

15451


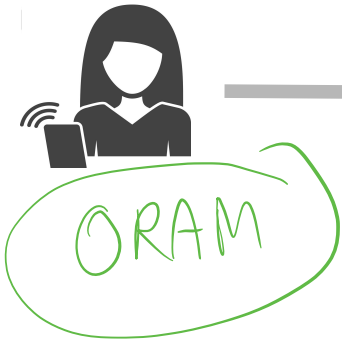
Oblivious RAM

aka. Elaine's favorite data structure


Elaine Shi



Motivating example: Cloud outsourcing




Access patterns to even encrypted data leak sensitive information.



**Liver
problem**

**Heart
problem**



Access patterns to even encrypted data leak sensitive information.

Access patterns of **binary search** leaks the rank of the number being searched.

```
func search(val, s, t)
  mid = (s + t)/2
  if val < mem[mid]
    search (val, 0, mid)
  else search (val, mid+ 1, t)
```

Access pattern leakage through





```
if (secret variable)
    read mem[x]
else
    read mem[y]
```

Recovering JPEG images through coarse-grained access patterns

Original



Recovered



Can we **provably** defeat
access pattern leakage

and **preserve efficiency**




Oblivious RAM (ORAM)

→ Random Access Machine



is an algorithmic technique that provably “encrypts” access patterns




Oblivious RAM (ORAM)

