




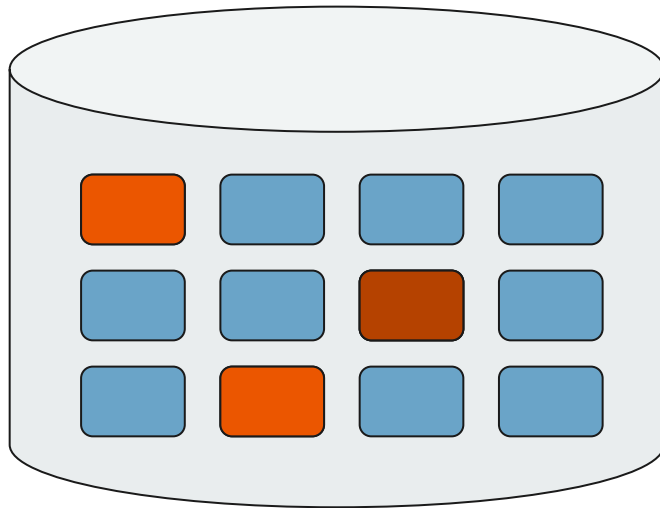




PanORAMa: Oblivious RAM with Logarithmic Overhead

Sarvar Patel, Giuseppe Persiano, [Mariana Raykova](#), Kevin Yeo

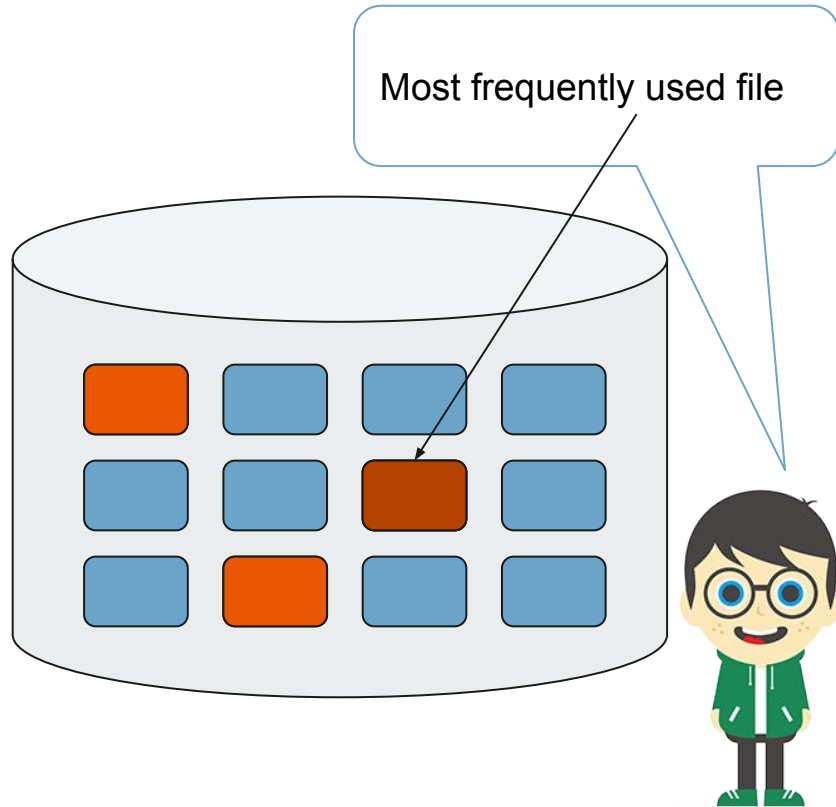









Bob knows what files I am accessing




Most frequently used file





Bandwidth is expensive



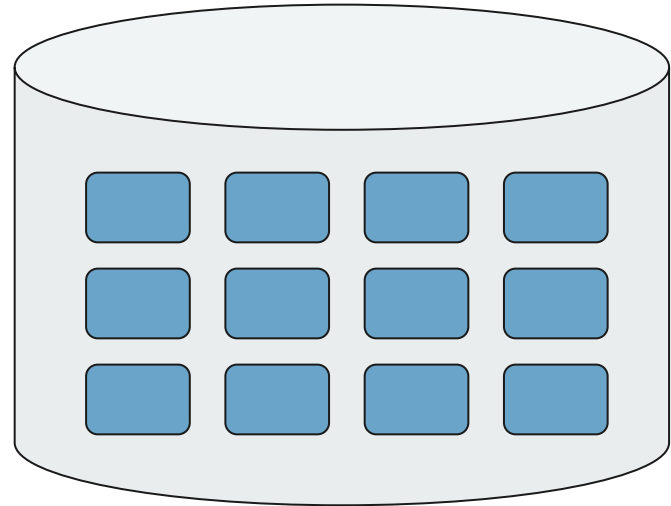
Retrieve the whole database



Oblivious RAM [GO'96]



Access pattern hiding



How Efficient Can an ORAM Construction be?



ORAM Lower Bound



ORAM Lower Bound

- Goldreich-Ostrovsky'96
 - Lower bound $O(\log_c N)$ blocks for database of N blocks and client memory of C blocks



ORAM Lower Bound

- Goldreich-Ostrovsky'96
 - Lower bound $O(\log_c N)$ blocks for database of **N blocks** and client memory of **C blocks**
 - Caveats:
 - Server **only moves and retrieves** blocks (does not do any computation)
 - Constructions with **statistical security**
 - Constructions that **work for any block size**
 - The client can have **oracle access to a private random function**



ORAM Lower Bound

- Goldreich-Ostrovsky'96
 - Lower bound $\Omega(\log_c N)$ blocks for database of **N blocks** and client memory of **C blocks**
 - Caveats:
 - Server **only moves and retrieves** blocks (does not do any computation)
 - Constructions with **statistical security**
 - Constructions that **work for any block size**
 - The client can have **oracle access to a private random function**
- Boyle-Naor'16
 - Formalized the “**balls and bins**” model
 - Evidence for the hardness of extending the lower bound beyond the ball and bins model
 - **Reduction from sorting circuits to ORAM**



ORAM Lower Bound

- Goldreich-Ostrovsky'96
 - Lower bound $\Omega(\log_c N)$ blocks for database of **N blocks** and client memory of **C blocks**
 - Caveats:
 - Server **only moves and retrieves** blocks (does not do any computation)
 - Constructions with **statistical security**
 - Constructions that **work for any block size**
 - The client can have **oracle access to a private random function**
- Boyle-Naor'16
 - Formalized the “**balls and bins**” model
 - Evidence for the hardness of extending the lower bound beyond the ball and bins model
 - **Reduction from sorting circuits to ORAM**
- Larsen-Nielsen'18
 - Lower bound extended to **computational online ORAM model (only block uploads/downloads)**





PanORAMa:

1

2

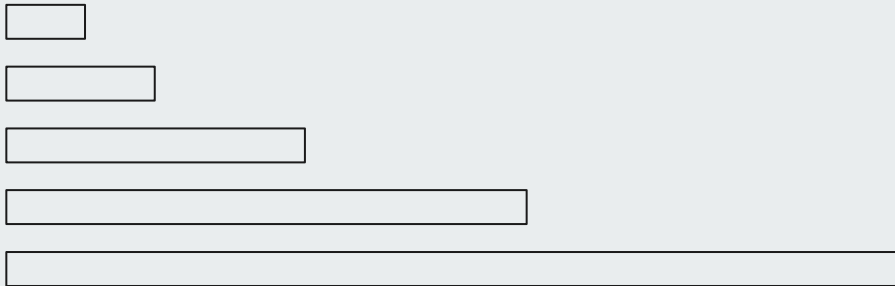
3

4

5




PanORAMa:



- New Oblivious RAM Construction
 - Improved asymptotic communication
 - $O(\log N \cdot \log \log N)$ blocks
 - Block size $\Omega(\log N)$
 - Can be instantiated in the balls and bins model
 - Follows the hierarchical paradigm



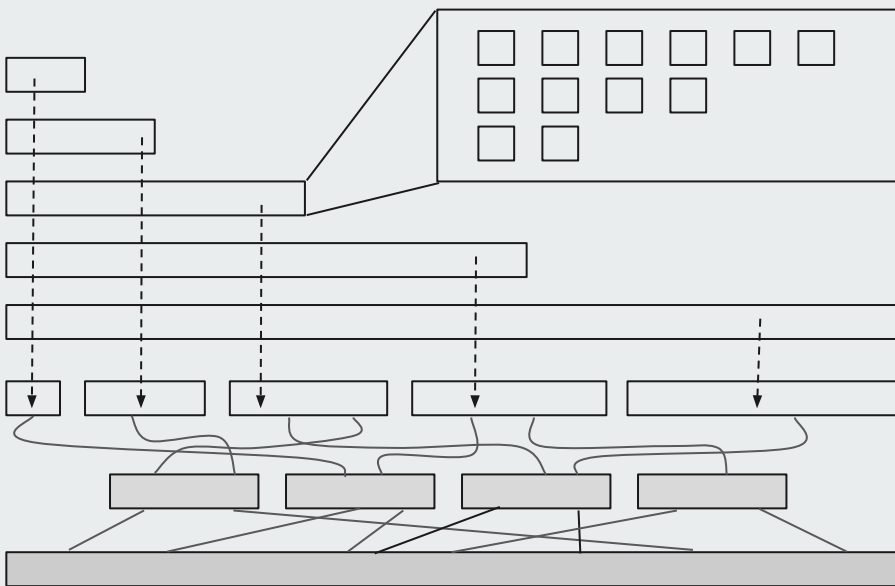
PanORAMa:



