

Flexway O-Sort: Enclave-Friendly and Optimal Oblivious Sorting

Tianyao Gu
Carnegie Mellon University
Oblivious Labs Inc.

Yilei Wang
Alibaba Cloud

Afonso Tinoco
Carnegie Mellon University
Oblivious Labs Inc.

Bingnan Chen
HKUST

Ke Yi
HKUST

Elaine Shi
Carnegie Mellon University
Oblivious Labs Inc.

Abstract

Oblivious algorithms are being deployed at large scale in real world to enable privacy-preserving applications such as Signal’s private contact discovery. Oblivious sorting is a fundamental building block in the design of oblivious algorithms for numerous computation tasks. Unfortunately, there is still a theory-practice gap for oblivious sort. The commonly implemented bitonic sorting algorithm is not asymptotically optimal, whereas known asymptotically optimal algorithms suffer from large constants.

In this paper, we construct a new oblivious sorting algorithm called flexway o-sort, which is asymptotically optimal, concretely efficient, and suitable for implementation in hardware enclaves such as Intel SGX. For moderately large inputs of 12 GB, our flexway o-sort algorithm outperforms known oblivious sorting algorithms by $1.32\times$ to $28.8\times$ when the data fits within the hardware enclave, and by $4.1\times$ to $208\times$ when the data does not fit within the hardware enclave. We also implemented various applications of oblivious sorting, including histogram, database join, and initialization of an ORAM data structure. For these applications and data sets from 8GB to 32GB, we achieve $1.44 \sim 2.3\times$ speedup over bitonic sort when the data fits within the enclave, and $4.9 \sim 5.5\times$ speedup when the data does not fit within the enclave.

1 Introduction

Oblivious algorithms [4, 6, 20, 21, 29, 32, 38, 44, 48] allow a trusted client (e.g., a secure hardware enclave) to securely outsource data to an untrusted storage (e.g., memory or disk), such that as the client accesses the data, the addresses visited leak nothing about the secret data or query. Oblivious algorithms have been deployed at a large scale in the real world. For example, Signal, the encrypted messenger app, relies on an oblivious key-value store to enable private contact discovery [43]. Specifically, users send their encrypted contacts to the Signal server running in a hardware enclave (also called TEE, short for Trusted Execution Environment),

and the enclave securely fetches matching entries from an oblivious key-value store, such that the access patterns do not leak which users’ entries are matched.

Oblivious sorting [2, 4, 24, 31, 38] is known to be a particularly important building block in the design of oblivious algorithms. Prior works [6, 12, 16, 25, 32, 34, 38, 46] showed that we can obliviously realize numerous computation tasks using oblivious sorting, such as histogram, page rank, database joins and group-by queries, list ranking, tree computations with Euler tour, tree contraction, graph algorithms such as breadth-first search, connected components, and minimum spanning tree/forest. More generally, any computational task represented as a streaming Map-Reduce algorithm or in the GraphLab programming model has an efficient oblivious implementation using oblivious sorting [25, 32, 34]. Oblivious sorting is also employed in Oblivious RAM (ORAM) [20] (a generic compiler that compiles any program to an oblivious counterpart), either for initialization, or as a load balancer for parallelization [10, 14].

Although oblivious sorting has been explored in various prior works, the practical state of the art is still unsatisfactory. Bitonic sort [5], due to its simplicity, is a go-to scheme in many implementations [3, 16, 33, 50]; however, its asymptotical work $O(N\log^2 N)$ is suboptimal, where N denotes the input size. On the other hand, the well-known AKS sorting network [2] and subsequent improvements [24, 36] achieve an optimal $O(N\log N)$ work, but their construction suffer from astronomically large constants. To bridge this significant theory-practice gap, some recent works [4, 23, 38, 42] proposed oblivious sort constructions that achieve either $O(N\log N)$ or $O(N\log N \text{polylog log } N)$ work; however, their practical performance are still far worse than bitonic sort.

In this paper, we explore the following important question:

Can we achieve *concretely efficient, enclave-friendly* oblivious sorting with *optimal* $O(N\log N)$ work?

1.1 Our Results and Contributions

Our work takes a step forward in designing practically efficient and asymptotically optimal oblivious sorting algorithms. Moreover, we show that our oblivious sorting algorithm is enclave-friendly from both an algorithmic perspective and with empirical implementation and evaluation. We summarize our contributions below.

Concretely efficient and optimal oblivious sorting. We propose a new oblivious sorting algorithm called *flexway oblivious sort* (or *flexway o-sort* for short) which achieves $(2.23 + o(1))N \log N$ work (measured in terms of the number of compare-and-exchange operations).

Enclave-friendliness. Our flexway o-sort algorithm is optimal in both of the following scenarios:

1. **[EPC \geq data]:** when the input array fits within the enclave’s protected memory region (also called EPC, short for Enclave Page Cache);
2. **[EPC $<$ data]:** when the input array does not fit within the EPC memory.

In the former setting, we want to optimize the *computational overhead* (also called *work*) of the algorithm. In the latter setting, we also need to minimize the number of *page swaps*. Specifically, when the enclave fetches (encrypted) data from insecure memory or disk, it needs the operating system’s help to perform a page swap. The page swap is a heavy-weight operation since it involves context switching as well as encrypting and decrypting the pages being swapped in and out. Existing benchmarking results [46] showed that the cost of a 4 KB page swap can be $66\times$ more expensive than moving a 4 KB page within the EPC memory. The vast majority of oblivious algorithms [4, 20, 21, 28, 32, 44, 48] in the literature focus on optimizing work but not the page swap overhead. Earlier works [46] showed that if we directly employ such an algorithm for hardware enclaves, then page swaps will likely become the dominant overhead when the data does not fit within the EPC memory. Also, it is known that $\frac{N}{B} \log \frac{M}{B}$ number of page swaps is necessary for any (even non-oblivious) sorting algorithm [1, 17].

Although recent enclave technologies such as SGX v2 and TDX has supported larger secure memory [11, 15], external-memory efficiency remains crucial for deploying oblivious sorting for several reasons:

- **Limited RAM Space:** Although Intel’s server-grade processors now support up to 1TB of EPC, many privacy-preserving systems are also limited by the amount of physical memory. For a private blockchain, e.g., high RAM requirements would significantly hinder participation.

- **Resource Allocation:** Sorting are often performed as maintenance tasks. In Signal’s usecase, e.g., if oblivious sorting is applied to scale the database, it should not consume memory needed by the latency-sensitive query service.
- **Enclave Creation Time:** The creation time of an SGX enclave increases linearly with the EPC size; e.g., creating a 64GB EPC enclave takes about 150 seconds in SGX V2 [15]. Therefore, it is desirable to minimize the EPC size for faster enclave creation.
- **Communication in Distributed Clusters:** when sorting is performed on a distributed clusters, page swap overheads translate into communication costs between nodes.

To the best of our knowledge, our flexway o-sort is the **first concretely efficient algorithm that achieves optimality in both dimensions**. Specifically, our flexway o-sort algorithm achieves $(2.23 + o(1))N \log N$ work and $(3 + o(1))\frac{N}{B} \log \frac{M}{B}$ number of page swaps where B denotes the page size, and M denotes the size of the EPC memory.

Open-source implementation. We implemented our algorithm and evaluated their concrete performance. The core algorithm implementation (counting both oblivious sorting algorithms) has 1,600 lines of code. Although our implementation uses Intel SGX, the algorithm design should work for any common hardware enclave architecture. Our implementation has been made open source at <https://zenodo.org/records/14629454>.

Evaluation. We show that our algorithm achieves significant speedup over prior oblivious sorting algorithms for both the [EPC \geq data] and [EPC $<$ data] scenarios. For an array of 12 GB size, our speedup over prior algorithms is depicted in the following table where for the [EPC $<$ data] scenario, we adopt an EPC size of 128 MB which is the same as SGXv1:

Scheme	Our speedup	
	EPC \geq data	EPC $<$ data
Randomized Shellsort	$28.8\times$	$208\times$
Bitonic (non-recursive impl.)	$5.49\times$	$38.4\times$
Bitonic (recursive impl.)	$1.32\times$	$4.10\times$
Multi-way bucket o-sort [38]	$14.3\times$	$12.4\times$

We also implemented various applications that rely on oblivious sorting, including histogram, ORAM initialization, and database join. We measure the end-to-end application performance when using our flexway o-sort, and compare it with (recursive) bitonic sort as a baseline. For these applications, we achieve $1.44 \sim 2.3\times$ speedup when the data fits within the enclave, and $4.9 \sim 5.5\times$ speedup when the data does not fit within the enclave.

Additional results. As a byproduct, we also construct an *oblivious shuffler*, which randomly permutes an input array without leaking the permutation. Earlier works showed that the oblivious shuffler is also a versatile primitive in oblivious algorithms and ORAM schemes [27, 40, 41]. We give more detailed evaluation results on oblivious shuffler in Section 5.4.

1.2 Technical Highlights

Starting point: multi-way bucket o-sort. Ramachandran and Shi [38] constructed multi-way bucket o-sort, which is asymptotically optimal in both work and page swaps, but unfortunately suffers from astronomical constants. Their construction reduces the task of oblivious sorting N elements to an oblivious p -way MergeSplit, which can be viewed as a special sorting algorithm where the keys come from a small domain $\{0, 1, \dots, p-1\}$. More precisely, a p -way MergeSplit accomplishes the following:

- **Input:** p bins each of size Z . Each bin contains *real elements* each marked with a key from $\{0, 1, \dots, p-1\}$ and *fillers*. The frequency of each distinct key in the input is bounded by Z .
- **Output:** Route the real elements to p destination bins depending on their key, and each output bin is padded with fillers to its capacity Z .

Ramachandran and Shi [38] showed that if we can achieve oblivious MergeSplit in $O(n \log n)$ cost where $n = pZ$, $p = \log N$, and $Z \in \text{poly} \log N$, then we can get oblivious sorting optimal in both work and number of page swaps.

Our key idea. We observe that Ramachandran and Shi’s framework still works if we use a slightly smaller $p \in \Theta(\sqrt{\log N})$. Our main contribution is to construct a concretely efficient $\Theta(\sqrt{\log N})$ -way MergeSplit algorithm that achieves $O(m \log m)$ work. To get this, we need some new algorithmic tricks. Notably, we show that if one can construct an oblivious Euler tour algorithm for a graph with p vertices, we can achieve p -way MergeSplit efficiently. To solve the oblivious Euler tour problem efficiently, we applied a *packing* trick. Specifically, as long as $p \in O(\sqrt{\log N})$, we can pack the adjacency matrix of the graph in $O(1)$ memory words. This allows us to obliviously access any entry in the adjacency matrix in $O(1)$ cost under a standard word-RAM model [19], and hence obtain an Euler tour obliviously in linear time. We refer the readers to Section 3.4 for the details of our novel MergeSplit algorithm.

