

# Onion ORAM: Constant Bandwidth ORAM with Server Computation

**Chris Fletcher**

Joint work with:

Ling Ren, Marten van Dijk, Sridhar Devadas

- **State of the art schemes**

- Bandwidth:  $O(\log N)$
- Client storage:  $O(1)$  ( Path ORAM =  $O(\log N)$  )
- Server storage:  $O(N)$

- **Is “optimal” ORAM possible?**

**$O(1)$  bandwidth,  $O(1)$  client storage,  $O(N)$  server storage**

- **Goldreich-Ostrovsky lower bound [1987, 1996]**

Given a program that runs in  $T$  time and an  $N$  block ORAM with  $O(1)$  client storage, the program+ORAM must run in  $\Omega(T \log N)$  time


- **$\Omega(T \log N)$  doesn't mean  $\Omega(T \log N)$  bandwidth!**

- **Example: Outsourced storage (Honest but curious)**

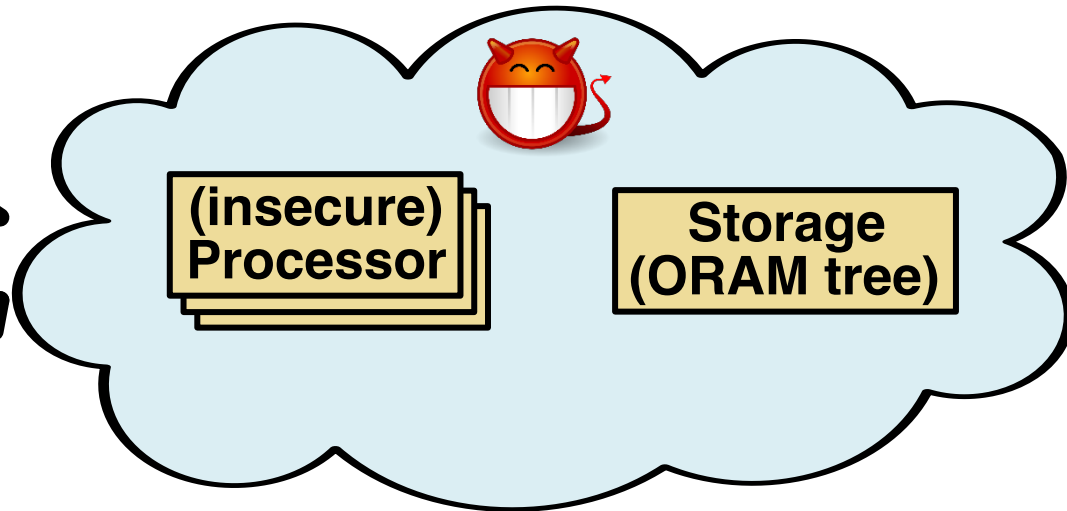
Client

Server

Trust boundary



Arbitrary messages



limited storage, compute

- “Read X, Y, Z, return  $F(X, Y, Z)$ ”
- Message stream must be oblivious

- XORing reads [Dautrich et al.], PIR+ORAM [Mayberry et al.]
- **XOR + Ring ORAM**
  - Permuted buckets  $\rightarrow$  one real block touched / read
  - $B, d_1, d_2, d_3, \dots$
  - $E(B, r), E(0, r_1), E(0, r_2), E(0, r_3) \dots$
  - Server sends:  $E(B, r) \oplus E(0, r_1) \oplus E(0, r_2) \oplus E(0, r_3) \oplus \dots$
  - Client computes:  $E(0, r_1) \oplus E(0, r_2) \oplus E(0, r_3) \oplus \dots$
- Both schemes make read bandwidth  $O(\log N) \rightarrow O(1)$
- Does not help on evictions!