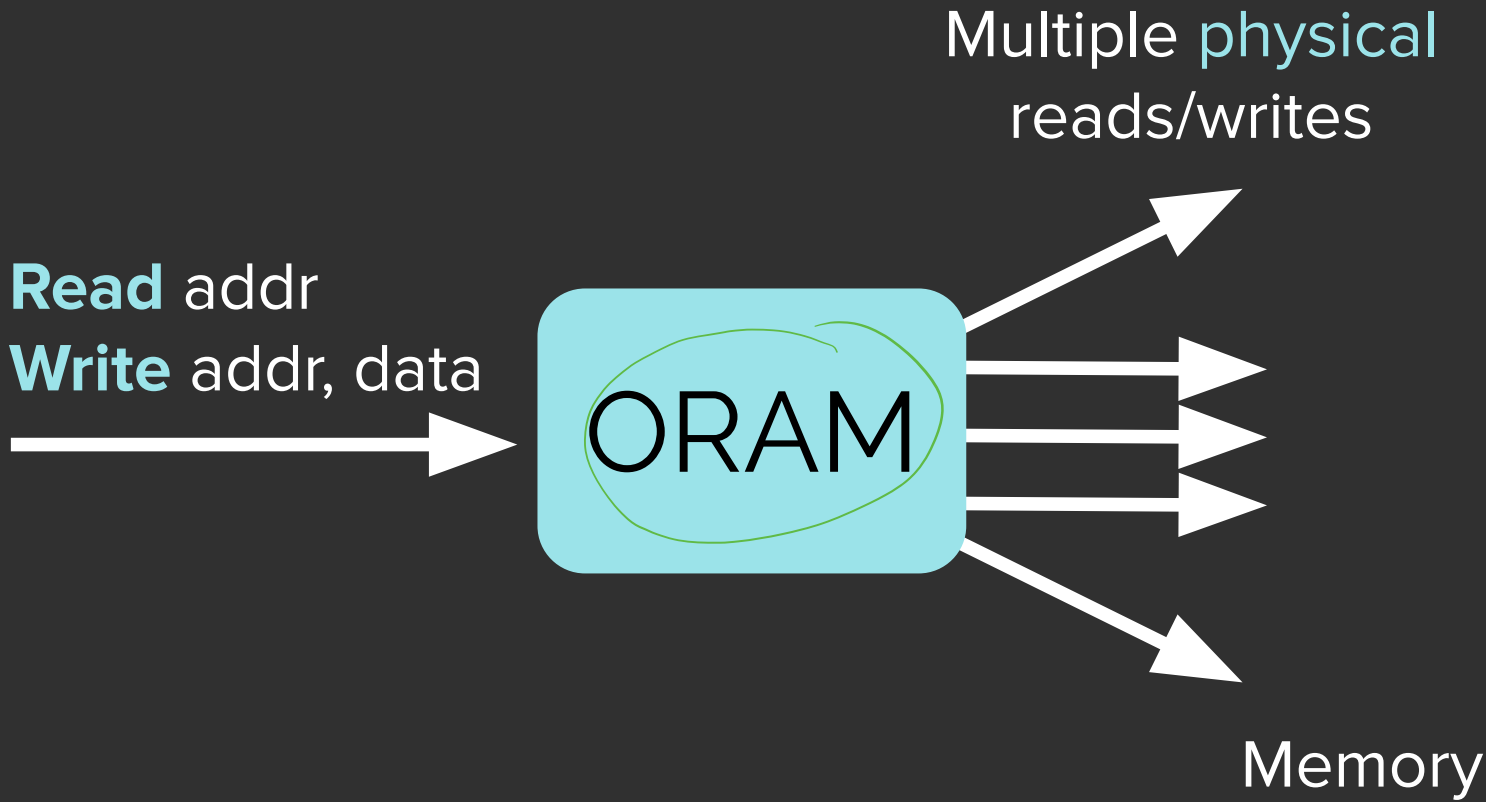


is an algorithmic technique that provably “encrypts” access patterns

 Permutation

 Shuffling






Read addr
Write addr, data



Multiple **physical**
reads/writes




**Security: physical accesses
independent of input requests**

Memory

Defining Security

Request
sequence 1




ORAM



Access Pattern
Distribution 1

Defining Security


Request
sequence 1



ORAM



Access Pattern
Distribution 1




ORAM



Access Pattern
Distribution 2

Defining Security


Request
sequence 1



ORAM



Access Pattern
Distribution 1




ORAM



Access Pattern
Distribution 2

Request
sequence 2

Request sequence 1




Access Pattern Distribution 1

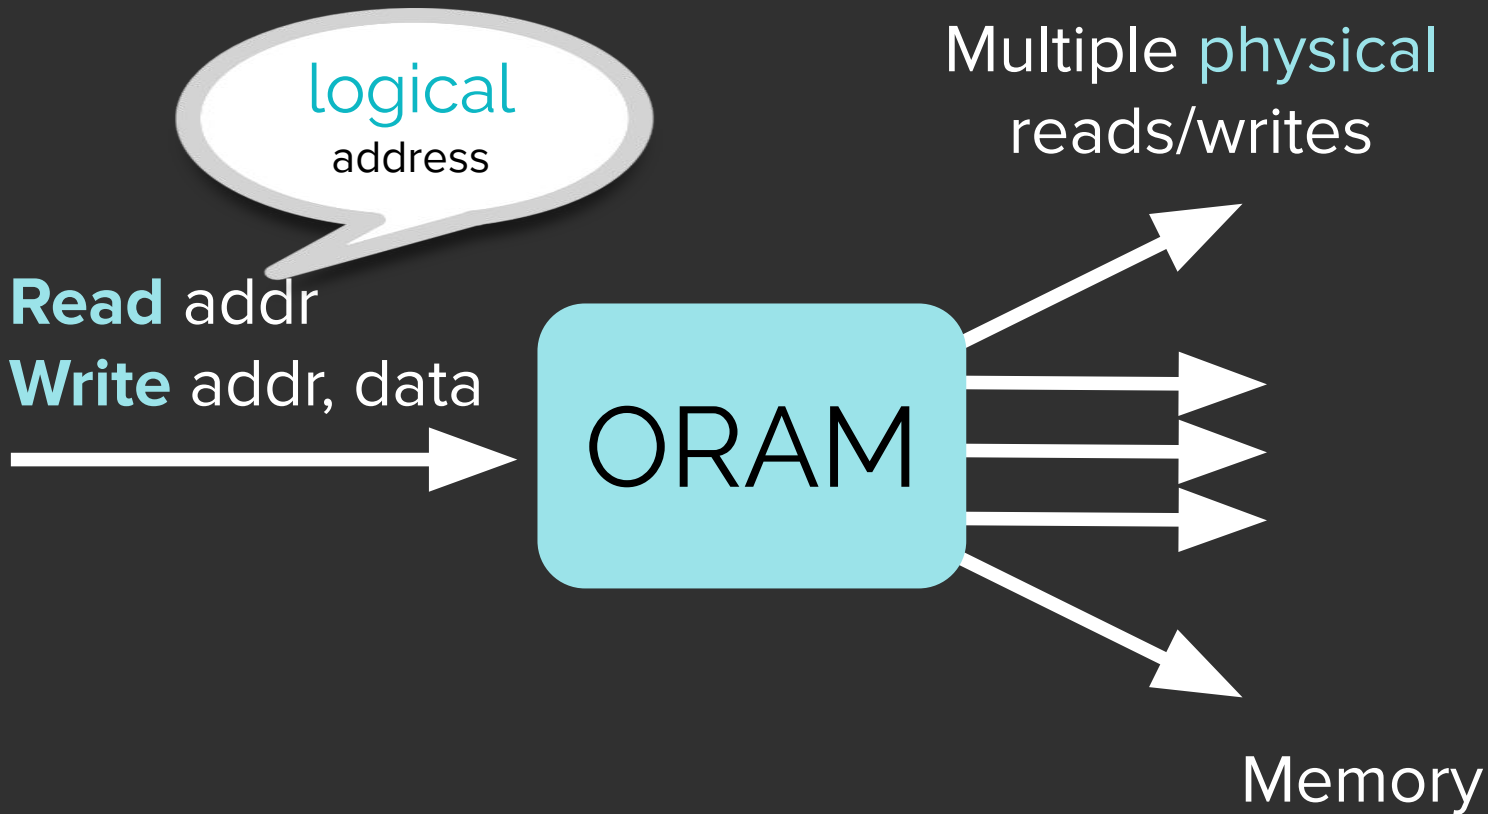
Holds for any 2 request sequences of equal length