Other constant-factor optimizations. Simply plugging in our new MergeSplit into the multi-way bucket o-sort framework allows us to significantly improve the concrete performance relative to the original multi-way bucket o-sort [38]; however, it is not enough for achieving the tight

constants claimed earlier. To get the promised result with $(2.23 + o(1))N \log N$ work and $(3 + o(1))\frac{N}{B} \log_{\frac{M}{B}} \frac{N}{B}$ page swaps, we introduce various additional optimizations.

Notably, while previous works use a uniform number of ways throughout the butterfly, we use a flexway butterfly construction that allows different number of ways in different layers (hence the name *flexway o-sort*). This approach ensures that the algorithm can adapt to all input sizes with only $(1 + o(1)) \times$ worst-case overhead. While the original butterfly network construction of Ramachandran and Shi [38] requires out-of-place data movements, we adopt an in-place variant of the butterfly network instead. Furthermore, we avoid the matrix transposition operations needed in the original multi-way bucket o-sort [38] through a careful choice of parameters. We describe various additional optimizations in subsequent technical sections as well as the appendices.

2 Preliminaries

2.1 Threat Model

We assume that the server uses secure hardware enclaves, such as Intel SGX, to ensure computational integrity. While we use Intel SGX as a test platform, our algorithmic constructions are compatible with various hardware enclave technologies. In this paper, we assume the hardware enclave itself is secure, as explored in complementary research on designing provably secure trusted hardware [18, 45, 49].

Due to the limited capacity of an enclave’s secure memory, a page swap mechanism is required to manage encrypted pages between external storage and internal memory. Even with sufficient secure memory, a malicious operating system could still revoke access rights to an enclave page, causing a page fault [8]. During page swaps, the OS could observe page-level access patterns and tamper with the pages. Additionally, the OS might monitor fine-grained memory accesses within the enclave through cache-timing attacks [7, 39]. Therefore, our algorithm must be fully oblivious [40] (also referred to as doubly oblivious [33]), ensuring that no secret information is leaked through page-level or EPC memory access patterns.

2.2 Background on Bucket Oblivious Sort

Our flexway o-sort algorithm builds on the bucket o-sort framework proposed by Asharov et al. [4], and a multi-way variant introduced by Ramachandran and Shi [38].

The “oblivious shuffle + non-oblivious sort” framework.

We depict the blueprint of the bucket oblivious sort algorithm [4] in Figure 1. The algorithm first *oblivious shuffle* (or *o-shuffle* for short) the input array, i.e., randomly permuting the input array without leaking the permutation, and then employ a comparison-based, non-oblivious sorting algorithm on

Table 1: **Comparison with sorting algorithms from prior works.** The bounds assume sufficiently large N and adopt the standard tall cache assumption: $B \geq \log^2 N$ and $M \geq B^2$. For the multi-way bucket o-sort [38], C_{AKS} and C_{SPMS} denote the large constants associated with the AKS sorting network [2] and the SPMS sorting algorithm [13]. For all the algorithms listed in the table, the work is dominated by the element-wise exchanges. A page read plus a page write is counted as one page swap.

Algorithm	Work (# Exchanges)	Page Swaps	Notes
Theoretically optimal [1, 17]	$\Omega(N \log N)$	$\Omega(\frac{N}{B} \log_M \frac{N}{B})$	
Prior works			
Bitonic (non-recursive) [5]	$(\frac{1}{4} + o(1))N \log^2 N$	$(\frac{1}{2} + o(1))\frac{N}{B} \log^2 N$	
Bitonic (recursive) [5, 16]	$(\frac{1}{4} + o(1))N \log^2 N$	$(\frac{1}{2} + o(1))\frac{N}{B} \log^2 \frac{N}{M}$	
Randomized Shell Sort [†] [23]	$24N \log N$	$(24 + o(1))N \log \frac{N}{M}$	// non-negl. failure prob.
Bucket o-sort [4]	$(1 + o(1))N \log N \log^2 \log N$	$(4 + o(1))\frac{N}{B} \log \frac{N}{B}$	
Multi-way bucket o-sort [38]	$C_{AKS} \cdot N \log N$	$C_{SPMS} \cdot \frac{N}{B} \log_M \frac{N}{B}$	// asymptotically opt.
Our result			
Flexway o-sort	$(2.23 + o(1))N \log N$	$(3 + o(1))\frac{N}{B} \log_M \frac{N}{B}$	// asymptotically opt.

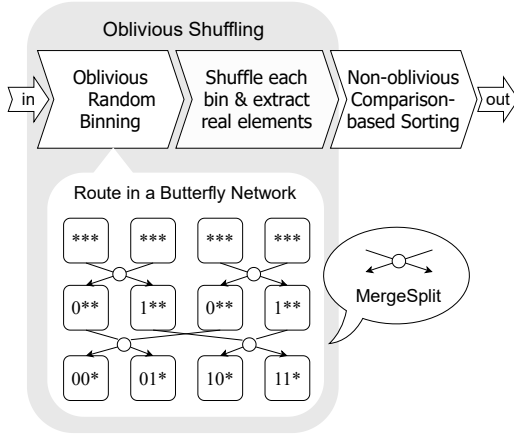


Figure 1: Procedure of Bucket Oblivious Sort.

the shuffled array. Asharov et al. [4] proved the obliviousness of this paradigm.

Oblivious random binning: core primitive in o-shuffle. A core primitive for realizing an o-shuffle is called an *oblivious random binning*, which addresses the following problem: given an input array of size N , we want to obviously place each element randomly into one of $(1 + \epsilon)N/Z$ bins, where $Z = \text{polylog } N$ is the size of each bin and $\epsilon > 0$ is a constant called the slack factor. For simplicity, prior works [4, 38] set $\epsilon = 1$. Asharov et al. [4] showed that we can realize an o-shuffle in the following way.

O-shuffle:


1. Run random binning on the input array.
2. Obviously shuffle the elements within each bin.
3. Extract all real elements and drop the fillers.

Note that Step 3 leaks how many real elements fall within each bin. Asharov et al. [4] showed that this leakage is safe as the number of real elements in each bin can be simulated.

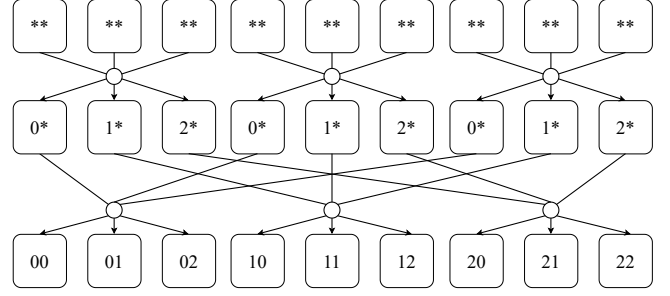
Realizing oblivious random binning. We now describe the oblivious random binning algorithm by Asharov et al. [4]. Their construction uses a butterfly network with $(\log N \text{Bin}) + 1$ layers, where each layer has $N \text{Bin} = (1 + \epsilon)N/Z$ bins of capacity Z . The algorithm works as follows:

Oblivious random binning:

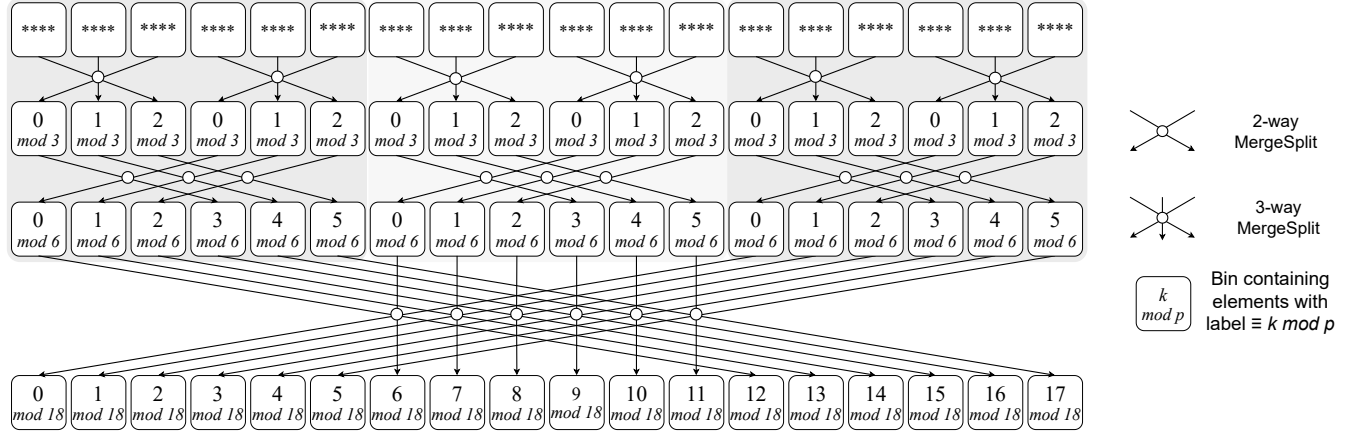
- Initially, a random label of $\log N \text{Bin}$ bits is assigned to each input element. The label denotes which final-layer bin the element goes to.
- Place exactly $Z/(1 + \epsilon)$ input elements in each bin at layer 0, and pad the bins to a full capacity with fillers.
- All elements are routed to their destination bins in the final layer along the butterfly network, through a sequence of MergeSplit operations as defined below. When routing from layer ℓ to layer $\ell + 1$, each element uses the ℓ -th bit of its label to decide the direction to go for the MergeSplit.



(a) 2-way butterfly network for NBin = 8. Uses binary labels.



(b) 3-way butterfly network for NBin = 9. Uses ternary labels.



(c) A $3 \times 2 \times 3$ flexway butterfly network for NBin = 18. The MergeSplit function routes elements based on the modulo of their labels.

Figure 2: Comparison of (a) the two-way butterfly network from [4], (b) the multi-way variant from [38], and (c) the flex-way variant proposed in this work. The flex-way butterfly network introduces fewer restrictions on the number of bins and is restructured to enable in-place MergeSplit (see section 4.2 for details).