- **New Oblivious RAM Construction**
 - **Improved asymptotic communication**
 - $O(\log N \cdot \log \log N)$ blocks
 - Block size $\Omega(\log N)$
 - Can be instantiated in the balls and bins model
 - Follows the hierarchical paradigm
- **New Oblivious Hash Table Construction**
 - Obliviousness for non-repeating queries
 - **Efficient initialization from random array**
 - Amortized query communication complexity
 - $O(\log N + \text{poly}(\log \log \lambda))$



PanORAMa:



- **New Oblivious RAM Construction**
 - **Improved asymptotic communication**
 - $O(\log N \cdot \log \log N)$ blocks
 - Block size $\Omega(\log N)$
 - Can be instantiated in the balls and bins model
 - Follows the hierarchical paradigm
- **New Oblivious Hash Table Construction**
 - Obliviousness for non-repeating queries
 - **Efficient initialization from random array**
 - Amortized query communication complexity
 - $O(\log N + \text{poly}(\log \log \lambda))$
- **New Multi-Array Shuffle Algorithm**
 - Efficient shuffle for **input with entropy**
 - Shuffle multiple sorted arrays
 - $O(N \log \log \lambda + N \log N \log \lambda)$
 - Not too many very small arrays



Construction Paradigms



- **Hierarchical ORAMs**
 - Worst case \neq Average case
 - Computation assumption beyond encryption in most cases



- **Tree ORAMs**
 - Worst case = Average case
 - Encryption - the only computational assumption


Construction Paradigms



- Hierarchical ORAMs
 - Worst case \neq Average case
 - Computation assumption beyond encryption in most cases



- Tree ORAMs
 - Worst case = Average case
 - Encryption - the only computational assumption



Construction Paradigms



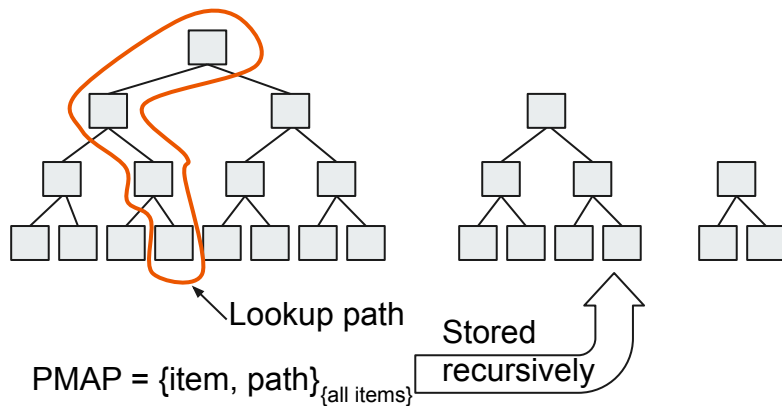
- Hierarchical ORAMs

- Worst case \neq Average case
- Computation assumption beyond encryption in most cases



- Tree ORAMs

- Worst case = Average case
- Encryption - the only computational assumption





Construction Paradigms

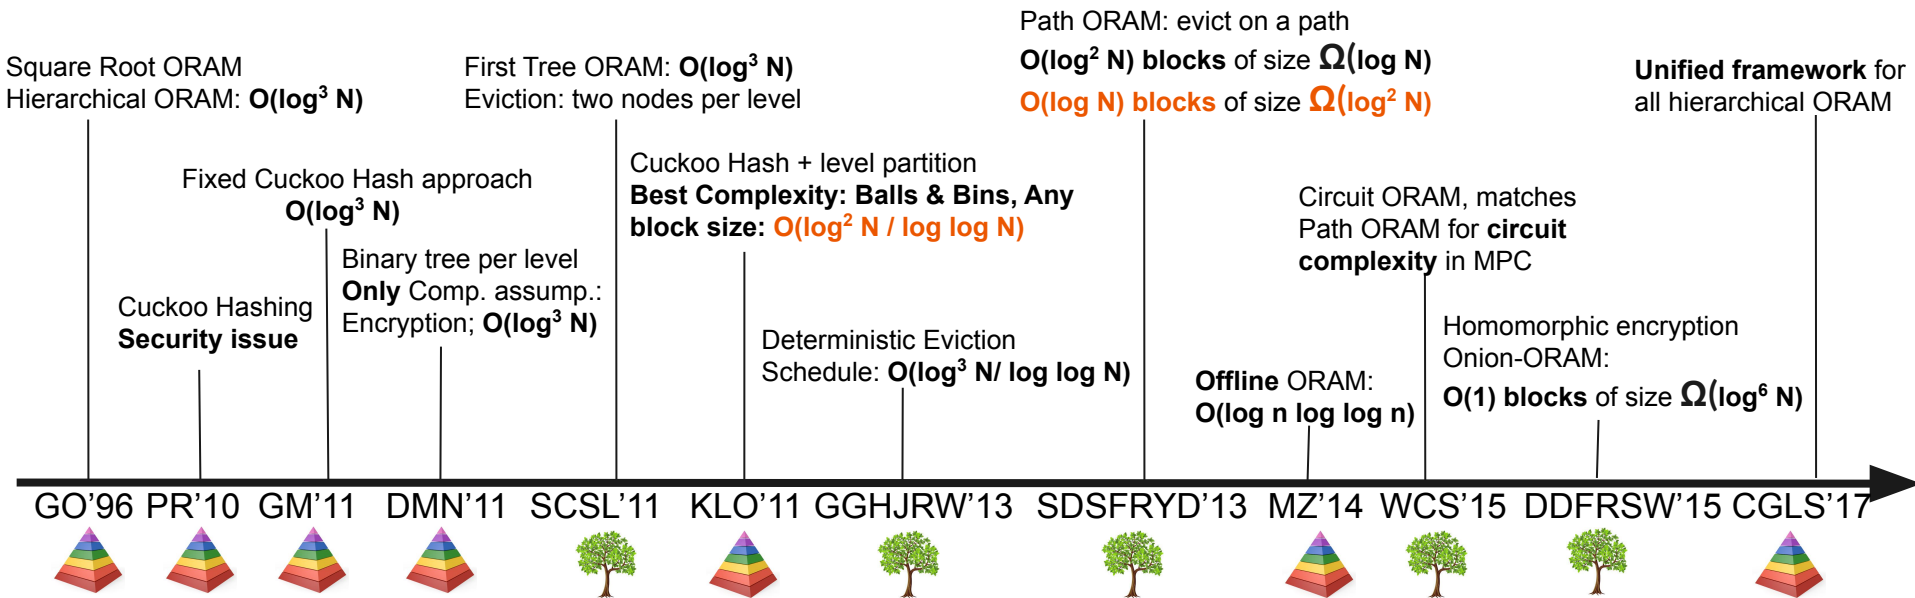


- **Hierarchical ORAMs**
 - Worst case \neq Average case
 - Computation assumption beyond encryption in most cases



- **Tree ORAMs**
 - Worst case = Average case
 - Encryption - the only computational assumption

Timeline and Complexity



The Hierarchical ORAM Paradigm



Hierarchical Construction

Level

1



2



⋮

i



⋮

Level i has capacity for **all items** assigned to **levels 1 to i** at any moment (2^i items)