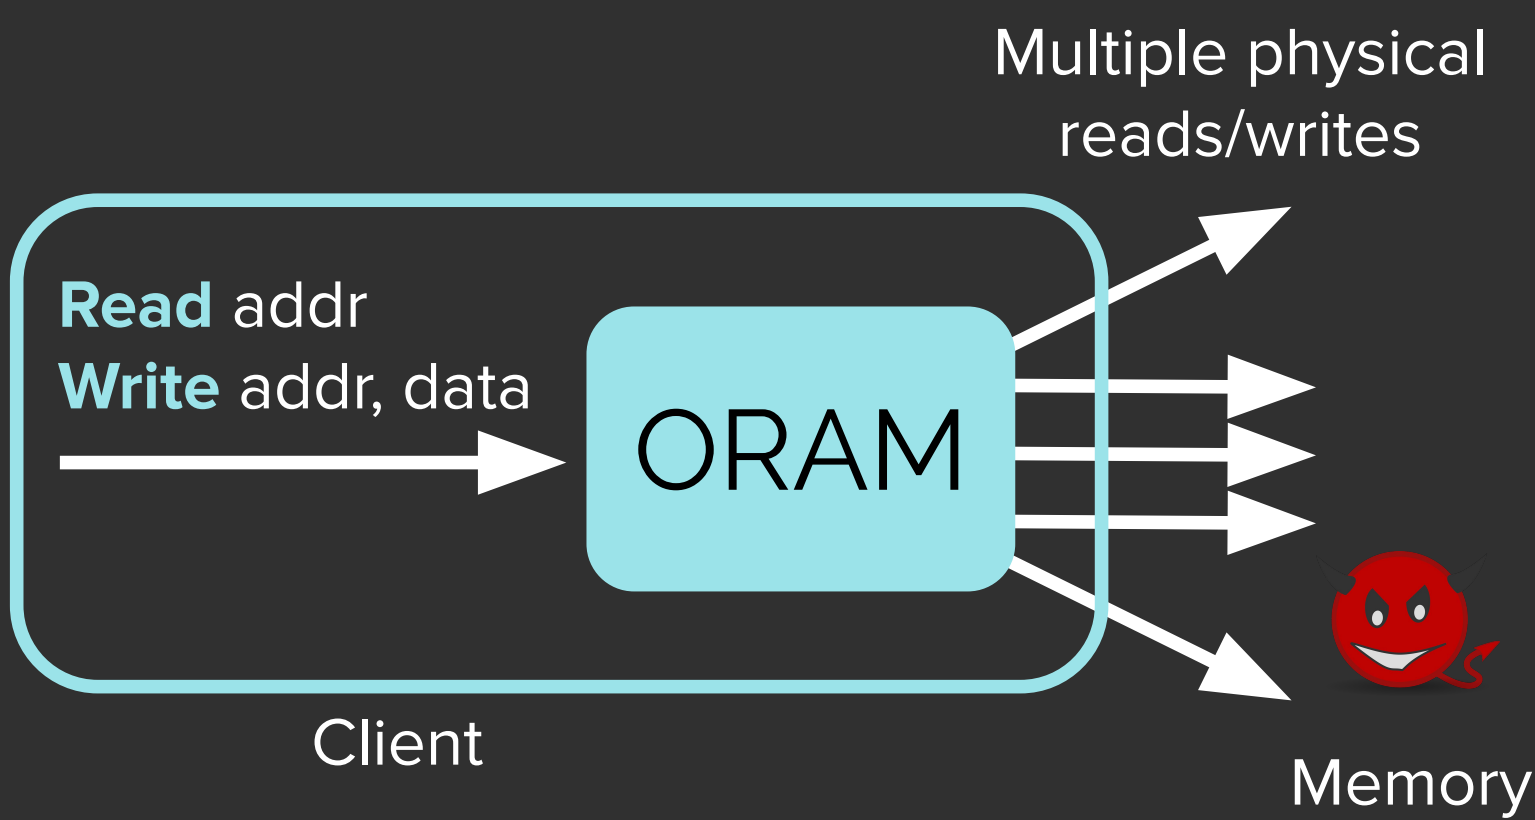


Request sequence 2



Access Pattern Distribution 2






Assumptions

- All blocks are encrypted *minimal storage unit*
- Client always reads a block, reencrypts it (possibly updating with new contents), and writes back

1 2 3 4 5

2 2 2 2 2

Strawman: permute blocks in memory




Memory


~~CPU~~
Client

Strawman: provides one-time security!

e.g., leaks frequency, co-occurrence

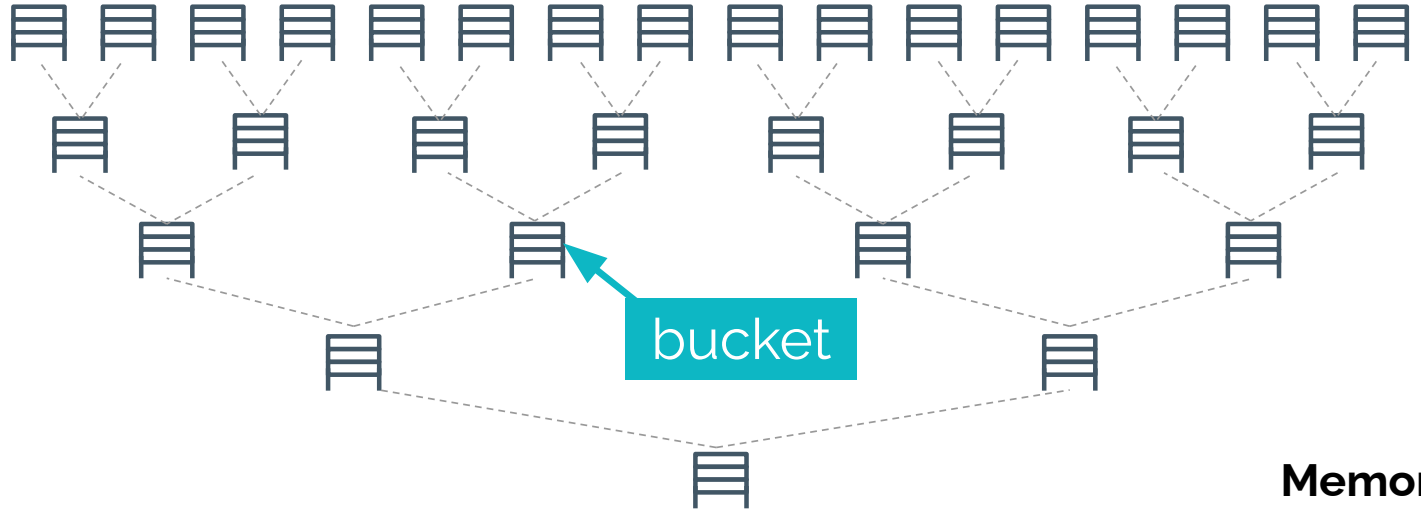


Blocks must move around in memory




A tree-based paradigm for ORAMs

[SCSL'11]



~~CPU~~
Client


Each bucket stores real and dummy blocks



Memory

CPU

Path invariant: every block mapped to a random path




Position
map



block x

~~CPU~~
Client

Reading a block is simple!



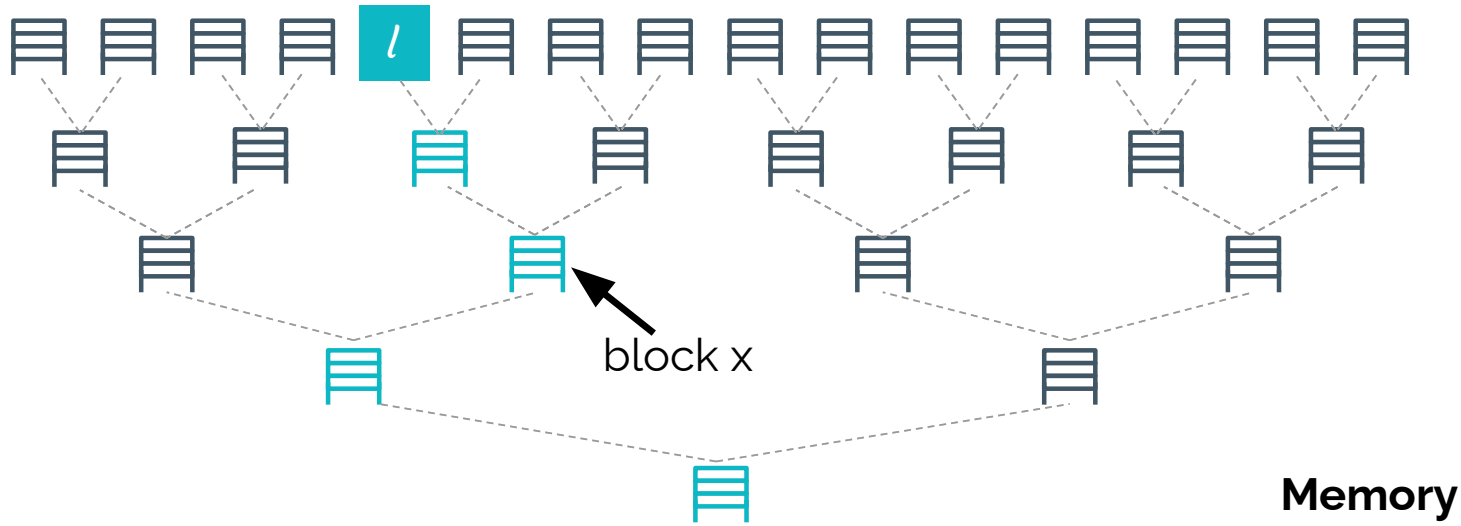
Position
map



block x

CPU

After being read, **block x** must relocate!