In expectation, the number of real elements that land in each bin is exactly $Z/(1 + \epsilon)$. Earlier work [4] proved that if $Z \in \text{polylog} N$, the probability that any bin overflows is negligible, so a simulator assuming no overflow produces access patterns statistically indistinguishable from real-world execution.

Core subroutine MergeSplit. In the butterfly network, a pair of bins at level ℓ route to a pair of bins at level $\ell + 1$ — this is accomplished through a MergeSplit operation.

The MergeSplit operation takes two input arrays (i.e., bins) each of size Z , where each array contains real elements tagged with a key 0 or 1, as well as some fillers. MergeSplit then routes each real element to either the left or the right output bin based on its key, and each output bin is padded with fillers to its capacity Z . For obliviousness, the algorithm’s access patterns should not leak where each element is going.

A multi-way variant. Ramachandran and Shi [38] introduced a p -way variant of the bucket o-sort algorithm, reducing the depth of the butterfly network from $O(\log(N/Z))$ to

$O(\log_p(N/Z))$. Each bin is connected to p bins in the next layer, and the MergeSplit operation is extended to handle p input and p output bins. Specifically, the MergeSplit routes each element at level ℓ to one of the p output bins at level $\ell + 1$ based on the ℓ -th digit of its label in base- p . By setting $p = \log N$, Ramachandran and Shi [38] gave a theoretical construction of oblivious sorting with optimal work. Figure 2b shows an example of p -way butterfly network where $p = 3$.

Key challenge: an efficient oblivious MergeSplit? The reason why the bucket o-sort [4] and the subsequent multiway variant [38] are inefficient is because they lack an efficient MergeSplit algorithm. Specifically, bucket o-sort uses bitonic sort to realize the MergeSplit. As bitonic sort requires $O(n \log^2 n)$ work on an input size n , the resulting oblivious sorting suffer from $O(N \log N (\log \log N)^2)$ work. While we can replace bitonic sort with a more efficient compaction algorithm with $O(n \log n)$ work [22, 40], the resulting oblivious sorting still requires $O(N \log N \log \log N)$ work.

By contrast, the multiway bucket o-sort algorithm [38] uses a $\log N$ -way butterfly network, and applies AKS sorting [2]

to realize the $\log N$ -way MergeSplit with $O(n \log n)$ work. This approach achieves asymptotic optimality for oblivious sorting (i.e. $O(N \log N)$ work), but suffers from astronomical constants due to the expander graphs in AKS.

It is also possible to use linear-work oblivious compaction [37] to realize the MergeSplit in the original (2-way) bucket o-sort [4]. This also achieves optimality in work but again known linear-work oblivious compaction [37] relies on expander graphs and suffers from enormous constants.

3 A New Multi-Way MergeSplit

3.1 Overview

Overview of our flexway o-sort. Our new flexway o-sort algorithm follows the same overall framework as multi-way bucket o-sort [4, 38]. However, unlike earlier works [38] that use either a 2-way or $\log N$ -way butterfly network, we set the number of ways to be $p \in \Theta(\sqrt{\log N})$. Both choices $p = \log N$ and $p \in \Theta(\sqrt{\log N})$ results in a butterfly network of $\Theta(\log_p N) = \Theta(\log N / \log \log N)$ layers.

The reason why we choose $p \in \Theta(\sqrt{\log N})$ (as opposed to $\Theta(\log N)$ like in earlier work [38]) is because it allows us to pack the adjacency matrix of a graph with p vertices into $O(1)$ memory words. This *packing trick* is later used to implement a more efficient p -way MergeSplit algorithm.

Overview of our p -way MergeSplit. As mentioned, Ramachandran and Shi’s multi-way bucket o-sort [38] is inefficient because they lacked an efficient multi-way MergeSplit. One of our main contributions is to devise a practical oblivious p -way MergeSplit algorithm that achieves $O(n \log n)$ work with input size $n = pZ$, where $p \in \Theta(\sqrt{\log N})$ and $Z \in 2^{O(\sqrt{\log N})}$. As shown in Figure 3, our novel p -way MergeSplit algorithm works as follows. Suppose we have an input array consisting of real elements and fillers, and each real element is marked with a key from $\{0, 1, \dots, p-1\}$, denoting which way it wants to go. Now, perform the following steps:

1. *Preprocess:* We first preprocess the input array and tag each filler with a key from $\{0, 1, \dots, p-1\}$, such that at the end, there is an equal number of elements with each key $k \in \{0, 1, \dots, p-1\}$.
2. *Interleave:* We run a recursive algorithm called Interleave to rearrange the input array such that in the output array, the elements have interleaving keys $0, 1, 2, \dots, p-1, 0, 1, 2, \dots, p-1, \dots$, and so on.
3. *Transpose:* Finally, rearrange elements into output bins by transposing the result of the previous step, so that the keys have the form $\underbrace{0, 0, \dots, 0}_Z, \underbrace{1, 1, \dots, 1}_Z, \dots, \underbrace{p, \dots, p}_Z$.

In the remainder of the section, we describe how each step is performed. Specifically, the core of our new multi-way MergeSplit is the Interleave subroutine, a recursive algorithm that calls another building block called Balance. The Balance subroutine is also the technical highlight of our algorithm, since it expresses the problem of rearranging an input array in a specific way as an Euler tour problem for a small graph with $p \in O(\sqrt{\log N})$ vertices. By leveraging the fact that the adjacency matrix of this graph can be packed in $O(1)$ memory words, we can make the algorithm highly efficient.

3.2 Balance

As mentioned, our p -way MergeSplit algorithm relies on a central building block called Balance.

Syntax. The Balance algorithm takes as input an array A containing n elements, each marked with a key from $\{0, 1, \dots, p-1\}$. It is promised that each distinct key appears an even number of times. The goal of Balance is to rearrange the array A such that

- each key appears the same number of times in the left and right half of the array; and
- the last element of A is the same as the input array.

The requirement of preserving the last element of the array is used in later building blocks when the array length is not a power of 2 — see the Permute algorithm in Section A.

Constraints. $p \in O(\sqrt{\log N})$.

Intuition: translate Balance into an Euler-tour on the exchange graph. First, we pair up elements in the left and right halves of the array A , specifically $A[i]$ and $A[n/2 + i]$ for each $i \in \{0, \dots, n/2 - 1\}$. Our goal is to exchange some of these pairs to make the array balanced.

We can equivalently think of this as a graph problem. We create a multi-graph G (called the exchange graph) with p vertices, one for each key in $\{0, 1, \dots, p-1\}$. If an element with the key u is paired with an element with the key v , we draw an edge between the vertices u and v . Observe that the resulting multi-graph may have parallel edges and self-loops. Further, since each key appears an even number of times in the input array, every vertex in G has an even number of incident edges (each self-loop counts twice).

Now, our goal is to orient the edges (i.e., assign a direction to each edge), such that every vertex has the same in-degree and out-degree. In particular, if there is a directed edge (u, v) , it means that during one encounter with the pair u and v , we should rearrange them such that u appears on the left and v appears on the right. Clearly, if every vertex has the same in-degree and out-degree, it means that the corresponding key appears the same number of times in the left and right halves.

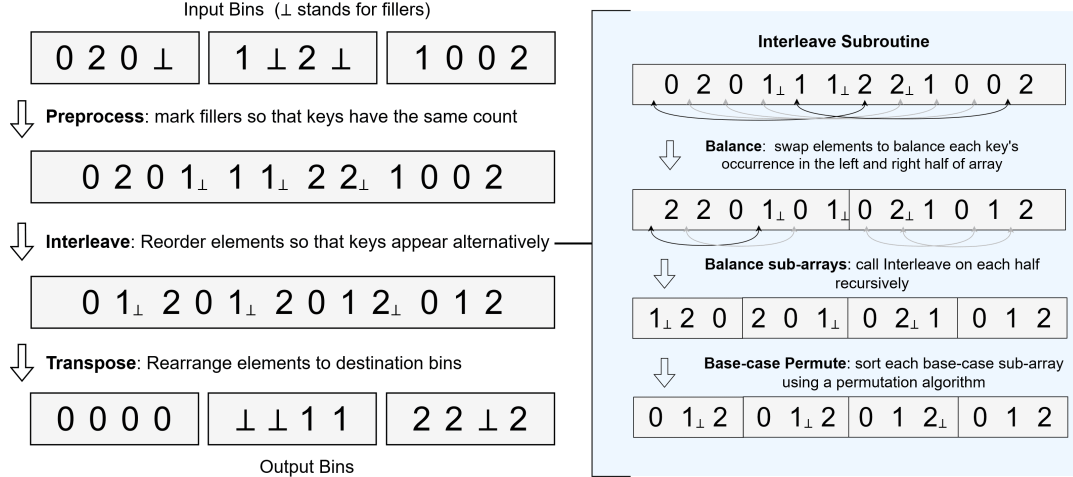


Figure 3: **Left:** An illustration of our MergeSplit construction. Input elements are marked with keys 0, 1, 2, and we use \perp to denote filler elements. **Right:** The Interleave subroutine, which in turn relies on the Balance subroutine as a building block.

To achieve this, we find an Euler tour in the exchange graph G and then orient the edges accordingly. To make the algorithm oblivious and efficient, we preprocess G to compress it into a simple graph. In particular, the preprocessing removes parallel edges and self-loops as they occur (see lines 4-6 of Algorithm 1). More specifically, if there are r edges between two vertices u and v , we do the following preprocessing:

- *Case 1: r is even.* In this case, we can assign $r/2$ of these edges one direction and the remaining $r/2$ the opposite direction. We can prune all these r edges, i.e. there is no edge left between u and v .
- *Case 2: r is odd.* In this case, we can assign $(r-1)/2$ of these edges one direction, and $(r-1)/2$ of them the opposite direction. We can prune these $r-1$ edges, such that there is only one edge left between u and v whose direction remains to be assigned.

With the above pre-processing, the pruned graph G always has 0 or 1 edge remaining between every pair of vertices. Therefore, the total number of edges is at most $p^2 \in O(\log N)$, and the adjacency matrix hence fits in $O(1)$ memory words. Moreover, every time we prune a pair of parallel edges or a self-loop, we reduce the degree of the endpoint(s) by 2, so all vertices still have even degree after the preprocessing. At this moment, we find an Euler tour to orient the remaining edges (see lines 7-12 of Algorithm 1).