$\log N - 1$



$\log N$





Hierarchical Construction: Search

Level

1



Linear scan

2



⋮

i



⋮

$\log N - 1$



$\log N$





Hierarchical Construction: Search

Level

1



Linear scan

2



Look up **searched item**

⋮

i



Look up **searched item**

Item found ✓

⋮

log N - 1



log N





Hierarchical Construction: Search

Level

1



Linear scan

2



Look up **searched item**

⋮

i



Look up **searched item**

Item found ✓

⋮

$\log N - 1$



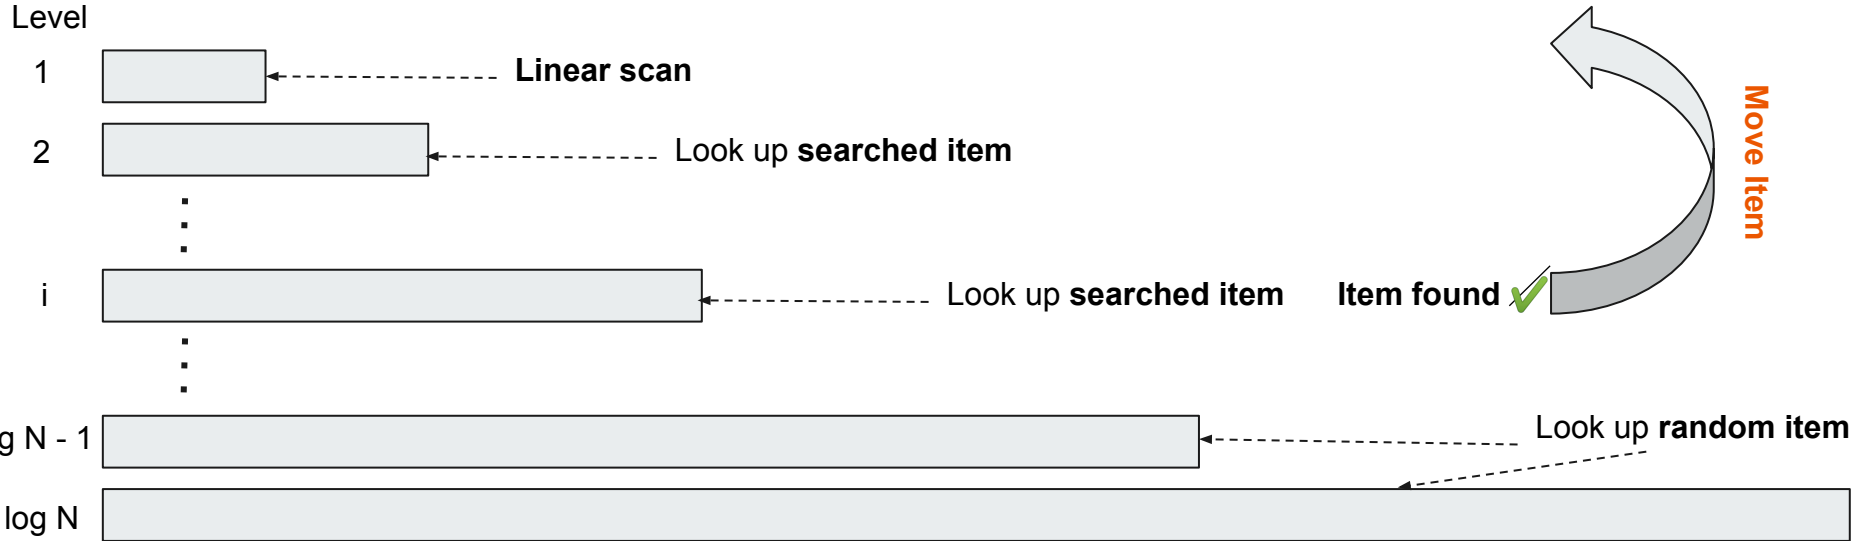
Look up **random item**

$\log N$



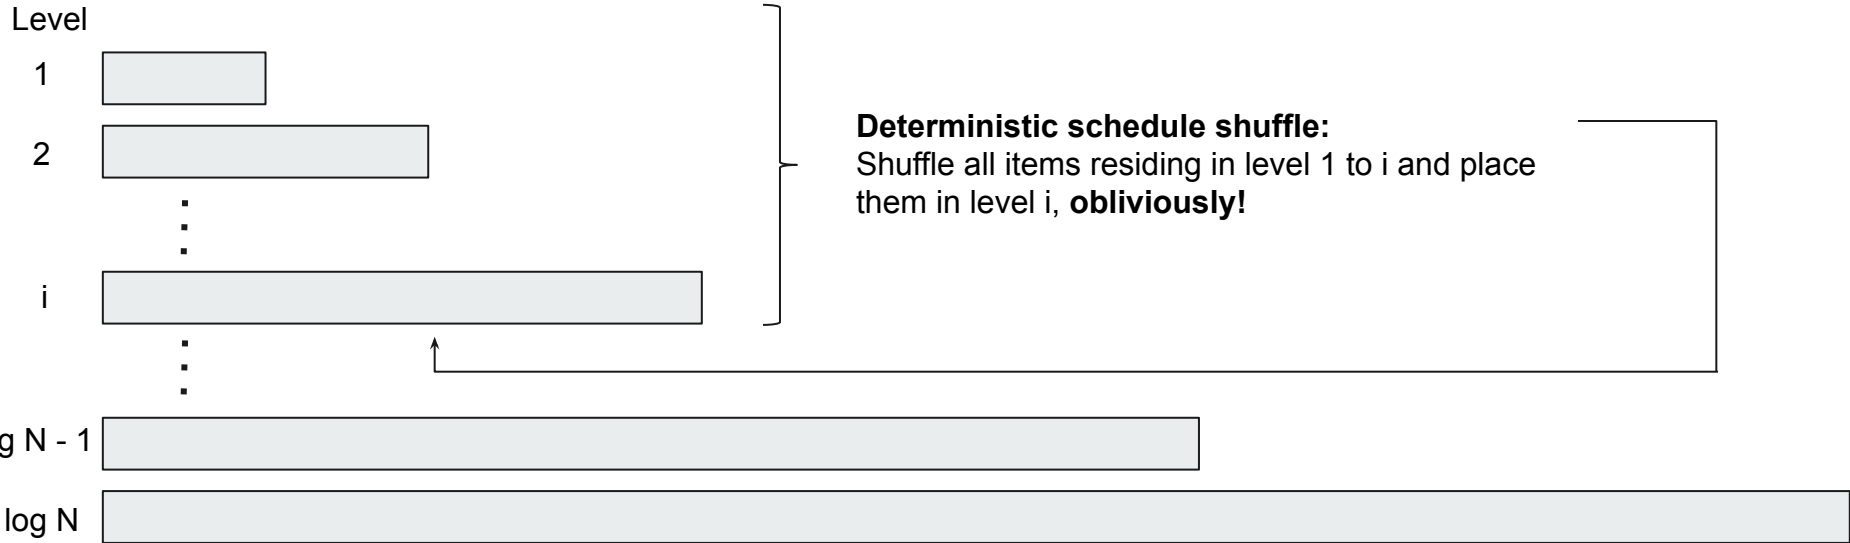


Hierarchical Construction: Search





Hierarchical Construction





Oblivious Hash Table [CGLS'17]

Level

1



2



⋮

i



⋮

$\log N - 1$



$\log N$



- Instantiate each level with
oblivious hash table (OHT)
- **Access pattern hiding for non-repeating queries**

Efficiency costs:


- **Query cost**
- **Oblivious initialization**
 - During shuffle



Existing Oblivious Hash Table Constructions




Existing Oblivious Hash Table Constructions



GO'96:


- Use pseudo-random function to match items to level buckets
- Query: retrieve item bucket $O(\log n)$
- Oblivious initializations: **several oblivious sorts** $O(n \log n)$

Existing Oblivious Hash Table Constructions



GO'96:

- Use pseudo-random function to match items to level buckets
- Query: retrieve item bucket $\mathbf{O(\log n)}$
- Oblivious initializations: **several oblivious sorts $\mathbf{O(n \log n)}$**




GM'11:

- Cuckoo hash table
- Query: $\mathbf{O(1)}$
- Oblivious initialization: oblivious sort $\mathbf{O(n \log n)}$



Existing Oblivious Hash Table Constructions




KLO'12:

- Level size n : **$\log n$** Cuckoo hash tables; each shuffle creates a new one
- Query all Cuckoo tables: **$O(\log n)$**
- Oblivious initialization: $\log n$ oblivious sorts on $n/\log n$ items per n queries: **$O(\log n)$**




Oblivious Two Tier Hash Table [CGLS'17]

Oblivious Two Tier Hash Table [CGLS'17]



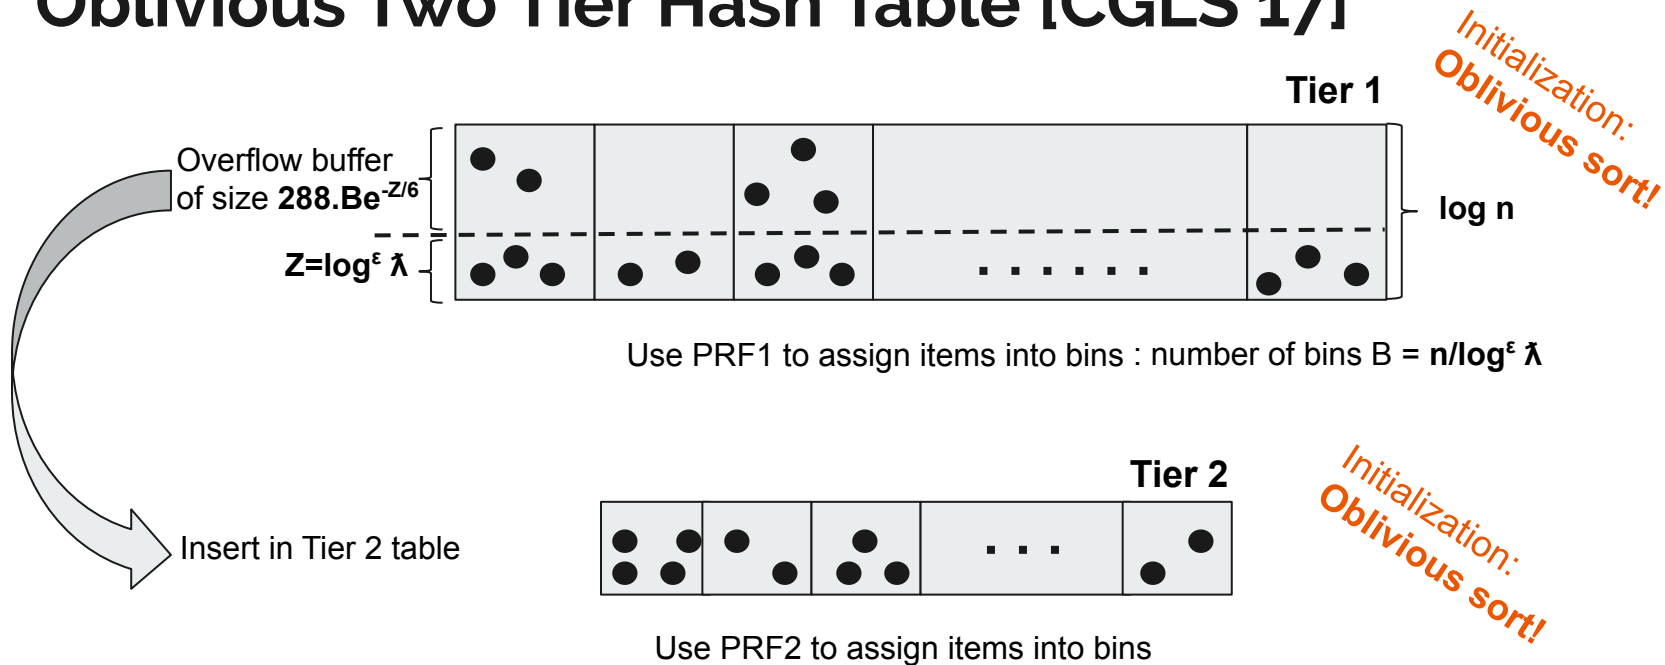
Use PRF1 to assign items into bins : number of bins $B = n/\log^\epsilon \lambda$

