

# Oblivious Sorting and Queues

Johannes Schneider

University of Liechtenstein, Vaduz, Liechtenstein

---

## Abstract

We present a deterministic oblivious LIFO (Stack), FIFO, double-ended and double-ended priority queue as well as an oblivious mergesort and quicksort algorithm. Our techniques and ideas include concatenating queues end-to-end, size balancing of multiple arrays, several multi-level partitionings of an array. Our queues are the first to enable executions of pop and push operations without any change of the data structure (controlled by a parameter). This enables interesting applications in computing on encrypted data such as hiding confidential expressions. Mergesort becomes practical using our LIFO queue, ie. it improves prior work (STOC '14) by a factor of (more than) 1000 in terms of comparisons for all practically relevant queue sizes. We are the first to present double-ended (priority) and LIFO queues as well as oblivious quicksort which is asymptotically optimal. Aside from theoretical analysis, we also provide an empirical evaluation of all queues.

*Keywords:* sorting, queues, complexity, oblivious algorithms, privacy preserving, computation on encrypted data, secure computing, fully homomorphic encryption, secure multi-party computation

---

## 1. Introduction

Advances in computing on encrypted data such as Fully Homomorphic Encryption (FHE) and secure multi-party computation (SMC) might make outsourcing computation securely practically feasible. Memory access must also be secured. For example, accessing the  $i$ -th element of an array of length  $n$  needs  $O(1)$  operations on RAM machines. But for a program running on encrypted data, the same access mechanism reveals access patterns. The knowledge of seemingly simple access patterns can help to disclose sensitive information such as stock trading patterns [17] or encryption keys [10]. A simple solution requires to access all array elements requiring  $O(n)$  instead of  $O(1)$  time. Oblivious RAM (ORAM) secures memory access more efficiently using multiple parties. Often relying on more than one party is not desirable. Current solutions for oblivious data structures also do not hide (high level) operations, which makes them unsuitable for omnipresent 'if-then-else' statements with private conditions and queue access in branches. Evaluating a confidential expression, keeping data as well as the expression itself secret, becomes straight forward using our LIFO queue and known techniques for computing on encrypted data. Such a scenario is important for cloud computing, ie. a cloud provider might host data for customers, which run their own analytics functionality. The customers wish to keep their data and algorithms private – in case of industrial automation an algorithm often means a mathematical expression on time-series sensor data.<sup>1</sup> To summarize, the main contributions are:

1. We present oblivious LIFO, FIFO and double-ended (priority) queues. The amortized overhead of an operation on the LIFO queue is  $O(\log n)$  in the maximal length  $n$  of the queue. Prior LIFO queues (based on priority queues [22]) required  $O(\log^2 n)$ . For a wide range of applications such as the producer-consumer problem in a streaming context our FIFO queue has only  $O(\log n)$  overhead which improves prior work [22] by a factor  $\log n$ . We are the first to introduce double-ended queues. Our double-ended queue needs  $O(\log^2 n)$ .
2. We are the first to derive oblivious data structures to support push and pop operations that might not alter the stored elements (depending on a parameter).

---

*Email address:* johannes.schneider@uni.li (Johannes Schneider)

<sup>1</sup>In fact, a request from industry motivated this feature.

3. Our deterministic mergesort algorithm improves on [9] for all relevant list sizes, eg. by two orders of magnitude for sorting of 10 billion elements.
4. We state the first oblivious quicksort algorithm. It is asymptotically optimal. The Monte Carlo algorithm succeeds with high probability, ie.  $1 - 1/n^c$  for an arbitrary constant  $c$ .

### 1.1. Overview of Technique

We structure the array representing the queue in subarrays (SA) of increasing size. A SA might be itself a queue. SAs are organized into parts that are merged and split if they are shifted between different SAs. Moving of elements between SAs can cause some of the push and pop operations to require linear run-time in the maximal queue length. But the time is amortized across many operations so that the average overhead is only (poly)logarithmic. Moving of parts between SAs happens based on the number of pops and pushes. It is not dependent on the data held in the queue. We develop a deterministic calling pattern that does not require knowing the number of stored elements in a queue. This allows to hide the number of operations together with another idea: We permit the pop and push of a special (empty) element that does not alter the number of stored elements in the data structure. Put differently, this disguises whether an operation on the data structure changed the stored elements or not. Furthermore, to ensure efficient access to both ends of a queue, eg. as needed for FIFO and double ended queues, we concatenate two ends of a (LIFO) queue.

### 1.2. Outline

We first discuss our model and some notation (Section 2). The main data structures are given in Section 3 (Stack) with a detailed explanation of core ideas and analysis, Section 4 (FIFO) and Section 5 (double-ended queue). Detailed case studies are given in Section 9 after explaining the technique thoroughly. This includes an explanation how obliviousness (and operation hiding) helps in securing code. Performance evaluation can be found in Section 11.

## 2. Preliminaries and Limitations

We assume knowledge of an upper bound on the maximal number of elements  $n$  that can be kept in the data structure, ie. a queue is represented by an array of fixed size  $n$ . This assumption is common for oblivious data structures. Adjusting the size of the data structure exactly to the actual number of elements is impossible since our goal is to conceal the number of elements contained in the queue. Our queues support two operations: Push (allowing empty elements) and Pop (allowing conditional popping). For obliviousness we proved an analogous definition as [6]. Essentially, obliviousness implies that memory access patterns are the same for any input.

**Definition 1.** *A data structure is oblivious if the sequence of memory access only depends on the number of push and pop operations. A sorting algorithm is oblivious if the sequence of memory accesses is the same regardless of the input.*

We use a special (empty) element “ $\emptyset$ ” also denoted by a dash ‘-’ indicating that an element in the queue is unused. Its bit representation must be different from any data item stored in the queue. We use variants of compare and exchange operations. The simplest form takes as input a binary bit  $b$  and two variables  $A$  and  $B$ . It assigns  $A := B$  if the bit  $b$  is 1, otherwise  $A$  is not changed, ie. it computes  $A := b \cdot B + (1 - b) \cdot A$ . The compare-exchange-and-erase  $CmpExEr(b, A, B)$  performs a compare and exchange as described and, additionally, it might erase  $B$ , ie. it sets variable  $B$  to  $\emptyset$  if  $b$  is 1 and leaves it unchanged otherwise (see PseudoCode  $CmpExEr$  in Algorithm 1). For the analysis we distinguish between input sensitive operations involving parameters of the push and pop elements as well as data of the queue and operations that do not directly depend on any input data (but potentially on the number of operations). The motivation is that for secure computation these distinctions are meaningful, since the former correspond to (slower) operations on encrypted data. For our algorithms input sensitive operations always dominate the time complexity – even when using non-encrypted data. They are split into elementary operations (+, -, ·), called  $E-Ops$ , and comparisons  $C-Ops$ , which are composed of elementary operation. The distinction is motivated since comparisons are used to measure performance of sorting algorithms. For encrypted operations, comparisons might have different time complexities, eg. for SMC such as [19] it is not clear how to perform a comparison in less than  $\Omega(n_b \cdot E - Ops)$  time, where  $n_b$  is the number of bits of a compared number.

### 3. LIFO (Stack)

70 For a Last-In-First-Out (LIFO) queue (also called Stack) a pop operation returns the most recently pushed element onto the data structure. To ensure obliviousness we access the same array elements independent upon the data contained in the queue. Our queue always accesses the first element. A newly pushed element is stored in the first position of the array. This implies that upon every insertion, we must shift elements to the right to avoid overwriting of a previously inserted element. It is easy to shift the entire queue to the right but this requires linear run-time. To  
 75 improve efficiency, we logically split the array representing the queue into subarrays (SAs) of exponentially growing size. We only shift parts of size at most  $2^k$  of a SA after every  $2^k$  push or pop operations.

LIFO Queue A with 3 subarrays  $S(i)$ :  $A = [S(0) \ || \ S(1) \ || \ S(2)]$   
 Subarray  $S(i)$  with 4 parts  $P(i,j)$ :  $S(i) = P(i,0) \ | \ P(i,1) \ | \ P(i,2) \ | \ P(i,3)$   
 Part  $P(i,j)$  with  $2^i$  elements:  $P(i,j) = E(i,j,0) \ E(i,j,1) \ E(i,j,2) \ \dots \ E(i,j,2^i-1)$   
 Legend: "||" separates subarrays, "|" parts and " " elements

Elements 1-12 pushed to ( $\Rightarrow$ ) LIFO Queue:

```

1 => [ 1| -| -| -|| - -| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
2 => [ 2| 1| -| -|| - -| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
3 => [ 3| 2| 1| -|| - -| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
4 => [ 4| 3| -| -|| 2 1| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
5 => [ 5| 4| 3| -|| 2 1| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
6 => [ 6| 5| -| -|| 4 3| 2 1| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
7 => [ 7| 6| 5| -|| 4 3| 2 1| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
8 => [ 8| 7| -| -|| 6 5| 4 3| 2 1| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
9 => [ 9| 8| 7| -|| 6 5| 4 3| 2 1| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
10 => [10| 9| -| -|| 8 7| 6 5| 4 3| 2 1|| - - - -| - - - -| - - - -| - - - -]
11 => [11|10| 9| -|| 8 7| 6 5| 4 3| 2 1|| - - - -| - - - -| - - - -| - - - -]
12 => [12|11| -| -||10 9| 8 7| 6 5| - -|| 4 3 2 1| - - - -| - - - -| - - - -]

```

Elements popped from ( $\Leftarrow$ ) LIFO Queue:

```

12 <= [ -|11| -| -||10 9| 8 7| 6 5| - -|| 4 3 2 1| - - - -| - - - -| - - - -]
11 <= [10| 9| -| -|| - -| 8 7| 6 5| - -|| 4 3 2 1| - - - -| - - - -| - - - -]
10 <= [ -| 9| -| -|| 8 7| 6 5| - -| - -|| 4 3 2 1| - - - -| - - - -| - - - -]
9 <= [ 8| 7| -| -|| - -| 6 5| - -| - -|| 4 3 2 1| - - - -| - - - -| - - - -]
8 <= [ -| 7| -| -|| 6 5| - -| - -| - -|| 4 3 2 1| - - - -| - - - -| - - - -]
7 <= [ 6| 5| -| -|| - -| - -| - -| - -|| 4 3 2 1| - - - -| - - - -| - - - -]
6 <= [ -| 5| -| -|| - -| - -| - -| - -|| 4 3 2 1| - - - -| - - - -| - - - -]
5 <= [ 4| 3| -| -|| - -| 2 1| - -| - -|| - - - -| - - - -| - - - -| - - - -]
4 <= [ -| 3| -| -|| 2 1| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
3 <= [ 2| 1| -| -|| - -| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
2 <= [ -| 1| -| -|| - -| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]
1 <= [ -| -| -| -|| - -| - -| - -| - -|| - - - -| - - - -| - - - -| - - - -]

