

# MELBOURNE SHUFFLE ON OBLIVIOUS STORAGE

---

Presented by: Yi Qian

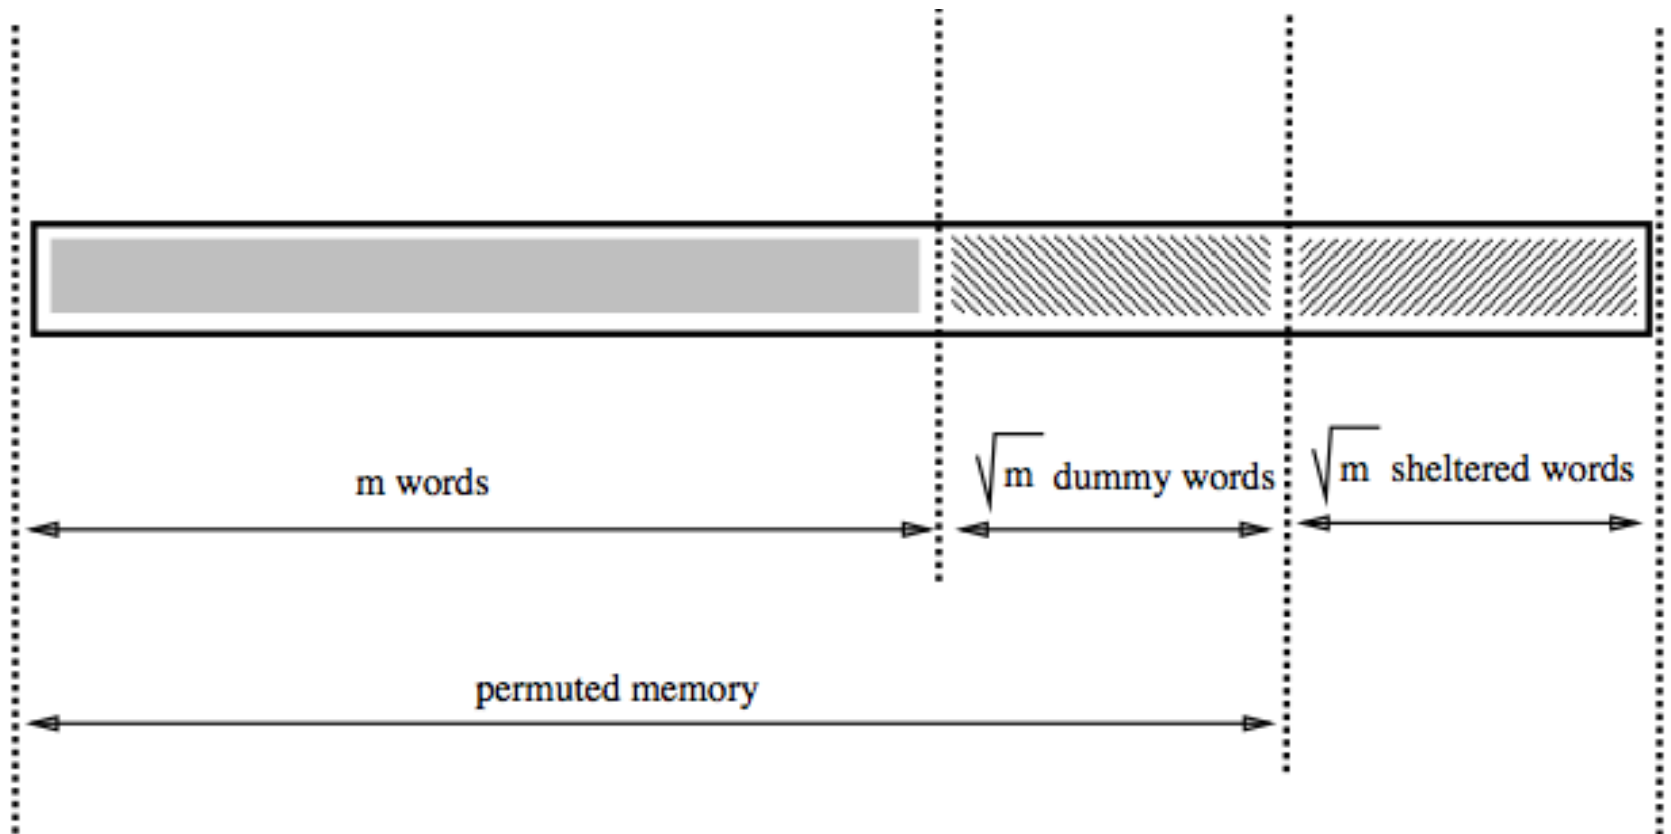
# ORAM

- A client wishes to store data at a remote untrusted
- Data access pattern can reveal very sensitive information to the untrusted server
- The goal of ORAM is to completely hide the data access pattern :
- 1) which data is being accessed; 2) how old it is (when it was last accessed); 3) whether the same data is being accessed (linkability); 4) access pattern (sequential, random, etc); or 5) whether the access is a read or a write.

# Naïve & Square Root Solutions

- Naive: For each access (write/read), download all data
- **Square Root: Covered by Prof already**
- Access overhead:  $O(\sqrt{n} (\log n)^2)$
- Client Memory (Message size):  $O(1)$
- Server Storage :  $O(n + \sqrt{n})$

# Square Root Solutions



# Square Root Solutions

- 1. A permutation  $P$  over the integers  $[1, n + \sqrt{n}]$  is uniformly selected and the content of the permuted memory is obviously relocated.
- 2. For each access with (location-index  $i$ ):
  - 2a. Scan through the entire shelter in a sequential order to find an item whose location-index is  $i$ . If the item is not found in the shelter, access the actual location  $P(i)$  in the permuted area. If found, access the next dummy item.
  - 2b. Scan again through the entire shelter in a sequential order, and write back the updated (re-encrypted) item of location-index  $i$ , in the next available location.