Oblivious Two Tier Hash Table [CGLS'17]




Use PRF1 to assign items into bins : number of bins $B = n / \log^\epsilon \lambda$


Oblivious Two Tier Hash Table [CGLS'17]



Oblivious Two Tier Hash Table [CGLS'17]



If not found, Retrieve bin PRF2(query)
else, Retrieve random bin.



*Amortized complexity:
 $O(\log^2 N / \log \log N)$*

PanORAMa Overview

- *Leverage entropy reuse to shuffle more efficiently*



PanORAMa Hierarchical Construction

Level

1



2



⋮

i



⋮

$\log N - 1$

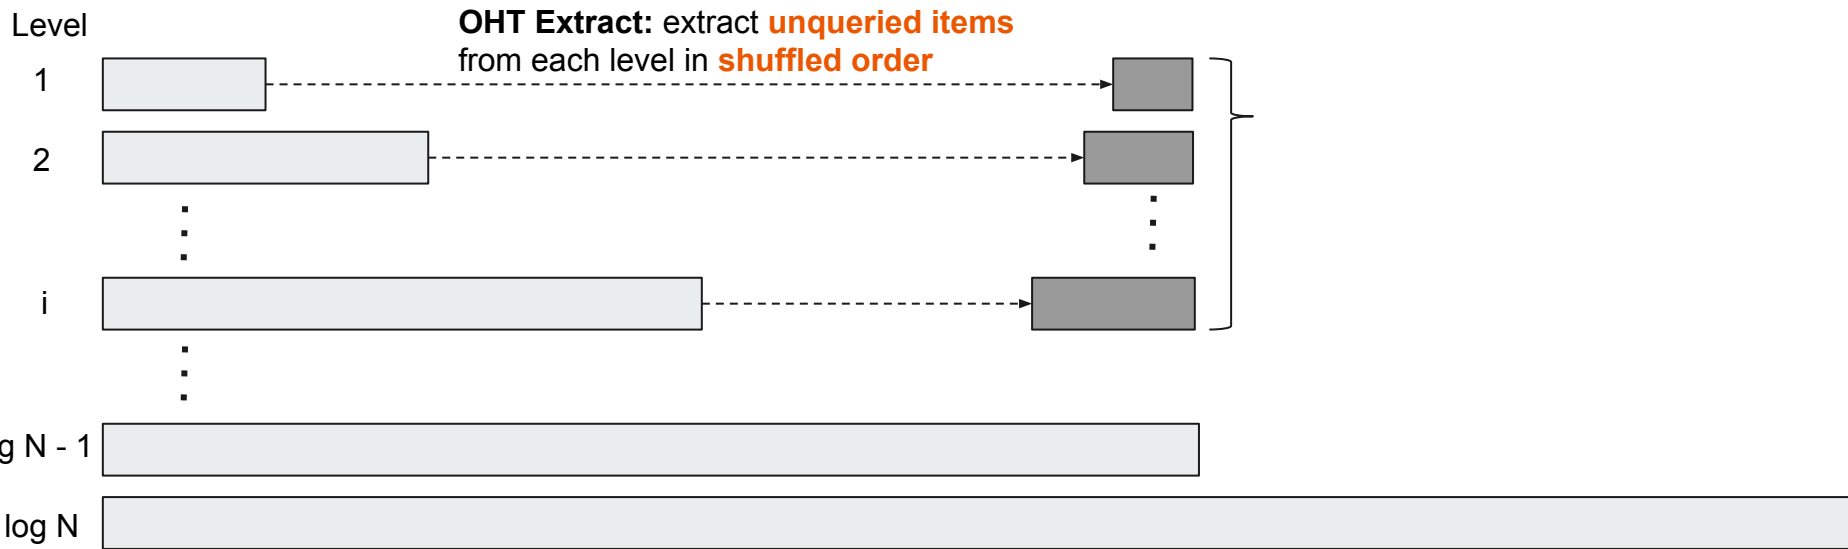


$\log N$



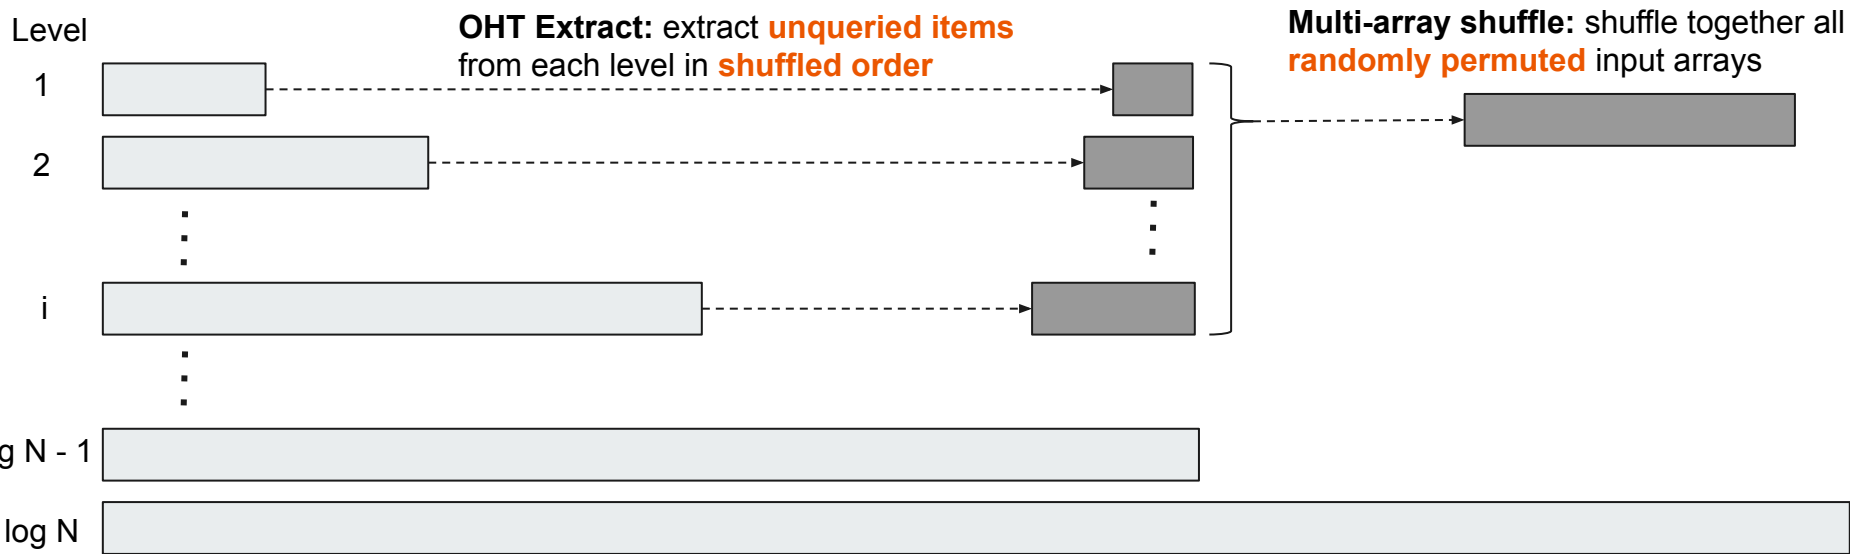


PanORAMa Hierarchical Construction






PanORAMa Hierarchical Construction





PanORAMa Hierarchical Construction

*No Oblivious Sorting
on a whole large level!*



Oblivious Hash Table

- *Oblivious initialization in $o(n \log n)$ leveraging input entropy*



Oblivious Hash Table

- **Definition.** Oblivious Hash Table (OHT):
 - **OHT.Init** - permutes input
 - **OHT.Build** - builds OHT from permuted input
 - **OHT.Lookup** - execute a query
 - **OHT.Extract** - extracts an array that contains unqueried item in random order
- **Security.**
 - **Access hiding: non-repeating** query sequences
 - Extract output **indistinguishable from random permutation**




