ORAM Lower bound [Larsen, Nielsen '18]

[GO '96]

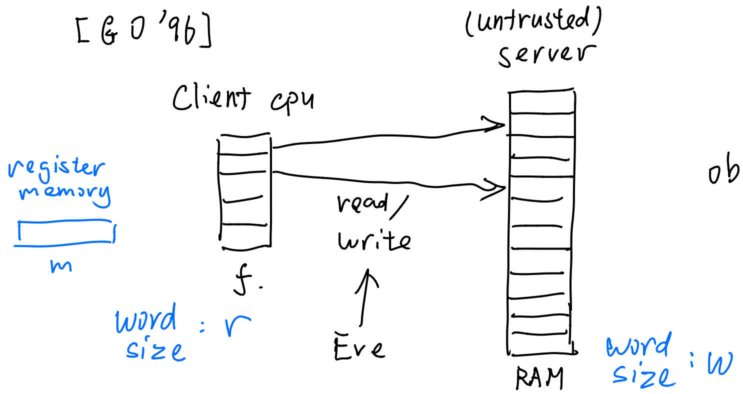

obfuscate access pattern.

N accesses take M access to ORAM

(bandwidth) overhead $= \frac{M \cdot w}{N \cdot r}$

Upper bounds

[GO '96] poly log N

[Stefanov, van Dijk, Shi, Chan, Fletcher, Ren, Yu, Devadas '13]

$O(\log N)$ but $w = O(\log^2 N)$ $r = w^2$ Path ORAM

[PPRY '18] PanORAMa $O(\log N \log\log N)$ $w = O(\log N) = r$

[AKLNPS '18] OptORAMa $O(\log N)$ $w = O(\log N) = r$

Lower bound

[GO '96] $\Omega(\log N)$

- "balls and bins" : the algorithm can't read the contents
- statistically secure : unbounded adversary

GO '96, and many constructions followed are only computationally secure.

[Boyle, Naor '16] Is there an ORAM Lower bound?

(only public randomness)

Thm. Suppose there is a circuit sorting n words with w-bits, with size $o(n w \log n)$, then there exists offline ORAM compiler with overhead $o(\log N)$

(offline) ORAM lower bound → efficient sorting circuits

× balls-bins

online ORAM?

[Larsen, Nielson '18] Yes, there is an ORAM Lower Bound!

$\Omega(\log N)$ for online ORAM
- any algorithm
- computationally secure
- any block size w

$\Omega(\log(Nr/m))$ r: word size for client

m: total memory bits for client

$O(1)$ blocks, $r \le m \le n^{1-\varepsilon} \Rightarrow \log n$

Array maintenance problem for dynamic array

$\begin{cases} \text{write}(i, \text{data}), \ \text{data} \in \{0,1\}^r \\ \text{read}(i) \end{cases}$ $i \in [n]$

re-use data structure lower bounds

U: updates Q: queries   cell probe model

all write(i, data) all read(i) [Pătraşcu, Demaine '13]

Idea:

1) ops: $\underbrace{w(1,1) \quad w(2,2)}_{N/2} \quad \underbrace{r(1) \quad r(2)}_{N/2}$

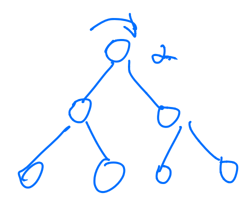

"access prev. accessed cells" must happen $N/2$

2) ops: $w(0,0)$ $w(0,0)$ $r(0)$ $r(0)$ (counting only, easy)

also need $N/2$ accesses o.w. security breach

3) $\underbrace{w(1,1) \ r(1)}_{N/4 \quad N/4}$ $\underbrace{w(2,2) \ r(2)}_{N/4 \quad N/4}$

$N/4$ accesses $N/4$ accesses

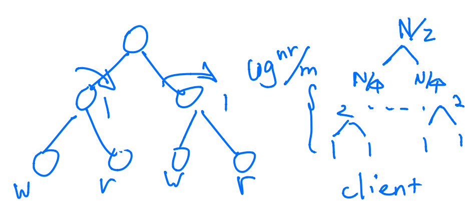

adversary distinguishes this

$$\begin{cases} w(0) & r(0) & w(0) & r(0) \\ w(0) & w(0) & r(0) & r(0) \end{cases}$$

$N/2$ "transfers"

$N/4$, then $N/4$ "transfers"

memory m

[PD '16]

Oblivious Cell Probe

Complexity: amortized over M operations

probes in expectation over $r \in \{0,1\}^\ell$ uniform

Security: $y := (op_1, \dots, op_M)$ op sequence

$A(y) := (A(op_1), \dots, A(op_M))$ probe sequence

$y \neq z$ $A(y), A(z)$ $\times$ distinguished w/ prob. $> 1/4$

in poly time $\log|U| + \log|Q| + w$

Correctness: fail prob. $< 1/3$

Thm. $\exists\, y = (op_1, \dots, op_N)$ $op_i \in U \cup Q$ s.t.

assuming security holds, takes $\Omega(N \log(Nr/m) \cdot r/w)$

(bandwidth overhead Mw/Nr) $\log(Nr/m)$

for $r \le m \le N^{1-\varepsilon}$ $\log N$

$$y := \underbrace{w(0,0)\ r(0)\ \cdots\ w(0,0)\ r(0)}_{M = 2n \text{ operations}}$$