**CSAIL**

---

**Can we make evictions  
 $O(1)$  Bandwidth?**

- **Example: Ring ORAM**

- ORAM on server is encrypted under FHE scheme  $E^{\text{FHE}}$
- Reads

**Client**

**Server**

**leaf,  $E^{\text{FHE}}(\text{address})$**



**$E^{\text{FHE}}(d) = \text{Select}(E^{\text{FHE}}(\text{address}), \text{Path}(\text{leaf}))$**

**$\text{RemoveBlock}(E^{\text{FHE}}(\text{address}), \text{Path}(\text{leaf}))$**

**$E^{\text{FHE}}(d)$**



- Evictions

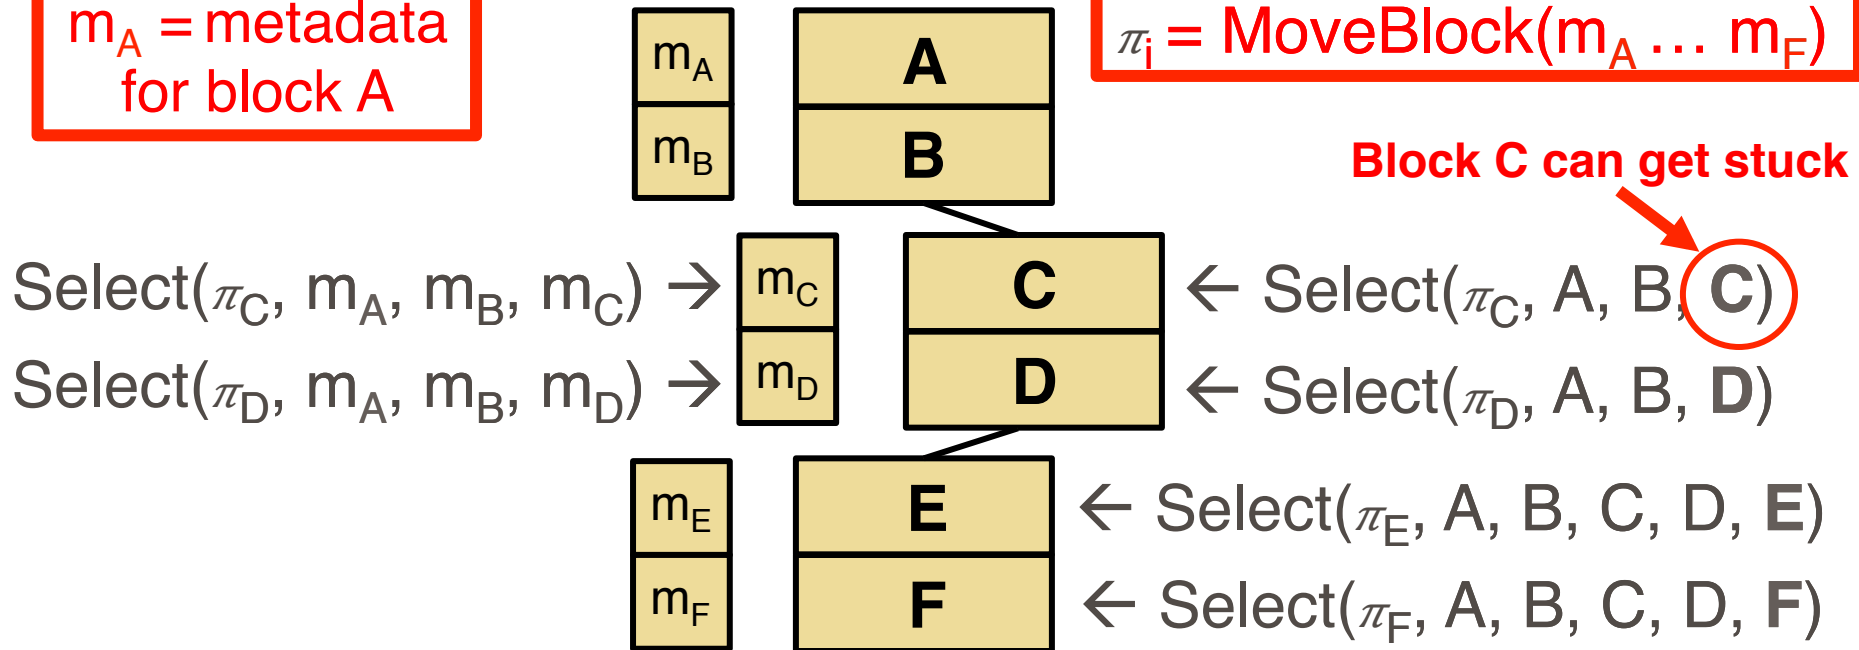
**$\text{Path}(\text{leaf}_g)' = \text{EvictPath}(\text{Path}(\text{leaf}_g))$**

**Read bandwidth is  $O(1)$ , no bandwidth for evictions!**

$$\text{Path}(\text{leaf}_g)' = \text{EvictPath}(\text{Path}(\text{leaf}_g))$$

$m_A$  = metadata  
for block A

$\pi_i = \text{MoveBlock}(m_A \dots m_F)$



- Only **Select()** touches *blocks*



Server computation:  $\text{polylog}(N)$



Bootstrap to manage noise [Apon et al., Mayberry et al.]



