



City Research Online

City St George's, University of London

Citation: Aksenov, V., Koval, N., Kuznetsov, P. & Paramonov, A. (2024). Memory Bounds for Concurrent Bounded Queues. In: UNSPECIFIED (pp. 188-199). ACM. ISBN 979-8-4007-0435-2 doi: 10.1145/3627535.3638497

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/36665/>

Link to published version: <https://doi.org/10.1145/3627535.3638497>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Memory Bounds for Bounded Queues

Anonymous Author(s)

Abstract

Concurrent data structures often require additional memory for handling synchronization issues in addition to memory for storing elements. Depending on the amount of this additional memory, implementations can be more or less *memory-friendly*. A *memory-optimal* implementation enjoys the minimal possible memory overhead, which, in practice, reduces cache misses and unnecessary memory reclamation.

In this paper, we discuss the memory-optimality of non-blocking *bounded queues*. Essentially, we investigate the possibility of constructing an implementation that utilizes a pre-allocated array to store elements and constant memory overhead, e.g., two positioning counters for `enqueue(...)` and `dequeue()` operations. Such an implementation can be readily constructed when the ABA problem is precluded, e.g., assuming that the hardware supports LL/SC instructions or all inserted elements are distinct. However, in the general case, we show that a memory-optimal non-blocking bounded queue incurs linear overhead in the number of concurrent processes. These results not only provide helpful intuition for concurrent algorithm developers but also open a new research avenue on the memory-optimality phenomenon in concurrent data structures.

1 Introduction

Developing concurrent algorithms is complex, and limiting synchronization is crucial in achieving high performance. Mutual exclusions allow building concurrent implementations on top of standard sequential algorithms at the expense of scalability. *Non-blocking* synchronization has the potential to improve scalability, but it often requires allocating extra memory. The classic Michael-Scott queue [13] serves as an illustrative example, allocating a new node for each enqueue and, therefore, using a significant amount of memory for node objects; this design results in increased memory usage and cache misses. Modern queues [5, 14, 22] incorporate several elements into a single node, significantly reducing memory allocation and creating a more compact memory structure that mitigates cache misses. We consider these solutions more *memory-friendly*, as they allocate less additional memory. The intuition is that the more memory-friendly the implementation is, the better and more predictable performance it usually provides under both high and low contention. This is especially important for general-purpose solutions in standard libraries of programming languages and popular frameworks.

In this paper, we raise the question of *memory-optimality*: **what is the absolute minimum of memory overhead a concurrent data structure must incur?** We narrow our

attention to *bounded queues*, ubiquitous in resource management systems and task schedulers. Bounded queues allow for a straightforward definition of memory overhead: the amount of memory that must be allocated on top of the fixed memory required for storing the queue elements.

A straightforward sequential bounded queue implementation maintains an array of size C (the queue capacity) equipped with two counters that track the total numbers of `enqueue(...)` and `dequeue()` invocations; their values modulo C point to the next working slots of the operations. See




Figure 1. Sequential bounded queue on top of an array of C slots and two counters.

Figure 1 for an illustration. This implementation utilizes exactly C array slots to store elements and two memory locations for the counters, producing the $\Theta(1)$ memory overhead.

Can we build a *concurrent* bounded queue with the same memory footprint? A trivial solution that utilizes coarse-grained locking incurs constant overhead required for the lock but is inefficient under high load. To the best of our knowledge, most fast and scalable concurrent queues are *non-blocking*, e.g., [14, 15, 17, 22]. While the classic Michel-Scott queue [13] algorithm allocates a new node for each element, these modern solutions pack multiple elements into a single node, making the implementation more memory-friendly. Nevertheless, the memory overhead remains linear in the number of elements.

For *bounded queues*, it is natural to utilize a pre-allocated array of C slots to store elements. This raises an interesting question: is it possible to design an algorithm where the memory overhead depends on neither the queue capacity nor the number of concurrent processes? Can we architect a non-blocking bounded queue using just an array for elements and two positioning counters for `enqueue(...)` and `dequeue()`, similarly to the sequential implementation?

Our contributions. We first observe that the primary challenge in building a concurrent bounded queue is the ABA problem. Intuitively, it may lead to `dequeue()` removing an element from the middle of the queue, assuming it is the first one. When the ABA problem is precluded, e.g., when all inserting elements are distinct or when using the LL/SC synchronization primitives instead of CAS, we show that it is possible to design an algorithm with *constant* memory overhead (Section 2).

In the general case, we prove that any *obstruction-free* [7]¹ bounded queue implementation must use additional $\Omega(T)$

¹Obstruction-freedom is known to be the weakest non-blocking progress condition.

memory locations for synchronization, where T is the number of processes (Section 3). To show that, we construct a non-linearizable execution for any algorithm that utilizes fewer memory locations. The lower bound appears to be tight: we present an algorithm with $O(T)$ memory overhead.

Practical impact. Our experience shows that many engineers have tried to build a concurrent bounded queue using only a pre-allocated array of the queue capacity and positioning counters and failed with that. Our work shows the impossibility of constructing such an algorithm and can potentially save a tremendous amount of time and mental energy for those willing to do that.

Notably, most modern *unbounded* concurrent queues pack multiple elements into fixed-capacity segments and try to reuse them [14, 15, 17]. These segments, also known as *ring buffers*, essentially are bounded queues with slightly relaxed semantics, allowing enqueue($\dots$) to fail spuriously and close the segment for further additions. To reuse the segments, they equip each slot with a 64-bit epoch value, leading to $\Theta(C)$ memory overhead, where C is the segment capacity. Our results do not prove but provide an intuition that it is impossible to implement such a reusable ring buffer structure with constant memory overhead.

Theoretical impact. This work opens a wide research avenue on memory-friendly and memory-optimal concurrent computing. It provides theoretical value by setting strict bounds on memory overhead in dynamic data structures and enabling intuition on the minimal memory overhead algorithms can achieve.

2 Memory Overhead: The Intuition

We start by developing intuition on memory-friendliness and the challenges in attaining constant memory overhead. To illustrate the latter, we present solutions that work under specific assumptions, thereby shedding light on the fundamental obstacles.

2.1 Intuitive Memory-Friendly Algorithm

We first present a *memory-friendly* bounded queue algorithm. While it is not *memory-optimal*, which we further show in Section 3, it can be tuned to be more or less memory-friendly, thus, developing an intuition on memory-friendliness.

