

A Generic Local-Spin Fetch-and- ϕ -based Mutual Exclusion Algorithm*

James H. Anderson
Department of Computer Science
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-3175
Email: anderson@cs.unc.edu

Yong-Jik Kim
Tmax Soft Research Center
272-6 Seohyeon-dong, Seongnam-si
Gyeonggi-do, Korea 463-824
Email: jick@tmax.co.kr

November 2005, Revised November 2006

Abstract

We present a generic *fetch-and- ϕ* -based local-spin mutual exclusion algorithm, with $O(1)$ time complexity under the remote-memory-references time measure. This algorithm is “generic” in the sense that it can be implemented using any *fetch-and- ϕ* primitive of rank $2N$, where N is the number of processes. The *rank* of a *fetch-and- ϕ* primitive is a notion introduced herein; informally, it expresses the extent to which processes may “order themselves” using that primitive. This algorithm breaks new ground because it shows that $O(1)$ time complexity is possible using a wide range of primitives. In addition, by applying our generic algorithm within an arbitration tree, one can easily construct a $\Theta(\max(1, \log_r N))$ algorithm using any primitive of rank r , where $2 \leq r < N$.

Keywords: Fetch-and- ϕ primitives, local spinning, shared-memory mutual exclusion, theory of concurrent algorithms, time complexity.

*Work supported by NSF grants CCR 9972211, CCR 9988327, ITR 0082866, and CCR 0208289. This work was presented in preliminary form at the 23rd IEEE International Conference on Distributed Computing Systems [3].

1 Introduction

Recent work on shared-memory mutual exclusion has focused on the design of algorithms that minimize interconnection-network contention through the use of *local spinning*. In local-spin algorithms, all busy waiting is by means of read-only loops in which one or more “spin variables” are repeatedly tested. Such spin variables must be either locally cacheable or stored in a local memory module that can be accessed without an interconnection network traversal. The former is possible on cache-coherent (CC) machines, while the latter is possible on distributed shared-memory (DSM) machines. As explained later, it is generally more difficult to design local-spin algorithms for DSM machines than for CC machines.

In this paper, several results concerning the time complexity of local-spin mutual exclusion algorithms are presented. The notion of time complexity assumed is that given by the *remote-memory-references (RMR) measure* [4]. Under this measure, an algorithm’s time complexity is defined as the total number of remote memory references (*i.e.*, references that require an interconnection network traversal) required in the worst case by one process to enter and then exit its critical section once. An algorithm may have different RMR time complexities on CC and DSM machines, because on CC machines, variable locality is dynamically determined, while on DSM machines, it is statically determined.

The main focus of this paper is mutual exclusion algorithms implemented using *fetch-and- ϕ* primitives. A *fetch-and- ϕ* primitive is characterized by a particular function ϕ (which we assume to be deterministic), accesses a single variable atomically, and has the effect of the following pseudo-code, where *var* is the variable accessed.

```
fetch-and- $\phi$ (var, input)
  old := var;
  var :=  $\phi$ (old, input);
  return(old)
```

In this paper, we distinguish between *fetch-and- ϕ* primitives that are comparison primitives and those that are not. A *comparison primitive* conditionally updates a shared variable after first testing that its value meets some condition; examples include *compare-and-swap* and *test-and-set*.¹ Non-comparison primitives update variables unconditionally; examples include *fetch-and-increment* and *fetch-and-store*.

In recent work [2], we established a time-complexity lower bound of $\Omega(\log N / \log \log N)$ remote memory references for any N -process mutual exclusion algorithm based on reads, writes, or comparison primitives. In contrast, several constant-time algorithms are known that are based on noncomparison *fetch-and- ϕ* primitives [5, 7, 12]. This suggests that noncomparison primitives may be the best choice to provide in hardware, if one is interested in implementing efficient blocking synchronization mechanisms.