Memory

Position map




CPU

block x



Pick a new random path and move x there



update position map

Memory


Position map



CPU

block x

Where on the new path can we write **block x** ?




Position
map



block x

CPU

Can we write it to the leaf?



Memory

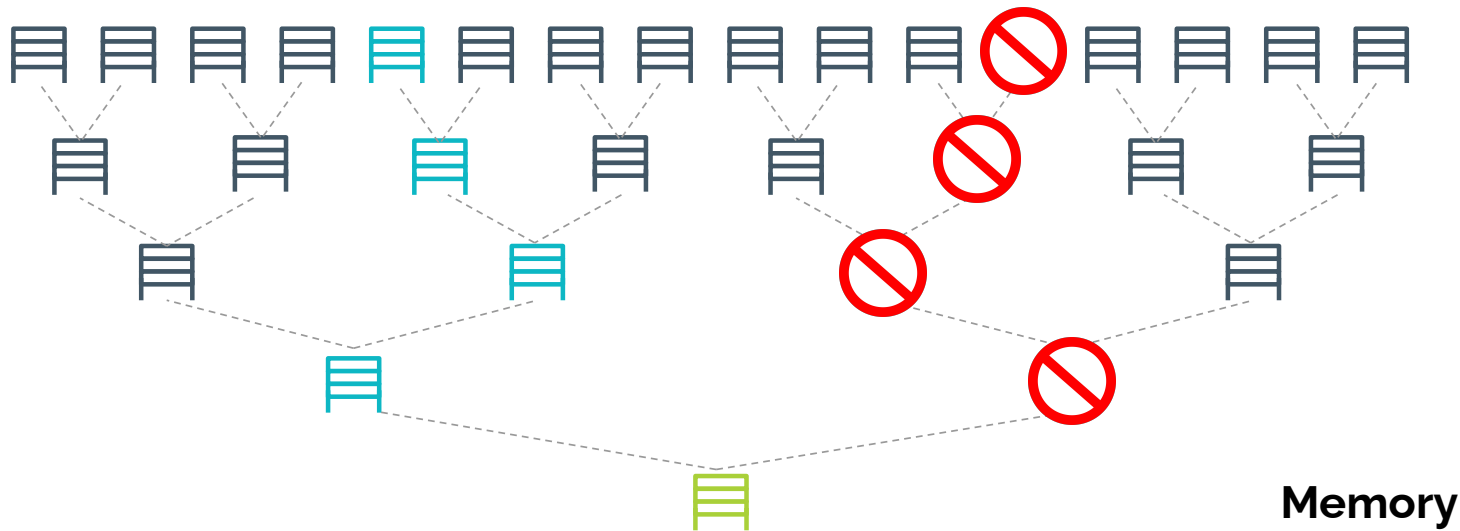
Position
map



block x

CPU

Writing to any non-root bucket leaks information



Position
map

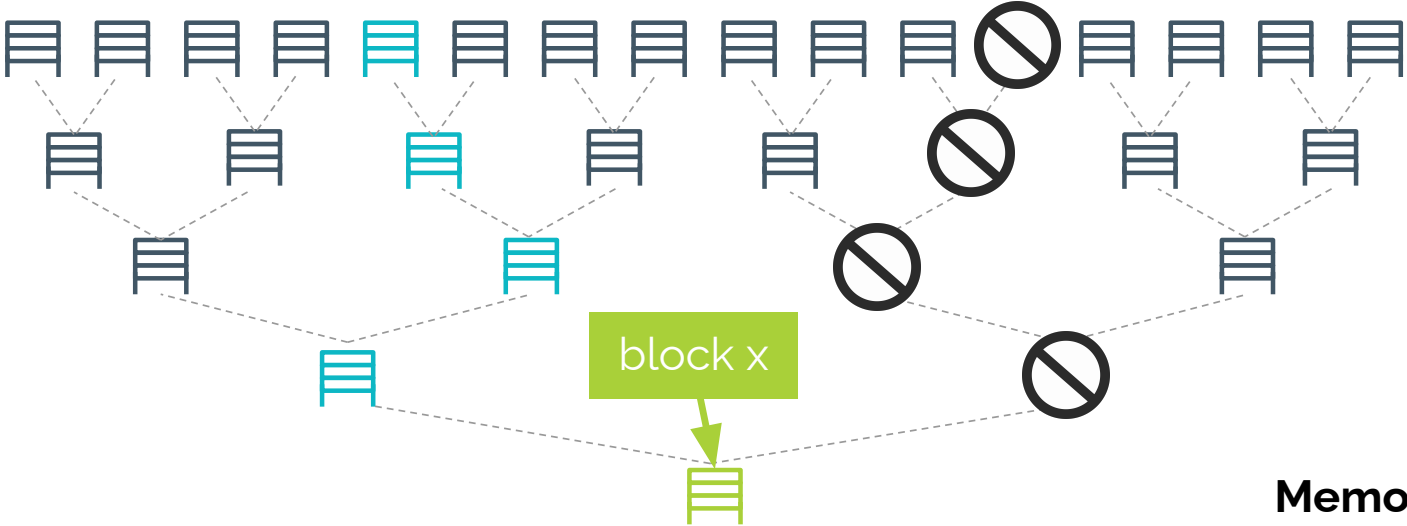


block x

Memory

CPU

Write it to the root!