The idea is to build a bounded queue on a conceptually infinite array with head and tail counters directly pointing to the array slots. Similar to the Michael-Scott queue design [13], enqueue($\dots$) first installs the element in the next empty slot, followed by the tail counter increment. Similarly, dequeue() begins by extracting the first element, incrementing its tail counter after that. Listing 1 presents the algorithm. As each enqueue-dequeue pair manipulates a unique array slot, the same element cannot be installed into the same slot, so the ABA problem is naturally eliminated.

```

1 fun enqueue(e: E): Bool = while(true) {
2   // Read the counters snapshot
3   t := tail; h := head
4   if t != tail: continue
5   // Is the queue full?
6   if t == h + C: return false
7   // Try to insert the element
8   done := CAS(&a[t],  $\perp$ , e)
9   // Increment the counter
10  CAS(&tail, t, t + 1)
11  if done: return true // finish on success
12 }
13
14 fun dequeue(): E? = while(true) {
15   // Read the counters snapshot
16   t := tail; h := head
17   if t != tail: continue
18   // Is the queue empty?
19   if h == t: return  $\perp$ 
20   // Try to extract an element
21   e := a[h]
22   done := e !=  $\perp$  && CAS(&a[h], e,  $\perp$ )
23   // Increment the counter
24   CAS(&head, h, h + 1)
25   if done: return e // finish on success
26 }

```

Listing 1. Memory-friendly bounded queue of capacity C implemented on an infinite array.

Infinite array. We suggest implementing the infinite array as a concurrent linked list of fixed-size segments, following the approach proposed earlier for synchronous queues and channels in Kotlin Coroutines [11, 12]. All segments are marked with a unique id and follow each other; Figure 2 illustrates the high-level structure. To access the i -th cell in the infinite array, we have to find the segment with the id = i / K (adding it to the linked list if needed) and go to cell $i \% K$ in it.

Memory overhead. Each segment has a memory overhead of $\Theta(1)$ to store its id and a pointer to the next segment. However, we still have to address the memory reclamation problem: once a segment no longer contains queue elements, it should be recycled, and the memory it occupies should be




Figure 2. High-level structure of the memory-friendly bounded queue based on a conceptually infinite array implemented via a concurrent linked list of fixed-size segments.

released. To make the memory overhead bounded, we suggest reusing segments by applying the technique to reclaim descriptors [1]. This approach requires $\Theta(T)$ additional segments on top of the ones that store values, where T is the number of concurrent processes. With storing K elements in each segment, the overall solution exhibits $\Theta(C/K + T \cdot K)$ memory overhead. Though not memory-optimal, this simple algorithm can be tuned to be more or less memory-friendly by adjusting the segment size K , achieving the minimum of $\Theta(T \cdot \sqrt{C})$ when choosing $K = \sqrt{C}$.

2.2 Constant Overhead with Distinct Elements

We now describe an algorithm that uses only $O(1)$ additional memory under a few reasonable assumptions. First, we require all inserting elements to be distinct, which is common in practice. This eliminates the *ABA problem* [8]. The second assumption is that the system is provided with an unlimited supply of versioned $\perp$ (null) values, so we can replace unique elements with unique $\perp$ -s on extractions. This can be achieved by stealing one bit from addresses (values) to mark them as $\perp$ -s and using the rest to store the $\perp$ version.

```

1 fun enqueue(e: E): Bool = while(true) {
2   // Read the counters snapshot
3   t := tail; h := head
4   if t != tail: continue
5   // Is the queue full?
6   if t == h + C: return false
7   // Try to insert the element
8   round := t / C; i := t % C
9   done := CAS(&a[i],  $\perp_{\text{round}}$ , e)
10  // Increment the counter
11  CAS(&tail, t, t + 1)
12  // Finish on success
13  if done: return true
14 }
15
16 fun dequeue(): E? = while(true) {
17   // Read counters+element snapshot
18   t := tail; h := head; e := a[h % C]
19   if t != tail: continue
20   // Is the queue empty?
21   if t == h: return  $\perp$ 
22   // Try to extract the element
23   round := h / C + 1; i := h % C
24   done := e !=  $\perp_{\text{round}}$  && CAS(&a[i], e,  $\perp_{\text{round}}$ )
25   // Increment the counter
26   CAS(&head, h, h + 1)
27   // Finish on success
28   if done: return e
29 }

```

Listing 2. Bounded queue algorithm with $O(1)$ additional memory that requires elements to be distinct and an unlimited supply of versioned $\perp$ (null) values.

Essentially, both `enqueue()` and `dequeue()` (1) read the counters in a snapshot manner, (2) try to perform a “round-valid” update (`enqueue()` replaces $\perp_{\text{round}}$ with the element while `dequeue()` does the opposite, where $\text{round} = \text{counter} / C$), and (3) increase the counter via CAS. The pseudocode is presented in Listing 2.

2.3 Constant Overhead with LL/SC

Another way to avoid the ABA problem is to use the LL/SC (*load-link/store-conditional*) ABA-immune synchronization primitives instead of CAS. In this algorithm (Listing 3), both `enqueue()` and `dequeue()` begin by reading a snapshot of the counters and the corresponding array slot — this part remains unchanged. However, to update the slot state, we now use the SC primitive instead of CAS. By inserting the LL instruction before reading the snapshot, we can ensure that the update fails if the cell has been changed.

While LL/SC are not available in most programming languages, they can be used when programming for ARM, PowerPC, and some other architectures. The algorithm also shows that LL/SC are conceptually more powerful than CAS.

```

1 fun enqueue(e: E): Bool = while(true) {
2   // Read the counters snapshot
3   t := tail; h := head
4   state := LL(&a[t % C])
5   if t != tail: continue
6   // Is the queue full?
7   if t == h + C: return false
8   // Try to insert the element
9   done := state ==  $\perp$  && SC(&a[t % C], e)
10  // Increment the counter
11  if LL(&tail) == t: SC(&tail, t + 1)
12  // Finish on success
13  if done: return true
14 }
15
16 fun dequeue(): E = while(true) {
17   // Read the counters snapshot
18   h := tail; h := head
19   e := LL(&a[h % C])
20   if t != tail: continue
21   // Is the queue empty?
22   if t == h: return  $\perp$ 
23   // Try to extract the element
24   done := e !=  $\perp$  && SC(&a[h % C],  $\perp$ )
25   // Increment the counter
26   if LL(&head) == d: SC(&head, h + 1)
27   // Finish on success
28   if done: return e
29 }

```

Listing 3. Bounded queue algorithm with $O(1)$ additional memory via LL/SC. This is a modification of the algorithm in Listing 2; the changes are highlighted.