Constant-time local-spin mutual exclusion algorithms that use noncomparison primitives have been proposed by T. Anderson [5], Graunke and Thakkar [7], Mellor-Crummey and Scott [12], Craig [6], and Landin and Hagersten [11]. In each of these algorithms, blocked processes wait within a “spin queue.” A process enqueues itself by using a *fetch-and- ϕ* primitive to update a shared “tail” pointer; a process’s predecessor (if any) in the queue is indicated by the primitive’s return value. A process in the spin queue waits (if necessary) until released by its predecessor. Although these algorithms follow a common strategy, they vary in the primitives used and the progress properties ensured. Some important attributes of each algorithm are listed below.

- T. Anderson’s algorithm uses *fetch-and-increment* and requires an underlying cache-coherence mechanism for spins to be local. Thus, it has $O(1)$ RMR time complexity only on CC machines.
- Graunke and Thakkar’s algorithm uses *fetch-and-store*. This algorithm also requires an underlying cache-coherence mechanism and thus has $O(1)$ RMR time complexity only on CC machines.
- Craig [6] and Landin and Hagersten [11] independently proposed the same algorithm, which is based on *fetch-and-store*. While Landin and Hagersten considered only CC machines, Craig presented constant-time variants of the algorithm for both CC and DSM machines.

¹ *compare-and-swap* and *test-and-set* are ordinarily defined to return a boolean condition indicating if the comparison succeeded. In this paper, we instead assume that each returns the accessed variable’s original value, as in [8]. It is straightforward to modify any algorithm that uses the boolean versions of these primitives to instead use the versions considered in this paper.

- Mellor-Crummey and Scott actually presented two variants of their algorithm, one that uses *fetch-and-store*, and a second that uses both *fetch-and-store* and *compare-and-swap*. In both, spins are local on both CC and DSM machines. However, the *fetch-and-store* variant is not starvation-free, and hence actually has unbounded RMR time complexity. The variant that also uses *compare-and-swap* is starvation-free and has $O(1)$ RMR time complexity on both CC and DSM machines.

The existence of these algorithms gives rise to a number of intriguing questions regarding mutual exclusion algorithms. Can $O(1)$ mutual exclusion algorithms be devised for both CC and DSM machines using primitives other than *fetch-and-increment* and *fetch-and-store*? Is it possible to automatically transform a local-spin algorithm for CC machines so that it has the same RMR time complexity on DSM machines? Given that the $\Omega(\log N / \log \log N)$ lower bound mentioned above applies to algorithms that use comparison primitives, we know that there exist *fetch-and- ϕ* primitives that are not sufficient for constructing $O(1)$ algorithms. For such primitives, what is the most efficient algorithm that can be devised? Can we devise a *ranking* of synchronization primitives that indicates the singular characteristic of a primitive that enables a certain RMR time complexity (for mutual exclusion) to be achieved? Such a ranking would provide information relevant to the implementation of *blocking* synchronization mechanisms that is similar to that provided by Herlihy’s wait-free hierarchy [8], which is relevant to *nonblocking* mechanisms.²

Contributions of this paper. While we are not yet able to fully answer all of these questions, we do take some initial steps towards their resolution in this paper. We begin by proposing a ranking of *fetch-and- ϕ* primitives. Informally, a primitive of rank r has sufficient symmetry-breaking power to linearly order up to r invocations of that primitive. Based on this notion of a rank, we then present a generic N -process *fetch-and- ϕ* -based local-spin mutual exclusion algorithm that has $O(1)$ RMR time complexity on both CC and DSM machines. This algorithm is “generic” in the sense that it can be implemented using any *fetch-and- ϕ* primitive of rank $2N$. Our generic algorithm breaks new ground because it shows that $O(1)$ RMR time complexity is possible using a wide range of primitives, on both CC and DSM machines.

We present our generic algorithm by first giving a variant that is designed for CC machines. We then present a transformation, which may be of more general interest, that can be used to replace spin loops that are local on CC machines with ones that are local on DSM machines. This transformation is then applied to our algorithm. This transformation makes use of an underlying two-process mutual exclusion algorithm and is correct as long as the added two-process algorithms are invoked safely (*i.e.*, by at most two processes at any time). We show that our DSM algorithm is correct by showing that, within it, these invocations are safe.