```

Figure 1: Pushes and pops onto a LIFO queue

More formally, a queue is implemented as an array  $A$  that is split into  $s$  subarrays (SA)  $S_i$  growing exponentially in size with  $i$ . The total length  $n$  of the array is  $n := \sum_{i=0}^{s-1} |S_i|$ . Each SA  $S_i$  itself is partitioned into  $q$  parts  $P_{i,0}, P_{i,1}, \dots, P_{i,q-1}$  of equal size  $|P_{i,j}| = |S_i|/q$ . The size of a part varies for different SAs. We denote the  $k$ -th element in  $P_{i,j}$  by  $E_{i,j,k}$ . Figure 1 shows the structure of a queue.  
 80

#### 3.1. Push, Pop and Shifting

We explain the shifting procedure shown in Figure 1 for a sequence of push operations. We always push an element onto the first position in the array  $A$  (or pop an element from there). After every modification of the queue,

Steps for a Push of 6 onto LIFO Queue:

```
Original queue:
[ 5| 4| 3| -|| - -| 2 1| - -| - -|| - - - -| - - - -| - - - -| - - - -]
after ShiftPartsRight (Sublist 0):
[ -| 5| 4| 3|| - -| 2 1| - -| - -|| - - - -| - - - -| - - - -| - - - -]
after E(0,0,0)=6
[ 6| 5| 4| 3|| - -| 2 1| - -| - -|| - - - -| - - - -| - - - -| - - - -]
final queue after executing emptyTwoParts (Sublist 0):
[ 6| 5| -| -|| 4 3| - -| 2 1| - -|| - - - -| - - - -| - - - -| - - - -]
```

Executing emptyTwoParts (Sublist 0):

```
Original queue:
[ 6| 5| 4| 3|| - -| 2 1| - -| - -|| - - - -| - - - -| - - - -| - - - -]
after ShiftPartsRight (Sublist 1):
[ 6| 5| 4| 3|| - -| - -| 2 1| - -|| - - - -| - - - -| - - - -| - - - -]
after shifting P(0,2) to P(1,0):
[ 6| 5| -| 3|| 4 -| - -| 2 1| - -|| - - - -| - - - -| - - - -| - - - -]
final queue after shifting of parts P(0,3) to P(2,0):
[ 6| 5| -| -|| 4 3| - -| 2 1| - -|| - - - -| - - - -| - - - -| - - - -]
```

Figure 2: Steps for pushing an element onto a LIFO queue

---

## Algorithm 1 LIFO

---

### Initialization(Number of SAs $s$ with $s \geq 1$ )

$q := 4$  {number of parts per SA}  
 $E_{i,j,k} := \emptyset, \forall i \in [0, s-1], j \in [0, q-1], k \in [0, 2^i - 1]$   
 $n_{pu} := n_{po} := 0$  {counter for pushes and pops}

### CmpExEr(b,A,B)

$A := b \cdot B + (1-b) \cdot A$  {Exchange A, B based on b}  
 $B := (1-b) \cdot B + b \cdot \emptyset$  {Delete B based on b}

### ShiftPartsRight(SA i,doOp)

$empty\&DoOp := doOp$  if  $E_{i,0,0} \neq \emptyset$  else 0  
**for** SA  $j := q-1$  **to** 1 **do**  
    $doShift := empty\&DoOp$  if  $E_{i,j,0} = \emptyset$  else 0  
   **for** Element  $k := 0$  **to**  $|S_i|/q - 1$  **do**  
     CmpExEr(doShift,  $E_{i,j,k}, E_{i,j-1,k}$ )

### EmptyTwoParts(SA i)

$isFull := 1$  if  $\bigwedge_{j=0}^{q-1} (E_{i,j,0} = \emptyset)$  else 0  
 ShiftPartsRight(SA i+1, isFull)  
**for** SA  $j := q-2$  **to**  $q-1$  **do**  
    $o := (j-q+2) \cdot |S_i|/q$  {offset for last 2 parts}  
   **for** Element  $k := 0$  **to**  $|S_i|/q - 1$  **do**  
     CmpExEr(isFull,  $E_{i+1,0,k+o}, E_{i,j,k}$ )

### Push(Element x)

$n_{pu} := MoveBetweenSAs(n_{pu}, EmptyTwoParts)$   
 $b := 1$  if  $x \neq \emptyset \wedge E_{i,0,0} \neq \emptyset$  else 0  
 $E_{0,0,0} := x$  if  $x \neq \emptyset$  else  $E_{0,0,0}$

### MoveBetweenSAs(nOps,Operation Op)

$mi := 0$  {Find maximal SA to empty/refill}  
**while**  $(nOps + 1) \bmod 2^{mi+1} = 0$  **do**  
    $mi := mi + 1$   
**for** SA  $i := \min(mi, s-2)$  **to** 0 **do**  
   Apply Operation  $Op$  on SA  $i$   
**return**  $(nOps + 1) \bmod 2^{\max(0, s-2)}$

### ShiftPartsLeft(SA i)

$full\&DoOp := 1$  if  $E_{i,0,q-1} \neq \emptyset$  else 0  
**for** SA  $j :=$  **to**  $q-2$  **do**  
    $doShift := full\&DoOp$  if  $E_{i,j,0} \neq \emptyset$  else 0  
   **for** Element  $k := 0$  **to**  $|S_i|/q - 1$  **do**  
     CmpExEr(doShift,  $E_{i,j,k}, E_{i,j+1,k}$ )

### RefillTwoParts(SA i)

$isEmpty := 1$  if  $\bigwedge_{j=0}^{q-1} (E_{i,j,0} \neq \emptyset)$  else 0  
 ShiftPartsLeft(SA i+1)  
**for** SA  $j := q-2$  **to**  $q-1$  **do**  
    $o := (j-q+2) \cdot |S_i|/q$  {offset for last 2 parts}  
   **for** Element  $k := 0$  **to**  $|S_i|/q - 1$  **do**  
     CmpExEr(doShift,  $E_{i,j,k}, E_{i+1,0,k+o}$ )

### Pop(doPop)

$result := E_{0,0,0}$  if doPop else  $\emptyset$   
 $E_{0,0,0} := \emptyset$  if doPop else  $E_{0,0,0}$   
 $n_{po} := MoveBetweenSAs(n_{po}, refillTwoParts)$   
**return** result

---

we modify (some) SAs to ensure that there is space for further pushes in the first SA. We shift elements to the right. Shifting is only done on a part level, ie. either we shift all elements of a part or none. We perform frequent shifts to overwrite empty small parts near the beginning of the array and less frequent shifts are conducted for larger parts situated towards the end of the array. We shift parts within a SA but also move parts between SAs, ie. either we merge two parts into one or we split a part into two parts. The subroutines for a push shown in Algorithm 1 are discussed next.

ShiftPartsRight and EmptyTwoParts: ShiftPartsRight shifts elements from one part to the next part (on the right) within a SA. It avoids overwriting of filled parts by checking if the part to be overwritten is indeed empty. To this end, we only check if the first position of a part is empty. No parts are moved, if the first part of the SA is empty. A parameter indicates whether shifting should take place or not. This is necessary to enable executions of push operations that do not modify the queue. If the parameter is false, ie. zero, then no elements are moved. The order of shifting is from back to front, ie. elements of the second to last part are shifted to the last part (given it is empty), then the third to last part is shifted to the second to last (if empty) and so on. EmptyTwoParts empties the last two parts of a SA  $i$  by merging them to form the first part of SA  $i + 1$ . It first empties the first part in SA  $i + 1$  by doing a ShiftPartsRight. Emptying only takes place if all parts of SA  $i$  are full and SA  $i + 1$  is not completely full. Without this condition for  $q > 2$  a full part would be (continuously) shifted towards the right for repeated insertions of the empty element  $\emptyset$ . This would lead to empty SAs followed by (partially) full SAs. As a consequence for pop operations we would have to undo the shifting (or search the entire array).

Push: A push operation first ensures that the first position of the array is empty. Then, it inserts the pushed element at the first position. A push and its suboperations are illustrated in Figure 2.

MoveBetweenSAs: Restructuring is done after every operation starting from some initial SA (down) to the very first SA in the beginning of the queue. The (index of the) initial SA depends on the number of operations and not the number of actual elements in the queue, which we wish to disguise. Parts of a SA are moved to the next SA, once a SA is full. It might seem reasonable to move all parts of a full SA to the next. However, for alternating pushes and pops this might trigger large performance penalties since parts are continuously moved back and forth between SAs. To disguise the number of elements in the queue (and thus parts), we access all parts in the same deterministic manner for any sequence of pushes of fixed length. Since we allow pushes of a special (empty) element that has no impact on the number of stored elements, the number of operations (as an indicator for the actual number of elements contains) is not exact. We assume that the array grows at a maximal rate, ie. every push is done using a non-empty element. Since we always empty two parts of a SA, we must create space in a SA by moving elements, whenever a sequence of operations could have resulted in the filling of two parts of that SA. For example, every push potentially fills one part in the first SA, since they are of size one. Thus, we would empty the first SA after every second push operation. For SA  $i$  with parts of size  $2^i$ , we would move two parts to the next SA after every  $2^{i+1}$  operations. But this approach fails for an arbitrary interleaving of operations pops and pushes of empty and non-empty elements. For example, for the following sequence of pushed elements 1, 2, 3,  $\emptyset$ , 4, 5, the algorithm would attempt to empty the first SA after having pushed 1, 2 and again after 1, 2, 3,  $\emptyset$ . The first SA contains 1, 2, 3 and misses one element to be full. Thus, the SA would not be emptied and two more elements could be (attempted) to be pushed onto the SA before trying to empty it again, but the SA becomes full after pushing one more element. Therefore, we perform restructuring operations more frequently, ie. for SA  $i$  we execute EmptyTwoParts after every  $2^i$  operations (rather than after  $2^{i+1}$ ). The last SA that can be emptied is the second to last, ie. the one with index  $s - 2$ . The restructuring is done in Algorithm MoveBetweenSAs which executes for a push operation EmptyTwoParts on all parts as described. It takes as input the counter of the current operations and returns the next value for the counter, which is (usually) the counter incremented by 1. However, once the maximal possible SA has been shifted the operation counter is reset to zero, eg. for  $s = 5$  the counter is reset after  $2^{s-2} = 8$  operations. The sequence of maximal SA indexes where parts might be moved to the next SA is a repetition of the sequence 0, 1, 0, 2, 0, 1, 0, 3.