**CSAIL**

---

**\*Discuss later: Does the previous scheme achieve optimal Bandwidth/storage?**

**Do we need bootstrapping?  
Do we need FHE?**



- **Additive-HE (e.g., Paillier)**

- Addition:  $E^{\text{AHE}}(a) \oplus E^{\text{AHE}}(b) = E^{\text{AHE}}(a + b)$

- Scalar multiplication:  $E^{\text{AHE}}(a) \otimes c = E^{\text{AHE}}(ca)$

- **Select from (X, Y):**

$$E^{\text{AHE}}(0) \otimes X \oplus E^{\text{AHE}}(1) \otimes Y = E^{\text{AHE}}(0+Y) = E^{\text{AHE}}(Y)$$

- **$Y = E^{\text{AHE}}(\text{plaintext})$**

- **Select op  $\rightarrow E^{\text{AHE}}(E^{\text{AHE}}(\text{plaintext}))$**

- Client decrypts twice

- (Possible) ciphertext blowup per layer

**Block gets extra layer of encryption**

- **Layers(output) =  $\max(\text{Layers}(\text{Block}_i) : \text{Blocks}) + 1$**

# ORAM Read + Additive-HE



- ORAM encrypted using 1 layer of  $E^{\text{AHE}}$  (abbreviated  $E$ )