Memory

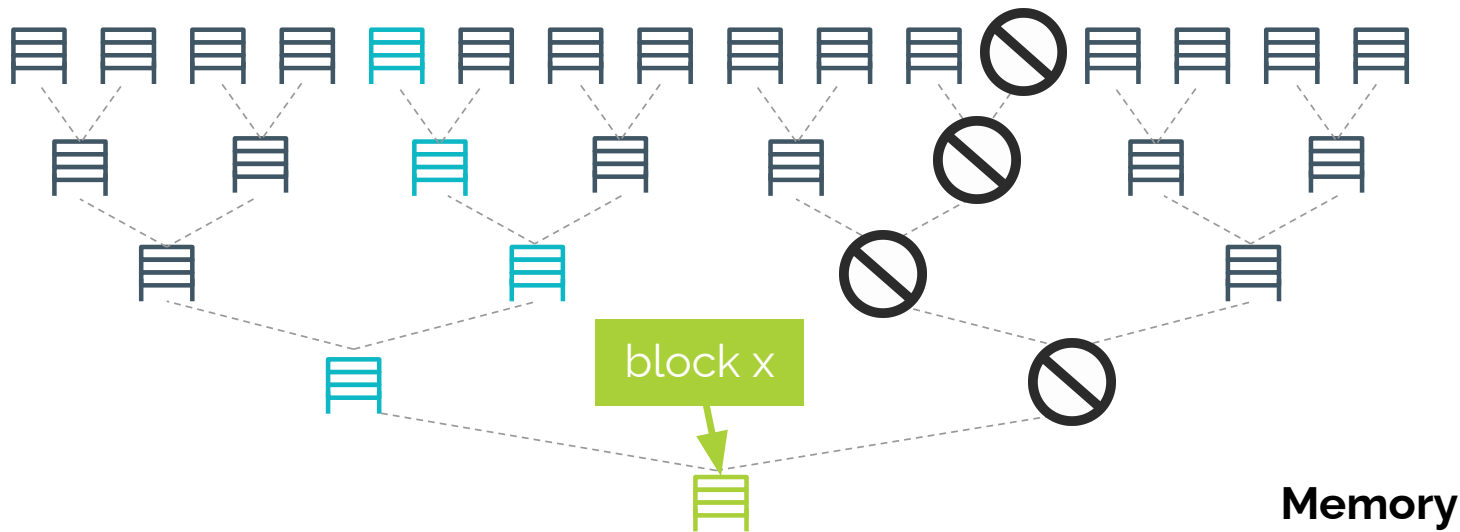
Position map



block x

CPU

Security: every request, visit a **random** path that has **not** been revealed

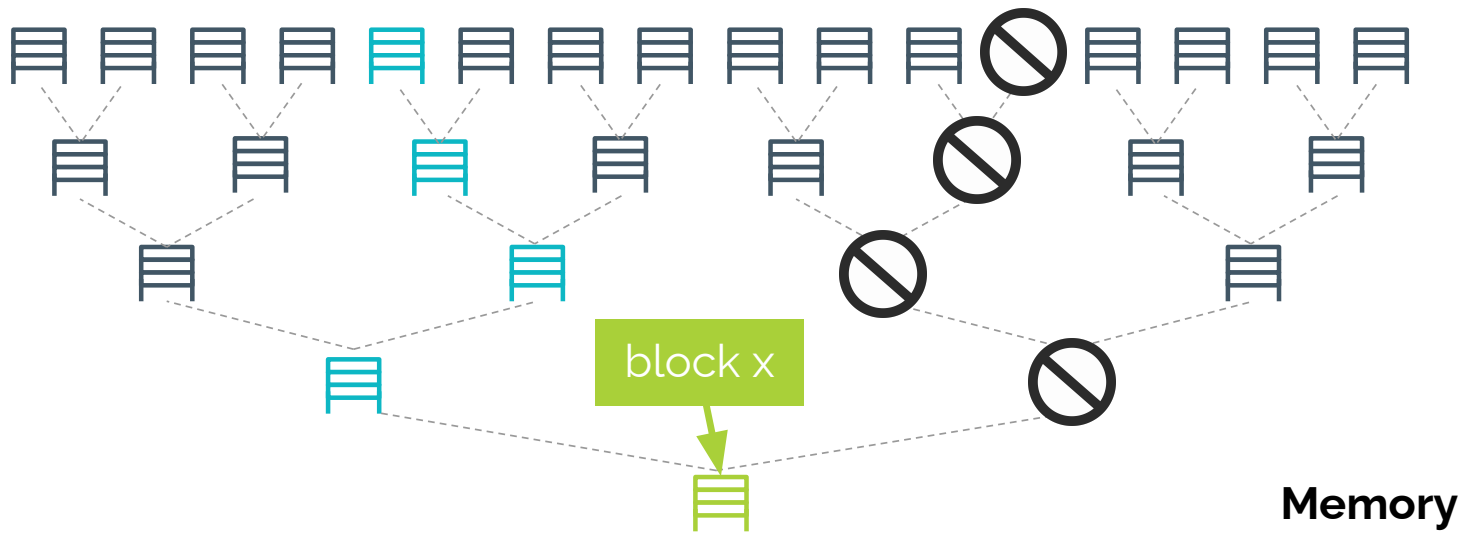


Position
map



CPU

Problem?