Pop: The pop operation (and its subroutines) behave analogously to push but reverse. ShiftPartsLeft shifts parts of a SA within the SA towards the beginning. In contrast to ShiftPartsRight, we do not need a parameter to indicate whether we actually perform the operation or not. ShiftPartsLeft only shifts a part, if the first part is empty. RefillTwoParts moves the first part of SA  $i + 1$  to the beginning of SA  $i$ . One full part in SA  $i + 1$  corresponds to two full parts in SA  $i$ . As for emptying of parts and right shifts, no non-empty parts are overwritten.

### 3.2. Analysis

135 **Theorem 1.** *The LIFO queue is oblivious.*

*Proof.* According to Definition 1 we require that memory accesses are independent of the input. (They are allowed to be dependent on the number of operations.) None of the procedures in Algorithm 1 accesses memory cell dependent on an input value, ie. all loop-conditions do not depend on the input and any conditional access to memory cells of the form ‘cell0:=a if cell1=x else b’ can be expressed as multiplications (Section 2).  $\square$

140 We analyze push and pop operations with respect to time complexity (Theorem 2) and correctness (Theorem 3). In the worst case a single operation might take  $\Omega(n)$ , where  $n$  is the maximal length of the queue. We prove that on average, the time is only logarithmic in  $n$ .

**Theorem 2.** *For the LIFO queue a pop and push operation requires amortized  $O(\log n)$  time, ie.  $14q \log(n/q)$  E-Ops and  $8q + 2$  C-Ops.*

145 The proof uses that two parts of SA  $i$  of length  $2 \cdot 2^i$  are refilled (emptied) after every  $2^i$  push (pop) operations. Since there are  $O(\log n)$  SAs we get time  $\sum_{i=0}^{\log n} 2^{i+1}/2^i = O(\log n)$ .

*Proof.* SA  $i$  is refilled (emptied) after every  $2^i$  pop (push) operations. After refilling (emptying) all SAs from index  $s-2$  to 0, ie. after  $2^{s-2}$  pop (push) operations, we start over by considering SA 0 only. The average run-time increases up to the point, where SA  $s-2$  is considered. Thus, it suffices to compute the average number of operations for a sequence of  $2^{s-2}$  pop (push) operations. We analyze pop operations by counting of E-Ops followed by C-Ops. CmpExEr needs 7 E-OPs (2 additions, 1 subtraction, 4 multiplications). ShiftsPartsLeft for SA  $i$  needs  $2^i \cdot 7(q-1)$ . RefillTwoParts on SA  $i$  performs one shift in list  $i+1$  and moves one part of it, yielding  $2^i \cdot 7(q-1) + 2^i \cdot 7q = 7 \cdot 2^i \cdot q$ . Since RefillTwoParts on SA  $i$  is called after every  $2^{i-1}$  pops, on average refilling of SA  $i$  contributes by  $7 \cdot 2^i \cdot q/2^{i-1} = 14q$  E-Ops. By definition we have  $n = \sum_{i=0}^{s-1} |S_i| = \sum_{i=0}^{s-1} q \cdot 2^i = q \cdot (2^s - 1)$  yielding  $s = (\log(n/q)) + 1$ . Summing over all SAs gives

$$\sum_{i=0}^{s-2} 14q = 14q(s-1) = 14q \log(n/q)$$

The analysis of C-Ops is analogous. CmpExEr contains zero comparisons. In ShiftPartsLeft we perform one comparison (line 1) and one in each of the  $2^i \cdot (q-1)$  iterations. A refill of SA  $i$  takes  $2q$  comparisons ( $q$  to compute isEmpty in RefillTwoParts (line 1) and  $q$  within ShiftPartsLeft). Therefore, the number of comparisons becomes  
150  $\sum_{i=0}^{s-2} 2q/2^{i-1} \leq 8q$ . Adding two C-Ops due to lines 1-2 in Algorithm Pop completes the proof for pop. The push operation is analyzed in the same manner.  $\square$

**Lemma 1.** *Each part  $P_{i,j}$  can only be in one of two states: empty (all elements being  $\emptyset$ ) or full (no elements being  $\emptyset$ ).*

This follows since we modify either all or none of the elements of a part.

155 *Proof.* Initially, all parts are empty. Parts of the first SA can only be full or empty, since they contain at most one element. Parts of SA  $i > 0$  are only modified due to mergers, splits and shifting. Right or left shifting of a part within a SA is done for entire parts. The part overwritten is an exact copy of the part being shifted. The part being shifted becomes empty, ie. all elements are set to the empty element. When all parts of SA  $i$  are full, the last two parts of a SA, ie.  $P_{q-2,i}$  and  $P_{q-1,i}$ , each of size  $2^i$  are shifted to the next SA, ie. to become the first  $P_{0,i+1}$  of size  $2^{i+1}$ . This part  
160 is filled completely. A filled part in SA  $i$  (see procedure RefillTwoParts) split into two parts of the same size, yields two full parts in SA  $i-1$ .  $\square$

**Theorem 3.** *The LIFO queue works correctly.*

We show that no elements are overwritten and no empty elements are returned if the array is non-empty since we refill and empty parts of SAs sufficiently often.

165 *Proof.* In Algorithm 1 no parts are overwritten if the first element of a part is non-empty – see definition and usage of variables *full&doOp*, *empty&doOp* in *ShiftPartsLeft/Right*; *isFull*, *isEmpty* in *Empty/RefillTwoParts*; line 2 of push with  $E_{i,0,0} \neq \emptyset$ . Since all elements of a part are either the empty element or differ from it (Lemma 1), checking the first element suffices to avoid overwriting of non-empty parts.

We first show that there is no interleaving of empty and non-empty SAs. Let the  $t$ -th SA be the largest SA such that at least one part is full. All SAs  $i < t$  contain at least one non-empty part. An arbitrary sequence of pushes cannot reduce the number of full parts in a SA below two. This follows since we only empty two parts of a SA if all four parts are full. An arbitrary sequence of pops cannot completely empty a SA except the last, since SA  $i$  being of size  $q \cdot 2^i$  is refilled with elements from SA  $i + 1$  after every  $2^i$  pops (see *MoveBetweenSAs* in Algorithm 1).

170 Next we show that there is no interleaving of SAs with some non-empty parts and SAs with only full parts. *EmptyTwoParts* executes on SA  $i$  before it executes on SA  $j < i$ . Upon execution there are two possibilities: Either no or two parts are moved to SA  $i + 1$ . In the first case at most 3 parts are full in SA  $i$  and thus, we could insert one more part, in the second case the SA is full and two parts are emptied. Either way, it suffices to empty SA  $i$  after two parts in SA  $i - 1$  (might) have been filled. Since this corresponds to  $2^{i-1}$  elements, our choice of calling *EmptyTwoParts*  $i$  after every two  $2^{i-1}$  operations suffices (see *MoveBetweenSAs* in Algorithm 1). Therefore, not all parts of a SA can be full, if there is space in a larger SA. For refilling parts an analogous argument applies.  $\square$

#### 4. FIFO

A First-In-First-Out (FIFO) queue needs fast access to the first and the last element. We use an array of LIFO queue variants of increasing lengths, ie. each SA of the FIFO queue is itself a LIFO queue. Each LIFO queue stores elements in ‘reverse’ order, meaning the first element to be popped in the LIFO queue is the oldest element the LIFO queue contains. In this way we can efficiently access the oldest element of each LIFO queue. The array structure is visualized in Figure 3. Each LIFO queue matches on SA.

For a pop operation the LIFO queue with largest index that is non-empty is identified. Then an element is popped from that queue. To make the algorithm oblivious we execute a pop operation on every LIFO queue within the FIFO queue. We start from the back and pop an element from each LIFO queue, ie. SA, until the first non-empty LIFO queue has been identified. For the remaining queues we execute pops using a parameter to indicate that, in fact, no element should be popped. The key point is that independent of the value of the parameter the same memory cells are accessed.

185 *PopperQueue:* A LIFO queue offers more functionality than is needed, since we do not push elements in the front but only pop them except for the first queue, which is just a single element. Opposed to a LIFO queue, we can therefore refill a SA completely. We reduce the number of parts from four to two. Using more parts per SA is slower since we must shift the same elements multiple times rather than moving them less often in bigger chunks, ie. larger SAs. We can reuse most LIFO procedures (Algorithm 1) without modification, ie. *ShiftsPartLeft*, *RefillTwoParts* and *Pop*. We call this LIFO variant “*PopperQueue*”. It is a special case of the LIFO queue from Section 3. It has the same (asymptotic) properties, but it is roughly a factor of two faster, since it uses less parts and therefore requires less shifts within a SA, ie. compare Theorem 2 for  $q = 4$  (LIFO) and  $q = 2$  (*PopperQueue*).

Due to the more involved array organization of a FIFO queue, the emptying of parts and refilling of parts needs careful attention. It is not possible to concatenate two parts to get a larger part without extra processing, ie. two arrays (of *PopperQueues*) placed after each other generally do not yield an array representing a larger *PopperQueue* with a valid structure. The concatenation could give partially filled parts. For example, assume that there are two queues with one SA and two parts, eg.  $[1|-]$  and  $[-|4]$ , naive concatenation yields  $[1|-||-4|-]$  having the partially filled part  $|-4|$ . Furthermore, we have to ensure a correct ordering of the elements within LIFO queues when moving elements between them.

190 If one last part of the *PopperQueue* stored in SA  $i$  is full, we move  $2^i$  elements from queue  $i$  to the very last part of queue  $i + 1$ . We pop one element after the other from queue  $i$  and put it directly into the last part of queue  $i + 1$ , ie. the element of the  $j$ -th pop is put at the  $j$ -th position of the last part. At this point the whole queue  $i + 1$  (except the last part that was just inserted) might be empty which would cause subsequent calls of pop on queue  $i + 1$  to fail. Therefore, we attempt to shift elements from the last part of the last SA of LIFO queue  $i + 1$  consisting of newly inserted elements up to the first SA of queue  $i + 1$ .