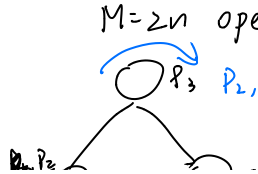

P_2, P_4 information transfer

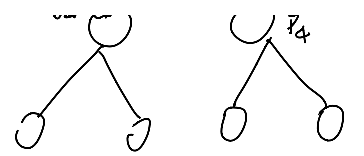

$P_1=0 \quad P_2=0 \quad P_3=0 \quad P_4=0$ Probe sequence

$P_1=5 \quad P_2=2 \quad P_3=2 \quad P_4=5$

$T(y)$ denote the tree for y

$T(z)$ denote the tree for $z := w(i_1, d_1)\ r(i_2) \cdots w(i_{2n-1}, d_{2n-1})\ r(i_{2n})$

Fix v, depth d.

Def $P_v(y)$ # probes assigned to v in tree y

assuming $<1/32$ fail prob. "information transfer" ← large for y

Lemma 1: $\mathbb{E}(|P_v(y)|) = \Omega(nr/w2^d)$ for $d \le \frac{1}{2}\log\frac{nr}{m}$

then $\mathbb{E}(|A(y)|) \geq \sum_v \mathbb{E}|P_v(y)| = \Omega(n \log(nr/m) \cdot r/w)$

Take Z_v random op sequence in the form of z

Lemma 2: Assume $< 1/32$ fail. prob. $\exists$ universal constant C

$$\mathbb{P}(P_v(Z_v) \geq Cnr/w2^d) \geq 1/2$$

information transfer is large for random z on each node v

Consequence: $\exists\ z$ s.t. lem 2 happens

then $\mathbb{E}|P_v(y)| \geq C/4\ nr/w2^d$

$$\left.\begin{array}{l} \text{o.w. } \mathbb{P}(P_v(y) \geq Cnr/w2^d) \leq 1/4 \\ \&\ \mathbb{P}(P_v(z) \geq Cnr/w2^d) \geq 1/2 \end{array}\right\} \text{gap}$$

adversary can reconstruct T_a given $a \in \{y, z\}$

and observe transfer on nodes

distinguishes y, z w/ prob. $> 1/4$

Proof of Lemma 2.

Assume o.w. $\mathbb{P}(|P_v(Z_v)| \geq \frac{1}{100} nr/w2^d) < 1/2$

Encoding argument

uniform, independent r-bit string

of data left subtree of v $d_j, \dots, d_{j+n/2^{d+1}-1}$

entropy, Shannon's source coding thm.

any encoding must use

$nr/2^{d+1}$ bits

in expectation conditioned on R.

Encode Now, if $P_v(Z_v) \ge \cdots$ or error is large

0 + encode directly

if $P_v(Z_v) < \cdots$ and error is small, write 1

use $\mathcal{D}$ write down all contents related to $P_v(Z_v)$

and the state of $\mathcal{D}$ before it.

also all errors

shorter than $nr/2^{d+1}$

Decode

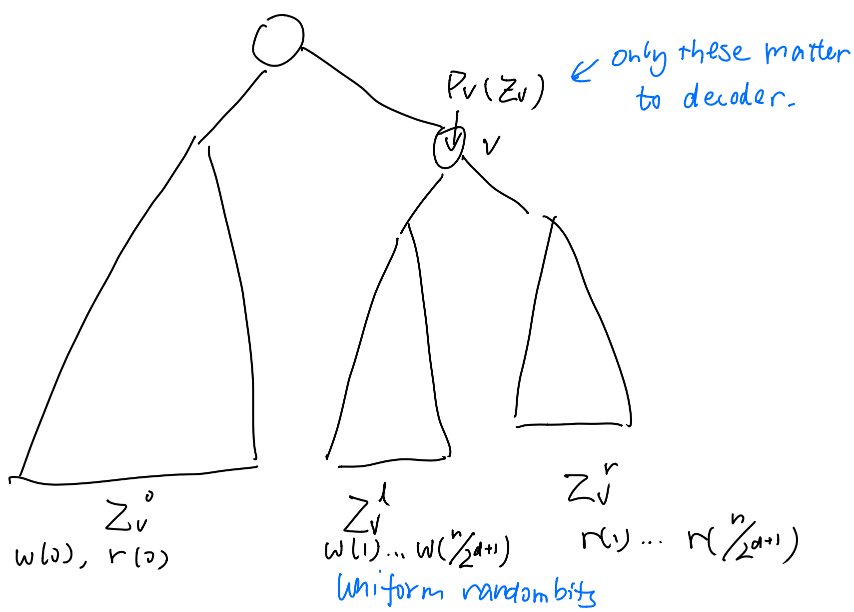

Bob can simulate the entire tree

and retrieve all reads, thus all $d_j \dots d_{j+n/2^{d+1}-1}$

Analysis

$c<1$

$$1 + \frac{3}{4} nr/2^{d+1} + \frac{1}{4} \cdot cnr/2^{d+1} < nr/2^{d+1}$$

- online ORAM $\Omega(\log(nr/m))$