Client


Server

leaf



Decrypt metadata

Compute  $\pi = E(0), E(0), \dots E(1), \dots E(0)$




Write updated metadata

Compute  $E(E(d)) = \text{Select}(\pi, \text{Path(leaf)})$



Decrypt  $d = E(E(d))$

Update  $d \rightarrow d'$

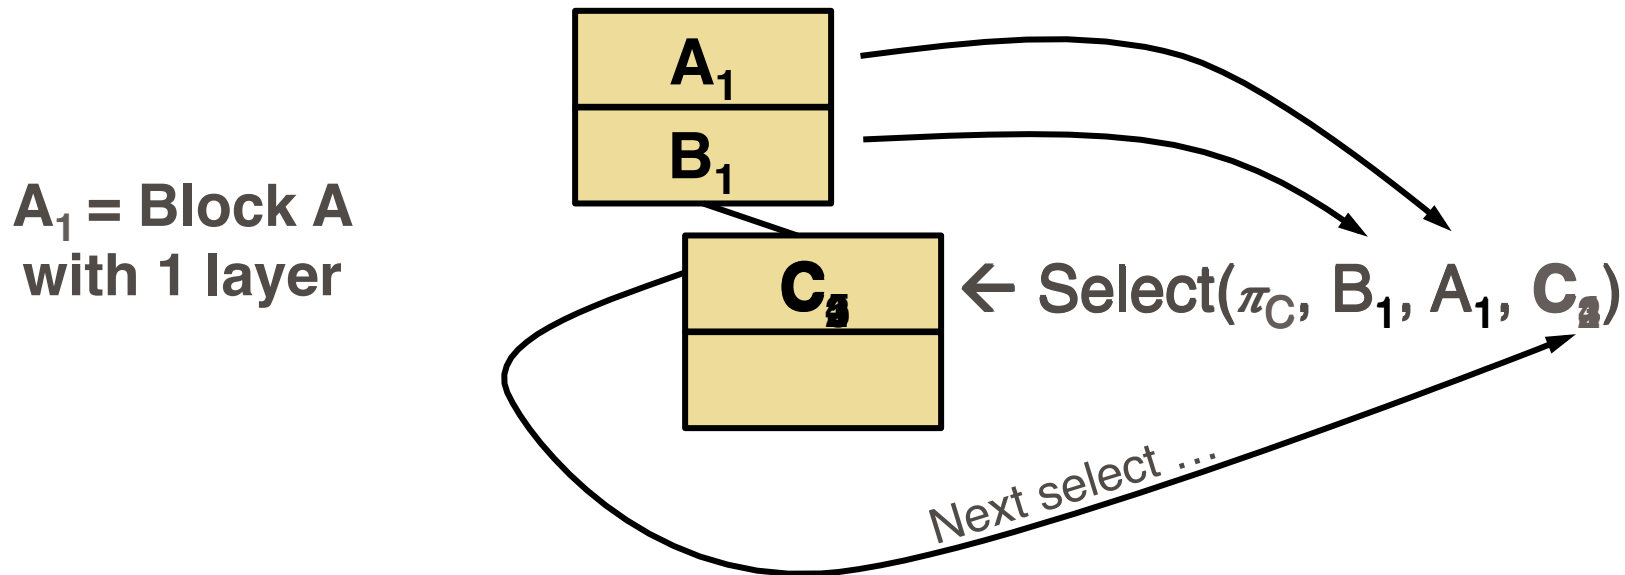


$\text{Path(leaf)[root].append}(E(d'))$



- Problem: Continuous reshuffling  $\rightarrow$  Unbounded layers
- Reason: Blocks can get stuck in buckets after evictions

$$\text{Layers}(\text{output}) = \max(\text{Layers}(\text{Block}_i) : \text{Blocks}) + 1$$



$O(T)$  evictions  $\rightarrow$  Slot with C gets  $O(T)$  layers






**CSAIL**

---

**ORAM with  $O(1)$  bandwidth,  
 $O(1)$  client storage,  
 $O(N)$  server storage**

**...with only additive-HE**

Design our ORAM eviction algorithm such that **buckets are guaranteed to be empty regularly**

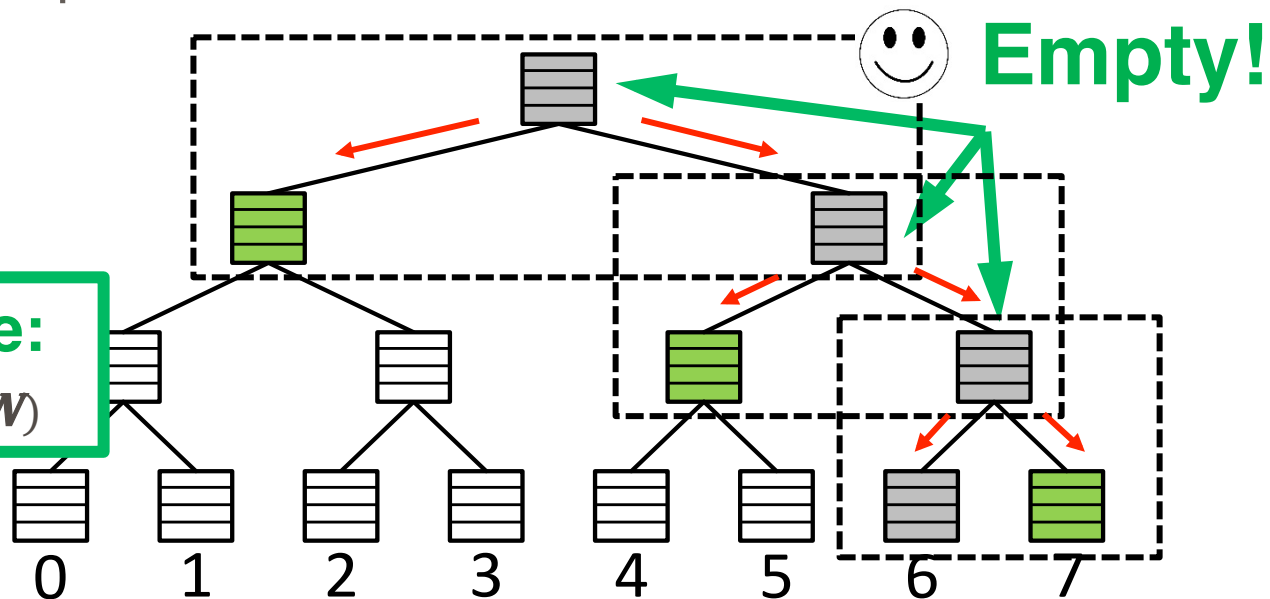


Design our ORAM eviction algorithm such that  
**buckets are guaranteed to be empty regularly**

1. Evict over reverse-lexicographic order of paths
2. Also evict to sibling buckets
3. Set  $Z, A$  s.t.  $\Pr[\text{bucket overflow}] = \text{negl}(\text{security parameter})$
4. Evict to 1 bucket triplet at a time

Example:  
evict to leaf 6

Informal guarantee:  
Max layers =  $\mathcal{O}(\log N)$



**Theorem:  $Z \geq A, N \leq A * 2^{\uparrow L} - 1$**

$$\rightarrow \Pr[\text{bucket overflow}] = e^{\uparrow - (2Z - A) \uparrow 2} / 6A$$

