store DB in unmodified form

$$\begin{array}{l} \text{Comm. } S \\ \text{online time } T \end{array} \Rightarrow S \cdot T \geq \sqrt{2}(n).$$

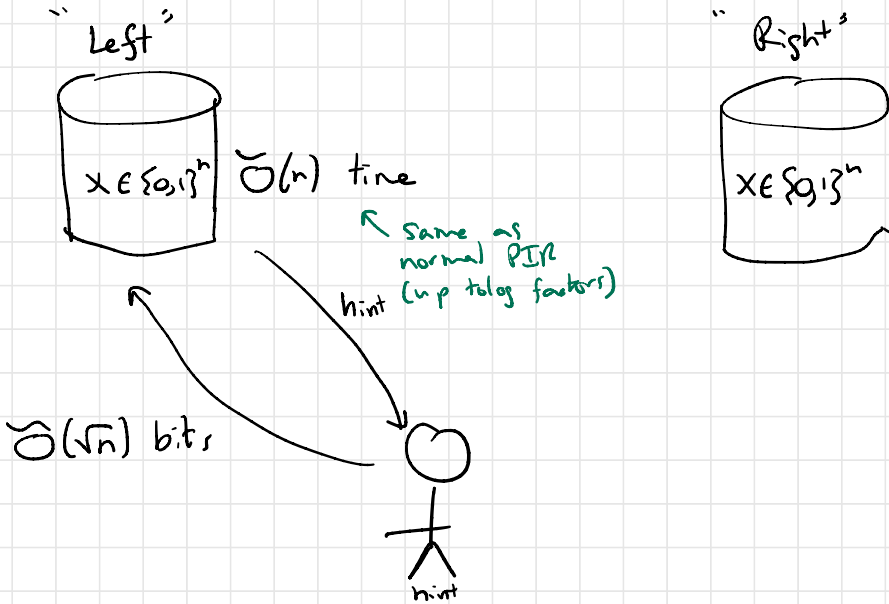
\Rightarrow This \sqrt{n} scheme is optimal w.r.t. these parameters.

- * Can get info-theoretic security.
- * Can get perfect correctness
- * Can reduce client running time.
- * Can reduce online communication.
- !
- !

See paper for
many of these
extensions.

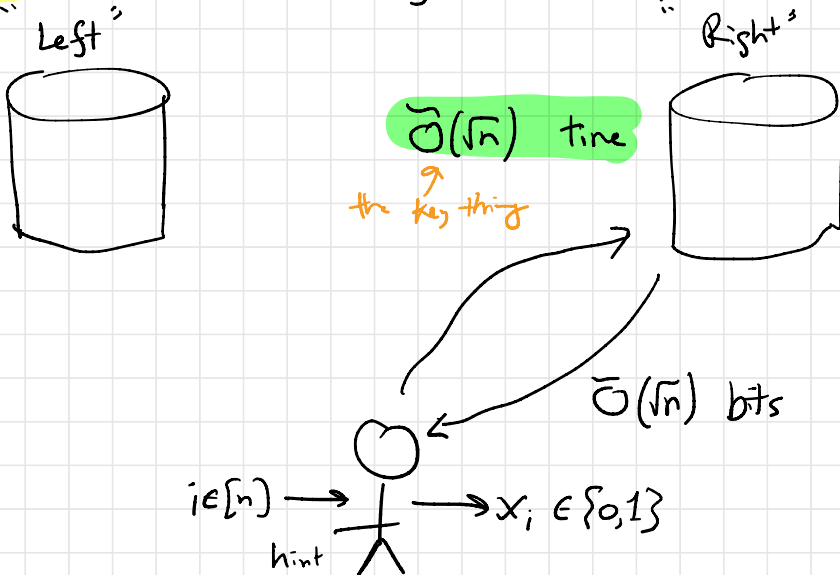
OFFLINE

(think: overnight batch job)




ONLINE

(think: when you browse web)



PIR Security & correctness properties are as before.

