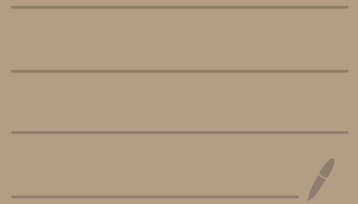


Privacy Enhancing Technologies

lecture 6: ORAM



Oblivious RAM

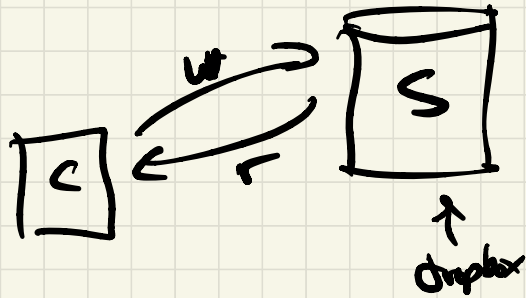
see Moodle

- PIR
 - ORAM
- } close to production
still very interesting
research developments

ORAM: privately read & write to
some database

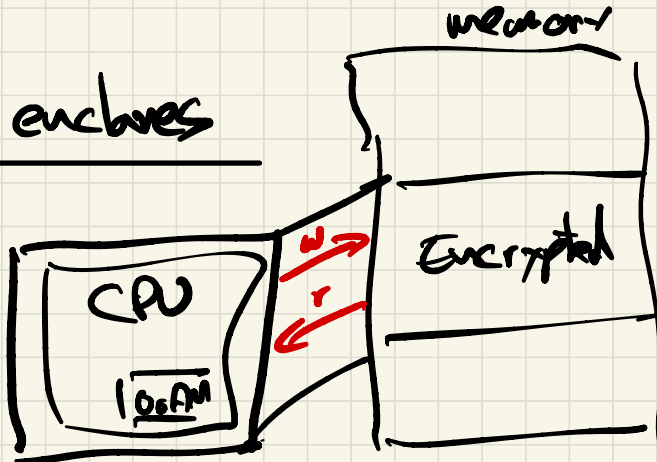
Examples

Private Dropbox

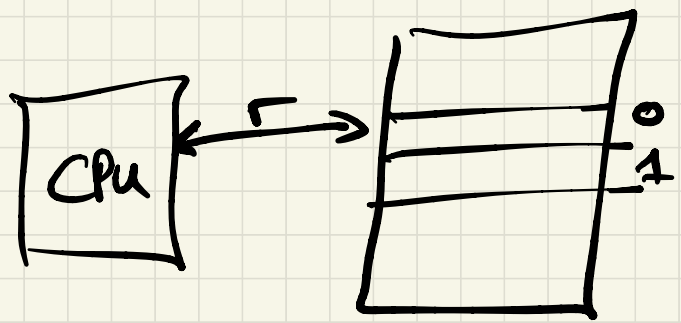


Hardware encloses

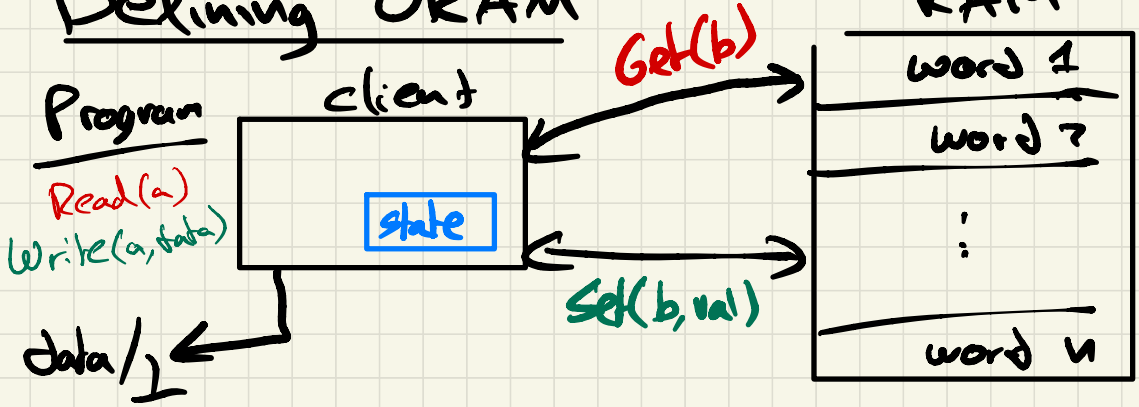
Intel
SGX



```
if (secret): ← encrypted
    read mem[0] ← encrypted
else
    read mem[1] ← encrypted
```



Defining ORAM



Correctness : The program behaves correctly (as if it was directly accessing RAM)

Privacy : the access patterns to the RAM leak nothing about the program's read & writes