Oblivious Hash Table

- **Definition.** Oblivious Hash Table (OHT):
 - **OHT.Init** - permutes input
 - **OHT.Build** - builds OHT from permuted input
 - **OHT.Lookup** - execute a query
 - **OHT.Extract** - extracts an array that contains unqueried item in random order
- **Security.**
 - **Access hiding: non-repeating** query sequences
 - Extract output **indistinguishable from random permutation**
- **Oblivious bin** = mini OHT
 - OHT that is instantiated on small input size $O(\text{polylog } N)$
 - We can use **oblivious sorting** without hurting efficiency

Oblivious Bin

- Oblivious Cuckoo Bin



- Cuckoo hash
- Oblivious sort to build and extract
 - Add **n dummies** in Build
 - Extract n items in Extract

- “Dynamic” Bin


- Items need to be added continuously in non-amortized manner
- Smallest ORAM level
- Size: $O(\log^7 n)$
- Use existing oblivious ram constructions, e.g. Goodrich, Mitzenmacher [GM'11]




Oblivious Hash Table



Oblivious Hash Table




Oblivious Hash Table




Oblivious resampling of bucket loads: $\Pr[\text{new load} > \text{cutoff}] < \text{negl}(\lambda)$ & $\Pr[\#\text{items} < \text{new load}] < \text{negl}(\lambda)$

Oblivious Hash Table




Oblivious resampling of bucket loads: $\Pr[\text{new load} > \text{cutoff}] < \text{negl}(\lambda)$ & $\Pr[\#\text{items} < \text{new load}] < \text{negl}(\lambda)$

Oblivious Hash Table



Oblivious resampling of bucket loads: $\Pr[\text{new load} > \text{cutoff}] < \text{negl}(\lambda)$ & $\Pr[\#\text{items} < \text{new load}] < \text{negl}(\lambda)$







Oblivious Hash Table: Create Oblivious Bins



⋮




- Items from Cuckoo stash are
- (encrypted) indexed with their source bucket
 - merged with items used to instantiate the last level



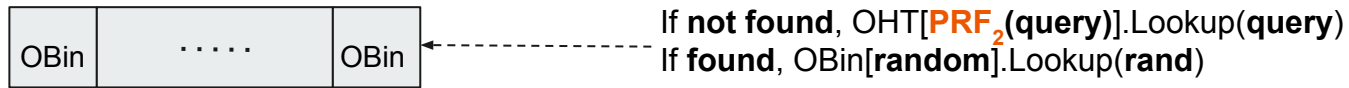

Oblivious Hash Table: Query




⋮



Oblivious Hash Table: Query



⋮





Oblivious Hash Table: Extract




⋮



Move items with origin in the Cuckoo stashes back to their corresponding levels

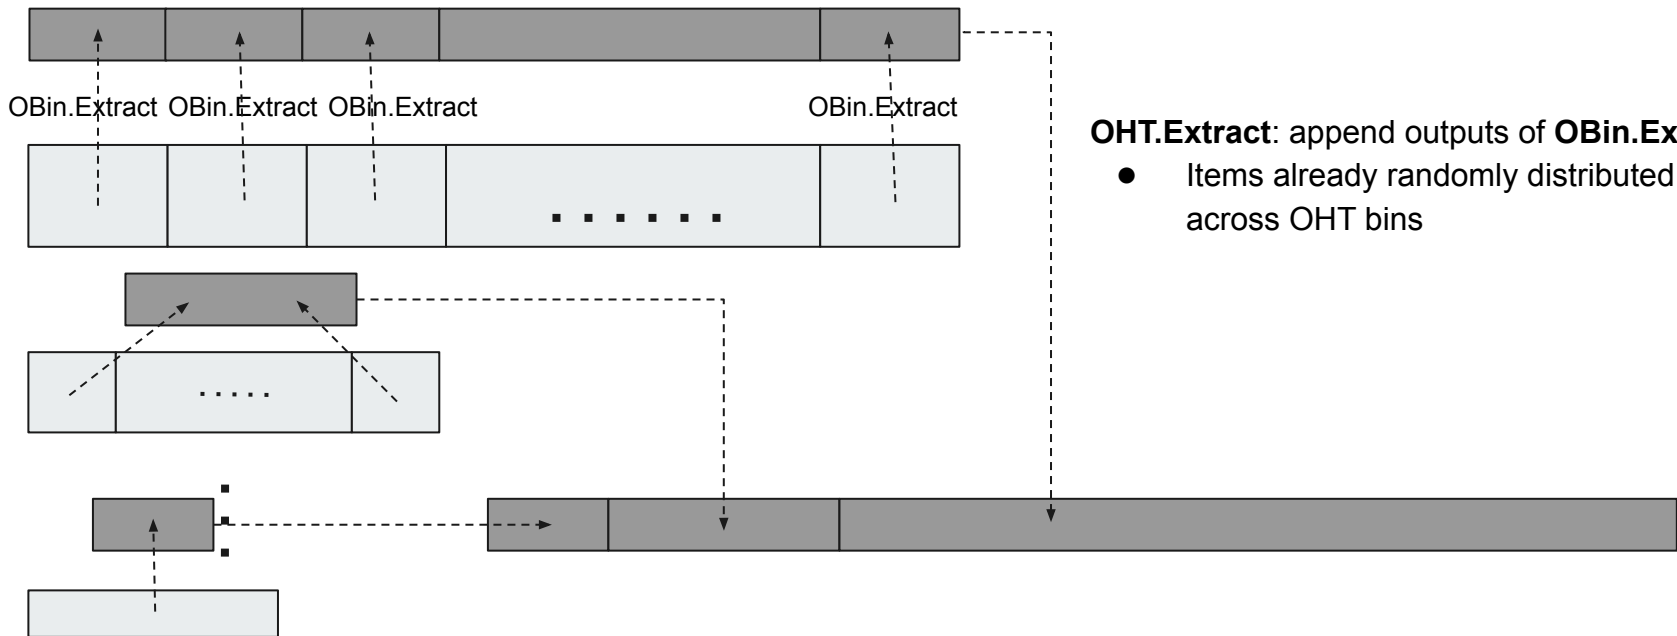


Oblivious Hash Table: Extract





Oblivious Hash Table: Extract





Oblivious Hash Table

Amortized Communication Complexity
over N accesses for OHT on N items:

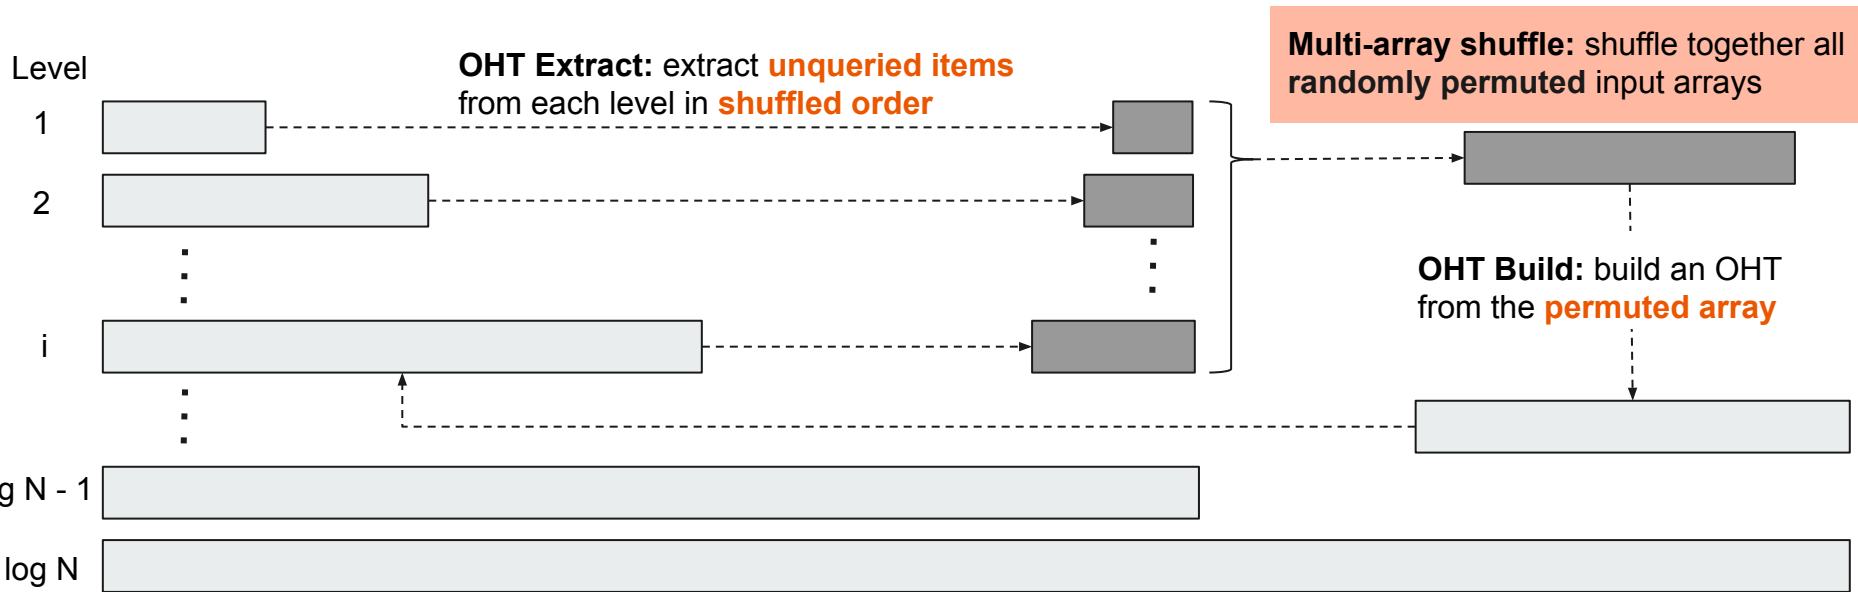
- OCuckooBin: $O(\log N + (\log \log \lambda))$