Both the CC and DSM variants of our basic algorithm have $\Theta(N)$ space complexity (for N processes), which is the same as the prior constant-time algorithms cited above. In the DSM case, $\Theta(N)$ is asymptotically optimal because each process needs a dedicated spin variable, as explained later. The term *space complexity* refers to the number of variables used. In addition, variable sizes are of importance. In our algorithms, some of the variables that are employed have sizes that are dependent on the generic *fetch-and- ϕ* primitive that is used; all other variables require $O(\log N)$ bits.

By applying our generic algorithm within an arbitration tree, one can easily construct a $\Theta(\max(1, \log_r N))$ algorithm using any primitive of rank r , where $2 \leq r < N$. For the case $r = \Theta(N)$, this algorithm is clearly asymptotically time-optimal. However, as shown in the second author’s Ph.D. dissertation [9], there exists a class of primitives with constant rank for which $\Theta(\max(1, \log_r N))$ is *not* optimal. This is shown by presenting a $\Theta(\log N / \log \log N)$ algorithm that can be implemented using any primitive that meets an additional condition. This algorithm is quite complicated, and therefore is not presented in this paper, due to space constraints. This additional condition arises due to the need to *reset* a variable that is repeatedly updated by *fetch-and- ϕ* primitive invocations. In the generic algorithm presented in this paper, variables are reset using simple writes. In the $\Theta(\log N / \log \log N)$ algorithm presented in [9], the *fetch-and- ϕ* primitive used must be of rank at least three and also be *self-resettable*, which means that the primitive itself can be used to reset a variable that has

²Herlihy’s hierarchy is concerned with *computability*: a primitive (or object) X is stronger than a primitive (or object) Y if X can be used to implement Y (in a non-blocking manner) but not vice versa. The ranking suggested here is not concerned with computability, but rather time complexity. Nonetheless, both rankings provide information concerning the usefulness of primitives. Herlihy’s hierarchy indicates which primitives should be supported in hardware if one is interested in implementing nonblocking algorithms; the proposed ranking indicates which primitives should be supported in hardware if one is interested in implementing scalable spin locks.

been updated using that primitive, *i.e.*, it is not necessary to perform resets using simple write operations. In the $\Theta(\log N / \log \log N)$ algorithm, this self-resetable feature is used (in addition to sometimes resetting with simple writes), with a resulting asymptotic improvement in time complexity for primitives of rank $o(\log N)$. As explained in [9], it follows from the $\Omega(\log N / \log \log N)$ lower bound mentioned above that this algorithm is time-optimal for certain self-resetable primitives of constant rank.

Organization. The rest of this paper is organized as follows. In Section 2, we present needed definitions. Then, in Section 3, we present our generic algorithm. A formal proof of correctness for two versions of the generic algorithm is presented in Appendices A and B. We end the paper with concluding remarks in Section 4.

2 Definitions

In the mutual exclusion problem, each process cycles through four code sections, termed “noncritical,” “entry,” “critical,” and “exit” sections, respectively. A process may halt within its noncritical section but not within its critical section. Furthermore, no variables (other than program counters) accessed within a process’s entry or exit section may be accessed within its critical or noncritical section. The objective is to design the entry and exit sections so that the following requirements hold.³

- **Exclusion:** At most one process executes its critical section at any time.
- **Starvation-Freedom:** If some process is in its entry (exit) section, then that process eventually executes its critical (noncritical) section.

We hereafter let N denote the number of processes in the system, and assume that each process has a unique process identifier in the range $0, \dots, N - 1$.