Oblivious Euler-tour algorithm for small graphs. Without the obliviousness requirement, there is a standard Euler-tour algorithm with $O(p^2)$ overhead where p^2 is the maximum number of edges for a graph with p vertices. Unfortunately, the standard Euler-tour algorithm is not oblivious. Our key insight is that we can have an oblivious version of the standard Euler-tour algorithm, as long as the adjacency matrix of the

graph can fit in $O(1)$ memory words, that is, the number of vertices $p \in O(\sqrt{\log N})$. In this way, we can obviously access each entry of the adjacency matrix by invoking $O(1)$ word-level operations supported by the RAM. Further, to ensure obliviousness, we also need to make sure that the Euler-tour algorithm does not abort prematurely which may leak the number of edges in the graph. In our oblivious implementation, we make sure that the loop always iterates for a fixed number of iterations (see line 8 of Algorithm 1).

Detailed algorithm description. We give a full description of our Balance algorithm in Algorithm 1. Basically,

- *Construct pre-processed graph.* Lines 4-6 construct the simple graph G where parallel edges and self-loops have been pruned. This part requires $O(n)$ numerical computation and no exchange.
- *Euler tour.* Lines 7-12 finds an Euler tour using depth-first search (DFS), and the orientation of the edges are stored in another directed simple graph D whose adjacency matrix can also be packed into $O(1)$ words. Starting from vertex 0, we traverse G through unvisited edges until reaching a dead end. The unvisited edge can be indexed using the least significant bit (LSB) operation. As each vertex has an even degree, we are guaranteed to return to the starting vertex. We then move to the next vertex and repeat the process until all the edges are visited. To make the search oblivious, we always pad the number of iterations to the worst-case, which equals the number of vertices plus the number of edges. In other words, we perform fake operations after all vertices have been visited. This part requires $O(p + \min(p^2, n)) = O(\min(\log N, n))$ numerical computation as long as $p \in O(\sqrt{\log N})$, and no exchange.

Algorithm 1 Balance(A, p)

Input: Input array A contains n elements each marked with a key in $\{0, 1, \dots, p-1\}$. Let k_i denote the key of the i -th element. Each key occurs even times.

Output: A is rearranged such that each key appears the same number of times in the left and the right half. Also, the last element of A should not change.

// All if conditionals use fake accesses to ensure that the access patterns are identical for both branches.

```
1:  $n \leftarrow |A|, m \leftarrow n/2$ 
2: if  $n = 2$  then return
3: Construct a simple graph  $G$  and directed graph  $D$ , both
   with  $p$  vertices numbered from 0 to  $p-1$  and no edge
   initially. Represent each graph with an adjacency matrix
   and pack it in a word.
4: for  $i \leftarrow 0$  to  $m-1$  do
5:   if  $(k_i, k_{i+m}) \in G$  then prune  $(k_i, k_{i+m})$  of  $G$ 
6:   else if  $k_i \neq k_{i+m}$  then add  $(k_i, k_{i+m})$  to  $G$ .
7: Start the walk at vertex  $t \leftarrow 0$ .
8: for  $p + \min(m, \frac{1}{2}p(p-1))$  iterations do
9:   if  $\exists v$  such that  $(t, v) \in G$  then
10:     Add  $(t, v)$  to  $D$  and delete it from  $G$ .
11:      $t \leftarrow v$ 
12:   else  $t \leftarrow \min(t+1, p-1)$ 
13: For all  $0 \leq u < v < p$ , add directed edge  $(u, v)$  to  $D$  if
   there is no edge between  $u$  and  $v$  in  $D$ .
14: if  $(k_{m-1}, k_{n-1}) \notin D$  then reverse all edges in  $D$ .
15: Reverse the edge between  $k_{m-1}$  and  $k_{n-1}$  in  $D$ .
16: for  $i \leftarrow 0$  to  $m-2$  do
17:   Exchange  $A[i]$  and  $A[i+m]$  if  $(k_i, k_{i+m}) \notin D$ .
18:   Reverse the edge between  $k_i$  and  $k_{i+m}$  in  $D$ .
```

- *Conditional exchange of elements.* Lines 13-18 exchange the element pairs based on the orientation of their corresponding edge in D . Some pairs may not have any corresponding edge in D because they appear even number of times and all the edges got pruned. Therefore, we add an edge in D between any pair that is not directly connected (the direction can be arbitrary, see line 13). Every time a conditional exchange occurs between two elements with keys u and v , we reverse the edge in D between vertices u and v in D (see line 18), so that next time u and v will be arranged in the opposite order. This effectively guarantees that the pruned edges between u and v are assigned to either direction the same number of times. This part takes $O(n)$ numerical computation and $n/2 - 1$ exchanges.

As an optimization, we can avoid exchanging the last ele-

ment pair by conditionally negate the adjacency matrix and use a reversed Euler tour (see line 14-15 of Algorithm 1).

Lemma 3.1 (Computational overhead of Balance). *Algorithm 1 incurs $O(n)$ numerical computation and $n/2 - 1$ exchanges.*

Lemma 3.2 (Obliviousness of Algorithm 1). *The memory access patterns of Algorithm 1 are deterministic and depend only on the length of the input array and the parameter p but not the contents of the array.*

Proof. As mentioned, the adjacency matrices of G and D can each be packed into $O(1)$ words. In this way, accessing an entry in the adjacency matrices requires $O(1)$ word-level operations. Further, all if conditionals use fake accesses to make the access patterns for both branches the same. Hence, lines 4-6 where we construct the pre-processed graph G have fixed access patterns. Lines 7-12 where we find the Euler tour iterate for a fixed number of times and the access patterns within each iteration of the loop are fixed. Similarly, Lines 13-18 where we perform the actual conditional exchanges also enjoy fixed access patterns. \square

3.3 Interleave

Syntax. Interleave receives an input array containing $n = pZ$ elements marked with keys in $\{0, 1, \dots, p-1\}$, each key appearing exactly Z times. The output is a rearranged array such that the i -th element has key $i \bmod p$. For $p = 3$ and $Z = 4$ as an example, the output elements should have the key sequence 0, 1, 2, 0, 1, 2, 0, 1, 2, 0, 1, 2.

Constraints. Z is a power of two, $p \in O(\sqrt{\log N})$.

Intuition. The basic idea behind Interleave is to recursively balance the number of each key on the left and right half of the array. At the base case, each key appears exactly once, and we use a permutation network to arrange them in order.

Detailed algorithm. Algorithm 2 shows the Interleave procedure. At line 3 we call Balance so that each key appears $Z/2$ times on the left and $Z/2$ times on the right. As Z is a power of two, we can call Interleave on both halves recursively.

For the base case, there are p elements with distinct keys in $\{0, 1, \dots, p-1\}$. In Section A, we define the Permute algorithm to sort these elements obliviously in $O(p \log p)$ time for $p \in O(\sqrt{\log N})$.

Lemma 3.3 (Computation cost of Interleave). *Algorithm 2 incurs $O(n \log n)$ numerical computation and no more than $\frac{1}{2}n(\log n + \log p)$ exchanges.*






Figure 4: An example of the Interleave algorithm. The algorithm recursively balances the input and runs permutation at the base case. In 4a, the dotted lines mean the two elements will not be swapped. In 4b and 4c, solid and dashed edges jointly denote the exchange graph G before pre-processing, and the solid edges denote the graph G after pre-processing.

Lemma 3.4 (Obliviousness of Algorithm 2). *The memory access patterns of Algorithm 2 are deterministic and depend only on the length of the input array and the parameter p but not the contents of the array.*

We defer the proof of Lemma 3.4 to Section B.2.

3.4 Multi-way MergeSplit

Syntax. A p -way MergeSplit takes in p input bins each containing Z elements. Each element is either a real or a filler element. Every real element has a key in $0, 1, \dots, p-1$, and a payload string. The goal of the MergeSplit function is to redistribute the real elements so that all elements with key j appear in the j -th output bin. All output bins are padded with fillers to a maximum capacity of Z . If any bin overflows (i.e., if any key appears more than Z times in the input), the algorithm simply aborts. For the special case $p = 2$, this is exactly the MergeSplit primitive used in the original bucket o-sort by Asharov et al. [4].

Algorithm 2 Interleave(A, p)

Input: The input array A contains $n = pZ$ elements each marked with a key from $\{0, 1, \dots, p-1\}$. Each key appears exactly Z times and Z is a power of two. We require $p \in O(\sqrt{\log N})$.

Output: A is rearranged so that the i -th element has key $i \bmod p$.

```

1:  $n \leftarrow |A|$ 
2: if  $n = p$  then: PERMUTE( $A$ ); return
3: BALANCE( $A, p$ )
4: INTERLEAVE( $A[0 : \frac{n}{2} - 1], p$ )
5: INTERLEAVE( $A[\frac{n}{2} : n - 1], p$ )

```

Constraints. $p \in O(\sqrt{\log N})$ and $Z \in 2^{O(\sqrt{\log N})}$.¹ We require the bin size Z to be a power of 2, but do *not* require p to be a power of 2. This weakened precondition provides flexibility for constructing the butterfly network in Section 4.

Obliviousness requirement. For obliviousness, we require that for any input where each key does not appear more than Z times, the memory access patterns of the algorithm are deterministic and the same.

Intuition. Our new MergeSplit algorithm has a preprocessing step to obviously mark every filler element also with a key so that each distinct key appears exactly Z times. Next, we call Interleave to rearrange elements into chunks of size p where each chunk contains elements with keys ordered from 0 to $p-1$. Finally, we create the bins by extracting the corresponding elements from each chunk.

Detailed algorithm. Algorithm 3 shows the MergeSplit procedure. The preprocessing involves two linear passes over the keys. The first pass counts the occurrences of each distinct key. If any key appears more than Z times, we detect a bin overflow and abort. The second pass marks the fillers and ensures that each key appears Z times. We make both passes oblivious and achieve linear runtime using a packing trick similar to that in algorithm 1. Note that representing all p counters requires $p \cdot \log Z \in O(\log N)$ bits. Using bitwise CPU instructions such as shifting and bitwise AND, we can pack all the counters into $O(1)$ memory words and update a counter obliviously in constant time.

Lemma 3.5 (Computation cost of MergeSplit). *Algorithm 3 incurs $O(pZ \log(pZ))$ numerical computation and no more than $pZ(\frac{1}{2} \log Z + \log p + 1)$ exchanges.*

¹Our flexway butterfly sorting network later chooses $Z = \log^c N$ for a constant $c > 1$ to get both negligible in N failure probability and optimal computational overhead.