Oblivious Multi-Array Shuffle

- *Random shuffle in $o(n \log n)$ leveraging input entropy: independently sorted input arrays*




PanORAMa Hierarchical Construction






Oblivious Multi-Array Shuffle






Oblivious Multi-Array Shuffle





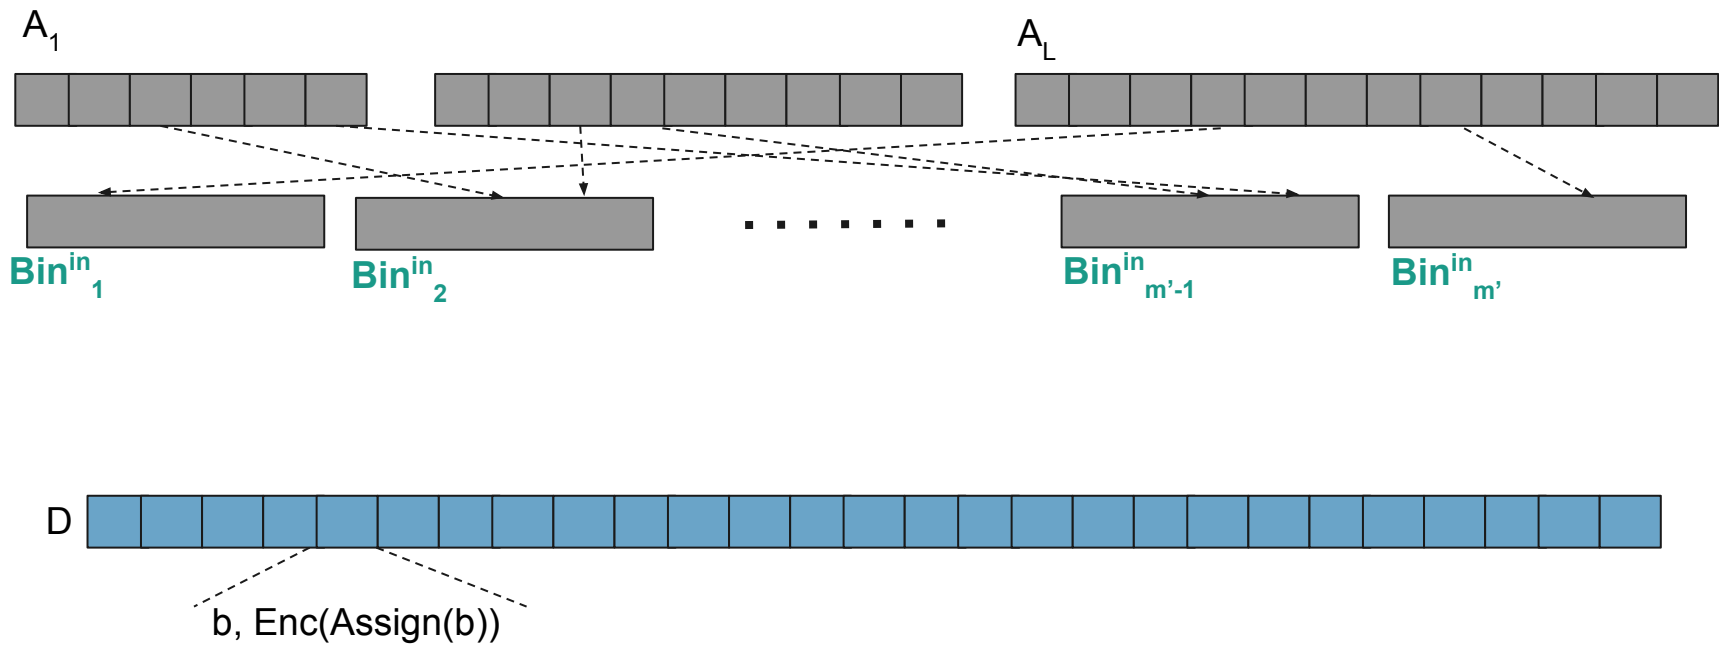
Oblivious Multi-Array Shuffle




$b, \text{Enc}(\text{Assign}(b))$ → **Obliviously Sample**



Oblivious Multi-Array Shuffle



Oblivious Multi-Array Shuffle





Bin Shuffle

Bin_i^{in}



For each j in $[L]$, Bin_i^{in} contains more elements from A_j than the number of elements from A_j assigned to $\text{Bin}_i^{\text{out}}$