2.4 $\Theta(T)$ Overhead with DCSS

Another way to eliminate the ABA problem is to use the Double-Compare-Single-Set (DCSS) synchronization primitive. DCSS(&A, expectedA, updateA, &B, expectedB) checks that the values located by addresses A and B are equal to expectedA and expectedB, respectively, updating A to updateA and returning true if the check succeeds, and returning false otherwise.

We use DCSS to atomically update the slot and check that the corresponding counter has not been changed. If DCSS fails, the algorithm helps to increment the counter and restarts the operation. The rest is the same: getting a counters snapshot (including the slot state for dequeue()), checking that the queue is not full/empty, and trying to update the slot state, incrementing the operation counter at the end. Listing 4 presents the pseudocode.

```

1 fun enqueue(e: E): Bool = while(true) {
2   // Read the counters snapshot
3   t := tail; h := head
4   if t != tail: continue
5   // Is the queue full?
6   if t == h + C: return false
7   // Try to insert the element
8   done := DCSS(&a[e % C],  $\perp$ , e, &tail, t)
9   // Increment the counter
10  CAS(&tail, t, t + 1)
11  // Finish on success
12  if done: return true
13 }
14
15 fun dequeue(): E = while(true) {
16   // Read the counters snapshot
17   t := tail; h := head
18   e := DCSS_Read(a[h % C])
19   if t != tail: continue
20   // Is the queue empty?
21   if t == h: return  $\perp$ 
22   // Try to extract the element
23   done := DCSS(&a[d % C], e,  $\perp$ , &head, h)
24   // Increment the counter
25   CAS(&head, h, h + 1)
26   // Finish on success
27   if done: return e
28 }

```

Listing 4. Bounded queue algorithm with $O(T)$ overhead via recyclable DCSS descriptors. This is a modification of the algorithm in Listing 2, the changes are highlighted.

The straightforward DCSS [6] implementation uses *descriptors*, which can be recycled in a way so that only $2T$ of them are required [1], thus, incurring $\Theta(T)$ additional memory. In short, a descriptor is a special object that “announces” an operation, in this case DCSS, which allows it to be *helped* by

other processes. Assuming the queue stores references, the distinction between *values* (enqueued elements) and *descriptors* can be implemented by either stealing a marker bit from addresses or using a language construction like `instanceof` in Java to check whether the object is a descriptor.

2.5 Values and Metadata

Notably, the algorithms with the “distinct elements” assumption (Subsection 2.2) and with DCSS (Subsection 2.4) require an ability to store elements and metadata ($\perp$ -s and DCSS descriptors) in the same array slots. In practice, to distinguish metadata from data, we either have to steal a bit from values, which results in a smaller universe of possible values, or use an Object wrapper, which incurs $\Omega(C)$ memory overhead. Besides, many programming languages do not allow storing primitive and Object types in the same memory locations, let alone bit stealing.

3 Memory Overhead: The Lower Bound

Recall that the memory overhead of a bounded queue refers to the extra memory necessary beyond that used to store the elements. We now prove that the minimal memory overhead of a linearizable non-blocking bounded queue is linear in the number of concurrent threads and show that this lower bound is tight by presenting a matching algorithm.

3.1 The Lower Bound: Proof Overview

We proceed by contradiction. Suppose that there exists an obstruction-free linearizable bounded queue of capacity C that employs $X < \frac{T}{24} - 1$ memory locations beyond the C memory locations to store the queue elements; T is the number of concurrent threads. (“24” here is a constant that simplifies the proof without affecting the desired $\Theta(T)$ lower bound.) We construct a non-linearizable execution with such an implementation, showing that any obstruction-free linearizable bounded queue must incur $\Theta(T)$ memory overhead.

To construct a non-linearizable execution, we first catch $X + 3$ threads (recall that X is the memory overhead) in a state when they are about to update different memory locations and pause these threads. For simplicity, we assume that these threads will perform writes to the memory locations and never update them to null; we explain how to support updates via CAS and null values later.

Then, carefully ensuring that the queue is empty, we successfully add C elements $x_1, x_2, \dots, x_C$ into the queue — following dequeue()s would have extracted these elements in the same order. However, we do not extract the elements now. Instead, we choose a thread caught right before updating a memory location, which stores an element x_i where $1 < i < C$, and is about to put y there. Intuitively, this write will remove v_i from the middle of the queue and put y instead. As we have caught $X + 3$ threads, and the queue uses X additional memory locations, at least three threads are about




Figure 3. A thread is about to remove v_i from the middle of the queue and put y instead.

to update the location where the queue elements are stored and one of them must point to a middle element. Figure 3 illustrates the resulting state.

We now replace the value v_i with y by invoking “write” in the caught thread. Assuming the bounded queue implementation does not differentiate elements, the state is now the same as when filling the queue with $v_1, v_2, \dots, v_{i-1}, y, v_{i+1}, \dots, v_C$ elements. Therefore, the following C dequeue()’s successfully extract these elements.

Intuitively, this execution is non-linearizable, as we have managed to put an element in the middle of the queue, observing “old” elements $v_1, v_2, \dots, v_{i-1}$, then a “new” element y , which could not be installed concurrently as the queue is bounded, and followed by “old” elements $v_{i+1}, \dots, v_C$.

Note that this is a very simplified version of the proof dedicated to creating an intuition on constructing a non-linearizable execution. To complete the proof, we need to be extremely careful with catching the threads, extending executions, and supporting CAS operations and null ($\perp$) values. The following subsections describe the model we use, the reasonable algorithm assumptions that simplify the proof, and the full proof itself.

3.2 Model

Threads and atomic operations. We consider a system of T processes (threads) $p_1, \dots, p_T$ communicating by invoking *primitives* on shared memory. Besides plain reads and writes, we assume Compare-and-Set primitive, denoted as CAS(&a, old, new), that atomically checks the value at address a and, if it equals old, replaces it with new and returns true, returning false otherwise.

While we do not support Fetch-and-Add implicitly, it does not break the proof when not manipulating memory locations where elements are stored. In particular, using Fetch-and-Add for updating counters is allowed.

The LL/SC (load-link/store-conditional) synchronization primitives, however, are not allowed, as they are ABA-immune and enable bounded queue implementation with constant memory overhead, as discussed in Subsection 2.3.