Def : • An operation $op \in \{ \text{read}(\cdot), \text{write}(\cdot) \}$
• $\text{Access}(op)$: set of RAM access made for op

Security/Privacy

$$\text{let } O = \{op_1, op_2, \dots, op_k\}$$

$$O' = \{op'_1, op'_2, \dots, op'_k\}$$

read(0)

write(7, "hello")

$$\{Access(op_1), \dots, Access(op_k)\} \quad \mathcal{N}^c$$

$$\{Access(op'_1), \dots, Access(op'_k)\}$$

"Trivial", or naïve solutions (RAM with n words)

(- local storage : n)

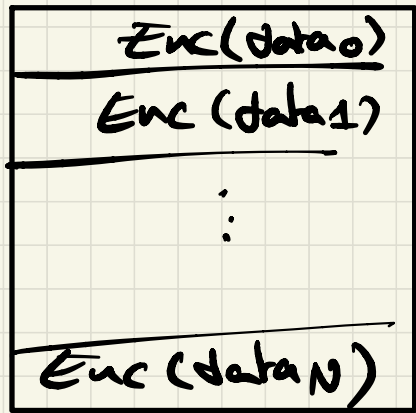
(- access overhead per op. : O)

↳ put all the RAM in the client storage

(- local storage : $O(\lambda)$ bits + $O(1)$ words

(- overhead : n RAM accesses per op.)

read(i)
write(i, data) } }



- read all the RAM one-by-one

⊖ re-encrypt each word

- and re-write to the RAM

What we can do :

- n^ϵ local storage (for constant $\epsilon > 0$)

- $O(\log n)$ accesses per op.
↑
 ϵ

PIR VS ORAM

Private reads from
public DB

Private reads & writes
to private DB/RAM

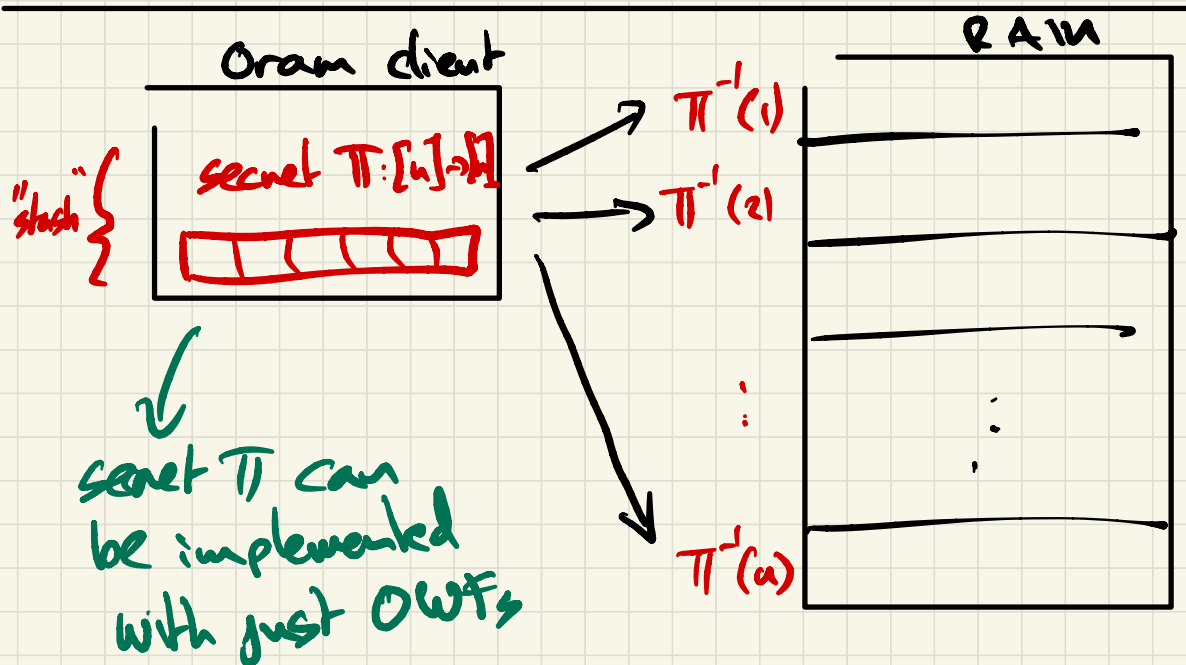
- DB is public & static
- Many clients \leftrightarrow 1 server
- client is stateless
- You need PK crypto
in single-server setting

- DB is private
and changes w.
every op
- 1 client \leftrightarrow 1 server
- client must have
state
- Can be built
from just OUP?