$\text{Bin}_i^{\text{out}}$





Bin Shuffle

Bin_i^{in}




Add dummies with each array index

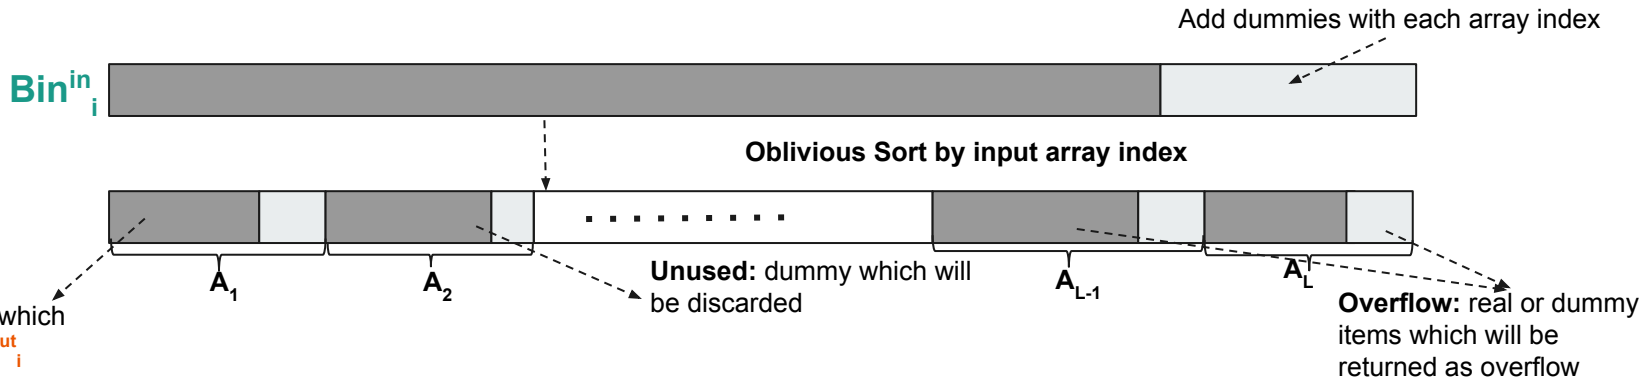
$\text{Bin}_i^{\text{out}}$



Bin Shuffle

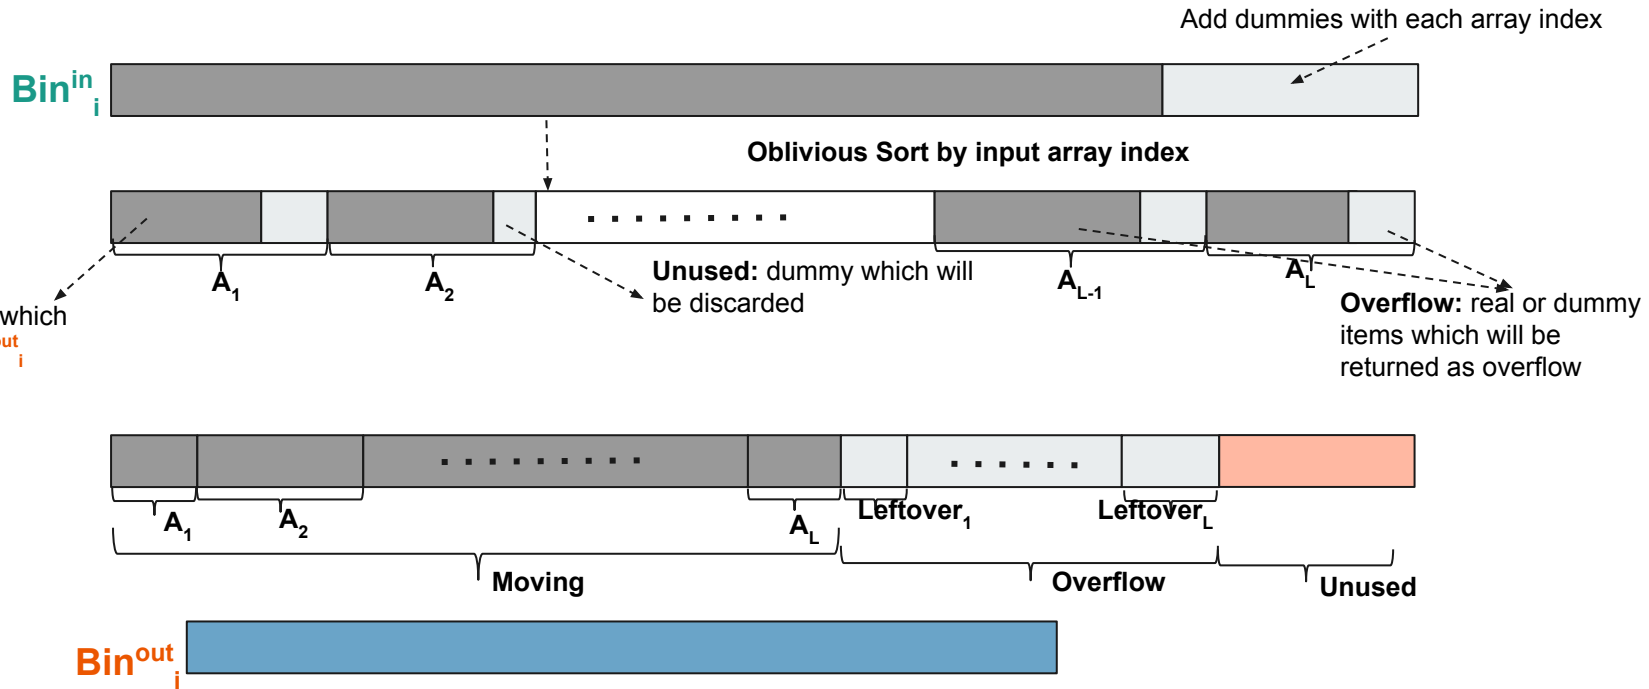


Bin Shuffle




$\text{Bin}^{\text{out}}_i$ [empty bar]

Bin Shuffle

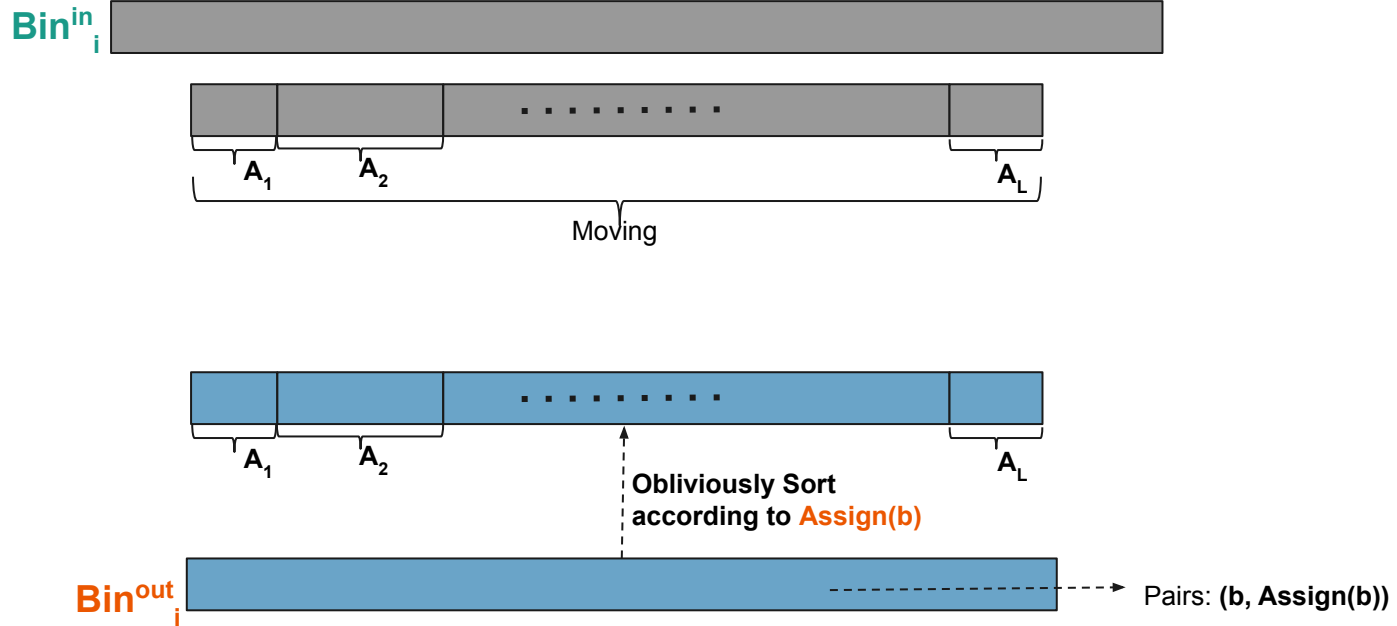




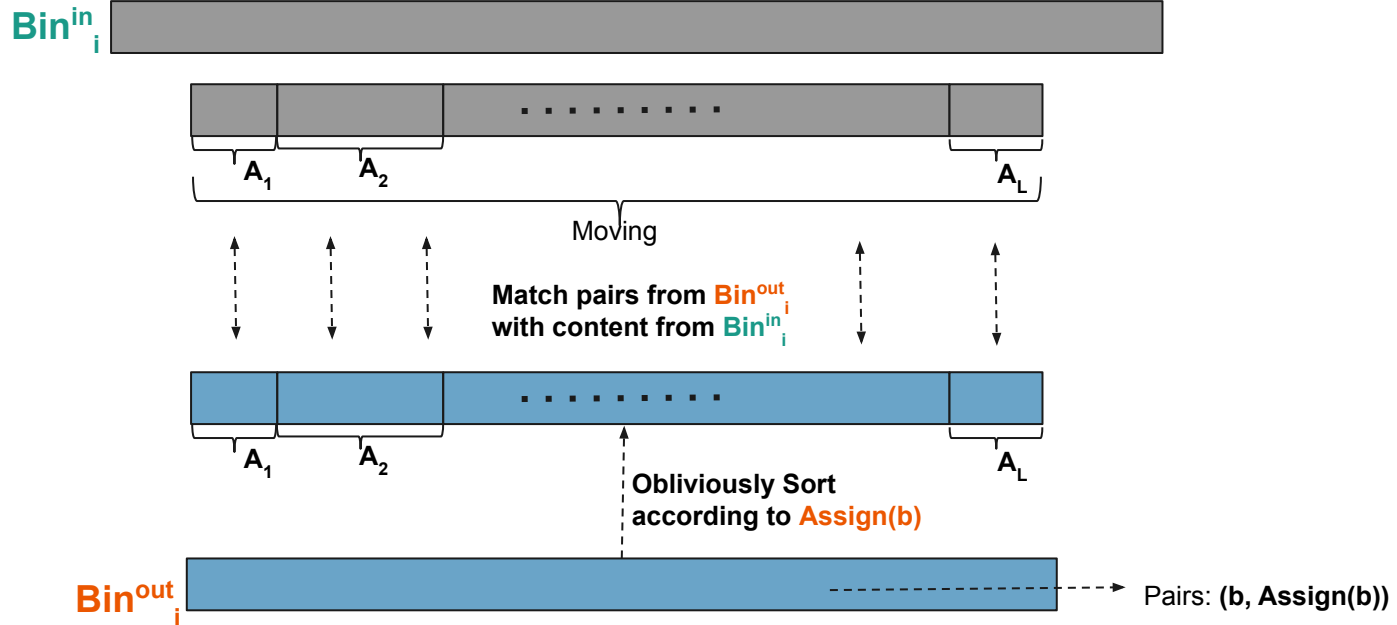
Bin Shuffle



Bin Shuffle



Bin Shuffle



Bin Shuffle

Binⁱⁿ_i



(b, Assign(b) = 1)

(b', Assign(b') = 2)