~~•  $Z = A = \theta(\log N) \omega(1) \Rightarrow \Pr[\text{bucket overflow}] = N^{\uparrow - \omega(1)}$~~

*Note:  $N = \text{poly}(\text{security parameter})$*

- **Asymptotics w/o server computation**

- Bandwidth =  $\mathcal{O}(\log^{\uparrow 2} N) \omega(1)$  blocks
- Client storage =  $\mathcal{O}(\log N) \omega(1)$  blocks
- Server storage =  $\mathcal{O}(N)$  blocks



- **Not competitive w/o server computation**


- **Same as previous proposal**

- Client sends leaf
- Server sends metadata
- Client sends  $\pi = E(0), E(0), \dots E(1), \dots E(0)$
- Server sends block


Assume layers  $\rightarrow$  ciphertext blowup

- **Simple scheme factoring in layers**

- Elements of  $\pi$  have 1 layer
- Pad blocks on path to  $S = \text{Max}(|\text{Block}_i| : \text{Blocks})$  bits
- Split each padded block into  $C$  chunks s.t.  $S / C = \text{Plaintext}(\pi \downarrow i) = P$








$\text{Path}(\text{leaf})[i] =$   $i^{\text{th}}$  triplet on path  
 $\text{Path}(\text{leaf})[i].\text{dest}[j] =$   $j^{\text{th}}$  block in  $i^{\text{th}}$  triplet's dest. bucket

- **Useful properties:**

1. At eviction start: non-leaf sibling buckets are empty
2. At eviction end: non-leaf destination buckets are empty

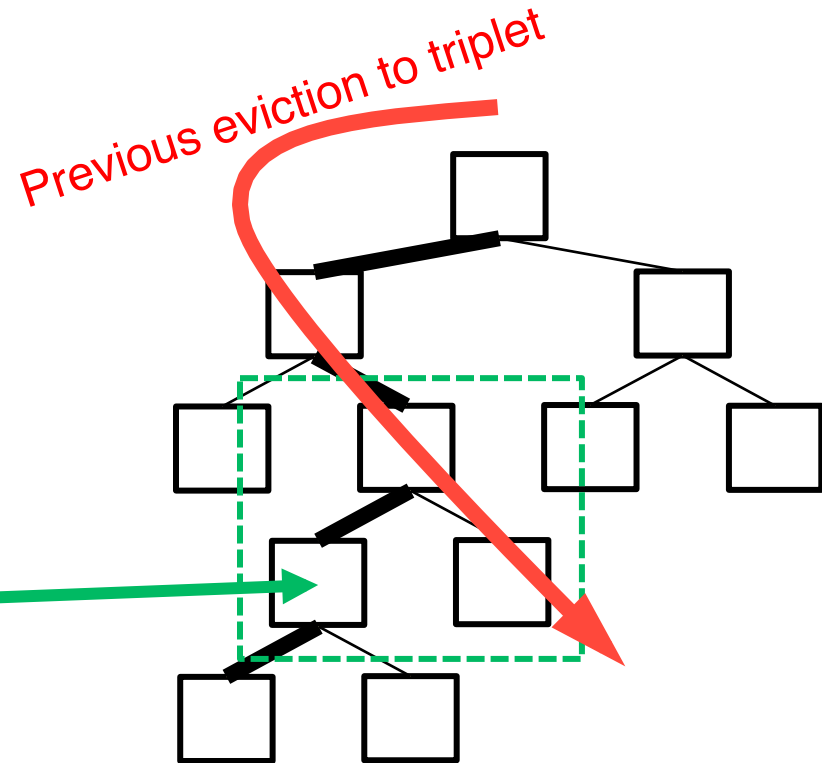


- Blocks get stuck in the leaves
- Non-leaves empty at regular intervals

- Analyze: Layers on destination bucket at start of select

Key intuition:  
destination bucket was  
sibling on last eviction

# Layers? →



**Theorem: buckets at level  $k < L$  have  $\leq c * k + 1$  layers**

- $c$  is constant,  $c=1$  in our final scheme

**Client**

**Server**

$\text{leaf}_g$  (eviction path) known by server


$E(\text{metadata for Path}(\text{leaf}_g))$



Compute  $\Pi = \{\pi \downarrow 0 \dots \pi \downarrow Z * L\}$

( $|\pi \downarrow i| = O(Z)$  encrypted coefficients)