Construction: Offline Phase



② Server sends parity of DB bits in each set


$$h_1 = \sum_{i \in S_1} x_i \pmod{2}$$

$$\vdots$$

$$h_{\sqrt{n}} = \sum_{i \in S_{\sqrt{n}}} x_i \pmod{2}$$

① Choose a random partition of $\{1, \dots, n\}$ into \sqrt{n} sets $S_1, \dots, S_{\sqrt{n}}$, each of size \sqrt{n} .

$S_1, \dots, S_{\sqrt{n}}$
(compress using pseudo-randomness)




③ Client stores his and S_i 's as hint.

Construction: Online Phase


① Find set S_j st. $i \in S_j$.

② Construct set S'_j *see next page
send to server



③ Reply w/ parity of DB bits in set.

$$a \leftarrow \sum_{i \in S'_j} x_i$$



$h_1, \dots, h_{\sqrt{n}}$

$S_1, \dots, S_{\sqrt{n}}$

④ output $x_i \leftarrow h_j \oplus a$.

How client constructs S' ...

Flip a coin w/ $p_{\text{heads}} = \frac{\sqrt{n}-1}{n}$ $\left[\frac{1}{\sqrt{n}} \right]$

↳ If heads, choose random $i \leftarrow S_j \setminus \{i\}$
Set $S' \leftarrow S_j \setminus \{i\}$
Output "fail".

↳ If tails, set $S' \leftarrow S_j \setminus \{i\}$.

Why this works...

Correctness: With probability $1 - p_{\text{heads}} \approx 1 - \frac{1}{\sqrt{n}}$...

$S' = S_j \setminus \{i\}$, so client outputs

$$h_j \oplus a = \left(\sum_{\ell \in S_j} x_\ell \right) + \left(\sum_{\ell \in S'} x_\ell \right) \pmod{2}$$

$$= \left(\sum_{\ell \in S_j} x_\ell \right) + \left(\sum_{\ell \in S_j} x_\ell \right) + x_i \pmod{2}$$

$$= x_i.$$

⇒ Scheme fails w.p. $\approx \frac{1}{\sqrt{n}}$.

Repeat the whole thing λ times in parallel to ensure that 1 executor succeed w.p. $\geq 1 - \left(\frac{1}{\sqrt{n}} \right)^\lambda \geq 1 - 2^{-\lambda}$.

Security

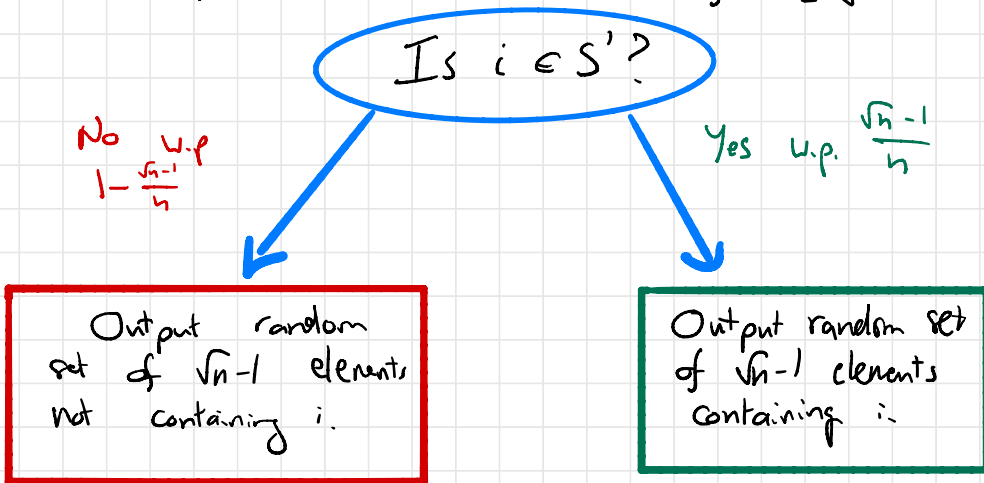
(Left server)

Sees only the random partition S_1, \dots, S_n
↳ independent of client's index: ✓

(Right server)

Claim: Set S' is a set of $\sqrt{n}-1$ random elements chosen w.o. replacement from $[n]$.