Bounded queue data structure. The *Bounded Queue* (BQ) data structure parameterized with type Type and capacity C stores elements of type Type and provides two operations:

- enqueue(x : Type) (or $\text{enq}(x)$): if the size of the queue is less than C , the operations adds value x and returns true; otherwise, the operation returns false.
- dequeue() (or $\text{deq}()$) retrieves the oldest element from the queue, returning $\perp$ if the queue is empty.

Here, $\perp$ is a special null value, available in most programming languages for references. Also, we omit type Type when discussing BQ.

Implementation. Informally, an *implementation* of a bounded queue BQ is a distributed algorithm A consisting of automata $A_1, \dots, A_T$. Given an operation invoked on BQ by p_i , A_i specifies the sequence of primitive operations on the shared memory that p_i needs to execute to compute a response. An *execution* of an implementation of BQ is a sequence of *events*: invocations and responses of high-level operations on BQ , as well as primitives applied by the processes on the shared-memory locations and the responses they return so that the automata A_i and the primitive specifications are respected. An operation is *complete* in a given execution if a matching response follows its invocation. An incomplete operation can be *completed* by inserting a matching response after invocation. We only consider *well-formed* executions: a process never invokes a new high-level operation before its previous operation returns. An operation o_1 *precedes* operation o_2 in execution α (we write $o_1 \leq_\alpha o_2$) if the response of o_1 precedes the invocation of o_2 in α .

For simplicity, we only consider *deterministic* implementations: all shared-memory primitives and automata $A_1, \dots, A_T$ exhibit deterministic behavior. Our lower bound can be, however, easily extended to randomized implementations.

Concurrent execution. A concurrent execution α of a bounded queue is *linearizable* [2, 9] if all complete operations in it and a subset of incomplete ones can be put in a total order S such that (1) S respects the sequential specification of the bounded queue, and (2) $\leq_\alpha \subseteq \leq_S$, i.e., the total order respects the *real-time* order on operation in α . Informally, every operation o of α can be associated with a *linearization point* put within the operation’s *interval*, the fragment of α between the invocation of o and the response of o . An implementation is *linearizable* if any its execution is linearizable.

Progress conditions. In this paper, we focus on *non-blocking* implementations that, intuitively, do not involve any form of locking: a failure of a process does not prevent other processes from making progress. We can also talk about non-blocking implementations of specific operations.

Popular liveness criteria are lock-freedom and obstruction-freedom [8]. An implementation of an operation is *lock-free* if, in every execution, at least one process is guaranteed to complete every such operation it invokes in a finite number of its own steps. An implementation of an operation is *obstruction-free* if it guarantees that if a process invokes this operation and, from some point on, takes sufficiently many steps *in isolation*, i.e., without contending with other

processes, then the operation eventually competes. We say that an implementation is obstruction-free (resp., lock-free) if all its operations are obstruction-free (resp., lock-free). Lock-freedom is a strictly stronger liveness criterion than obstruction-freedom: any lock-free implementation is also obstruction-free, but not vice versa [8].

3.3 Algorithm Restrictions

Atomic operations. Recall the model supports read, write, and Compare-and-Set operations. For simplicity, we do not support Fetch-and-Add, but it does not break our proof when not manipulating memory locations where elements are stored. In particular, using Fetch-and-Add for updating counters is allowed.

Values and metadata. We restrict our attention to algorithms that assume a clear separation between *value-locations*, used exclusively to store queue elements, and *metadata-locations*, used to store everything else (e.g., counters or descriptors). Importantly, we impose no restrictions on the number of bits that can be stored in memory locations for both metadata and data, allowing metadata-locations to hold either value or metadata. Notably, we require that values and metadata be indistinguishable — metadata without context can be seen as a value. Making values and metadata distinguishable would require stealing a bit, automatically imposing $\Theta(C)$ memory overhead. Therefore, we find this restriction reasonable.

Value-independence. Furthermore, we assume that the algorithm is agnostic to *values*: intuitively, the values passed to the operation only reflect the value-locations in the memory and not the steps performed by the algorithm. We call such algorithms *value-independent*. A value-independent algorithm must be *non-creative*, *enqueue-value-independent* and *dequeue-value-independent*. Intuitively, (1) a enqueue or dequeue operation can only use the values stored in value-locations or provided as an argument, (2) it does not matter what is enqueued, x or y —the only difference is that y 's should be stored instead of x 's in the memory, and (3) if the only difference in the memory is x and y in one value-location, a dequeue operation must return y instead of x . We present the formal definitions below. To the best of our knowledge, all the state-of-the-art queue implementations satisfy these assumptions.

Definition 3.1. A BQ algorithm A is *non-creative* if each time an operation op uses as an argument in a write and CAS or returns a value v , then op previously read v in a value-location or received v as an argument (if op is *enqueue*).

Definition 3.2. Let S be a memory state reached by some execution of a BQ algorithm A . Suppose that $a_1, \dots, a_k$ and v are distinct values and $a_1, \dots, a_k$ do not appear in value-locations in S . Let a process p be *idle* in S (not executing any operation). Let S_1 be the memory state reached from S by

p sequentially executing $\text{enq}(a_1), \dots, \text{enq}(a_i), \dots, \text{enq}(a_k)$ and S_2 be the memory state of A reached from S by p sequentially executing $\text{enq}(a_1), \dots, \text{enq}(v), \dots, \text{enq}(a_k)$ (i.e., instead of a_i , v is inserted). A is *enqueue-value-independent* if S_1 is identical to S_2 , except that every value-location storing a_i in S_1 stores v in S_2 .

Definition 3.3. Let S be a memory state reached by some execution of a BQ algorithm A . Suppose that there is a value x that appears only in one value-location in S and let process p be idle in S . Let process p sequentially apply $\text{dequeue}()$ operations as long as they are successful, i.e., until the queue becomes empty, and suppose that exactly one of these $\text{dequeue}()$ operation op returns x . A is *dequeue-value-independent* if when we replace x with y in S and replay $\text{dequeue}()$ operations by the process p , we obtain an execution of A in which op now returns y .

Definition 3.4. A BQ implementation A is *value-independent* if it is *non-creative*, *enqueue-value-independent*, and *dequeue-value-independent*.

3.4 Auxiliary Notions

Definition 3.5. In the proof, we construct an execution where some processes are *paused* just before performing $\text{CAS}(l, x, y)$ operation. Such processes are called *poised*, and the location l is said to be *covered* by that process.

For simplicity, we treat $\text{l.write}(y)$ as $\text{CAS}(l, *, y)$ with “*” matching any value.