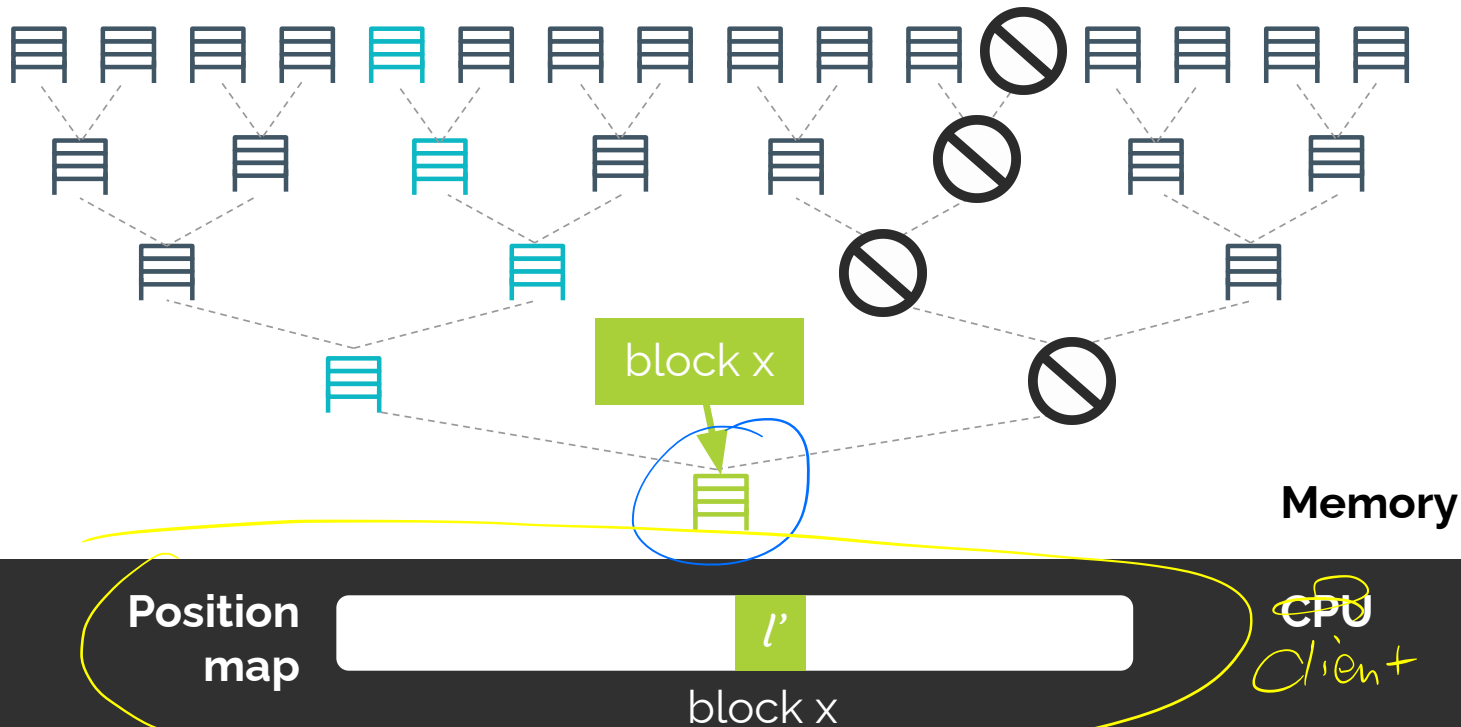
Position
map



block x

CPU

Problem: root will overflow



Further improvements


Remove position map

Resolve overflow

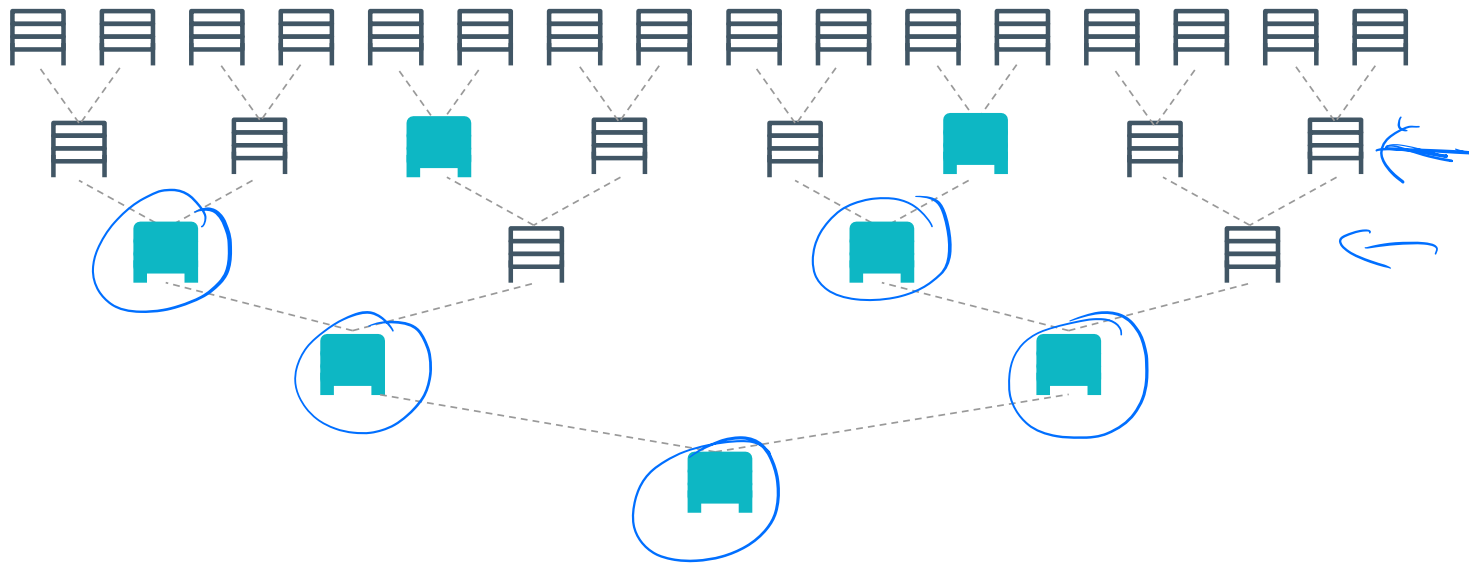
A **background eviction** process percolates blocks upwards




A background eviction process percolates blocks upwards




Every request: pick 2 random buckets per level to evict



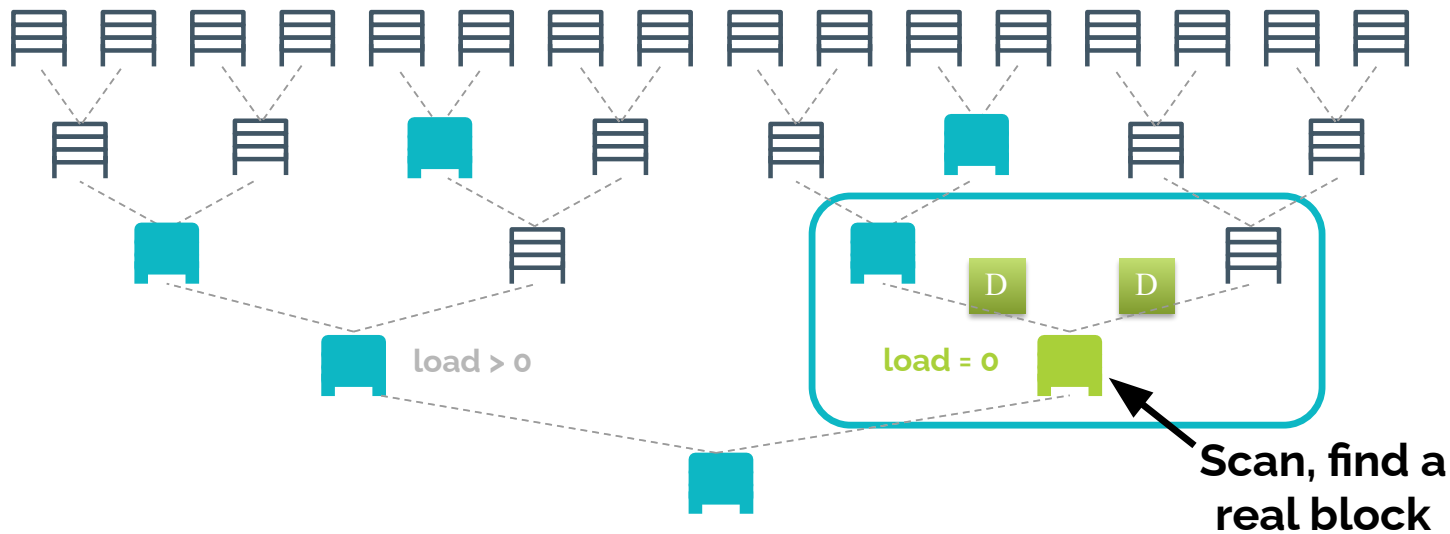
Every request: pick 2 random buckets per level to evict