Algorithm 3 p -way MergeSplit for $p \in O(\sqrt{\log N})$

Input: $A := A_0 || A_1 || \dots || A_{p-1}$ where $p \in O(\sqrt{\log N})$. For $j \in \{0, \dots, p-1\}$, each bin A_j has size Z where Z is a power of 2 and $Z \in 2^{O(\sqrt{\log N})}$. Each A_j contains real and filler elements; and each real element has a key from $\{0, \dots, p-1\}$.

Output: p bins denoted $A' = A'_0 || A'_1 || \dots || A'_{p-1}$. We want to route all real elements in the input with key k to A'_k , padded with fillers to a size of Z . Output Abort if any bin overflows.

- 1: **for** $k \leftarrow 0$ **to** $p-1$ **do** // Preprocess
- 2: $C_k \leftarrow$ Count of real elements marked with key k .
- 3: Abort if any $C_k > Z$.
- 4: Mark the next $Z - C_k$ fillers in A with key k .
- 5: Interleave(A, p)
- 6: **for** $k \leftarrow 0$ **to** $p-1$ **do** // Transpose
- 7: $A'_k \leftarrow [A[k], A[p+k], \dots, A[p(Z-1)+k]]$

* The preprocessing can be implemented obliviously with $O(1)$ linear scans by packing the counts in $O(1)$ words.

Lemma 3.6 (Obliviousness of Algorithm 3). *As long as the input promises that each key appears no more than Z times, then the memory access patterns of Algorithm 3 are deterministic and depend only on the length of the input array and the parameter p but not the contents of the array.*

We defer the proof of Lemma 3.6 to Section B.2.

4 Flexway O-Sort

4.1 Basic Algorithm and External-Memory Efficiency

As mentioned, our basic algorithm is the following: we use the same framework of multi-way bucket o-sort [38] with the following changes: 1) we set the number of ways $p \in \Theta(\sqrt{\log N})$, and 2) we adopt a new p -way MergeSplit algorithm described in Section 3.4.

This gives us an oblivious random binning algorithm, from which we can easily construct an oblivious random permutation. To get oblivious sorting, we need to run a non-oblivious comparison-based sort after the oblivious random permutation. The work of Ramachandran and Shi [38] adopts the SPMS algorithm which is theoretically optimal in terms of both work and number of page swaps, but concretely inefficient. In our work, we use the external merge-sort algorithm [30] to realize the non-oblivious sort.

External-memory efficiency. The above basic algorithm achieves $O(N \log N)$ work. At this moment, we explain why

this algorithm can achieve an optimal number of page swaps as observed by Ramachandran and Shi [38]. The reason is that the butterfly network structure enjoys good *locality* — see Figure 5 (left). Instead of working on the MergeSplit operations layer by layer, we can work on them k layers at a time: each time we fetch a batch of bins into the enclave and perform k layers of MergeSplit operations among those bins. Then, we perform a *matrix transposition* operation to bring together bins that will be grouped in the next k layer's MergeSplit operations. In Figure 5 (left), we show a 2-way butterfly network (rather than multi-way) for clarity, and we use both $k=1$ and $k=2$ due to imperfect rounding. Standard analysis shows that if the batch size is $\Theta(M)$, the above approach achieves $O(N/B \cdot \log_{M/B}(N/B))$ number of page swaps [38], which is optimal even for non-oblivious sort [1].

4.2 Other Optimizations

So far, we have significantly improved the constants relative to Ramachandran and Shi [38] while maintaining the asymptotical optimality. However, to achieve the constants claimed in Table 1, we need several additional tricks.

New rounding technique and flexible ways. Prior works [4, 38] insist both the bin size Z and the number of ways each MergeSplit performs to be powers of two. Consequently, they have to round the number of elements in the butterfly network to the next power of two, leading to a potential $2\times$ overhead. We instead do not impose the power-of-two constraint on the number of ways, and moreover, allow a non-uniform number of ways at different levels of the butterfly network. More concretely, we have the following problem. Recall that initially, the number of bins per level is $\text{NBin} := (1 + \epsilon)N/Z$. Our goal is to round NBin up to some integer that can be expressed as $\text{NBin}^* := p_1 \times p_2 \times \dots \times p_L$, and moreover:

1. For every $\ell \in [L]$, $\lfloor \sqrt{\log N} \rfloor / 2 \leq p_\ell \leq \lfloor \sqrt{\log N} \rfloor$;
2. $\text{NBin}^* = (1 + o(1))\text{NBin}$.

The first condition ensures that the number of ways is not too small and hence the total work is $O(N \log N)$. The second condition makes sure that the rounding introduces only $(1 + o(1))\times$ overhead. An example of such a flexway butterfly network is shown in Figure 2c. In Section B.1, we show how to find a solution for any NBin .

Our flexway butterfly also requires a corresponding modification in our external-memory implementation. Specifically, the number of layers k we fetch in a batch will also be non-uniform at different depths in the butterfly. At any point of time, we always pick largest k such that the batch size can maximally utilize the enclave's EPC memory (of size B).

In-place butterfly and saving matrix transposition. As shown in Figure 5 (left), the original multi-way bucket o-sort [38] suffers from several concrete inefficiencies. First,

MergeSplit operations are performed *not in-place*, i.e., the outputs cannot be written to the input bins. This requires extra working buffer in the enclave’s EPC memory, and thus reduces the effective M parameter. Second, in between every k layers of MergeSplit operations, the original multi-way bucket o-sort relies on a *matrix transposition* algorithm to rearrange bins, such that bins that are grouped in the next k layers of MergeSplit become close together. This matrix transposition increases the concrete overhead.

To address these overheads, our new algorithm shown in Figure 5 (right) introduces the following optimizations:

1. *Avoid matrix transposition.* We set the bin size to be at least a page size, that is, 4 KB. In this way, we can skip the matrix transpose and directly fetch the relevant bins from discontinuous memory locations into the enclave.
2. *In-place operations.* By indexing the buckets in a reverse lexicographical order, our butterfly network supports in-place reads and writes both within the enclave and for accesses to external memory.

First, within the enclave, whenever a MergeSplit operates on a set of bins, we write the outputs to the same set of bins. Second, when the enclave fetches a set of bins from the external memory and performs a batch of MergeSplit operations, the output bins are written back to the same locations in external memory.

The in-place optimization improves the memory utilization by $2\times$ and hence reduces the page swap overhead.

Solver for optimal concrete parameters. Given the enclave EPC memory size M , the page size B , the input size N , and the desired security parameter, we need to choose several internal parameters, including the bucket size, number of buckets per layer, batch size, and the number of ways at each level of the butterfly network. We implemented a solver to search for the parameters that minimizes the overhead of rounding, and maximally balances the computation and the page swap overhead.

Table 2 lists the optimal butterfly network parameters determined by our solver for various input sizes. It also lists the concrete constants for element-wise exchanges and page swaps in our o-sort algorithm. Compared with the theoretically calculated constants of Table 1, here, the constants for element-wise exchanges are larger, but the constants for page swaps are smaller — this is by intention since the solver attempts to balance these two sources of overheads and minimize the total running time under our testing setup.

5 Experimental Results

5.1 Experimental Setup

We evaluated the sorting algorithms on an Intel Xeon Platinum 8352S processor with 2.2 GHz base frequency and DDR4 RAM.

We evaluated both the $[EPC > \text{data}]$ and the $[EPC < \text{data}]$ settings. In the latter setting, we limit the EPC size to 128 MB which matches SGX v1. In both settings, we varied the number of elements and the size of the payload strings.

Baselines. We compare with the following baselines:

- *Prior oblivious sort algorithms.* We compare with bitonic sort [5], randomized shell sort [23], and multi-way bucket o-sort [38]. For the multi-way bucket o-sort, we did not implement the version in the original paper [38] due to its astronomical constants. Instead, we use the “oblivious bin placement” algorithm described in earlier work [10] for multi-way MergeSplit, and external-memory mergeSort for the comparison-based sort.
- *Non-oblivious baseline.* For the $[EPC > \text{data}]$ setting, we use `std::sort` as the non-oblivious baseline, and for the $[EPC < \text{data}]$ setting, we use external-memory MergeSort as the non-oblivious baseline.

5.2 Results on Oblivious Sorting

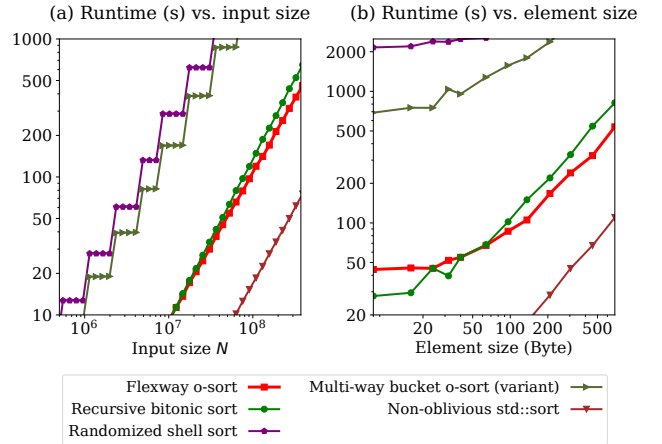


Figure 6: Comparing our sorting algorithm with prior works when $EPC \geq \text{data}$. Fig (a) shows the runtime in relation to the input size, where each element consists of an 8-byte key and a 120-byte payload. Fig (b) fixes an input size of 100 million and varies the size of each element.

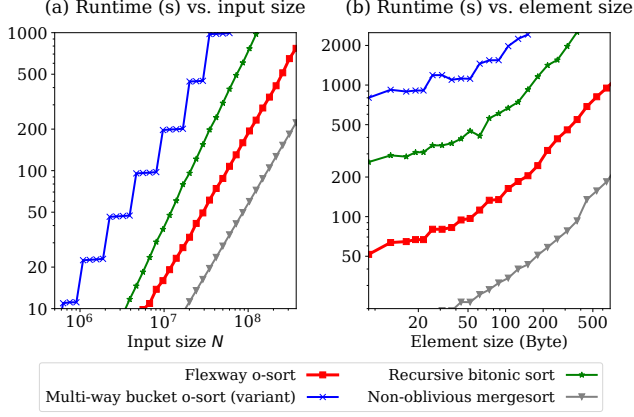


Figure 7: Comparing our sorting algorithm with prior works (128 MB EPC). Other parameters are the same as Figure 6.

The “EPC > data” setting. For an input of 100 million 128-byte elements, our flexway o-sort is $28.8\times$ faster than randomized shell sort [23], $14.3\times$ faster than multi-way bucket o-sort [38], and 32% faster than the recursive bitonic sort. As the input size expands, our speedup over bitonic sort increases due to our asymptotic enhancement. For example, with 1 billion elements each 200-bytes-wide, our flexway o-sort becomes 88% faster than the recursive bitonic sort.