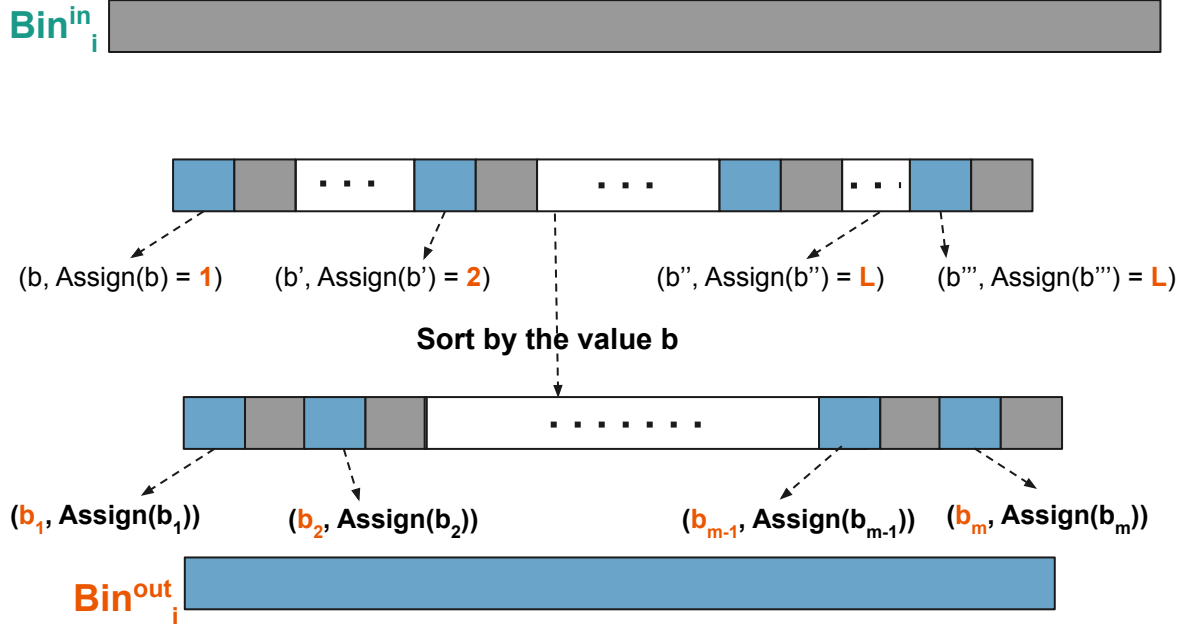
(b'', Assign(b'') = L)

(b''', Assign(b''') = L)

Bin^{out}_i

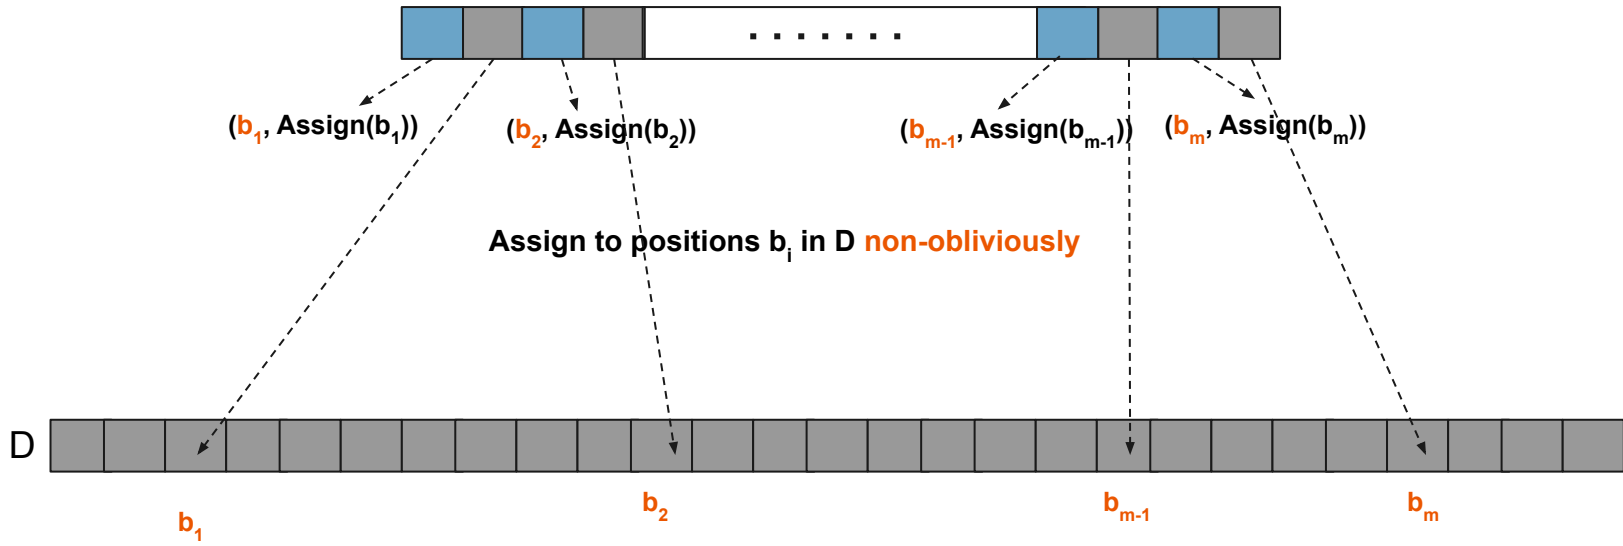


Oblivious Multi-Array Shuffle






Bin Shuffle



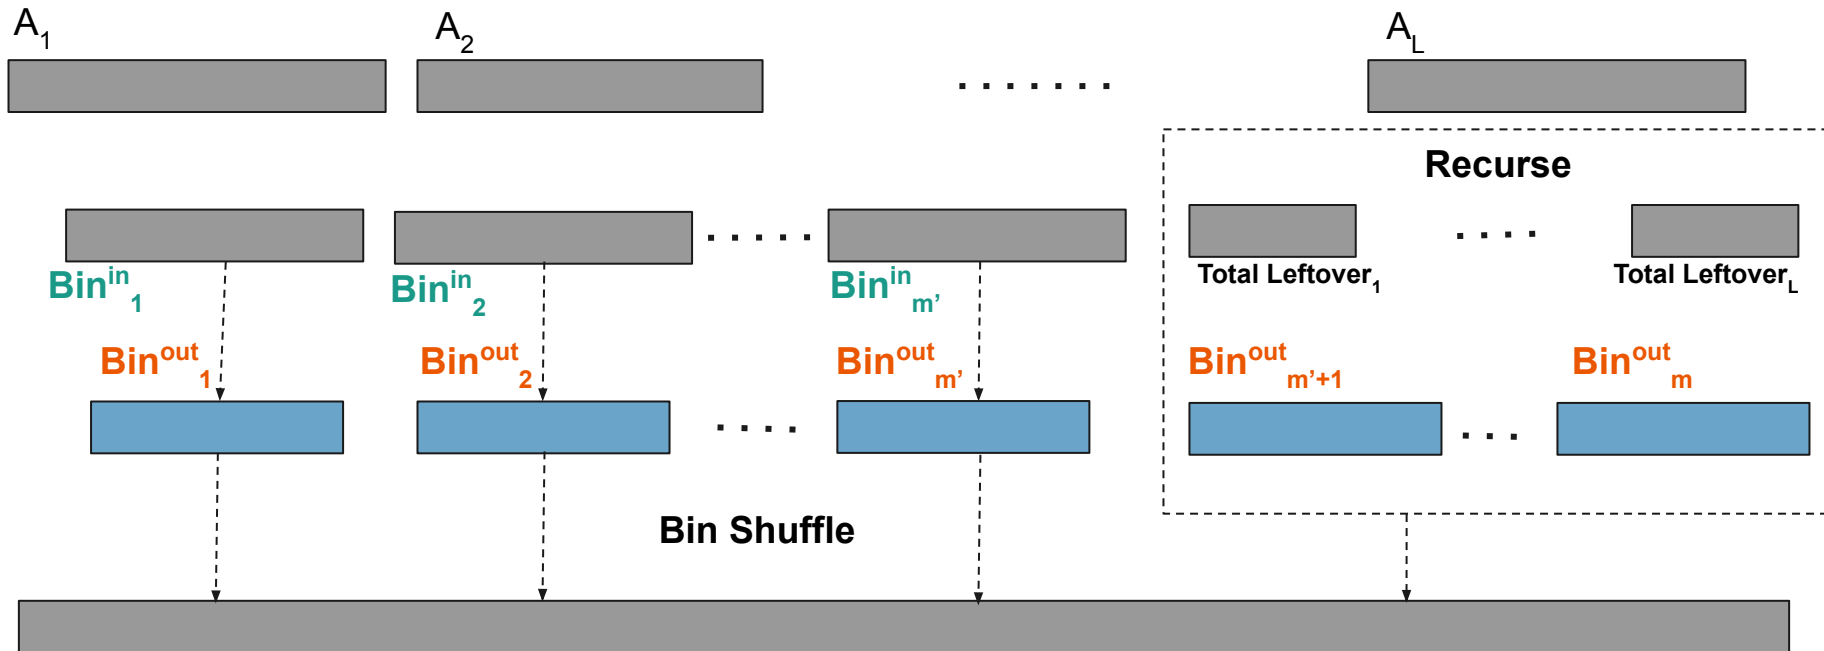


Oblivious Multi-Array Shuffle





Oblivious Multi-Array Shuffle





Ball and Bins Model

- We can instantiate the PanORAMa construction in the model where GO'96 proved $O(\log N)$ lower bound
 - Server does no computation on the data \Rightarrow satisfy “balls and bins” requirement
 - GO'96 allows client to oracle access to private random function \Rightarrow replace PRF
- PanORAMa complexity: $O(\log N \cdot \log \log N)$

Follow-up Work




ORAM with Logarithmic Complexity

- **OptORAMa: Optimal Oblivious RAM [AKLNPS18] (eprint 2018/892)**
 - Communication complexity: $O(\log N)$
 - Oblivious compaction: $O(N)$



Overview



- **PanORAMa**: new ORAM construction with improved asymptotic complexity
 - $O(\log N \cdot \log \log N)$ for block size $\Omega(\log N)$
- New Efficient Building Blocks
 - **Oblivious Hash Table**
 - **Oblivious Multi-Array Shuffle**

Thanks!

Questions?