$\Pi, E(\text{updated metadata for Path}(\text{leaf}_g))$



For triplet  $i$ :

$\text{Path}(\text{leaf}_g)[i].\text{sibling} = \text{Path}(\text{leaf}_g)[i].\text{src}$

For slot  $j$ :

$\text{args} = \{\text{Path}(\text{leaf}_g)[l].\text{dst}[j], \text{Path}(\text{leaf}_g)[l].\text{src}\}$

$\text{Path}(\text{leaf}_g)[i].\text{dst}[j] = \text{Select}(\pi \downarrow Z * i + j, \text{args})$

**Problem: layers in leaves are not bounded**

- **At end of each eviction...**

**Client**


**Server**

$$E(d_j) = \text{Path}(\text{leaf}_g)[L].\text{dst}[j]$$

For  $j = 0 \dots Z - 1$

Decrypt  $E(d_j)$  to a constant number of layers, yielding  $e_j$   
Compute  $E(e_j)$

$$E(e_j)$$



- **Layer theorem now applies to all levels**
- **Adds constant amortized bandwidth if  $Z \sim A$**



**CSAIL**

---

# Setting parameters

## Problem: each layer can add ciphertext blowup


- Layer bound =  $O(\log N)$
- Paillier (1999):  $n \rightarrow n^2$  (n = RSA modulus) ☹️
- Damgård-Jurik (2001):  $n^s \rightarrow n^{s+1}$  ☺️
  - s = free parameter
  - Strategy: set  $s \approx \log N$ ,  $\log N$  layers  $\rightarrow n^{s \log N} = n^{O(\log^2 N)}$
- Operations are like Paillier:
$$E(a) \oplus E(b) = E(a)E(b) \qquad E(a) \otimes b = E(a)^b$$
- Best attack: factor  $n$ , complexity  $\exp(|n|^{1/3} (\log |n|)^{2/3})$   
 $\therefore |n| = \theta(\log^3 N) \rightarrow$  defeat attacks w/ complexity  $N^{\omega(1)}$

- So far ... **Select** =  $\bigoplus_i \pi \downarrow i \otimes \text{Block} \downarrow i$       “trivial linear PIR”

- Each select adds 1 layer  
layer bound =  $\log N$
- $Z$  inputs  $\rightarrow |\pi| = Z * \text{layer bound} * |n| = \log^5 N \omega(1)$

- **Hierarchical PIR [Lipmaa 2005]**

- Multiplexer tree
- $Z$  inputs  $\rightarrow |\pi| = \log Z$  coefficients  
 $\rightarrow$  select adds  $\log Z$  layers
- $\therefore$  layer bound  $\uparrow = \log M \log \log N$
- $Z$  inputs  $\rightarrow |\pi| = \log Z * \text{layer bound} \uparrow * |n|$   
 $= \log^4 N \log^2 \log N$



At least better in theory 😊



- **Strategy: set  $|\Pi| = |\{\pi \downarrow 0 \dots \pi \downarrow Z * L\}| = O(B)$**
- **I.e.,  $\Pi$  contributes constant (amortized) bandwidth**
  
- **Let  $Z=A=\log N \omega(1)$**
- **$|\pi \downarrow read| = |n| * \log^2 N \omega(1)$  (n = RSA modulus)**
- **$|\pi \downarrow evict| = |n| * \log N \log^2 \log N$  (mux tree)**
- **$|\Pi \downarrow evict| = |n| * \log^2 N \log^2 \log N = \log^5 N \log^2 \log N$**
  
- **Final asymptotics:**
  - Block size  $B = \Omega(\log^5 N \log^2 \log N)$
  - Bandwidth =  $O(B)$
  - Client storage =  $O(B)$
  - Server storage =  $O(BN)$

- **Decrease block size**  $B = \Omega(k \cdot \log^2 N \log^2 \log N)$ 
  - Modern schemes w/o computation:  $B = O(\log^2 N)$
- **How?**
  - Server computation is  $O(\log^2 N) \omega(1)$  blocks - is  $O(\log N)$  possible?
  - Is there a suitable additive-HE scheme with  $k = o(\log^3 N)$ ?
- **Protect against malicious servers**
  - Server performs select incorrectly
- **Improve Garbled RAM schemes?**
  - Use ORAM as a blackbox
- **Parameterization for SWHE for Onion ORAM**
  - No bootstrapping needed