# Square Root Solutions

- Step 2 takes  $O(\sqrt{n})$ . The **bottleneck** is **step 1**.
- Step 1: oblivious shuffle.
  - Need to shuffle the entire data every  $\sqrt{n}$  accesses.
- Assign a random integer to each data. Then do Oblivious sorting.
- Oblivious sorting: sort an array non-adaptively
- Batcher's Sorting Network:  $O(n (\log n)^2)$

# Batcher's Sorting Network

- Normal sorting alg:  $O(n \log n)$ . Non-oblivious.
  - Do compare-and-switches between pairs of keys. Which pairs depends on the input array.
  - Batcher's Sorting: fixed indexes of compare-and-switches.
- $\text{Sort}(x_1, \dots, x_n)$ :
  - $\text{Sort}(x_1, \dots, x_{n/2})$ , then  $\text{Sort}(x_{n/2+1}, \dots, x_n)$ , and then  $\text{Merge}(x_1, \dots, x_n)$ .
- $\text{Merge}(x_1, \dots, x_n)$ :
  - $\text{Merge}(x_i, \text{ for } i \text{ odd})$ , then  $\text{Merge}(x_i \text{ for } i \text{ even})$ , and then  $\text{Comp}(x_2, x_3)$ ,  $\text{Comp}(x_4, x_5)$ ,  $\dots$ ,  $\text{Comp}(x_{n-2}, x_{n-1})$ .

$\text{Comp}(x_i, x_j)$  : compare and switch with key  $i$  and key  $j$

# Batcher's Sorting Network Example

- Given 2 7 6 3 9 4 1 8
- If we sort the first and second halves separately we obtain: 2 3 6 7 1 4 8 9
- Sorting the odd-indexed keys (2, 6, 1, 8) and then the even-indexed keys (3, 7, 4, 9) while leaving them in odd and even places respectively yields: 1 3 2 4 6 7 8 9
- This list is now almost sorted: doing a comparison switch between the keys in positions (2 and 3),
- (4 and 5) and (6 and 7) will in fact finish the sort.