Definition 3.6. If a value v was never used as an argument in the execution E , we say that v is *fresh* in E .

Let E be a finite execution of A , and p be an idle process in E . A *fill* procedure applied to E (or just *fill*) consists of a process p enqueueing C different values in isolation. We say that a fill procedure is *successful* if all C enqueues are successful and C different value-locations store these values.

An *empty* procedure (or just *empty*) consists of process p executing C dequeue operations in isolation. We say that an empty procedure is *successful* if all these C dequeue operations are successful.

We say that a fill-empty-procedures pair is *up-to-date* if both procedures are successful and dequeues from the empty procedure return the values enqueued by the fill.

Lemma 3.7. Every execution can be extended by one working process with an up-to-date fill-empty procedure at the end.

To prove Lemma 3.7, we need the following lemmas first.

Definition 3.8. We say that the queue is *logically empty* if it is in a state for which a new dequeue operation would return $\perp$ if executed in isolation.

Lemma 3.9. Consider some state of a system S : there are *poised* processes and a *logically empty* queue. If we consequently perform a *fill* procedure and an *empty* procedure on

the queue and one of these procedures is not successful, then one of the poised operations “took effect”, i.e., is linearized before any future operation.

Proof. We have a logically empty queue and we perform a fill procedure. Suppose this fill is not successful. By definition, there can be two cases. First, some enqueue operation is not successful then another enqueue operation succeeded (exactly, C enqueues succeed) – which can only be the one that was poised and, thus, it is linearized. Second, all enqueue operations are successful, but not all C new values are stored in the shared memory. Thus, there exists some poised dequeue operation that can return the non-stored value. But after the successful empty procedure, it would mean that this poised operation is linearized.

After the fill procedure, the logical state of the queue is full. Thus, C dequeues should succeed. If the empty procedure is not successful, thus, some old poised dequeue takes place. This means that it was already linearized before that moment. $\square$

Lemma 3.10. *Consider some state of a system: a logically empty queue with poised operations. If we perform a successful fill procedure followed by a successful empty procedure and dequeue operations of the empty procedure return not the exactly same values that were enqueued by the fill procedure, then one of the poised operations “took effect”.*

Proof. Let us look at the value v that was enqueued and was not dequeued by the empty procedure. In the end, the queue is empty; thus, the value v was dequeued by some operation. Since it was not returned by any operation from the empty procedure, it should be some poised operation. Thus, after the empty procedure, this poised operation is linearized. $\square$

By combining the last two lemmas, we construct the following result.

Lemma 3.11. *Consider some state of a system: a logically empty queue with poised operations. If a fill-empty-procedures pair applied to this state is not up-to-date, then one of the poised operations “took effect”, i.e., is linearized before any future operations.*

We are now ready to prove Lemma 3.7 and show that every execution can be extended by one working process with an up-to-date fill-empty procedure at the end.

Proof of Lemma 3.7. In our proof, a dedicated process repeatedly performs alternating fill and empty procedures, until the corresponding pair is up-to-date. By Lemma 3.11, the number of such iterations cannot exceed T , as in our proof we catch at most $T - 1$ concurrent operations $\square$

3.5 The Lower Bound: Complete Proof

We now present a complete proof of the memory overhead lower bound. Specifically, we prove Theorem 3.12 below.

Theorem 3.12. *Any obstruction-free linearizable value-independent implementation of Bounded Queue with capacity C uses $C + \lfloor \frac{T-24}{24} \rfloor$ value-locations, assuming that there exists at least $\Theta(T \cdot C)$ different values and $\frac{T}{2} < C$.*

Recall that the key idea is to construct a non-linearizable execution. Before diving deeper into details, we describe the execution we target to construct and prove that it is non-linearizable.

Lemma 3.13. *Let $k = 2X + 3$. Let E be an execution that results in a state that has k distinct value-locations $l_1, \dots, l_k$ such that for all i $1 \leq i \leq k$ we can replace x_i with y_i in l_i using poised CASes. Also, all x_i are distinct, and for all i , $x_i \neq y_i$, and x_i, y_i are non- $\perp$.*

Then we can extend E to a non-linearizable execution.

Proof. First, we extend E using one process so it ends with up-to-date fill-empty pair (by Lemma 3.7, this is possible). By the definition of a successful fill, there are C distinct values stored in memory after it, so, since we have $C + X$ value-locations in total, there are at least $C - X$ of these new values that are stored in a single location. As $k = 2X + 3$, after the fill, there are at least 3 locations among $l_1, \dots, l_k$ such that a value of some enqueue from fill is stored only in that location. Thus, there is l_i that stores a value of some enq from the fill procedure that is not a first enq , nor the last one, and this value is stored only in l_i . Using enqueue-value-independence we change the argument of the corresponding enqueue to x_i , reaching the state where after the fill l_i is the only location storing x_i and x_i is returned in the following empty procedure.

Therefore, if we replay this updated fill procedure we could perform poised CASes substituting x_i with y_i before the empty procedure. Thus, if at this point we replay the empty procedure, the dequeue that used to return x_i now returns y_i by dequeue-value independence.

By that, the last fill-empty-procedures pair have operations

$enq(v_1)$	$\dots$	$enq(x_i)$	$\dots$	$enq(v_C)$
$deq \rightarrow v_1$	$\dots$	$deq \rightarrow y_i$	$\dots$	$deq \rightarrow v_C$

We call a pair of type $\langle enqueue(v), dequeue \rightarrow v \rangle$ a *matching* pair and it is *true-matching* if in any given linearization, these enqueue and dequeue indeed match each other.

We claim that $\langle enq(v_1), deq \rightarrow v_1 \rangle$ and $\langle enq(v_C), deq() \rightarrow v_C \rangle$ are true-matching pairs in the constructed execution.

To see this, note that for a matching pair to be not true-matching there should be another enqueue whose value is dequeued. But v_1 and v_C are fresh values since the only non-fresh value among enqueues is x_i and by the construction, it is not the last, nor the first enqueue. Thus, there can not be other poised $enq(v_1)$ or $enq(v_C)$.

The question now is where to place $enq(y_i)$ to find a true-match $deq \rightarrow y_i$. It cannot be placed before $enq(v_1)$ or after

$enq(v_C)$ since the pairs for v_1 and v_C are true-matching and they cannot transpose due to FIFO property of the queue. So, it can only be placed between $enq(v_1)$ and $enq(v_C)$.