The “EPC < data” setting. When sorting 100 million 128-byte elements using a 128 MB EPC, our flexway o-sort is at least $12.4\times$ faster than the multi-way bucket o-sort algorithm [38], $38\times$ faster than a non-recursive implementation of bitonic sort, and $4.1\times$ faster than a recursive version of bitonic. At a larger data size of 1 billion with 200-bytes-wide elements, our flexway o-sort becomes $7.2\times$ faster than the recursive bitonic sort.

Performance of MergeSplit. Figure 8 shows a micro benchmark of the multi-way MergeSplit algorithm. Our new algorithm is approximately $18\times$ faster than the naïve approach using two bitonic sorts [5, 10], and $2\times$ faster than a multi-way instantiation from OrCompact [35, 40].

5.3 Applications

Table 3 compares the runtime of three applications implemented with our algorithm and the recursive bitonic sort:

- **Histogram:** obviously count frequency of 256-byte URLs.
- **ORAM initialization:** obviously initializing an ORAM tree. We implement the initialization algorithm described in EnigMap [46], assuming each data entry is 128 bytes.
- **DB join:** left join two tables based on an 8-byte ID. Apart from the ID, each row is 256 bytes long for both tables.




Figure 8: Throughput of our multi-way MergeSplit compared with prior methods. Each bin contains 4096 elements of 136 bytes, including the 8-byte label.

Table 3: Benchmark Results for Different Applications.

Application	Input size	Runtime (s) (128 MB EPC)		Runtime (s) (EPC \geq data)	
		Ours	Bitonic	Ours	Bitonic
Histogram	2^{23}	45.66	187.2	35.14	71.80
	2^{26}	447.5	2448	328.6	760.6
ORAM init	2^{23}	153.0	529.1	123.7	139.8
	2^{26}	1483	7400	1086	1568
DB join	2^{23}	247.2	715.1	132.1	172.7
	2^{26}	2038	9954	1206	1985

In the histogram application, with an input size of 2^{26} elements, our algorithm is $5.5\times$ faster than the baseline given a 128 MB EPC, and $2.3\times$ faster when the EPC is unlimited. Notably, the speedup surpasses the results in Figure 6 and Figure 7. This is because histogram requires costly oblivious comparisons between long URLs, and our algorithm incurs significantly fewer comparisons than bitonic sort.

For ORAM initialization, our algorithm achieves a $5.0\times$ speedup over the baseline with a limited 128 MB EPC, and a 44% acceleration when the EPC is unlimited.

For database Join, when both tables have 2^{26} rows, our algorithm is $4.9\times$ faster than the baseline with a 128 MB EPC, and 65% faster when the EPC is unlimited.

5.4 Results on Oblivious Shuffling

Since our flexway o-sort first obviously shuffles the array and then applies a non-oblivious comparison-based sort, we also obtain an oblivious shuffle as a by-product.

For shuffling 100 million elements of 128 bytes with no EPC limit, our flexway o-shuffle is $72.5\times$ faster than the Waks-on/Waks-off o-shuffle [41] and $2.3\times$ faster than OrShuffle [40]. When running in SGX with a 128 MB EPC, our flexway o-shuffle algorithm is $5.5\times$ faster than OrShuffle [40] and $16\times$ faster than the multi-way bucket o-shuffle [38].




Figure 9: Comparing our shuffling algorithm with prior works when EPC size \geq data size. Parameters are the same as Fig. 6.

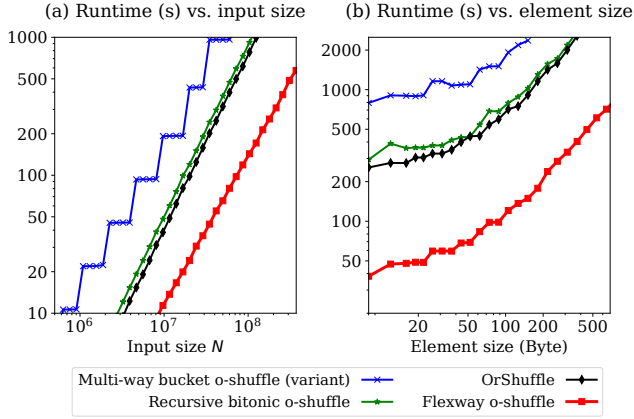


Figure 10: Comparing our shuffling algorithm with prior works (128 MB EPC). Parameters are the same as Fig. 7.

6 Concurrent Work

Concurrent work [35] implements oblivious sorting and shuffling in a distributed model using the same algorithmic framework as ours. Earlier work [9] observed that butterfly networks are well-suited not only for external-memory algorithms but also for distributed settings.

In comparison, our new MergeSplit building block is both asymptotically and concretely faster than theirs and can serve as a drop-in replacement in their framework. As shown in Figure 8, our MergeSplit algorithm is approximately $2\times$ faster than OrCompact [40], the algorithm employed in [35]. Since MergeSplit dominates the computation for large inputs, this improvement is expected to reduce the computational overhead in [35] by nearly 50%.

7 Conclusion

In this paper, we introduced flexway o-sort, a new oblivious sorting algorithm that is asymptotically optimal, concretely efficient, and well-suited for hardware enclaves. Future work will focus on parallelizing the algorithm to further improve performance and adapting the algorithm to circuits for protocols such as multi-party computation and fully-homomorphic encryption.

8 Ethics Considerations

This work introduces an algorithm designed to defend against existing attacks on secure enclave systems, aligning with privacy protection laws and contributing to the public good. All experiments were conducted on local test machines without involving any private or sensitive information. We do not anticipate any harm or adverse consequences resulting from the publication of this research.

9 Open Science

To promote transparency and facilitate further research, we have open-sourced our algorithm’s implementation, including the evaluation codes, at <https://zenodo.org/records/14629454>. We welcome feedback and contributions to improve the implementation and further the research in oblivious sorting and other oblivious algorithms.

References

- [1] Alok Aggarwal and S. Vitter, Jeffrey. The input/output complexity of sorting and related problems. *Commun. ACM*, 31(9):1116–1127, sep 1988.
- [2] M. Ajtai, J. Komlós, and E. Szemerédi. An $O(n \log n)$ sorting network. In *STOC*, 1983.
- [3] A. K. M. Mubashwir Alam, Sagar Sharma, and Keke Chen. Sgx-mr: Regulating dataflows for protecting access patterns of data-intensive sgx applications. *Proceedings on Privacy Enhancing Technologies*, 2021:5 – 20, 2020.
- [4] Gilad Asharov, T-H. Hubert Chan, Kartik Nayak, Rafael Pass, Ling Ren, and Elaine Shi. Bucket oblivious sort: An extremely simple oblivious sort. In *SOSA*, 2020.
- [5] Kenneth E. Batcher. Sorting networks and their applications. In *AFIPS*, 1968.
- [6] Marina Blanton, Aaron Steele, and Mehrdad Alisagari. Data-oblivious graph algorithms for secure computation and outsourcing. In *ASIA CCS*, 2013.

- [7] Ferdinand Brasser, Urs Müller, Alexandra Dmitrienko, Kari Kostiaainen, Srdjan Capkun, and Ahmad-Reza Sadeghi. Software grand exposure: SGX cache attacks are practical. In *11th USENIX Workshop on Offensive Technologies (WOOT 17)*, Vancouver, BC, August 2017. USENIX Association.
- [8] Jo Van Bulck, Nico Weichbrodt, Rüdiger Kapitza, Frank Piessens, and Raoul Strackx. Telling your secrets without page faults: Stealthy page Table-Based attacks on enclaved execution. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1041–1056, Vancouver, BC, August 2017. USENIX Association.
- [9] T.-H. Hubert Chan, Kai-Min Chung, Wei-Kai Lin, and Elaine Shi. MPC for MPC: secure computation on a massively parallel computing architecture. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPIcs*, pages 75:1–75:52. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [10] T.-H. Hubert Chan and Elaine Shi. Circuit opram: Unifying statistically and computationally secure orams and oprams. 2017.
- [11] Pau-Chen Cheng, Wojciech Ozga, Enriquillo Valdez, Salman Ahmed, Zhongshu Gu, Hani Jamjoom, Hubertus Franke, and James Bottomley. Intel tdx demystified: A top-down approach. *ACM Comput. Surv.*, 56(9), apr 2024.
- [12] Shumo Chu, Danyang Zhuo, Elaine Shi, and T.-H. Hubert Chan. Differentially oblivious database joins: Overcoming the worst-case curse of fully oblivious algorithms. In *2nd Conference on Information-Theoretic Cryptography, ITC 2021, July 23-26, 2021, Virtual Conference*, volume 199 of *LIPIcs*, pages 19:1–19:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [13] Richard Cole and Vijaya Ramachandran. Resource oblivious sorting on multicores. *ACM Trans. Parallel Comput.*, 3(4), mar 2017.
- [14] Emma Dauterman, Vivian Fang, Ioannis Demertzis, Natacha Crooks, and Raluca Ada Popa. Snoopy: Surpassing the scalability bottleneck of oblivious storage. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles, SOSOP ’21*, page 655–671, New York, NY, USA, 2021. Association for Computing Machinery.
- [15] Muhammad El-Hindi, Tobias Ziegler, Matthias Heinrich, Adrian Lutsch, Zheguang Zhao, and Carsten Binnig. Benchmarking the second generation of intel sgx hardware. In *Proceedings of the 18th International Workshop on Data Management on New Hardware, DaMoN ’22*, New York, NY, USA, 2022. Association for Computing Machinery.
- [16] Saba Eskandarian and Matei Zaharia. Oblidb: Oblivious query processing for secure databases. *Proc. VLDB Endow.*, 13(2):169–183, oct 2019.
- [17] Alireza Farhadi, MohammadTaghi Hajiaghayi, Kasper Green Larsen, and Elaine Shi. Lower bounds for external memory integer sorting via network coding. In *STOC*, 2019.
- [18] Andrew Ferraiuolo, Rui Xu, Danfeng Zhang, Andrew C. Myers, and G. Edward Suh. Verification of a practical hardware security architecture through static information flow analysis. In *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2017, Xi’an, China, April 8-12, 2017*, pages 555–568. ACM, 2017.
- [19] M. L. Fredman and D. E. Willard. Blasting through the information theoretic barrier with fusion trees. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing, STOC ’90*, page 1–7, New York, NY, USA, 1990. Association for Computing Machinery.
- [20] O. Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In *STOC*, 1987.
- [21] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *J. ACM*, 1996.
- [22] Michael T. Goodrich. Data-oblivious external-memory algorithms for the compaction, selection, and sorting of outsourced data. *CoRR*, abs/1103.5102, 2011.
- [23] Michael T. Goodrich. Randomized shellsort: A simple data-oblivious sorting algorithm. *J. ACM*, 58(6), dec 2011.
- [24] Michael T. Goodrich. Zig-zag sort: A simple deterministic data-oblivious sorting algorithm running in $O(N \log N)$ time. In *STOC*, 2014.
- [25] Michael T. Goodrich and Michael Mitzenmacher. Privacy-preserving access of outsourced data via oblivious RAM simulation. In *ICALP*, 2011.
- [26] C. A. R. Hoare. Algorithm 64: Quicksort. *Commun. ACM*, 4(7):321–, July 1961.