# Batcher's Sorting Network Complexity

- **S(n)**: comparisons of  $\text{sort}(x_1, \dots, x_n)$
- **M(n)** comparisons of  $\text{merge}(x_1, \dots, x_n)$
- **S(n) = 2S(n/2) + M(n)**,
- **M(n) = 2M(n/2) + n/2 - 1.**
- **S(2)=M(2)=1**
- Inductively, **M(n)  $\leq$  n/2 log n and S(n)  $\leq$  n/2 (log n)<sup>2</sup>**
  
- **Square root solution shuffle every  $\sqrt{n}$  accesses.** Hence, amortized overhead access is  **$(\sqrt{n}(\log n)^2)$**

# Melbourne Shuffle

Table 1: Comparison of data-oblivious sorting and shuffle algorithms over  $n$  items.

	<i>Randomized</i>	Private Memory	Message Size	External Memory	I/Os
Batcher's network [3]		$O(1)$	$O(1)$	$O(n)$	$O(n(\log n)^2)$
Batcher's network I		$O(\sqrt{n})$	$O(\sqrt{n})$	$O(n)$	$O(\sqrt{n}(\log n)^2)$
Batcher's network II [9]		$O(\sqrt{n})$	$O(\sqrt[8]{n})$	$O(n)$	$O(n^{7/8})$
AKS [1], Zig-zag sort [8]		$O(1)$	$O(1)$	$O(n)$	$O(n \log n)$
Randomized shellsort [7]	✓	$O(\sqrt{n})$	$O(\sqrt{n})$	$O(n)$	$O(\sqrt{n} \log n)$
Melbourne shuffle	✓	$O(\sqrt{n})$	$O(\sqrt{n})$	$O(n)$	$O(\sqrt{n})$
Melbourne shuffle ( $c \geq 3$ )	✓	$O(\sqrt[c]{n})$	$O(\sqrt[c]{n})$	$O(n)$	$O(c\sqrt[c]{n}^{c-1})$

# Melbourne Shuffle

- Takes advantage of allowing reasonable client memory and message size.
- The access overhead is measured w.r.t. the amortized number of I/O executions.
- Supported operations:
  - Read/write an element at location  $loc$
  - Read/write a list of elements at locations  $loc, \dots, loc+L-1$
  - Read/write a dist of elements at locations specified by:
    - $Loc_i + L_i-1$  for all  $i$ .

# Melbourne Shuffle Problem

- $(\text{Enc}(\pi(A)), \alpha) \leftarrow \text{Shuffle}(s, S, A, \pi)$
- Input:  $s$  is the secret key,  $S$  is the server storage,  $A$  is the input array of length  $n$  and  $\pi$  is a pseudorandom permutation.
- Return (1) the encryption of the permutation of  $A$  according to  $\pi$ ; (2) a transcript  $\alpha$  of the operations that transform  $\text{Enc}(A)$  to  $\text{Enc}(\pi(A))$  using space  $S$ .
- Every operation in  $\alpha$  must be oblivious


# Melbourne Shuffle Basic Algorithm

- Client memory / message size:  $O(\sqrt{n} \log n)$
- Server storage:  $O(n \log n)$
- Overhead access:  $O(\sqrt{n})$
  
- **Intuition:**
  - 1. Split input array  $A$  into  $\sqrt{n}$  buckets (each of  $\sqrt{n}$  elements)
  - 2. For every bucket do some kind of oblivious shuffle(according to  $\pi$ ).
  - 3. Clean up and return the  $\text{Enc}(\pi(A))$ .