A \sqrt{n} Construction

* local storage : $O(n)$ bits
 $+ O(\sqrt{n})$ words

* overhead per op : $O(\sqrt{n})$
(amortized)



- Read(i), Write(i , data) \Rightarrow read $\pi(i)$ & write $\pi(i)$
- store (i , data) in local stash

- read/write something that is already in the stash:
 - do the op on the stash
 - do a random read/rewrite on RAM

- this works until your stash is full ...

↳ say you evict (i) from stash

↳ next time you have to access (i) you leak to the RAM that you're re-accessing the same memory

- when the stash is full, we re-shuffle the entire RAM ☹

⇒ How to do this with little overhead?

Oblivious sorting

How to sort data with uniform access patterns

Bad example: merge sort

Inefficient example: Bubble sort

$e_1 \quad e_2 \quad e_3 \quad e_n \quad \dots \quad e_n$
↔ ↗
 $O(n^2)$ oblivious sort

Ideal: oblivious $O(n \log n)$ sort

Practical: Batched sort $O(n \log^2 n)$
oblivious

The protocol

- init:
- client inits a stash of \sqrt{n} words
 - RAM is init. to $\text{Enc}(0)$

while (true):

- client sample a permutation π
- client shuffles the RAM according to π
↳ empty the stash
- Process \sqrt{n} ops:
 - if addr is in stash: do a random RAM read and store
 - otherwise, read addr to stash
 - if write, do it on the stash

Per \sqrt{n} ops, overhead is:

- once: $O(n \log^2 n)$

- \sqrt{n} times: 1

\Rightarrow Amortized overhead per op.
is $\tilde{O}(\sqrt{n})$

\nwarrow ignoring $\log(n)$ factors

Tree-based ORAMs

↳ $O(\text{poly}(\log(n)))$ overhead

↳ Here: "Simple ORAM"

$$O(\log^4 n)$$

How this works

1) build a "bad" ORAM
with stash $\frac{n}{2}$ and
 $O(\log^3 n)$ overhead

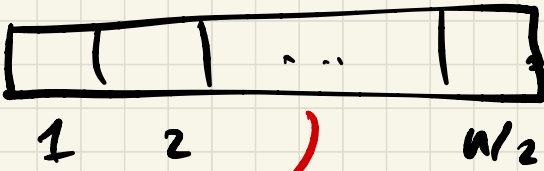
2) Recurse! Store the stash
in another ORAM

3) Repeat \log times
- constant stash
- overhead $O(\log n) \cdot O(\log^3 n)$

building a BAD ORAM

Client

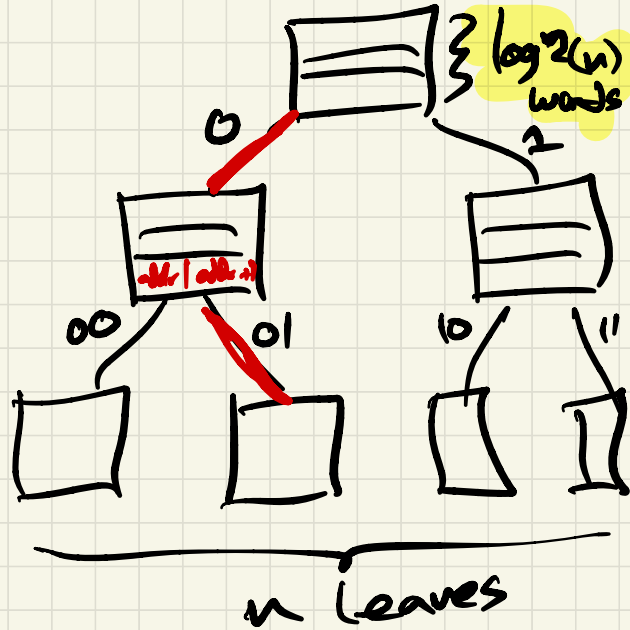
"position map"



$$b = \left\lfloor \frac{a \text{ addr}}{2} \right\rfloor$$

$$\boxed{\text{pos} = 01}$$

RAM



Invariant: The data at address

a is somewhere on the p th from the root to the leaf node indicated at position $\left\lfloor \frac{a}{2} \right\rfloor$ in the position map

How this works?

Read (addr):

- look up pos at element $\lfloor \frac{addr}{2} \rfloor$ in the position map
- read $\log^2(u) \cdot \log(u)$ elements of RAM from root to leaf pos
↳ [addr: val, addr+1: val']
- pick a new random leaf pos' and update position map $\lfloor \frac{addr}{2} \rfloor \rightarrow pos'$
- re-encrypt (addr, val) (addr+1, val') and insert this into root node
- if Overflow: walk down the path from root to pos' and "flush" all elements down the tree
↳ if this overflows, abort 😞