- [27] William Holland, Olga Ohrimenko, and Anthony Wirth. Efficient oblivious permutation via the waksman network. In *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, ASIA CCS '22, page 771–783, New York, NY, USA, 2022. Association for Computing Machinery.
- [28] Zahra Jafargholi, Kasper Green Larsen, and Mark Simkin. Optimal oblivious priority queues and offline oblivious RAM. In *SODA*, 2021.
- [29] Marcel Keller and Peter Scholl. Efficient, oblivious data structures for MPC. In *Asiacrypt*, 2014.
- [30] Donald E. Knuth. *The Art of Computer Programming, Volume III: Sorting and Searching*. Addison-Wesley, 1973.
- [31] Wei-Kai Lin, Elaine Shi, and Tiancheng Xie. Can we overcome the $n \log n$ barrier for oblivious sorting? In *SODA*, 2019.
- [32] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. Oblivm: A programming framework for secure computation. In *IEEE Symposium on Security and Privacy*, 2015.
- [33] Pratyush Mishra, Rishabh Poddar, Jerry Chen, Alessandro Chiesa, and Raluca Ada Popa. Oblix: An efficient oblivious search index. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 279–296, 2018.
- [34] Kartik Nayak, Xiao Shaun Wang, Stratis Ioannidis, Udi Weinsberg, Nina Taft, and Elaine Shi. Graphsc: Parallel secure computation made easy. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 377–394. IEEE Computer Society, 2015.
- [35] N. Ngai, I. Demertzis, J. Ghareh Chamani, and D. Papadopoulos. Distributed & scalable oblivious sorting and shuffling. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 156–156, Los Alamitos, CA, USA, may 2024. IEEE Computer Society.
- [36] M. S. Paterson. Improved sorting networks with $o(\log n)$ depth. In *Algorithmica*, 1990.
- [37] Enoch Peserico. Deterministic oblivious distribution (and tight compaction) in linear time. *CoRR*, abs/1807.06719, 2018.
- [38] Vijaya Ramachandran and Elaine Shi. Data oblivious algorithms for multicores. In *SPAA '21: 33rd ACM Symposium on Parallelism in Algorithms and Architectures, Virtual Event, USA, 6-8 July, 2021*, pages 373–384. ACM, 2021.
- [39] Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *ACM Conference on Computer and Communications Security (CCS)*, pages 199–212, New York, NY, USA, 2009. ACM.
- [40] Sajin Sasy, Aaron Johnson, and Ian Goldberg. Fast fully oblivious compaction and shuffling. In *ACM CCS*, 2022.
- [41] Sajin Sasy, Aaron Johnson, and Ian Goldberg. Waks-on/waks-off: Fast oblivious offline/online shuffling and sorting with waksman networks. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, CCS '23, pages 3328–3342, New York, NY, USA, 2023. Association for Computing Machinery.
- [42] Elaine Shi. Path oblivious heap: Optimal and practical oblivious priority queue. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*, pages 842–858. IEEE, 2020.
- [43] Signal. Technology deep dive: Building a faster oram layer for enclaves. <https://signal.org/blog/building-faster-oram/>, 2022.
- [44] Emil Stefanov, Marten van Dijk, Elaine Shi, Christopher W. Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. Path ORAM: an extremely simple oblivious RAM protocol. In *CCS*, pages 299–310. ACM, 2013.
- [45] Pramod Subramanyan, Rohit Sinha, Ilia A. Lebedev, Srinivas Devadas, and Sanjit A. Seshia. A formal foundation for secure remote execution of enclaves. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 2435–2450. ACM, 2017.
- [46] Afonso Tinoco, Sixiang Gao, and Elaine Shi. Enigmap : External-memory oblivious map for secure enclaves. In *Usenix Security*, 2023.
- [47] Abraham Waksman. A permutation network. *J. ACM*, 15(1):159–163, jan 1968.
- [48] Xiao Wang, Hubert Chan, and Elaine Shi. Circuit oram: On tightness of the goldreich-ostrovsky lower bound. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, page 850–861, New York, NY, USA, 2015. Association for Computing Machinery.
- [49] Danfeng Zhang, Yao Wang, G. Edward Suh, and Andrew C. Myers. A hardware design language for timing-sensitive information-flow security. In *Proceedings of the Twentieth International Conference on Architectural*

- [50] Wenting Zheng, Ankur Dave, Jethro G. Beekman, Raluca Ada Popa, Joseph E. Gonzalez, and Ion Stoica. Opaque: An oblivious and encrypted distributed analytics platform. In *NSDI*, 2017.

A Additional Algorithmic Details

At the base case of Interleave, there are p elements with distinct keys in $\{0, 1, \dots, p-1\}$, and we need to sort them in ascending order. While we can apply a concretely-efficient sorting network in practice given that p is small, it does not satisfy our requirement for an asymptotically $O(n \log n)$ MergeSplit algorithm. Instead, our construction applies a Waksman’s permutation network [47] to reorder the elements obliviously using no more than $p \log p$ exchanges. For p being power of two, Waksman [47] gave an algorithm to calculate the routing plan using a permutation matrix of size $p \times p$. Since $p \in O(\sqrt{\log N})$ and each entry of the matrix can be represented with one bit, we can make the algorithm oblivious by packing the permutation matrix into $O(1)$ words. However, as we want our butterfly network to have flexible ways, we need to emulate the permutation network for arbitrary p no more than $\sqrt{\log N}$. In this section, we describe how to achieve this goal by recursively applying our Balance algorithm.

Syntax. Permute takes in an array A with n elements whose keys are a permutation of $\{0, 1, \dots, n-1\}$. The goal of Permute is to rearrange the array A such that the i -th element has key i .

Constraints. We require $n \in O(\sqrt{\log N})$.

Intuition: when n is a power of 2. Our permutation network structure is the same as Waksman, but we propose an algorithm that can obliviously calculate the routing plan along the way using Balance as a primitive. For simplicity, we first assume n to be a power of two. As a preprocessing step, we modify each element’s key to be its original key modulo $n/2$. The set of keys hence becomes $\{0, 1, \dots, n/2-1\}$ and each distinct key appears exactly twice. Also, the number of distinct keys is no more than $\sqrt{\log N}$. With these preconditions, we can call Balance on the array, so that both the left and right half contain a suit of these new keys. We then call Permute on each half recursively. As a result, for $i \in \{0, 1, \dots, n/2-1\}$, both $A[i]$ and $A[i+n/2]$ have new key i , which means their original keys are i and $i+n/2$. To complete the permutation, we just need to conditionally exchange $A[i]$ and $A[i+n/2]$, putting the one with the original key i at the front.

Algorithm 4 Permute(A)

Input: The input array A contains n elements, where $n \in O(\sqrt{\log N})$. The i -th element has a key $\pi(i)$ along with a payload, where π is a permutation of set $\{0, 1, \dots, n-1\}$. Also, we let each element own a stack, which can be integrated into the same word storing the key.

Output: A is rearranged such that the i -th element has key i .

```

1:  $n \leftarrow |A|, m \leftarrow \lceil n/2 \rceil$ 
2: if  $n = 1$  then return
3: If  $n$  is odd, pad a filler with key  $n$  at the end of  $A$ .
4: for  $i \leftarrow 0$  to  $2m-1$  do
5:    $b \leftarrow \lfloor A[i].key / m \rfloor$ . Push  $b$  to  $A[i].stack$ .
6:    $A[i].key \leftarrow A[i].key \bmod m$ 
7: BALANCE( $A, m$ )
8: If  $n$  is odd, remove the last element of  $A$ .
9: PERMUTE( $A[0 : m-1]$ )
10: PERMUTE( $A[m : n-1]$ )
11: for  $i \leftarrow 0$  to  $n-1$  do
12:   Pop  $b$  from  $A[i].stack$ .
13:    $A[i].key \leftarrow A[i].key + b \cdot m$ 
14: for  $i \leftarrow 0$  to  $\lfloor n/2 \rfloor - 1$  do
15:   Obviously exchange  $A[i]$  and  $A[i+m]$ . Put the one
      with smaller key to the front.
```

Detailed algorithm: when n is not a power of 2. Algorithm 4 gives a full description of Permute and generalizes it to handle lengths that are non-powers of 2.

- *Handling lengths that are non-powers of 2.* If the current n is not even, we pad a filler element with key n at the end of the array (line 3). Since Balance does not change the last element, the filler is still at the end, and we may exclude it after performing the Balance operation. In an actual implementation, the filler can be imaginary and need not occupy any extra space.
- *Save and restore the original keys without extra space.* When we modify an element’s key k to the new key $k \bmod \lceil n/2 \rceil$, we need to save the original key k . This can be achieved without extra space. Conceptually, imagine that each element owns a stack, and when we modify k to the new key $k \bmod \lceil n/2 \rceil$, we push the bit whether $k \geq \lceil n/2 \rceil$ to its stack (lines 5-6). This way, we can recover the original key k after the recursion (lines 12-13). Since storing $k \bmod \lceil n/2 \rceil$ and whether $k \geq \lceil n/2 \rceil$ does not take up more bits than storing the original k , we can integrate the stack into the same word that stores the key.