But then we have $C + 1$ successful enqueues in a row, so there must be $deq \rightarrow v_1$ before $enq(v_C)$ which means that the pair $\langle enq(v_1), deq \rightarrow v_1 \rangle$ from the table is not a true-match, but we have proven it already.

Thus, $enq(x_i)$ corresponds to $deq \rightarrow y_i$ which makes the execution not linearizable. $\square$

We are finally ready to prove the main Theorem 3.12 and show that any obstruction-free linearizable value-independent implementation of Bounded Queue with capacity C uses $C + \lfloor \frac{T-24}{24} \rfloor$ value-locations, assuming that there exists at least $\Theta(T \cdot C)$ different values and $\frac{T}{2} < C$.

Proof of Theorem 3.12. The proof proceeds by contradiction. Namely, we pick the number of additional value-locations X small enough, though still linear in T , and show that for each implementation that uses at most X additional value-locations, we can build a non-linearizable execution. We prove that it is sufficient to pick $X = \lfloor \frac{T-24}{24} \rfloor$.

Our procedure to create non-linearizable execution is partitioned into four steps. During the first step, the queue is filled with fresh values, and $\frac{T}{2}$ processes are poised right before the CAS operation. In the second step, we just call empty and fill procedures in order to obtain fresh values in memory. The third and fourth steps are the most complicated. Depending on the BQ algorithm, we get three cases depending on poised operations. The goal of each case is to come up with an execution that brings the queue to the state required by Lemma 3.13. Finally, using this lemma we build a non-linearizable execution.

Now, we dive into the details. Our construction algorithm works in four steps.

First step. We try to extend an execution $\frac{T}{2}$ times with an up-to-date fill-empty pair using Lemma 3.7. For each pair, we require a new idle process. Though, for each try, we don't let a process finish an extension, rather we stop this process right before it is about to CAS at a not yet covered value-location from $\perp$. We say that we *catch* the process when we stop it right before performing a CAS. When we catch the process, we restart with another process using Lemma 3.7. We claim that using this procedure we can catch $\frac{T}{2}$ processes.

To see that, assume by the contrary that we do not catch $\frac{T}{2}$ processes. Thus, one of them managed to perform a successful fill. By the definition of successful fill, there must be C fresh values stored in the value-locations after it, so a process made successful attempts to CAS C value-locations and was not poised. This implies that all those locations are covered since all locations initially store $\perp$ and hence cannot be modified without being previously covered. But since one process covers only one location and $C > \frac{T}{2}$, C locations cannot be covered, a contradiction.

Second step. We extend an execution using a new idle process to end up with a successful fill. It is indeed possible by the trivial corollary of the Lemma 3.7: if we can ensure successful fill and empty, we can just stop after the fill.

Third step. We try to extend an execution $\frac{T}{2}$ times with an up-to-date fill-empty pair using Lemma 3.7. For each pair, we require a new idle process. Though, for each try, we do not let a process finish an extension, rather we stop this process right before it is about to CAS on a location that satisfies the following *catch criteria*:

1. The corresponding value-location was not yet covered by some CAS at Step 3;
2. This value-location is already covered by some CAS at Step 1;
3. The poised CAS must be from a non-bottom value different from the previous CASes poised during Step 3.

Fourth step. This step is to account for the case when there are few processes caught at Step 3.

If there are $\frac{T}{4} - 3X$ idle processes at this point, we consecutively extend the execution with a successful fill using each of them. It is indeed possible by the trivial corollary of the Lemma 3.7. We catch a process if it satisfies the following *catch criteria*:

1. The corresponding value-location is not yet covered by CAS at Step 4;
2. This value-location is already covered by CAS at Step 1;
3. The poised CAS must be from a non-bottom value (to possibly a bottom one) different from the previous CASes poised during Step 4.

Now, we explain how our poised operations provide a state desired by Lemma 3.13. This part becomes a little bit technical — we get *three* cases depending on how many CASes we poised during Step 3:

- Suppose that we poised less than $\frac{T}{4}$ processes.

After Step 2, the successful fill is performed, so C fresh distinct values are stored in C different locations. We have $\frac{T}{2}$ value-locations covered by CASes at Step 1, so, since we have $C + X$ value-locations in total, there are $\frac{T}{2} - 2X$ locations among covered that store unique values. Let us call the set of these locations that are not modified during Step 3 as L' . We claim that $|L'| > \frac{T}{4} - 2X$.

To see this, assume the opposite, namely that L' has at most $\frac{T}{4} - 2X$ value-locations. Then at least $(\frac{T}{2} - 2X) - |L'| = \frac{T}{4}$ of our chosen value-locations were modified. It means that during our procedure in Step 3, we made a CAS on these locations but did not catch it. This could happen only if one of the catch criteria of Step 3 was not satisfied. However, criteria (2) and (3) are always satisfied due to the way we chose locations: the chosen value-locations are already covered by some CAS from Step 1, and they stored non-bottom value

(they store some of C values from Step 2). Thus, the only criterion that could not be satisfied is (1), i.e., it is already covered by a CAS. But, we poised less than $\frac{T}{4}$ CASes in the main condition; thus, (1) is satisfied for some location – we could have poised at least one more CAS. Thus, $|L'|$ is greater than $\frac{T}{4} - 2X$.

We claim that during Step 4, we catch $\frac{T}{4} - 3X$ processes. Otherwise, one of them was able to complete a successful fill, which implies that C fresh values are stored, so at most X locations are left unmodified. Thus, at least $|L'| - X = \frac{T}{4} - 3X$ locations in L' must be overwritten – there will be attempts to successfully CAS and hence cover them. Let L be the set of value-locations covered by poised operations from Step 4. Finally, we apply Lemma 3.13 with l_i being locations from L .

We now need to show that the Lemma 3.13 requirements are satisfied. By the construction of L , all $l_i \in L$ are indeed distinct. For each $1 \leq i \leq |L|$ there is a poised CAS of the form $CAS(l_i, \perp, y_i)$ and a CAS of the form $CAS(l_i, x_i, mb_i)$. If $mb_i \neq \perp$, then we can change the value in l_i from x_i to mb_i . Otherwise, we can perform both CASes changing the value in l_i from x_i to y_i . We also need to show that all x_i are distinct. It follows from the fact that all values in L' are. Finally, $|L| = \frac{T}{4} - 3X \geq 2X + 3$, thus, we can choose $2X + 3$ locations from L and apply Lemma 3.13.