FIFO Queue F with 3 LIFO queues A(m):      F            = <L(0)> <L(1)> <L(2)>  
 LIFO Queue L(m) with m+1 subarrays S(i): L(m)    = [S(0) || ... || S(m-1)]  
 Subarray S(i) with 2 parts P(i,j):            S(i)        = P(i,0) | P(i,1)  
 Part P(i,j) with 2<sup>i</sup> elements:                P(i,j)      = E(i,j,0) E(i,j,1) E(i,j,2) ... E(i,j,2<sup>i</sup>-1)

Elements pushed to (=>) FIFO Queue

1 =>	<[ -  1]>	<[ -  -   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
2 =>	<[ 1  2]>	<[ -  -   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
3 =>	<[ -  3]>	<[ 1  2   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
4 =>	<[ 3  4]>	<[ 1  2   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
5 =>	<[ -  5]>	<[ 1  2   3 4  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
6 =>	<[ 5  6]>	<[ 1  2   3 4  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
7 =>	<[ -  7]>	<[ 1  2   3 4  5 6]>	<[ -  -   - -  - -   - - - -  - - - -]>
8 =>	<[ 7  8]>	<[ 5  6   - -  - -]>	<[ 1  2   - -  3 4   - - - -  - - - -]>
9 =>	<[ -  9]>	<[ 5  6   7 8  - -]>	<[ 1  2   - -  3 4   - - - -  - - - -]>
10 =>	<[ 9 10]>	<[ 5  6   7 8  - -]>	<[ 1  2   3 4  - -   - - - -  - - - -]>

Elements popped from (<=) FIFO Queue

1 <=	<[ 9 10]>	<[ 5  6   7 8  - -]>	<[ -  2   3 4  - -   - - - -  - - - -]>
2 <=	<[ 9 10]>	<[ 5  6   7 8  - -]>	<[ 3  4   - -  - -   - - - -  - - - -]>
3 <=	<[ 9 10]>	<[ 5  6   7 8  - -]>	<[ -  4   - -  - -   - - - -  - - - -]>
4 <=	<[ 9 10]>	<[ 5  6   7 8  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
5 <=	<[ 9 10]>	<[ -  6   7 8  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
6 <=	<[ 9 10]>	<[ 7  8   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
7 <=	<[ 9 10]>	<[ -  8   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
8 <=	<[ 9 10]>	<[ -  -   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
9 <=	<[10  -]>	<[ -  -   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>
10 <=	<[ -  -]>	<[ -  -   - -  - -]>	<[ -  -   - -  - -   - - - -  - - - -]>

Figure 3: Sequence of pushes and pops onto a FIFO queue



215 The push operation for the FIFO queue appends elements to the end of the very first LIFO queue. Since it is of length two, we shift the second element of it to the left and then set the second position to the newly inserted element.

**Corollary 1.** *For the FIFO queue a pop operation requires  $O(\log^2 n)$  and a push  $O(\log n)$  time on average.*

*Proof.* For a pop of the FIFO queue we do a pop for each of the PopperQueues giving  $\sum_{i=0}^s O(i) = O(s^2) = O(\log^2 n)$ . For a push we move blocks of size  $2^i$  from SA  $i$ , ie. PopperQueue  $i$ , to SA  $i + 1$  after every  $2^i$  operations, which needs time linear in the queue length. Summation gives  $\sum_{i=0}^s O(2^i/2^i) = O(s) = O(\log n)$ .  $\square$

#### 220 4.1. Fast FIFO (and Double-Ended Queues)

FIFO queues are often used as buffers to distribute peak loads across a longer timespan. Commonly, a producer pushes elements onto the queue continuously (as a stream), while a consumer repeatedly takes an element and processes it. Buffering always introduces some delay in processing. Thus, usually an additional delay is tolerable. A pop on the fast FIFO queue only returns an element given the queue has been filled partially, ie. it is at least half full. Our FIFO queue that has only amortized  $O(\log n)$  overhead rather than  $O(\log^2 n)$ . The idea is to use two queues “back to back”: one for popping and one for pushing. The two queues share the last part, ie. both treat this part as belonging to them. Thus, elements are pushed onto one of the queues and are continuously shifted to the right with newly inserted elements until they reach the queue for popping. A pop only returns an element after its last part of the last SA (shared with the pushing queue) has been filled. The same ideas also apply to double-ended queues.

230 For the Fast FIFO Queue (B2B Queue) the time complexity of a push and pop matches the corresponding operations for the LIFO queue.

**Corollary 2.** *For the B2B-FIFO Queue a pop and push operation require  $O(\log n)$ .*

## 5. Double-Ended Queue

235 A double-ended queue supports popping elements at the head and tail as well as prepending elements at the beginning and appending them at the end. We combine ideas for LIFO and FIFO queues. We use an array of queues (as for FIFO queues) to address the need to push elements to the head of the array and pop them from the tail. Since elements can also be pushed at the back, we use LIFO queues, ie. SA  $i$  of the double-ended queue is given by a LIFO queue with  $i + 1$  SAs (rather than PopperQueues). Pushing to the back requires identifying the last non-empty SA, ie. queue, as for popping from the back in the FIFO queue. However, we can only push the element onto the queue, if it is non-full, otherwise we push it onto the next queue. Popping elements from the front might trigger refilling of SAs. In turn, we have to move the newest elements of one SA to another. Identifying the newest elements of a LIFO queue (with elements sorted by age, ie. ascending insertion order) is cumbersome, since there is only efficient access to the oldest element. To reverse order, we remove all elements from the array (using a sequence of pops) and insert them into a temporary LIFO queue. This yields a queue sorted by newest to oldest elements. Then we move elements by popping them from the temporary queue to the queue to refill, ie. for queue  $i$  we move  $2^{i+1}$  elements. The remaining elements are pushed back onto the emptied queue (used to create the temporary LIFO queue).

**Theorem 4.** *Any operation on the double-ended queue has amortized time  $O(\log^2 n)$ .*

250 Operations are similar to the LIFO queue, except for refilling and emptying that needs an additional logarithmic factor due to the popping and pushing of elements rather than direct access in  $O(1)$ .

*Proof.* Pushing and popping to the front works the same as for LIFO queues except for the refilling and emptying of full SA. We require an additional logarithmic factor, since we cannot just copy elements of one SA, ie. queue, to another but we first pop them from the LIFO queue onto a temporary queue. More precisely, each element access using a pop requires amortized  $O(\log n)$  as shown in Theorem 2 rather than  $O(1)$ . Pushing and popping to the back requires executing a constant number of push and pop operations for all parts constituting LIFO queues. Since we have  $O(\log n)$  queues and each operation on a LIFO queue requires  $O(\log n)$  (see Theorem 2), a single push and pop operation requires  $O(\log^2 n)$ .  $\square$

## 6. Double-Ended Priority Queue

In this scenario, each data item has a priority. A double-ended priority queue can return either the data element with the smallest or largest priority. The queue structure is the same as for double-ended queue. We ensure that each SA, ie. LIFO queue, contains elements sorted in descending priority. When moving elements from one queue to another, ie. to empty a full queue or refill a queue, we first create one single sorted array containing all elements from both queues and then refill the smaller queue up to half of its capacity with the elements of smallest priority and put the other elements in the larger queue. The sorting can be done by merging both arrays.

Popping the element of minimum priority requires finding the smallest element in SA 0. Popping the element of maximum priority requires checking all parts, since we do not know which parts contain elements and which do not as well as which part contains the element with largest priority. More precisely, we first (peek) all parts and find the element and part with the maximum element. After that we perform a pop on the (first) queue containing the maximum element. This is done by executing a pop for all parts. The parameter of the pop operation, determining whether the operation indeed removes an element from the queue, must be set accordingly, ie. it is  $\emptyset$  for all but the queue containing the maximum element.

The restructuring is somewhat more involved. Upon a push that requires restructuring, eg. either refilling or emptying queue  $i$  we first create one sorted array in increasing order by merging both queues as done for ordinary mergesort (see also Section 7). We then refill SA  $i$  until it is half full with the smallest elements (in reversed order) and insert the remaining to the next SA (in reversed order).

**Theorem 5.** *Any operation on the double-ended priority queue has amortized time  $O(\log^2 n)$ .*

*Proof.* We discuss time followed by correctness. Pushing an element to the front (or popping the element of minimum priority) works the same as for LIFO queues except for the emptying and refilling of full SA. We require an additional logarithmic factor to move elements from queue  $i$  to queue  $i + 1$  (or the other way around), we create a temporary queue by repeatedly popping the element of maximum priority and adding it to the temporary queue. Each element access using a pop requires amortized  $O(\log n)$  as shown in Theorem 2 rather than  $O(1)$ . Moving the elements from the temporary queue onto the (new) queues  $i$  and  $i + 1$  has the same asymptotic time complexity. Therefore, push to the front need  $O(\log^2 n)$  time. Popping the maximum priority element requires executing a pop operation for all LIFO queues (plus restructuring). Since we have  $O(\log n)$  queues and each operation on a LIFO queue requires  $O(\log n)$  (see Theorem 2), a single push and pop operation requires  $O(\log^2 n)$ . Popping the maximum priority element requires executing a pop operation for all LIFO queues (plus restructuring). This requires  $O(\log^2 n)$ .

Correctness of a pop of maximum priority follows, since we maintain all queues in descending order of priority. Thus, the element of maximum priority is the first element in one of the queues. Since we consider the first elements of all queues and return the one of maximum priority, correctness follows. For the minimum we only investigate the first queue. Since upon every restructuring operation on queue  $i$  we keep the smallest half of both queues  $i$  and  $i + 1$  in queue  $i$ , it holds that after a restructuring all elements in SA  $i$  are smaller than any element in SA  $i + 1 > i$ . Using induction, we have that the smallest element is in SA 0.  $\square$

## 7. Oblivious Mergesort

Our oblivious mergesort algorithm (O-Mergesort) divides an unsorted array (or list) into SAs of one element. It repeatedly merges two arrays of equal length to obtain a new sorted array of double the length until there is only one array remaining. This array is sorted. To make the sorting procedure oblivious requires a queue that supports a conditional pop, ie. we pop the element of the array if it is smaller than another element. For short arrays (of length 1), we use a naive sort. Otherwise, two PopperQueues are merged by repeatedly comparing the first element of each queue  $A$  and  $B$  and appending the smaller one to the result array  $C$ . Note, that since  $A$  and  $B$  are sorted the element put into  $C$  is the smallest element in both arrays. We pop an element from the array which element we just appended to  $C$  – see Algorithm O-Merge 2.