Example. For example, Figure 11 depicts the network structure of the Permute algorithm for $n = 7$ (see Algorithm 4). The first two levels of the network are recursive calls of

Lemma 3.6: Obliviousness of Algorithm 3

Proof. For $p \in O(\sqrt{\log N})$ and $Z \in 2^{O(\sqrt{\log N})}$, it only requires $p \log Z \in O(\log N)$ bits to store all the counts in the preprocessing step. As described in the detailed algorithm, by packing the counts into $O(1)$ words, the preprocessing step have a fixed access pattern. By Lemma 3.4, the call to Interleave is also deterministic and depends only on the input size and p . The transposition at the end clearly enjoys fixed access patterns as well. \square

B.3 Analysis of Flexway O-Shuffle

Theorem B.1. *Let the bin size $Z \in \Omega(\log N(\log \log N)^3)$, and the slack factor $\epsilon \in \Theta(\frac{1}{\log \log N})$, the flexway o-shuffle algorithm described in Section 4 is oblivious.*

Proof. The proof is similar to earlier works [4, 38]. If there is no overflow, the input requirements of algorithms 1, 2, 3 will be satisfied, and the access pattern of the oblivious random binning process is deterministic and depend only on input length N , M , B , the size of each element, and the desired failure probability (which in turn decide how we choose the other parameters including Z , ϵ and the number of ways). After the oblivious random binning, we sort each bucket and reveal the number of real elements per bucket. Asharov et al. [4] shows that such leakage is safe because the number of real elements in each last-level bucket is simulatable without knowledge of the input array - it's just the bin loads of a balls-into-bins process. Therefore, it suffices to prove that no bin will receive more than Z elements except with negligible (in N) probability, meaning a simulator that assumes no overflow can produce access patterns statistically indistinguishable from real-world execution. This holds as long as Z is super-logarithmic in N due to the following reasoning.

Consider a fixed bin $A_j^{(\ell)}$ at level ℓ and index j . It can only receive real elements from P_ℓ initial bins, each filled with $Z/(1+\epsilon)$ real elements, where $P_\ell := \prod_{h=1}^{\ell} p_h$. An element with label τ reaches $A_j^{(i)}$ only when $\tau \equiv j \pmod{P_i}$, which occurs with a probability $1/P_i$. Since the labels are chosen independently, we can apply a Chernoff bound to show that $A_j^{(i)}$ overflows with a probability

$$P_{\text{overflow}} = \exp(-\Omega(\epsilon^2 Z)) \leq \exp(-\Omega(\log N \log \log N)) \\ = N^{-\Omega(\log \log N)}.$$

At each level, there are $(1+\epsilon)N/Z$ bins, and the butterfly network has a maximum of $\log((1+\epsilon)N/Z)$ levels (corresponding to the case where every level is two-way). By applying a union bound over all levels and all bins, we obtain the desired bound on the failure probability. \square

Theorem B.2. *Given bin size $Z \in \Theta(\log N(\log \log N)^3)$, flexway o-shuffle incurs $(2+o(1))N \log N$ exchanges and $O(N \log N)$ numerical computation.*

Proof. By Lemma B.1, the flexway butterfly network contains no more than $(\log O(N/Z))/\log p$ levels, and MergeSplit is called $(1+o(1))N/(pZ)$ times at each level. By Lemma 3.5, each MergeSplit operation incurs no more than $pZ(\frac{1}{2} \log Z + \log p + 1)$ exchanges.

Substituting $p \in \Theta(\sqrt{\log N})$ and $Z \in \Theta(\log N(\log \log N)^3)$, it requires the following number of exchanges for routing:

$$(1+o(1))N \frac{\log O(N/Z)}{\log p} (\frac{1}{2} \log Z + \log p + 1) \\ \in (2+o(1))N \log N.$$

Calling bitonic sort to shuffle each bin at the last level requires $\frac{1}{4}Z \log Z(\log Z + 1)$ exchanges. Since the size of each bin is $\text{polylog } N$ and the label length is $\Theta(\log N)$, the probability of label collision in each bin is $O(N^{-C})$ for some positive constant C . Therefore, we only need $1+o(1)$ trials in expectation. Since there are $\frac{(1+o(1))N}{Z}$ bins, the expected number of exchanges to sort the bins at the last level is

$$(1+o(1))N \frac{1}{4} \log Z(\log Z + 1) \in o(N \log N).$$

Finally, removing fillers takes $O(N)$ exchanges. Hence, the total number of exchanges is

$$(2+o(1))N \log N.$$

For numerical computation, labeling all the elements requires $O(N \log N)$ computation, as we assumed that each label has $O(\log N)$ bits and generating each random bit requires constant time. To extract the keys for MergeSplit, only one division and one modulus operation are required per element. Under the word-RAM model, we can determine the multiplicative inverse of all choices of the divisor $p \in O(\sqrt{\log N})$ at compile time, allowing each division and modulo operation to be done in constant time. Since there are $\Theta(N)$ elements and $\Theta(\frac{\log N}{\log \log N})$ levels, key extraction requires $o(N \log N)$ numerical computation in total. For both MergeSplit and bitonic sort, the amount of numerical computation is asymptotically upperbounded by the number of exchanges. Therefore, the time complexity of numerical computation is $O(N \log N)$. \square

Theorem B.3. *Given that $B \geq \log^2 N$, $M \geq B^2$, and let $Z \in \Theta(\log N(\log \log N)^3)$, the number of page swaps for flexway o-shuffle is*

$$((2+o(1)) \log_{\frac{M}{B}} \frac{N}{B} + 1) \frac{N}{B}.$$

Proof. By Lemma B.1, there are $(1+o(1))N/Z$ bins in total. Each MergeSplit performs at least $p_{\min} = \frac{1}{2}\sqrt{\log N}$ ways of partitioning and at most $p_{\max} = \sqrt{\log N}$ ways of partitioning. Hence, the total number of levels in the butterfly network can be upper-bounded as:

$$L_{\text{total}} \in \lceil \log_{p_{\min}} ((1+o(1))N/Z) \rceil \\ \subset (1+o(1)) \frac{\log N - \log Z}{1/2 \cdot \log \log N}$$

A single pass of batch execution can route all elements through at least:

$$L_{\text{batch}} = \lfloor \log_{p_{\max}} \frac{M}{Z} \rfloor \geq \frac{\log M - \log Z}{1/2 \cdot \log \log N} - 1$$

levels. Therefore, the number of passes is at most:

$$\begin{aligned} n_{\text{pass}} &= \lceil \frac{L_{\text{total}}}{L_{\text{batch}}} \rceil \\ &\in (1 + o(1)) \frac{\log N - \log Z}{\log M - \log Z - 1/2 \cdot \log \log N} + 1. \end{aligned}$$

Substituting $Z \in \Theta(\log N (\log \log N)^3)$, we obtain:

$$n_{\text{pass}} \in (1 + o(1)) \frac{\log N}{\log M - 3/2 \log \log N} + 1$$

Given $B \geq \log^2 N$, we can further derive:

$$n_{\text{pass}} \leq (1 + o(1)) \log_{\frac{M}{B}} \frac{N}{B} + 1.$$

Each pass requires $(1 + o(1)) \frac{N}{B}$ page swaps, except that in the first pass we only need to read $(1 + o(1)) \frac{N}{B}$ pages and similarly in the last pass we only need to write $(1 + o(1)) \frac{N}{B}$ pages. Between every neighboring passes, matrix transposition results in another $(1 + o(1)) \frac{N}{B}$ page swaps. Therefore, the total number of page swaps is

$$\begin{aligned} n_{\text{swap}} &\in 2(n_{\text{pass}} - 1)(1 + o(1)) \frac{N}{B} + (1 + o(1)) \frac{N}{B} \\ &= ((2 + o(1)) \log_{\frac{M}{B}} \frac{N}{B} + 1) \frac{N}{B}. \end{aligned}$$

□

B.4 Analysis of Flexway O-Sort

Theorem B.4. *Let the bin size $Z \in \Omega(\log N (\log \log N)^3)$, and the slack factor $\epsilon \in \Theta(\frac{1}{\log \log N})$. Our flexway o-sort algorithm described in Section 4 is oblivious.*

Proof. The proof resembles earlier works [4, 38], which showed that if we apply an oblivious shuffling algorithm and then any non-oblivious, comparison-based sort, the resulting algorithm is oblivious. □

Theorem B.5. *Given bin size $Z = \log N (\log \log N)^3$, flexway o-sort incurs $(2.23 + o(1))N \log N$ exchanges and $O(N \log N)$ numerical computation.*

Proof. When merging multiple sorted chunks in the external memory mergesort, it is sufficient to store pointers in the heap. This approach ensures that each element is copied only once during each pass. Consequently, the complexity of merging in the exchange model can be expressed as $O(N \log_{M/B} \frac{N}{B})$, which is $o(N \log N)$.

At the base case of the external mergesort, each quick-sort incurs $\frac{\ln 2}{3} n \log n \approx 0.23 n \log n$ exchanges in expectation [26]. Suppose that the batch size is M' , where $M' \leq N$, the number of exchanges incurred across all base case instances amounts to an expected value of $0.23 N \log M' \leq 0.23 N \log N$. By Theorem B.2, the expected number of exchanges required by flexway o-sort is

$$(1.23 + c + o(1))N \log N.$$

Since external memory mergesort uses $O(N \log N)$ numerical computation, the overall numerical computation of flexway o-sort is still $O(N \log N)$. □

Theorem B.6. *Given that $M \geq B^2$, $B \geq \log^2 N$, and let $Z \in \Theta(\log N (\log \log N)^3)$, the number of page swaps for flexway o-sort is*

$$((3 + o(1)) \log_{\frac{M}{B}} \frac{N}{B} + 1) \frac{N}{B}.$$

Proof. First, we show that during the last pass of shuffling, elements can be rearranged into $O(N/M)$ sorted chunks. Let M_{batch} represent the batch size at the last level.

Case 1: If $M_{\text{batch}} \geq M/2$, then there are at most $2(1 + o(1))N/M$ batches, with each batch producing a sorted chunk.

Case 2: If $M_{\text{batch}} < M/2$, then we can allocate a buffer of size $M/2$ in the enclave and always copy the shuffled elements first to this buffer. When the buffer is full, we sort all the $M/2$ elements and write them out as a chunk. Consequently, there can be at most $\lceil 2N/M \rceil$ sorted chunks.

Assigning a one-page buffer to each sorted chunk, in the worst-case, the required number of page swaps to merge all the chunks hierarchically is:

$$n'_{\text{swap}} \in (1 + o(1)) \frac{N}{B} \cdot (1 + \log_{\frac{M}{B}} O(\frac{N}{M})).$$

Considering $M \geq B^2$ and $B \geq \log^c N$, we have:

$$n'_{\text{swap}} \in (1 + o(1)) \frac{N}{B} \cdot \log_{\frac{M}{B}} \frac{N}{B}.$$

Combining this with Theorem B.3, we obtain the target expression. □