Scan, find a real
block, write to a child




Dummy eviction ensures security




Dummy eviction ensures security

Eviction process does not leak information



Thm: bucket size = $\log n$

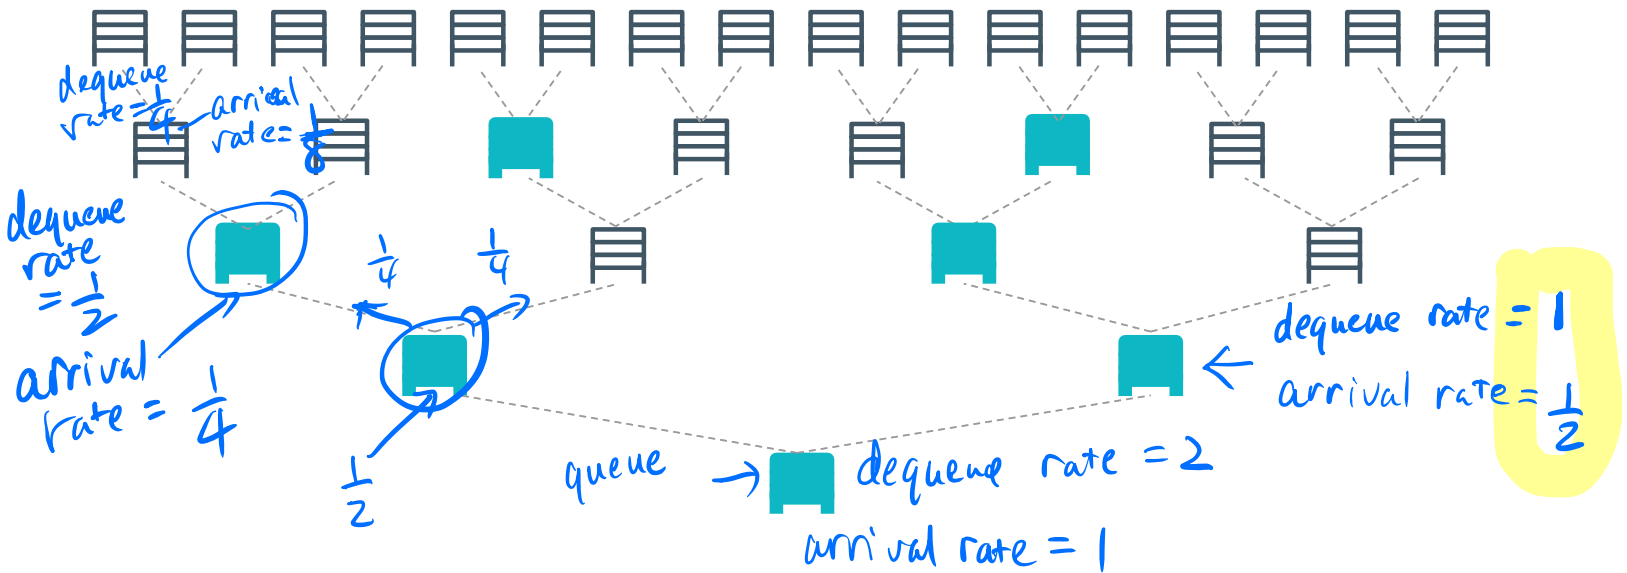


no overflow w.h.p.

[SCSL'11]


$O(\log n)$

$\uparrow 1 - \frac{1}{nc}$



Proof: use queuing theory and measure concentration bounds.

From 15259:




$$\Pr[\text{queue length} > R] \leq \exp(-C \cdot R)$$

in particular, suppose $R \geq \underbrace{C' \cdot \log N}_{\text{sufficiently large constant}}$ \Downarrow Suitable constant.

$$\Pr[\text{queue length} > R] \leq \frac{1}{N^c} \rightarrow \text{Constant.}$$

discrete version of Burke's Thm

Thm: bucket size = $\log n$ \Rightarrow no overflow w.h.p. [SCSL'11]