**Theorem 6.** *Sorting of an array of  $n$  elements requires at most  $85n \log n$  C-Ops and a total of  $n \cdot (3 + 560(\log n - 1) + 28 \log^2 n)$  E-Ops.*

---

**Algorithm 2 O-Merge**

---

**Input:** Sorted PopperQueue  $A$  and  $B$  of length  $l$

**Output:** Merged LIFO Queue  $C$

```
1: if  $l = 1$  then
2:    $b := 1$  if  $A[0] \leq B[0]$  else  $0$ 
3:    $C[0] := A[0] \cdot b + B[0] \cdot (1 - b)$ 
4:    $C[1] := B[0] \cdot b + A[0] \cdot (1 - b)$ 
5: else
6:    $eleA := A.pop(1)$  {Returns smallest element in  $A$ }
7:    $eleB := B.pop(1)$ 
8:   for  $k = 0$  to  $2 \cdot l - 2$  do
9:     {Set  $C[k]$  to the smallest element in  $A$  union  $B$  and remove the element}
10:     $b := 1$  if  $eleA \leq eleB$  else  $0$ 
11:     $C[k] := eleA \cdot b + (1 - b) \cdot eleB$ 
12:     $eleA := A.pop(b) \cdot b + (1 - b) \cdot eleA$ 
13:     $eleB := A.pop(1 - b) \cdot (1 - b) + b \cdot eleB$ 
14:   end for
15:    $C[2 \cdot l - 1] := eleA \cdot b + (1 - b) \cdot eleB$ 
16: end if
```

---

*Proof.* The merger of two arrays of size  $l$  each requires  $4l$  pop operations, each requiring 18 comparisons using Theorem 2 with  $q = 2$ . Additionally, we need one more comparison per iteration. This gives a total of 85l S-Ops for merging two arrays. In total we get the following bound  $\sum_{j=0}^{\log n - 1} 2^{\log n - j} \cdot 85 \cdot 2^j \leq 85n \log n$  S-Ops.

The naive sort of two arrays of size one comparing the two elements requires 5 E-Ops. In total there are  $n/2$  queues of length 1 yielding a total of  $5n/2$ . The merger of two arrays of size  $l > 1$  requires each  $4l$  pop operations, each requiring  $28(\log(l/2) + 1) = 28 \log l$  E-Ops.<sup>2</sup>, giving a total of  $112l \log l$ . Additionally, we need 5 E-Ops for each of the  $2l$  operations, giving a total of  $10l$  E-Ops. Overall we get  $l(10 + 112 \log l)$ . Overall, we get

$$\begin{aligned} 5n/2 + \sum_{j=1}^{\log n - 1} 2^{\log n - j - 1} \cdot 112 \cdot 2^j \cdot (\log(2^j) + 10) \\ = 5n/2 + \sum_{j=1}^{\log n - 1} n \cdot 56 \cdot (\log(2^j) + 10) \\ \leq 3n + 56n(\log^2 n/2 + 10(\log n - 1)) \\ = n \cdot (3 + 560(\log n - 1) + 28 \log^2 n) \end{aligned}$$

□

The analysis uses that we merge  $\frac{n}{2^i}$  arrays of length  $2^i$  and Theorem 2 to bound the time to merge two arrays. We improve on [9] by a factor of more than 1000 in terms of the number of comparisons, ie. C-Ops. Comparisons are often used to analyze sorting algorithms, since typically the total operations involved is proportional to the number of comparisons. In our case, this does not necessarily hold, since we only require one comparison for shifting a large number of elements. Therefore, the costs for shifting might dominate the costs for comparisons. To ensure a fair and objective comparison among algorithms we also analyzed the number of other operations, ie. E-Ops, since they are the dominant factor in our algorithm. With respect to the total number of operations O-Mergesort is asymptotically worse by a factor  $\log n$ . However, due to the extremely large constants used in the state-of-the-art [9] we use less operations for all practically relevant scenarios, ie. for arrays of length up to roughly  $2^{5300}$ . For illustration, when sorting 10 billion elements we need more than 100x less E-Ops. Furthermore, E-Ops (or XORs, ANDs) are generally less complex than comparisons, therefore in practice the speed-up might be even larger.

---

<sup>2</sup>We have  $\log(l/2) + 1$  rather than  $\log(l/2)$  using Theorem 2 with  $q = 2$  since we merge arrays of length  $l = 2^x$  but we only support mergers of arrays of length  $2^y - 1$ , thus we need  $y = x + 1 = \log(l/2) + 1$

## 8. Quicksort

Our oblivious quicksort algorithm (O-Quicksort) is a comparison-based divide and conquer algorithm. Small arrays of size at most  $4 \log^2 n$  are sorted using O-Mergesort. Larger arrays are recursively split into two smaller (sub)arrays. An array is split using a pivot element. All elements less or equal to the pivot are put in one array and all larger elements in the other array. Ideally, both arrays are of the same size. However, naive splitting likely leads to badly balanced arrays leading to  $O(n^2)$  run-time since an oblivious algorithm must treat both parts as potentially large. However, when choosing the median as pivot, it is possible to ensure that both arrays are of equal size. We compute an approximate median for all elements (Section 8.1). Unfortunately, choosing an approximate median still leaves some uncertainty with respect to the exact array lengths after the splitting. Therefore, in the partition process (Section 8.1), rather than swapping elements within one array, we create two arrays of fixed length, one for elements larger than the pivot and one for all other elements. Since the length of each of the two arrays must be fixed using conservative upper bounds, their sum of lengths exceeds the length of the array to be split. To get a single sorted array requires a special reunification of both arrays (Section 8.2). For simplicity, we assume that all elements to be sorted are distinct. This assumption is removed in Section 8.3.

### 8.1. Random Pivot Choice and Partition

Algorithm RandomPivot chooses several elements uniformly at random, sorts them and picks the median. By choosing the median of a sufficiently large sample of elements we ensure that the chances of a split resulting in very unbalanced arrays is small. We pick a fixed number of samples  $n_p$ , sort them, eg. using the O-MergeSort algorithm, and then pick the middle element  $l/2$  of the sorted array of length  $l$  as pivot.

---

#### Algorithm 3 RandomPivot

---

**Input:** Array  $A$  of length  $l$ , number of samples  $n_p$

**Output:** Pivot  $p$

- 1:  $P :=$  Set of  $n_p$  elements of  $A$  chosen uniformly at random
  - 2:  $SP :=$  Sorted samples  $P$  {eg. using O-MergeSort}
  - 3:  $p := SP[l/2]$  {Choose middle element (= Median) as pivot}
- 

For the partitioning the entire array is split into two arrays, one with all elements being smaller than the pivot and one with all elements being larger. The two arrays are given by two LIFO queues. We push elements that are smaller than the approximated median on one of the queues and the larger elements on the other queue. We discuss the case of duplicates in Section 8.3.

**Theorem 7.** *Algorithm RandomPivot returns a pivot  $p$  such that at least a fraction  $c_f = \frac{1}{2} \cdot \frac{1}{1 + \sqrt{13c(\log n)/n_p}} \geq 1/4$  of elements of an array of length  $n$  are larger than  $p$  and the same fraction is smaller than  $p$  with probability  $1 - 1/n^c$  for  $n_p \geq 13 \cdot c \cdot \log n$ .*

We obtain tail estimates using carefully applied Chernoff bounds.

**Theorem 8** (Chernoff Bound). *The probability that the number  $X$  of occurred independent events  $X_i \in \{0, 1\}$ , i.e.  $X := \sum_i X_i$ , is not in  $[(1 - c_0)\mathbb{E}[X], (1 + c_1)\mathbb{E}[X]]$  with  $c_0 \in ]0, 1[$  and  $c_1 \in ]0, 1[$  can be bounded by  $p(X \leq (1 - c_0)\mathbb{E}[X] \vee X \geq (1 + c_1)\mathbb{E}[X]) < 2e^{-\mathbb{E}[X] \cdot \min(c_0, c_1)^2/3}$ .*

Proof of Theorem 7:

*Proof.* The value of  $c_f$  is minimized, when  $n_p$  is smallest. Thus, the bound  $c_f \geq 1/4$  follows from substitution of the lower bound, ie.  $n_p = 13 \cdot c \cdot \log n$ , into  $c_f = \frac{1}{2} \cdot \frac{1}{1 + \sqrt{13c(\log n)/n_p}} = 1/2 \cdot \frac{1}{1 + \sqrt{1}} = 1/4$ . The theorem holds if the pivot does not stem from the  $c_f \cdot n$  smallest or largest elements. If we pick less than  $c_f \cdot n_p < n_p/2$  elements  $S \subseteq A$  from the  $c_f \cdot n$  smallest and less than  $c_f \cdot n_p < n_p/2$  elements  $L \subseteq A$  from the  $c_f \cdot n$  largest elements this will be the case. The reason being that the pivot  $p$  is the element at position  $n_p/2$  in the sorted sequence and, thus, it will not be from

360 the set of  $c_f \cdot n$  smallest or largest elements. We expect to pick  $c_f \cdot n_p$  elements  $S$  out of the  $c_f \cdot n$  smallest elements (and analogously for the largest), ie.  $E[|S|] = c_f \cdot n_p$ . We seek the smallest factor  $f > 1$  such that when exceeding the expectation by factor  $f$  the pivot is not chosen correctly. We have  $f \cdot c_f \cdot n_p = n_p/2$ , if  $f = 1/(2 \cdot c_f)$ . The probability that the expectation is exceeded by a factor  $f > 1$  or more is given using a Chernoff bound (see Theorem 8) by

$$\begin{aligned} \text{prob}(|S| > f \cdot E[S]) &< 1/2^{(f-1)^2/3 \cdot c_f \cdot n_p} \leq 1/2^{1/12 \cdot (f-1)^2 \cdot n_p} \text{ (Using } c_f \geq 1/4) \\ &= 1/2^{1/12 \cdot (1/(2 \cdot c_f) - 1)^2 \cdot n_p} \text{ (Using } f = 1/(2 \cdot c_f)) \\ &= 1/2^{1/12 \cdot (1 + \sqrt{13c \log n/n_p})^2 \cdot n_p} = 1/2^{13/12 \cdot c \cdot \log n} = 1/n^{13/12 \cdot c} \end{aligned}$$