- Suppose that in Step 3 at least $\frac{T}{4}$ processes are poised while more than $\frac{T}{8}$ of them CAS to non-bottom. Let L be the set of value-locations covered by poised CASes that CAS to non-bottom.

Now, we can apply Lemma 3.13 with l -s being locations from L .

Let us check Lemma's requirements are satisfied. Locations in L are distinct by the catch criteria. We can indeed substitute x_i with y_i in each l_i applying the corresponding CAS. x_i are non-bottom because of the catch criteria and y_i are non-bottom because of the current case assumption. All x_i are distinct because of the catch criteria. Finally, $|L| \geq \frac{T}{8} \geq 2X + 3$, thus, we can choose $2X + 3$ locations from L and apply Lemma 3.13.

- Suppose that in Step 3 at least $\frac{T}{4}$ processes are poised and at least $\frac{T}{8}$ of them CAS to bottom.

Here, we have at least $\frac{T}{8}$ pairs of the type $CAS(l_i, \perp, y_i)$ and $CAS(l_i, x_i, \perp)$.

Note that x -s are those values that were present in the memory after Step 2, whereas y -s are those introduced during Step 1. Recall that by non-creative property y -s are all arguments of enqueue operations.

Filling the queue with fresh values in Step 2 guarantees that at most X values out of those y -s are present at the start of Step 3. Since all x -s are pairwise distinct

by the catch criteria of Step 3, there can be no more than X values from y -s among them. We no longer consider CAS pairs where CAS is done from x which is equal to some y_i . However, there are at most X of the non-satisfying pairs and we still have at least $\frac{T}{8} - X$ good $CAS(l_i, \perp, y_i), CAS(l_i, x_i, \perp)$ pairs. Let $L = \{l_i\}$ be the set of value-locations covered by these pairs.

Finally, we can apply Lemma 3.13 with l_i being elements of L .

Requirements of the Lemma are satisfied. Indeed, l -s are distinct and x -s are distinct by the catch criteria of Step 3. For each l_i we can substitute x_i with y_i performing a pair of CASes. Moreover, x -s are distinct from y -s, since y -s are values from Step 1 and by the construction of L there are no values from Step 1. Finally, $|L| = \frac{T}{8} - X \geq 2X + 3$, thus, we can choose $2X + 3$ locations from L and apply Lemma 3.13. $\square$

3.6 The Upper Bound

Finally, we show that our lower bound is tight by presenting an algorithm that asymptotically matches the lower bound and uses $O(T)$ extra memory. Please note that our goal is to show the upper bound of the memory overhead rather than provide a practical and efficient implementation.

Algorithm overview. The key idea is to use descriptors for enqueue($\dots$) operations, with the possibility of reusing them [1], to match the desired memory overhead. In addition, we use an T -size “announcement” array that stores references to descriptors of “in-progress” enqueue($\dots$) invocations. Intuitively, a descriptor in this array declares an *intention* to perform an enqueue operation on some chosen cell. This enqueue succeeds if the enqueues counter has not been changed (which is similar to our DCSS-based solution above). The descriptor stores a target cell to write and an operation status in successful field. We say that a descriptor/operation/thread “covers” a cell if it intends to put an element there.

Thus, in total, we have $O(T)$ memory overhead that asymptotically matches the lower bound shown in the previous subsection.

Now, we explain how descriptors and the announcement array help us achieve a correct algorithm. As in the previous algorithms, a dequeue($\dots$) (resp., enqueue($\dots$)) operation starts with taking a snapshot of the counters in lines and checks if the queue is *empty* (resp., *full*).

The implementation of dequeue($\dots$) slightly differs from those in the algorithms from Section 2. To read an element, it goes through the announcement array to check for the ongoing enqueue operations on the target cell, which is calculated from the counters. If one exists, our dequeue takes the argument of that ongoing operation. Otherwise, it takes a value from the target cell. Finally, it tries to increment

the dequeues counter and returns the found element if the corresponding CAS succeeds.

As for `enqueue()`, it creates a special descriptor, announces it in our announcement array, and then tries to atomically apply the operation. Then, the operation increments the enqueues counter, possibly helping some concurrent operation. If the descriptor is successfully applied, the operation completes. Otherwise, the operation is restarted.

There can be only one descriptor in the successful state that can “cover” a given cell in the elements array, no other successful descriptor can point to the same cell. Thus, only the thread that started covering the cell is eligible to update it. When the thread with exclusive access for modifications finishes the operation, it “releases” the cell, so, another operation is able to “cover” it. However, when an enqueue operation finds a cell to be covered, it replaces the old descriptor (related to the thread that covers the cell now) with a new one and completes. The thread with the exclusive access to the cell helps this enqueue operation to put the element.

Summing everything up, our algorithm is lock-free (Appendix A.1) and it uses just $O(T)$ additional memory to C values in the bounded queue. Please note that we do not provide any guarantees about the execution time of operations on our BQ. One should expect the queue to answer the queries in $O(1)$, however, in our algorithm each operation should traverse through the whole announcement array leading to $\Theta(T)$ time. This leaves us with an open question whether it is possible to implement a memory-optimal queue that serves requests in $O(1)$ time.

Implementation details. Due to space constraints, we discuss the implementation details and prove the algorithm’s correctness in Appendix A.

4 Related Work

Memory efficiency has always been one of the central concerns in concurrent computing. Many theoretical bounds have been established on memory requirements of various concurrent abstractions, such as mutual exclusion [3], perturbable objects [10], or consensus [23]. However, it appears that minimizing memory overhead in dynamic concurrent data structures has not been in the highlight until recently. A standard way to implement a lock-free bounded queue, the major running example of this paper, is to use descriptors [16, 20] or additional meta-information per each element [4, 18, 19, 21]. The overhead of resulting solutions is proportional to the queue size: a descriptor contains an additional $\Omega(1)$ data to distinguish it from a value, while an additional meta-information is $\Omega(1)$ memory appended to the value by the definition.

The fastest queues, however, store multiple elements in nodes so the memory overhead per element is relatively small from the practical point of view [14, 17, 22], and we consider them as memory-friendly, though not memory-optimal.

A notable exception is the work by Tsigas et al. [19] that tries to answer our question: whether there exists a lock-free concurrent bounded queue with $O(1)$ additional memory. The solution proposed in [19] is still a subject to the ABA problem even when all the elements are different: it uses only two null-values, and if one process becomes asleep for two “rounds” (i.e., the pointers for `enqueue()` and `dequeue()` has made two traversals through all the elements), waking up it can incorrectly place the element into the queue. Besides resolving the issue, our algorithm, under the same assumptions, is shorter and easier to understand (see Subsection 2.2).