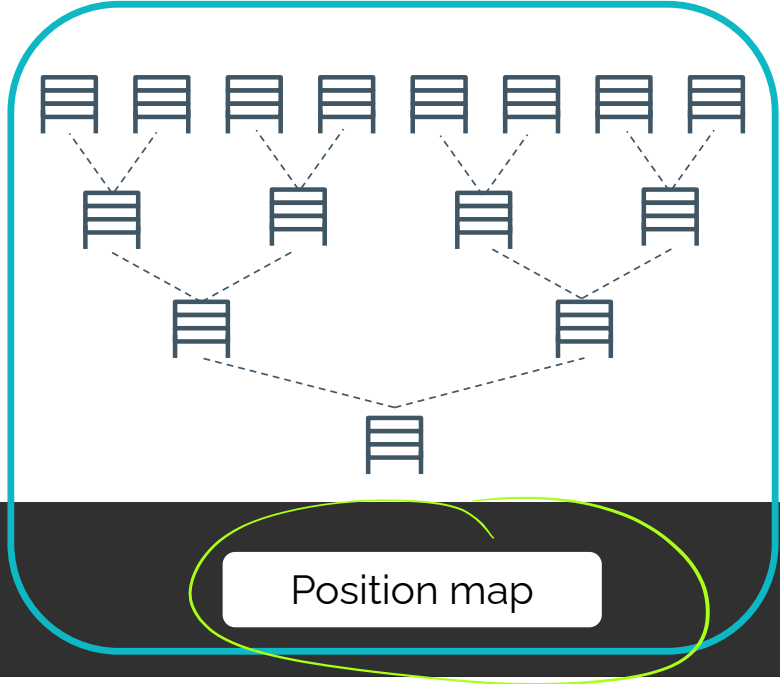
Every request incurs $O(\log^2 n)$ cost



Further improvements

Remove position map

Resolve overflow

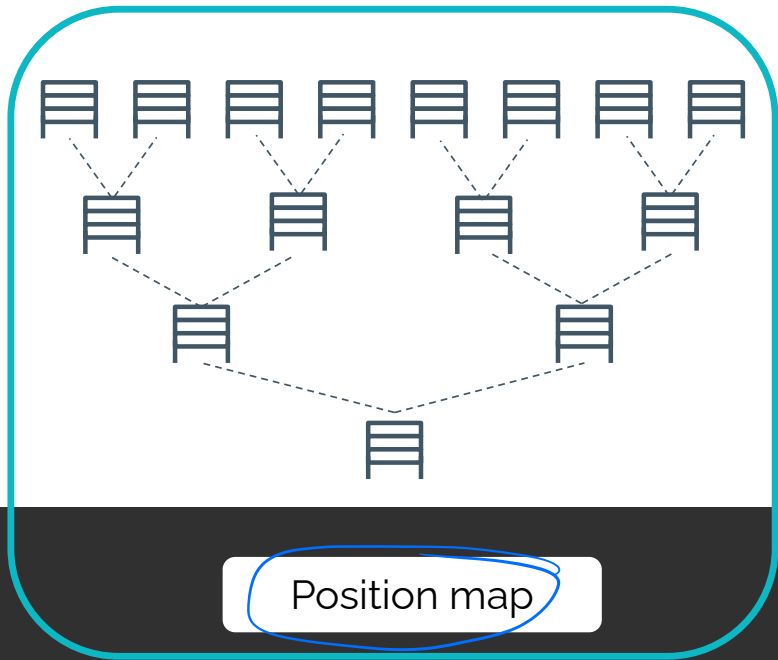



Position map


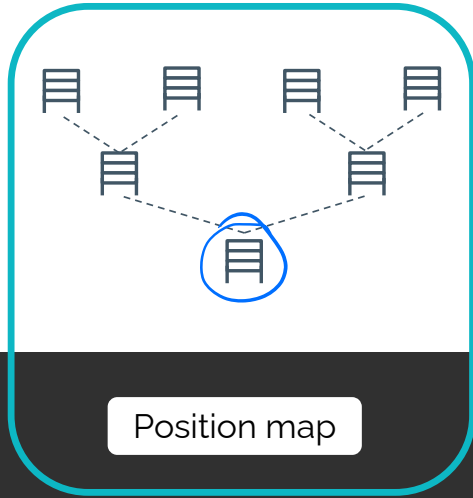


Store position map **recursively** in a smaller ORAM

block size $> 10 \log N$



Cost with eviction: $O(\log^3 n)$




1: $x \leftarrow \text{position}[a]$
2: $\text{position}[a] \leftarrow \text{UniformRandom}(0 \dots 2^L - 1)$

3: **for** $\ell \in \{0, 1, \dots, L\}$ **do**
4: $S \leftarrow S \cup \text{ReadBucket}(\mathcal{P}(x, \ell))$
5: **end for**

6: $\text{data} \leftarrow \text{Read block } a \text{ from } S$
7: **if** $\text{op} = \text{write}$ **then**
8: $S \leftarrow (S - \{(a, \text{data})\}) \cup \{(a, \text{data}^*)\}$
9: **end if**

10: **for** $\ell \in \{L, L - 1, \dots, 0\}$ **do**
11: $S' \leftarrow \{(a', \text{data}') \in S : \mathcal{P}(x, \ell) = \mathcal{P}(\text{position}[a'], \ell)\}$
12: $S' \leftarrow \text{Select min}(|S'|, Z) \text{ blocks from } S'$
13: $S \leftarrow S - S'$
14: $\text{WriteBucket}(\mathcal{P}(x, \ell), S')$
15: **end for**

16: **return** data



**Path
ORAM**

bucket size = 4

**Achieves $O(\log^2 n)$ cost
with recursion**

Perform eviction only on the read path



**Path
ORAM**

The **most aggressive** eviction algorithm

- Pack blocks as close to leaves as possible subject to path invariant

Blocks that cannot be evicted in **stash**


```
1:  $x \leftarrow \text{position}[a]$ 
2:  $\text{position}[a] \leftarrow \text{UniformRandom}(0 \dots 2^L - 1)$ 
3: for  $\ell \in \{0, 1, \dots, L\}$  do
4:    $S \leftarrow S \cup \text{ReadBucket}(\mathcal{P}(x, \ell))$ 
5: end for
6:  $\text{data} \leftarrow \text{Read block } a \text{ from } S$ 
7: if  $\text{op} = \text{write}$  then
8:    $S \leftarrow (S - \{(a, \text{data})\}) \cup \{(a, \text{data}^*)\}$ 
9: end if
10: for  $\ell \in \{L, L - 1, \dots, 0\}$  do
11:    $S' \leftarrow \{(a', \text{data}') \in S : \mathcal{P}(x, \ell) = \mathcal{P}(\text{position}[a'], \ell)\}$ 
12:    $S' \leftarrow \text{Select min}(|S'|, Z) \text{ blocks from } S'$ 
13:    $S \leftarrow S - S'$ 
14:    $\text{WriteBucket}(\mathcal{P}(x, \ell), S')$ 
15: end for
16: return  $\text{data}$ 
```

Path ORAM


**ACM CCS Best
Student Paper**
[SDS+'13]

**Achieves $O(\log^2 n)$ cost
with recursion**

```
1:  $x \leftarrow \text{position}[a]$ 
2:  $\text{position}[a] \leftarrow \text{UniformRandom}(0, \dots, 2^l - 1)$ 
3: for  $\ell \in \{0, 1, \dots, L\}$  do
4:    $S \leftarrow S \cup \text{ReadBucket}(\mathcal{P}(x, \ell))$ 
5: end for
6:  $\text{data} \leftarrow \text{Read block } a \text{ from } S$ 
7: if  $\text{op} = \text{write}$  then
8:    $S \leftarrow (S - \{(a, \text{data})\}) \cup \{(a, \text{data}^*)\}$ 
9: end if
10: for  $\ell \in \{L, L - 1, \dots, 0\}$  do
11:    $S^\ell \leftarrow \{(a', \text{data}') \in S : \mathcal{P}(x, \ell) = \mathcal{P}(\text{position}[a'], \ell)\}$ 
12:    $S^\ell \leftarrow \text{Select min}(|S^\ell|, Z) \text{ blocks from } S^\ell$ 
13:    $S \leftarrow S - S^\ell$ 
14:    $\text{WriteBucket}(\mathcal{P}(x, \ell), S^\ell)$ 
15: end for
16: return  $\text{data}$ 
```




**ACM CCS Best
Student Paper**
[SOS+13]



“Ascend” processor
running a variant of
our ORAM

~**2X** average overhead on
standard benchmarks
(ORAM not always on)



Thank you!