In the same manner we can compute  $\text{prob}(|L| > f \cdot E[L])$ . Therefore the probability of both events becomes for  $n$  sufficiently large:

$$\text{prob}((|L| \leq f \cdot E[L]) \wedge (|S| \leq f \cdot E[S])) \geq 1 - (\text{prob}(|L| > f \cdot E[L]) + \text{prob}(|S| > f \cdot E[S])) \geq 1 - 1/n^c$$

365 □




Figure 4: Merger of two subarrays within O-Quicksort

## 8.2. Sorting and End-to-End Array Merger

370 So far we have obtained well-balanced, but unsorted SAs. Since we do not have access to their exact lengths we use a conservative bound on their lengths given by the analysis of the partition process. O-Quicksort recurses on two separate arrays stemming from the partitioning. We sort the array of smaller elements  $B$  than the pivot in ascending order and the array of larger elements  $C$  in descending order. At the end, both arrays must be merged to get a single final array of sorted elements. This requires some care since we do not know their exact lengths. Due to the partitioning process we can bound the minimum length of  $B$  and  $C$  to be  $l/2 \cdot (1 + \epsilon)$  with  $\epsilon := \sqrt{13 \cdot c \cdot (\log n)/l}$ . We copy the elements (up to the guaranteed minimum length bound) to the final array, so that these elements appear sorted. This means we fill the final array with  $B$  from the left end towards the right and with  $C$  from the right end. The

---

**Algorithm 4** O-Quicksort

---

**Input:** Array  $A$  of length  $l$ , Sort ascending:  $asc$

**Output:** Sorted array  $A$

```
1:  $\epsilon := \sqrt{13 \cdot c \cdot (\log n) / l}$ 
2: if  $l \geq 4 \log^2 n$  then
3:    $B, C := \text{Partition}(A)$ 
4:   O-Quicksort( $B, l/2 \cdot (1 + \epsilon), True$ )
5:   O-Quicksort( $C, l/2 \cdot (1 + \epsilon), False$ )
6:   for  $k = 0$  to  $l/2 \cdot (1 - 2\epsilon)$  do
7:      $B[k] := B[k]$  {Copy elements from  $B$  to  $A$ }
8:      $B[l - k] := C[k]$  {Copy elements from  $B$  to  $A$ }
9:   end for
10:  for  $k = l/2 \cdot (1 - 2\epsilon)$  to  $l/2 \cdot (1 + \epsilon)$  do
11:    if  $B[k] \neq \nu_\emptyset$  then  $B[k] := B[k]$  {Copy elements from  $B$  to  $A$ }
12:    if  $C[k] \neq \nu_\emptyset$  then  $B[l - k] := C[k]$  {Copy elements from  $C$  to  $A$ }
13:  end for
14: else
15:    $B :=$  Sort using O-MergeSort or other alg. either ascending or desc. depending on  $asc$ 
16: end if
```

---

375 entire process is illustrated in Figure 4. The remaining elements are handled in the same fashion, but before setting an array element in  $A$  to be an element from  $B$  or  $C$ , we check whether the element in  $A$  is still empty.

**Theorem 9.** *O-Quicksort needs  $O(n \log n)$  C-Ops and E-Ops. It produces a correct sorting with probability  $1 - 1/n^c$  for an arbitrary constant  $c$ .*

380 The recurrences are somewhat involved, since the lengths of both arrays used for recursion exceeds the original length of the array being split. We conduct a staged analysis to obtain (asymptotically tight) bounds.

*Proof.* The complexity  $T(n)$  of O-Quicksort in terms of comparisons can be stated using a recurrence. For one call (ignoring recursion) to an array of length  $l > 4 \log^2 n$  we have that the complexity is given by partitioning the array being  $O(l)$  plus the reunification of both sorted arrays, ie. the copying of elements being also  $O(l)$ . Thus, we get a total of  $O(l) = c_0 \cdot l$  for some constant  $c_0$ . We obtain the following recurrence for an array of length  $l$  using  $\epsilon = \sqrt{a/l}$  with  $a := 13 \cdot c \cdot (\log n)$ :  
385

$$\text{First call: } T(l) = 2T(l/2 \cdot (1 + \sqrt{a/l})) + c_0 \cdot l$$

$$\begin{aligned} \text{Second call: } T(l/2 \cdot (1 + \sqrt{a/l})) &\leq 2T(l/4 \cdot (1 + \sqrt{a/l}) \cdot (1 + \sqrt{a/(l/2)})) + c_0 \cdot l/2 \cdot (1 + \sqrt{a/l}) \\ &\leq 2T(l/4 \cdot (1 + \sqrt{a/(l/2)}))^2 + c_0 \cdot l/2 \cdot (1 + \sqrt{a/l}) \end{aligned}$$

$$\text{Third call: } T(l/4 \cdot (1 + \sqrt{a/(l/2)}))^2 \leq 2T(l/8 \cdot (1 + \sqrt{a/(l/4)}))^3 + c_0 \cdot n/8 \cdot (1 + \sqrt{a/(l/4)})^3$$

$$i\text{-th call: } T(l/2^i \cdot (1 + \sqrt{a/(l/2^{i-1})}))^i \leq 2T(l/2^{i+1} \cdot (1 + \sqrt{a/(l/2^i)}))^{i+1} + c_0 \cdot l/2^i \cdot (1 + \sqrt{a/(l/2^i)})^{i+1} \quad (1)$$

Assume we start splitting the entire array  $A$  with  $l = n$ . The total number of operations (C-Ops and E-Ops) at recursion depth  $i$  is given by the additive term in Equation (1) multiplied by the number of calls to O-Quicksort being  $2^i$ , ie.

$$2^i c_0 \cdot n/2^i \cdot (1 + \sqrt{a/(n/2^i)})^{i+1} = c_0 \cdot n \cdot (1 + \sqrt{a/(n/2^i)})^{i+1}$$

390 The total operations for the first  $r := \log n - 8 \log \log n$  recursions is given by:

$$\sum_{i=0}^{r-1} c_0 \cdot n \cdot (1 + \sqrt{a/(n/2^i)})^{i+1} \leq c_0 \cdot n \cdot \sum_{i=0}^{r-1} (1 + \sqrt{a/(\log n^8)})^{\log n} \leq c_0 \cdot n \cdot \log n$$

After  $r$  recursions the size of the input sequence for the recursive calls is at most  $n/2^r \cdot (1 + \sqrt{a/(n/2^{r-1})})^r \leq 2 \cdot \log^8 n$  (for  $n$  sufficiently large). For another  $6 \log \log n$  recursions on an array of length  $2 \log^8 n$  the number of operations is bounded by:

$$c_0 \cdot 2 \log^8 n \cdot \sum_{i=0}^{6 \log \log n - 1} \left(1 + \sqrt{\frac{4 \log n}{(\log n)^2}}\right)^{6 \log \log n} = c_0 \cdot 2 \log^8 n \cdot \sum_{i=0}^{6 \log \log n - 1} \left(1 + \frac{2}{\sqrt{\log n}}\right)^{6 \log \log n} \leq 4c_0 \log^8 n$$

The size of the remaining arrays is  $4 \log^2 n$  using the same derivation as above using  $r$  recursions. To sort such an array using O-MergeSort requires  $O(\log^2 n \log \log n)$  C-Ops and  $O(\log^2 n (\log \log n)^2)$  E-Ops (see Theorem 6). There are  $2^{\log n - 2 \log \log n} = n / \log^2 n$  such arrays, giving a total of  $O(n \log \log n)$  C-Ops and  $O(n (\log \log n)^2)$  E-Ops. To obtain a correctly sorted queue all executions of RandomPivot must be successful. We perform at most  $\log n - 2 \log \log n$  recursions. Thus, in total there are at most  $n$  calls to RandomPivot, each succeeding with probability at least  $1 - 1/n^{c'}$  for an arbitrary constant  $c'$ . The probability that all succeed is at least  $(1 - 1/n^{c'})^n \geq 1 - 1/n^{c'-2}$ . Choosing  $c' = c + 2$  concludes the proof.  $\square$

### 8.3. Equal or Duplicate Elements

So far we focused on arrays of distinct elements. For non-distinct elements our algorithm can fail to compute balanced arrays in case the chosen median is not unique. In the most extreme case all elements are the same and the split would result in one empty array and one array being the same as the array to be split. Elements can always be made distinct by appending a unique number at the end, eg. by appending a counter to a array of elements  $(0, 0, 0)$ , we get  $(00, 01, 02)$ . We assign elements that are equal to the pivot  $p$  to both arrays such that their lengths maintain balanced. In a first phase we create two arrays  $B, C$  and maintain a counter  $l_p$  for the elements equal to  $p$  by distinguishing three cases for an element  $x$  that is compared to the pivot  $p$ , ie.  $x < p$ ,  $x > p$  and  $x = p$ . In the first case, we assign  $x$  to array  $B$  and increment the length counter of  $l_B$ . In the second case we assign  $x$  to  $C$  and increment the length counter  $l_C$  of  $C$ . In the third case, we increment just the counter  $l_p$ . In the second phase we distribute  $l_p$  copies of  $p$  to the arrays  $B$  and  $C$  such that their difference in length is as small as possible. We perform  $l$  iterations, where  $l$  is the number of elements in the array to be partitioned, ie.  $A$ . In each iteration we subtract one from  $l_p$ . If  $l_p$  is zero the arrays remain the same. Otherwise, if the lengths of  $l_B$  is less than  $l_C$ , we append a copy of the pivot to  $B$  and increment the length counter  $l_B$  otherwise we do the same for  $C$ . The (asymptotic) complexity remains the same.

## 9. Applications

### 9.1. Confidential Expressions

Confidential expressions, hiding data as well as operations on the data are a rather straight forward application of our LIFO queue with conditional operations as well as basic operations, e.g. addition and multiplication, from secure computing schemes such as fully homomorphic encryption (FHE) or secure multi-party computation (MPC). We first discuss the evaluation of non-confidential expressions. For brevity we only discuss the evaluation of expressions involving numbers, additions and multiplications. We focus on evaluating postfix expressions. When using encrypted values, the expression remains confidential. That is to say, despite computing the expression we do not learn anything about the expression except its length. The key to achieve this is the conditional push, ie. we execute a push operation but it only has an impact given that the element to be pushed is different from the special element  $\emptyset$  that is not appended to the stack. Our algorithm 5 requires linear run-time in terms of the number of elements in the expression (or, more precisely, in the bound we get for the length of the expression).