The tightest algorithm we found is the recent work by Nikolaev [15], that proposes a lock-free bounded queue with capacity C implemented on top of an array with $2C$ memory cells. While the algorithm manipulates the counters via Fetch-And-Add, it still requires descriptors, one per each ongoing operation. This leads to the additional overhead linear in T . Thus, the total memory overhead is $\Omega(C + T)$, while the lower bound is $\Theta(T)$ and does not depend on the queue capacity.

5 Discussion

In this paper, we proved that any non-blocking implementation of bounded queue incurs $\Omega(T)$ memory overhead, while showing that the bound is tight with a matching algorithm. In Section 2, we also presented a series of algorithms with *constant* memory overhead that work under several practical restrictions on the system or/and the application, along with a simple DCSS-based algorithm that matches the lower bound but requires an ability to store values and descriptors in the same array slot.

We find the following relaxations important to consider in future work: (1) the single-producer and single-consumer application restrictions, (2) the ability to store descriptors in value-locations (free in JVM and Go) or “steal” a couple of bits from addresses (free in C++) for containers of references, and (3) relaxing the object semantics, e.g., to probabilistic guarantees. Each of these cases corresponds to a popular class of applications, and determining the optimal memory overhead in these scenarios is appealing in practice.

It would also be interesting to consider the problem of memory overhead incurred by *wait-free* implementations since such solutions are typically more complicated.

Finally, in this paper, we only focused on bounded queues, as they allow a natural definition of memory overhead. One may try to extend the notion to the case of *unbounded* queues by considering the *ratio* between the amount of memory allocated for the metadata with respect to the memory allocated for data elements. Defining the bounds on this ratio for various data structures, such as pools, stacks, etc., remains an intriguing open question.

References

- [1] Maya Arbel-Raviv and Trevor Brown. 2017. Reuse, don't recycle: Transforming lock-free algorithms that throw away descriptors. In *31st International Symposium on Distributed Computing (DISC 2017)*.
- [2] Hagit Attiya and Jennifer Welch. 2004. *Distributed Computing. Fundamentals, Simulations, and Advanced Topics*. John Wiley & Sons.
- [3] James E. Burns and Nancy A. Lynch. 1993. Bounds on Shared Memory for Mutual Exclusion. *Inf. Comput.* 107, 2 (1993), 171–184.
- [4] Steven Feldman and Damian Dechev. 2015. A wait-free multi-producer multi-consumer ring buffer. *ACM SIGAPP Applied Computing Review* 15, 3 (2015), 59–71.
- [5] Anders Gidenstam, Håkan Sundell, and Philippas Tsigas. 2010. Cache-aware lock-free queues for multiple producers/consumers and weak memory consistency. In *PODC*. 302–317.
- [6] Timothy L Harris, Keir Fraser, and Ian A Pratt. 2002. A practical multi-word compare-and-swap operation. In *International Symposium on Distributed Computing*. Springer, 265–279.
- [7] Maurice Herlihy, Victor Luchangco, and Mark Moir. 2003. Obstruction-Free Synchronization: Double-Ended Queues as an Example. In *23rd International Conference on Distributed Computing Systems (ICDCS 2003)*, 19–22 May 2003, Providence, RI, USA. IEEE Computer Society, 522–529.
- [8] Maurice Herlihy and Nir Shavit. 2013. *The art of multiprocessor programming*. Morgan Kaufmann.
- [9] Maurice P Herlihy and Jeannette M Wing. 1990. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 12, 3 (1990), 463–492.
- [10] Prasad Jayanti, King Tan, and Sam Toueg. 2000. Time and Space Lower Bounds for Nonblocking Implementations. *SIAM J. Comput.* 30, 2 (2000), 438–456.
- [11] Nikita Koval, Dan Alistarh, and Roman Elizarov. 2019. Scalable FIFO Channels for Programming via Communicating Sequential Processes. In *European Conference on Parallel Processing*. Springer, 317–333.
- [12] Nikita Koval, Dan Alistarh, and Roman Elizarov. 2023. Fast and Scalable Channels in Kotlin Coroutines. In *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming* (Montreal, QC, Canada) (PPoPP '23). Association for Computing Machinery, New York, NY, USA, 107–118. <https://doi.org/10.1145/3572848.3577481>
- [13] Maged M. Michael and Michael L. Scott. 1996. Simple, Fast, and Practical Non-Blocking and Blocking Concurrent Queue Algorithms. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing, Philadelphia, Pennsylvania, USA, May 23–26, 1996*, James E. Burns and Yoram Moses (Eds.). ACM, 267–275.
- [14] Adam Morrison and Yehuda Afek. 2013. Fast concurrent queues for x86 processors. In *ACM SIGPLAN Notices*, Vol. 48. 103–112.
- [15] Ruslan Nikolaev. 2019. A Scalable, Portable, and Memory-Efficient Lock-Free FIFO Queue. In *DISC*.
- [16] Peter Pirkelbauer, Reed Milewicz, and Juan Felipe Gonzalez. 2016. A portable lock-free bounded queue. In *ICAPP*. 55–73.
- [17] Raed Romanov and Nikita Koval. 2023. The State-of-the-Art LCRQ Concurrent Queue Algorithm Does NOT Require CAS2. In *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming* (Montreal, QC, Canada) (PPoPP '23). Association for Computing Machinery, New York, NY, USA, 14–26. <https://doi.org/10.1145/3572848.3577485>
- [18] Niloufar Shafiei. 2009. Non-blocking array-based algorithms for stacks and queues. In *ICDCN*. 55–66.
- [19] Philippas Tsigas and Yi Zhang. 2001. A simple, fast and scalable non-blocking concurrent FIFO queue for shared memory multiprocessor systems. In *SPAA*. 134–143.
- [20] John D Valois. 1994. Implementing lock-free queues. In *ICPADS*. 64–69.
- [21] Dmitry Vyukov. [n.d.]. Bounded MPMC Queue. <http://www.1024cores.net/home/lock-free-algorithms/queues/bounded-mpmc-queue>.
- [22] Chaoran Yang and John Mellor-Crummey. 2016. A wait-free queue as fast as fetch-and-add. *ACM SIGPLAN Notices* 51, 8 (2016), 16.
- [23] Leqi Zhu. 2016. A tight space bound for consensus. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18–21, 2016*, Daniel Wicks and Yishay Mansour (Eds.). ACM, 345–350.