# Melbourne Shuffle Basic Algorithm

- Storage S split into 3 parts: I, T and O.
- I: for input A
- T: An encrypted temporary array for shuffle
- O: output array
  
- Algorithm includes two phases:
  - 1. Distribution phase: Put elements from every bucket of I to every bucket of T according to the permutation.
  - 2. Clean up and return the final output array

# Distribution Phase



```

1: max_elems  $\leftarrow p \log n$ 
2: buckets  $\leftarrow \sqrt{n}$ 
3: Up to date on phase: distribute elements of I into T
4: for  $id_I \in \{0, \dots, \text{num\_buckets} - 1\}$  do {read buckets of  $I$ }
5:   bucket $_M \leftarrow \text{getRange}(I, id_I \times \sqrt{n}, \sqrt{n})$ 
6:   rev_bucket $_M \leftarrow \text{empty\_map}()$  {Reverse map of bucket ids in  $T$  to elements}
7:   for  $e \in \text{bucket}_M$  do {Assign elements their bucket ids in  $T$ }
8:      $(x, v) \leftarrow \text{Dec}(e)$ 
9:      $id_T \leftarrow \lfloor \rho(x) / \sqrt{n} \rfloor$  {Bucket id of element  $(x, v)$  in  $T$  according to its location in  $O$ }
10:    rev_bucket $_M[id_T].\text{add}(\text{Enc}(x, v))$  {Collect elements of same bucket}
11:   end for
12:   {Can be done via a single putRangeDist for  $\sqrt{n}$  batches of size max_elems}
13:   for  $id_T \in \{0, \dots, \text{num\_buckets} - 1\}$  do {Distribute bucket $_M$  in buckets of  $T$ }
14:     if  $\text{size}(\text{rev\_bucket}_M[id_T]) > \text{max\_elems}$  then
15:       fail { $\rho$  moves more than  $p \log n$  elements from a bucket of  $I$  to a bucket of  $T$ }
16:     end if
17:     {Hide how many real elements go to  $T$ 's buckets by padding with encrypted dummies}
18:     rev_bucket $_M[id_T] \leftarrow \text{dummy\_pad}(\text{rev\_bucket}_M[id_T], \text{max\_elems})$ 
19:     {Write a batch of max_elems from every bucket of  $I$  to every bucket of  $T$ }
20:     putRange( $T, id_T \times \sqrt{n} \times \text{max\_elems} + \text{max\_elems} \times id_I, \text{rev\_bucket}_M[id_T]$ )
21:   end for
22: end for

```



# Clean up Phase

```
23: {Clean-up phase: clean T and write the result to O}
24: for  $id_T \in \{0, \dots, \text{num\_buckets} - 1\}$  do {read buckets of  $T$ }
25:    $\text{bucket}_M \leftarrow \text{getRange}(T, id_T \times \sqrt{n} \times \text{max\_elems}, \sqrt{n} \times \text{max\_elems})$ 
26:   {Decrypt the bucket, remove dummy, sort real elements using  $\rho$  and re-encrypt}
27:    $\text{bucket}_M \leftarrow \text{clean}(\text{bucket}_M)$ 
28:   {The distribution phase guarantees that  $\text{bucket}_M$  contains exactly  $\sqrt{n}$  elements}
29:    $\text{putRange}(O, id_T \times \sqrt{n}, \text{bucket}_M)$ 
30: end for
```

# Security & Oblivious



- The algorithm fails when there is a bucket in  $T$  contains
- $> p \log n$  elements from a bucket in  $I$ .
- Throw  $\sqrt{n}$  balls into  $\sqrt{n}$  bins. Pick  $p \geq \epsilon$  is sufficient to guarantee failure occurs with neg-probability
- Roughly: Every operation in the algorithm is oblivious.
- 1. Read data from every bucket of  $I$
- 2. Write to  $T$  with encrypted buckets of length  $\sqrt{n} p \log n$
- 3. Clean up phase is oblivious