To evaluate confidential expressions, all array elements of the input  $A$  must be encrypted as well as variables that depend on the array elements. Variables that indicate array lengths or the number of operations do not have to be encrypted.<sup>3</sup> To this end, one can use any of the known scheme for computing on encrypted data such as FHE or MPC, eg. [5] or [20]. These schemes provide basic operations such as addition and multiplication that allow to construct other operations such as comparisons, subtractions and more. However, certain operations like accessing an array

<sup>3</sup>In Algorithm 5 the values of  $n$  and  $s$  do not have to be encrypted. In the LIFO queue and its sub-procedures,  $n_{pu}$ ,  $n_{po}$ ,  $q, o, mi$  and  $s$  remain unencrypted. All other variables are encrypted.

---

**Algorithm 5 Case Study: PostFix expressions**

---

**Input:** LIFO Queue  $A$  of postfix symbols of length at most  $n$

**Output:** Result of evaluation

```
1:  $st := LIFO(s)$  {Choose number of SAs  $s$  such that array can hold at least  $n$  elements}
2: for  $i := 0$  to  $n - 1$  do
3:    $symb := A.pop(1)$ 
4:    $toPush := symb$  if  $symb$  is a number else  $\emptyset$ 
5:    $st.push(toPush)$ 
6:    $isAdd := 1$  if  $symb = '+'$  else  $0$ 
7:    $resAdd := st.pop(isAdd) + st.pop(isAdd)$ 
8:    $isMul := 1$  if  $symb = '*'$  else  $0$ 
9:    $resMul := st.pop(isMul) \cdot st.pop(isMul)$ 
10:   $toPush := isMul \cdot resMul + (1 - isMul) \cdot \emptyset$ 
11:   $toPush := isAdd \cdot resAdd + (1 - isAdd) \cdot toPush$ 
12:   $st.push(toPush)$ 
13: end for
14: return  $st.pop(1)$ 
```

---

element using an encrypted index might occur linear overhead in the length of the array. In our case, we manage to keep the asymptotic running time, since we only have to directly substitute additions, subtractions, multiplication and comparisons operations.

## 435 9.2. Stock Span Analysis

The Stock Span Problem is motivated by financial analysis of stocks. The span of a stock's price on day  $d$  is the maximum number of consecutive days until day  $d$ , where the price of the stock has been at most its price on  $d$ . The well-known text book solution is given in Algorithm 6 taking linear time in the number of prices  $n$ .<sup>4</sup> A straight forward solution gives a quadratic run-time algorithm due to the nested loops, ie. due to the worst case behavior of the inner loop. This renders the solution impractical for larger datasets. The total number of iterations (when being summed across all outer loop iterations) of the inner loop is only  $n$ . A single iteration of the inner loop could perform all  $n$  iterations for some inputs. To ensure obliviousness we would have to execute (asymptotically)  $n$  iterations of the inner loop for every execution of the outer loop. Furthermore, the code contains direct array access, eg.  $price[i]$ . In the obvious manner, this would also incur linear run-time overhead. However, it is possible to transform the nested loop by essentially changing the inner loop to an 'if'-conditional first without changing the number of iterations of the outer loop. Then we make the loop oblivious using a conditional expression `if-then-else`. Essentially, in Algorithm 6 we replace the `while` and `do` keyword in line 5 by an `if` and `then`. Lines 6 to 8 form the `else` part. We only show the final pseudocode after the translation of the 'if' into oblivious code in Algorithm 7. Since we must execute both branches of the `if` to keep the condition confidential, the algorithm requires that we can execute the 'pop' operation without impacting the data structure, ie. without actually performing a pop. This is supported by our data structure by using a special element in case the condition evaluates to true. The algorithm uses a peek operation, which returns the first element without removing it. It can be implemented using a combination of pop and push operation, eg.  $x := pop(1), push(x)$ .

## 10. Related Work

455 In 2009 Gentry [5] introduced a fully homomorphic encryption(FHE) scheme based on lattices. Since then the field of computing on encrypted data (and circuits) has evolved rapidly, as summarized in [15]. All approaches for FHE are based on either addition and multiplication or XOR and AND. Secure computations can also be carried out using multiple parties, such that no party learns a secret (if it does not collude with other parties). Secure multi-party computation was introduced three decades ago [25, 7] and is still subject to extensive research, eg. [20]. Both

---

<sup>4</sup>We adjusted it from <http://www.geeksforgeeks.org/the-stock-span-problem/>



---

**Algorithm 6 Case Study: Stock Span**

---

**Input:** LIFO Queue  $price$  of prices of length at most  $n$

**Output:** LIFO  $S$  with spans (in reverse order)

```
1:  $st := LIFO(s)$  {Choose number of SAs  $s$  such that array can hold at least  $n$  elements}
2:  $st.push(0)$ 
3:  $S.push(1)$  {Span first element is 1}
4: for  $i = 0$  to  $n - 1$  do
5:   while  $st.peek() \neq \emptyset \wedge price[st[0]] \leq price[i]$  do  $st.pop()$ 
6:    $span := i + 1$  if  $st.peek() \neq \emptyset$  else  $i - st[0]$ 
7:    $S.push(span)$ 
8:    $st.append(i)$ 
9: end for
```

---

---

**Algorithm 7 Case Study: Oblivious Stock Span**

---

**Input:** LIFO Queue  $price$  of prices of length at most  $n$

**Output:** LIFO  $S$  with spans (in reverse order)

```
1:  $st := LIFO(s)$  {Choose number of SAs  $s$  such that array can hold at least  $n$  elements}
2:  $pi := price.pop(1)$ 
3:  $st.push((0, pi))$ 
4:  $S.push(1)$  {Span first element is 1}
5:  $i := 0$ 
6: for  $k := 0$  to  $n - 1$  do
7:    $(sti, stp) := st.pop(1)$ 
8:    $popNext := 1$  if  $sti = \emptyset$  or  $sti \leq pi$  else 0
9:    $pi := price.pop(popNext) \cdot popNext + (1 - popNext) \cdot pi$ 
10:   $i := i + (1 - popNext)$ 
11:   $span := i + 1$  if  $st.peek() \neq \emptyset$  else  $i - sti$ 
12:   $span := \emptyset$  if  $popNext$  else  $span$ 
13:   $S.push(span)$ 
14:   $pushi := (i, pi)$  if  $(1 - popNext)$  else  $(\emptyset, \emptyset)$ 
15:   $st.push(pushi)$ 
16: end for
```

---

460 SMC and FHE can compute expressions beyond addition and multiplication such as comparisons. These operations could be used as black boxes to make our algorithms work on encrypted data. [5, 23] mentioned that circuit privacy is achievable by adding a (large) noise vector of the encrypted 0. The original work on SMC [25] already allowed to hide circuits using garbled circuits. Our approach also allows to achieve circuit privacy in a novel manner by hiding whether a certain operation really impacted the computation. Our work is not limited to circuits. “Oblivious RAM” (ORAM) [6] disguises access patterns by a client from a server using randomization. The original solution [6] is based on a hierarchy of buffers such that each level of the hierarchy consists of several blocks. One block per level is read and always written on the first level. For each level, some blocks are real and some are dummy containing random data. The original algorithm has been improved, eg. in a recent tree-based scheme [21] each data block is stored somewhere in the tree, ie. following a specific path from the root. After an access minor reordering involving the accessed elements takes place, potentially, resulting in some data being sent to the client. Some schemes, eg. [21, 24], trade-off performance (or memory consumption) and the risk for memory access errors. Oblivious data structures for ORAM covering arrays and queues are discussed in [24]. They make use of data access locality encountered for arrays and queues. A single access to a parent node in the position map returns pointers to multiple children. They hide whether the operation is a read or write to a memory cell. However, assuming one knows that a write must occur in a function, one knows that some memory cell is modified. We do not use traditional ORAM techniques. Furthermore, in our scenario, knowing that a certain operation is performed, ie. a pop or push, still gives no hint whether the data structure was modified or not.

Other work designed oblivious data structures particularly for SMC, eg. [12, 22]. The work [12] uses ORAM structures and secret sharing among parties to achieve obliviousness. In contrast, [22] presents a deterministic scheme for priority queues using the bucket heap concept for priority queues [3] coming with  $O(\log^2 n)$  overhead. Bucket heaps partition memory into buffers of size  $2^{2^{i+1}}$  and signal blocks of size  $2^{2^i}$  [3]. Buckets store actual elements, whereas buffers store overflowing elements. Once a buffer is full it is merged into a bucket. [22] adjusted this setting to use blocks of equal size. Our queue shares the common idea of organizing data in blocks of increasing size that is also found in other work, eg. [14]. We differ from prior work [14, 22, 3] in several aspects, eg. we perform a more fine-grained partitioning using multiple blocks, eg. in the view of [22] we introduce one more level of partitioning for buffer and data blocks. We have also come up with a deterministic oblivious sequence of restructuring operations to handle empty and full blocks rather than counting the number of elements in the queue, eg. [22]. In contrast to our work, prior work also does not hide the impact of an operation (ie. they do not hide the number of elements in a bucket), which is essential for securing control structures. Our fast B2B-FIFO queue introduces novel ideas such as block sharing not found in prior work.

The paper [1] shows how to compute the  $k$ -th ranked element for SMC. The paper [13] discusses median computation. Such operations might prove valuable also for sorting, eg. selecting the median element for quicksort. However, both protocols [1, 13] disclose the outcome of secure comparisons, which might require non-desirable client interaction and is not considered secure. The SMC protocol for sorting in [26] runs in constant rounds but needs to know the product of the range of numbers  $R$  and it has communication and computational complexity that is proportional to product of the range of numbers times the number of elements, ie.  $O(n \cdot R)$  (an improved version has  $O(n^2)$ ). To achieve constant rounds it relies on the evaluation of unbounded fan-in gates.