Here is one (funny) way to sample a set S' of $\sqrt{n}-1$ random elements from $[n]$.



This is exactly how client constructs S'

⇒ S' is a random set, indep of i .

⇒ Info theoretic security

Efficiency

* Need to repeat 1 times to drive failure prob $\rightarrow 0$.

* Offline cost

Server time : $1/n$ random reads
Comm : upload: 1 PRG seed
 : download: $1/\sqrt{n}$ parity bits
Storage : $1/\sqrt{n}$ parity bits \rightarrow seed

* Online time

Server : $1/\sqrt{n}$ random reads
Comm : $1/\sqrt{n}$ bits \leftarrow can reduce to $\text{poly}(1, \log n)$
Client time : ???

\leftarrow Not clear how to reduce.

Standard PR server $n/2$ XORs over DB
 in linear scan

Off/on PR affine server $1/n$ XORs in random order

$\rightarrow \sim 128x$ cost $1/2$ of 1

$\rightarrow \sim 10x$ cost $1/2$ of random reads



Extensions

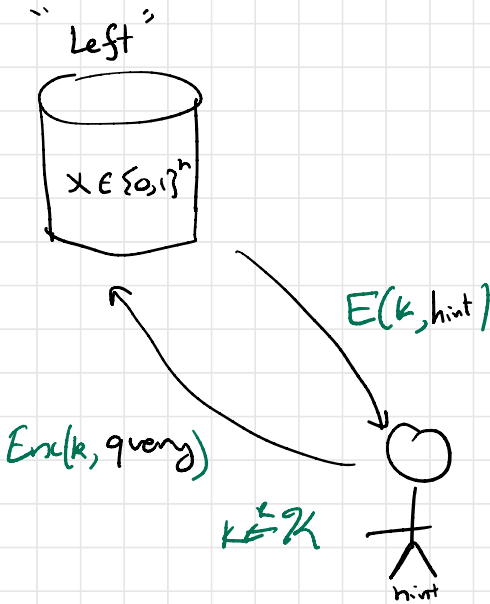
- Can reuse one hint for many subsequent online queries (requires tweaks)

⇒ Amortized server cost after q queries

$$\frac{\tilde{O}(n)}{q} + \tilde{O}(\sqrt{n}) \rightarrow \tilde{O}(\sqrt{n}) \text{ for } q \text{ large.}$$

- Works in single-server setting

OFFLINE



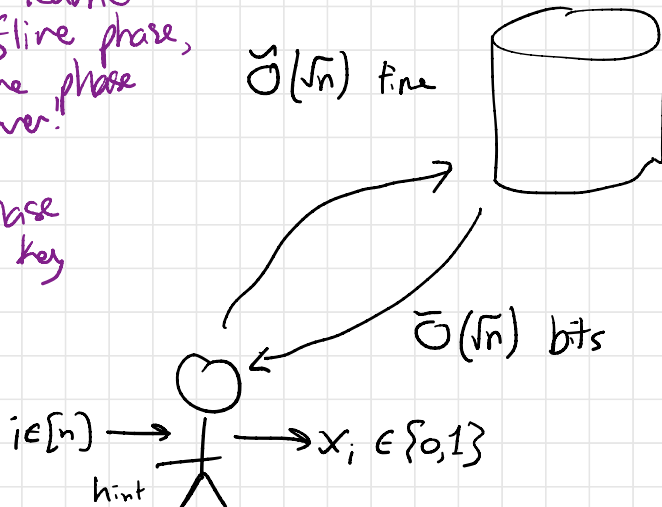
Idea: Offline phase happens "under encryption."

Can do w/ FHE or (worse parameters) lin. hom. ENC.

ONLINE

Since server learned nothing in offline phase, can run online phase with same server.

Also, online phase requires no pub key ops.



Moral of the story:

1. Can reduce online server time in PIR using preprocessing.

↳ Still new things to say about an old problem.

2. Still not clear whether these techniques will work in practice at scale.

3. Competition is no privacy.

↳ How can we ever outperform that?



This is the important question to think about as you are working on crypto + privacy tech.