# Optimized Melbourne Shuffle

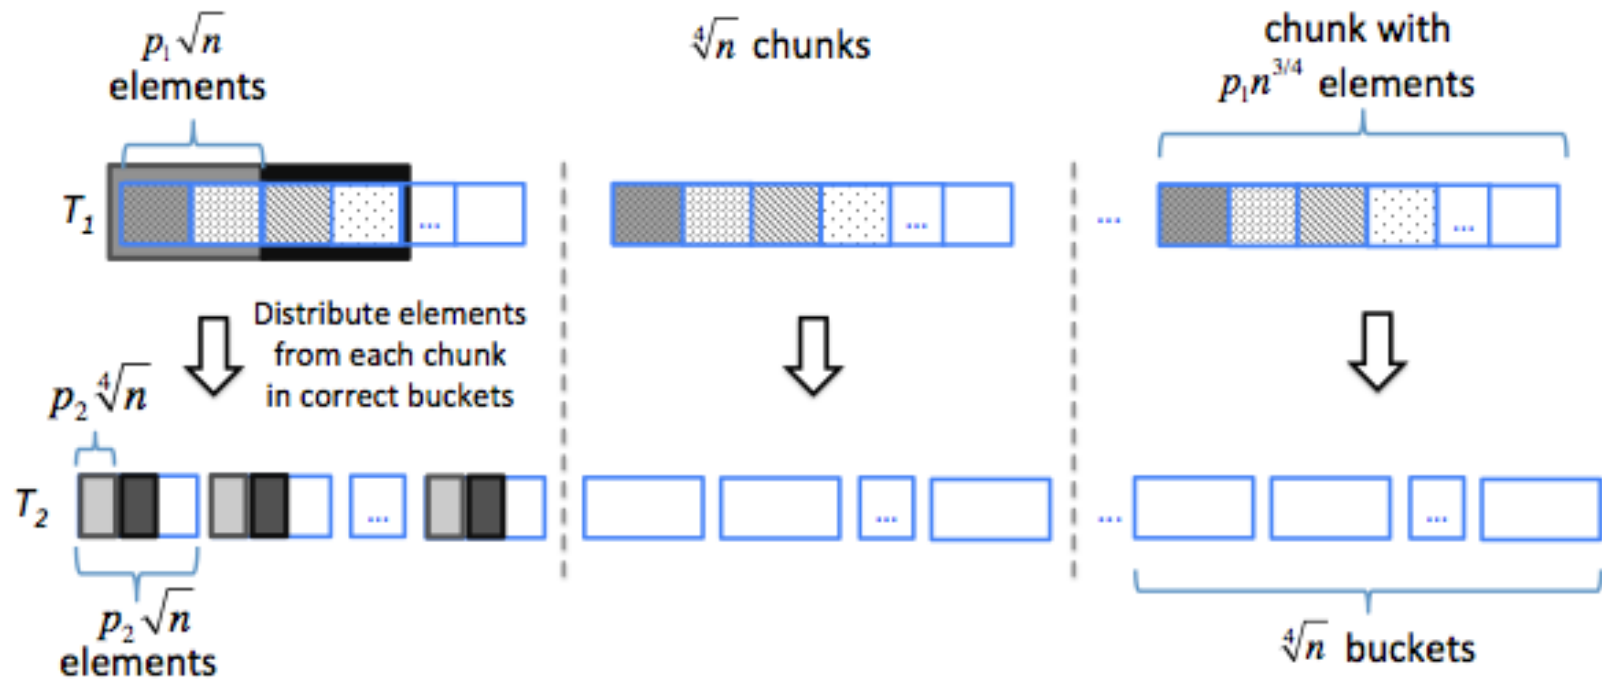
- Idea: split the distribution phase into two phases.
- 1. From I to T1. Again, split I into  $\sqrt{n}$  buckets. Distribute (according to the permutation) the elements in every buckets into chucks of length  $p_1 n^{3/4}$  in T1 (there are  $n^{1/4}$  such chucks and T1 is of length  $p_1 n$ ).
- 2. From T1 to T2. Distribute (according to the permutation) the elements in every chucks to the buckets of length  $p_2 \sqrt{n}$  in T2 (there are  $\sqrt{n}$  such buckets and T2 is of length  $p_2 n$ )

# I to T1

Cloud Storage



# T1 to T2



# Analysis

- $O(\sqrt{n})$  I/O access
- client memory / message size:  $\leq p_2\sqrt{n}$
- Server Storage:  $(p_1+p_2)n$
  
- Security & Oblivious: by similar arguments.