Sorting networks are naturally oblivious, since they use a fixed sequence of comparisons among elements in an array that is not related to the data stored in the array. They have been used for secure sorting in [11, 9]. The work [9] is based on a network with  $19600 \cdot n \log n$  comparators. A comparator can be implemented by a comparison yielding a bit  $b$  followed by an exchange of two elements  $A, B$ , ie.  $A := b \cdot B + (1 - b) \cdot A$  and  $B := b \cdot A + (1 - b) \cdot B$ . Therefore a comparator needs 7 E-Ops in addition to the comparison, yielding  $156800 \cdot n \log n$  operations. Though this is asymptotically optimal, it is of little practical relevance due to the number of comparators needed. Additionally, the depth (of the sorting network) is of order  $n \log n$ , which makes it non-parallelizable. Our algorithms improve on it for all relevant scenarios (see Section 7 for a detailed comparison). The oblivious randomized Shellshort algorithm [8] is asymptotically optimal in terms of the number of comparisons using several techniques such as permutation of the array as well as shaker and brick passes of the array.

Oblivious algorithms for geometric problems are presented in [4]. Algorithms for graphs incurring overhead up to linear factor (in the number of nodes) are given in [2]. Other work [18] based on ORAM designed data structures for maps. They allow for history independence, ie. different sequences of operations lead to indistinguishable (memory layouts of the physical) data structures.


## 11. Evaluation

We shed light on two aspects that are not immediate from the asymptotic analysis. First, on the one hand our oblivious data structures are more involved than using a naive oblivious implementation traversing the entire array for each operation, on the other hand we have shown that asymptotically it outperforms a naive implementation. The key question is, whether our oblivious queues outperform already for queues of small capacity or only for those with large capacity. Therefore, we compared our implementation against a simple ‘linear’ oblivious queue that accesses all elements (that could be stored) for each operation. Thus, the run-time is linear in the capacity. Second, how much slower is our array compared to a non-oblivious queue. We have shown that the asymptotic factors are of order  $O(\log n)$  and  $O(\log^2 n)$  depending on the queue type. Here, we give more precise insights.

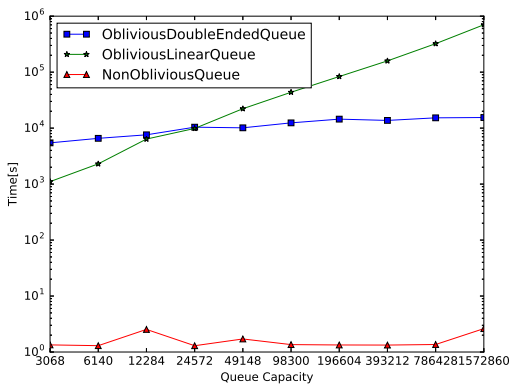
We implemented the benchmarks in Python. The evaluation was run on a machine equipped with an Intel 2.5 GHz quad-core CPU with 8 GB RAM on Windows 7. For the non-oblivious queue we ran 1 Mio. operations. For the oblivious linear queue, ie. the naive oblivious queue traversing the entire array for each operation, we attempted to run 1 Mio operations, but stopped after 1 hour if the computation was still ongoing and estimated the time it would take to compute 1 Mio. operations. For the oblivious data structures we executed  $\max(100000, 2 \cdot \text{capacity})$  operations, since the maximal run-time is achieved if we execute a multiple of the capacity. Each operation was chosen randomly among a push and pop operation. Due to the obliviousness it does not matter what parameters we use for the push and pop operation.




(a) FIFO Queues



(b) LIFO Queues



(c) Double-Ended Queues



(d) Double-Ended Priority Queues

Figure 5: Running Times results for 1 Mio. operations for our oblivious queues compared to linear oblivious and non-oblivious queues

The plots in Figures 5a, 5b, 5c and 5d show the run-times comparing all queue variants for increasing maximum

530 queue sizes for FIFO and LIFO queues. Qualitatively all queues behave similarly as predicted by the asymptotic  
analysis. For small queue sizes (LIFO and FastFIFO up to about 60, FIFO up to about 500) a simple linear oblivious  
queue has an edge over our more complex queues. For double-ended queues performance is somewhat worse, but  
simple linear queues are also outperformed for moderate queue sizes. With growing queue sizes the exponential gap  
535 becomes clearly visible between the linear oblivious queue and our implementations. The LIFO and fast FastFIFO  
queue are more than 100x faster for queues of capacity about 10000. For FIFO queues we reach the boundary of  
100x performance improvement for queues of capacity about 100000. Note, that it is not uncommon that a simple  
algorithm with bad asymptotic behavior outperforms more complex algorithms. For example, a naive bubble sort  
might also be faster than mergesort for small data sizes although the later is asymptotically much faster.

540 When it comes to overhead of our oblivious queue compared to the built-in Python queue (surrounded by a  
wrapper), which uses memory proportional to the actual array size, the asymptotic behavior is well-visible. Our  
LIFO and FastFIFO queue both have an asymptotic overhead of  $\log n$  compared to non-oblivious queues that directly  
accesses queue elements. This results in close to parallel lines in Figures 5b and 5b. The overhead is roughly a factor  
40 for queues of size 10000. For FIFO queues the asymptotic overhead is larger, ie.  $\log^2 n$ . The overhead is a factor  
545 of 200 for arrays of the same size. In the light of overhead that typically comes with secure computation, eg. FHE or  
SMC that can reach more than 5-6 orders of magnitude [16], our overhead is very modest. We also want to emphasize  
that to the best of our knowledge no prior work has compared against a non-oblivious implementation. We believe  
this is a very important benchmark.

## 12. Conclusions

550 We have presented oblivious queues accompanied by theoretical and practical investigation having only  
(poly)logarithmic overhead. Since queues are an essential part of everyday programming, we believe that they will  
play a major role for enabling computation using encrypted data, in particular with focus on expression hiding. Still,  
many more common data structures and operations have to be realized efficiently before any of the existing technolo-  
gies such as FHE and MPC become practical for a large range of applications

## 555 13. References

- [1] G. Aggarwal, N. Mishra, and B. Pinkas. Secure computation of the median (and other elements of specified ranks). *Journal of cryptology*, 23(3):373–401, 2010.
- [2] M. Blanton, A. Steele, and M. Alisagari. Data-oblivious graph algorithms for secure computation and outsourcing. In *Proceedings of the 8th ACM SIGSAC symposium on Information, computer and communications security*, pages 207–218. ACM, 2013.
- 560 [3] G. S. Brodal, R. Fagerberg, U. Meyer, and N. Zeh. Cache-oblivious data structures and algorithms for undirected breadth-first search and shortest paths. In *Algorithm Theory-SWAT 2004*, pages 480–492. 2004.
- [4] D. Eppstein, M. T. Goodrich, and R. Tamassia. Privacy-preserving data-oblivious geometric algorithms for geographic data. In *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 13–22. ACM, 2010.
- [5] C. Gentry. Fully Homomorphic Encryption Using Ideal Lattices. In *Proc. 41st ACM Symposium on Theory of Computing*, pages 169–178, 2009.
- 565 [6] O. Goldreich. Towards a theory of software protection and simulation by oblivious rams. In *Proc. of 19th Symposium on Theory of computing(STOC)*, pages 182–194, 1987.
- [7] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *Proc. of 19th Symp. on Theory of computing*, pages 218–229, 1987.
- 570 [8] M. T. Goodrich. Randomized shellsort: A simple oblivious sorting algorithm. In *Proceedings of the 21st Symposium on Discrete Algorithms(SODA)*, pages 1262–1277, 2010.
- [9] M. T. Goodrich. Zig-zag sort: A simple deterministic data-oblivious sorting algorithm running in  $o(n \log n)$  time. In *Proc. of 46th Symp. on Theory of Computing*, pages 684–693, 2014.
- [10] M. S. Islam, M. Kuzu, and M. Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In *NDSS*, volume 20, page 12, 2012.
- 575 [11] K. V. Jonsson, G. Kreitz, and M. Uddin. Secure multi-party sorting and applications. *IACR Cryptology ePrint Archive*, 2011:122, 2011.
- [12] M. Keller and P. Scholl. Efficient, oblivious data structures for MPC. pages 506–525, 2014.
- [13] Y. Lindell and B. Pinkas. Secure multiparty computation for privacy-preserving data mining. *Journal of Privacy and Confidentiality*, 1(1):5, 2009.
- 580 [14] J. C. Mitchell and J. Zimmerman. Data-oblivious data structures. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 25. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2014.

- [15] C. Moore, M. O’Neill, E. O’Sullivan, Y. Doroz, and B. Sunar. Practical homomorphic encryption: A survey. In *Int. Symp.on Circuits and Systems (ISCAS)*, pages 2792–2795, 2014.
- 585 [16] M. Naehrig, K. Lauter, and V. Vaikuntanathan. Can homomorphic encryption be practical? In *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*, pages 113–124. ACM, 2011.
- [17] B. Pinkas and T. Reinman. Oblivious RAM revisited. In *Advances in Cryptology (CRYPTO)*, pages 502–519. 2010.
- [18] D. S. Roche, A. J. Aviv, and S. G. Choi. A practical oblivious map data structure with secure deletion and history independence. *arXiv preprint arXiv:1505.07391*, 2015.
- 590 [19] J. Schneider. Lean and fast secure multi-party computation: Minimizing communication and local computation using a helper. *13th Int. Conf. on Security and Cryptography(SECRYPT)*, 2016.
- [20] J. Schneider. Lean and fast secure multi-party computation: Minimizing communication and local computation using a helper. *SECRYPT*, 2016, extended version: arXiv:1508.07690, <https://arxiv.org/abs/1508.07690>.
- [21] E. Stefanov, M. Van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas. Path oram: An extremely simple oblivious RAM protocol. In *Proc. of the SIGSAC conference on Computer & communications security*, pages 299–310, 2013.
- 595 [22] T. Toft. Secure data structures based on multi-party computation. In *Proceedings of the 30th annual symposium on Principles of distributed computing*, pages 291–292, 2011.
- [23] M. Van Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan. Fully homomorphic encryption over the integers. In *Advances in cryptology–EUROCRYPT 2010*, pages 24–43. 2010.
- [24] X. S. Wang, K. Nayak, C. Liu, T. Chan, E. Shi, E. Stefanov, and Y. Huang. Oblivious data structures. In *Proc. of the Conference on Computer and Communications Security*, pages 215–226, 2014.
- 600 [25] A. C.-C. Yao. How to generate and exchange secrets. In *Foundations of Computer Science(FOCS)*, 1986.
- [26] B. Zhang. Generic constant-round oblivious sorting algorithm for MPC. In *Provable Security*, pages 240–256. 2011.