We assume the existence of a *generic fetch-and- ϕ* primitive, as defined in Section 1. We will use “*Vartype*” to denote the type of the accessed variable *var*. (The accessed variable’s type is part of the definition of such a primitive.) For example, for a *fetch-and-increment* primitive, *Vartype* would be **integer**, and for a *test-and-set* primitive, it would be **boolean**. In our algorithms, we use \perp to denote the initial value of a variable accessed by a *fetch-and- ϕ* primitive (*e.g.*, if *Vartype* is **boolean**, then \perp would denote either *true* or *false*). We now define the notion of a “rank,” mentioned earlier.

Definition: The *rank* of a *fetch-and- ϕ* primitive is the largest integer r satisfying the following.

For each process p , there exists a constant array $\alpha_p[0..\infty]$ of input values, such that, for any integer value a_p , if p performs the sequence of *fetch-and- ϕ* invocations given by

for $i := a_p$ **to** ∞ **do** *fetch-and- $\phi(v, \alpha_p[i])$* **od**

on a variable v (of type *Vartype*) that is initially \perp (for some choice of \perp), then in any interleaving of these invocations by the N different processes, **(i)** any two invocations among the first $r - 1$ by different processes write different values to v , **(ii)** any two *successive* invocations among the first $r - 1$ by the same process write different values to v , and **(iii)** of the first r invocations, only the first invocation returns \perp .

A *fetch-and- ϕ* primitive has *infinite rank* if the condition above is satisfied for arbitrarily large values of r . \square

As our generic algorithm shows, a *fetch-and- ϕ* primitive with rank r has enough power to linearly order r invocations by possibly different processes unambiguously. Note that it is not necessary for the primitive to fully order invocations by the same process, since each process can keep its own execution history.

³Progress properties other than starvation-freedom may be of interest as well.

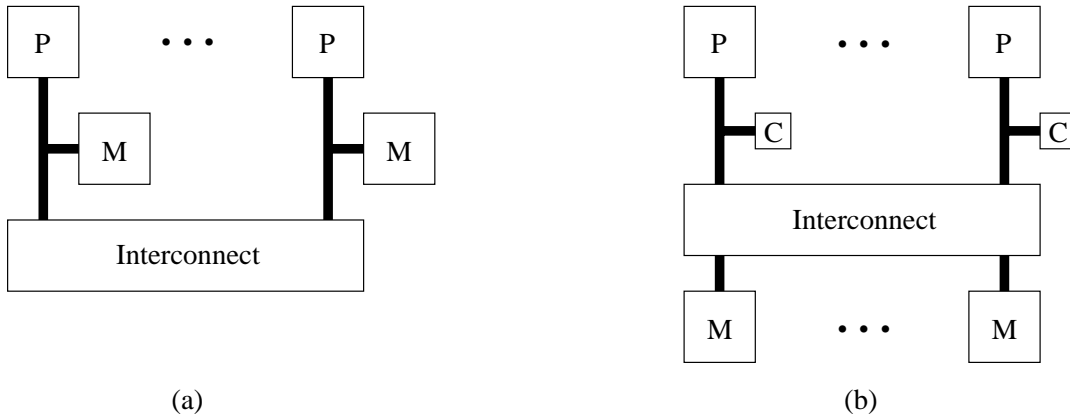


Figure 1: **(a)** DSM model. **(b)** CC model. In both insets, ‘P’ denotes a processor, ‘C’ a cache, and ‘M’ a memory module.

Examples. An r -bounded *fetch-and-increment* primitive on a variable v with range $0, \dots, r - 1$ is defined by $\phi(\text{old}, \text{input}) = \min(r - 1, \text{old} + 1)$. (In this primitive, the input parameter is not used, and hence we may simply assume $\alpha_p[j] = \perp$ for all p and j .) If v is initially 0, then any r consecutive invocations on v return distinct values, $0, 1, \dots, r - 1$. Moreover, any further invocation (after the r^{th}) returns $r - 1$, which is the same as the return value of the r^{th} invocation. Therefore, an r -bounded *fetch-and-increment* primitive has rank r , and an unbounded *fetch-and-increment* primitive has infinite rank.

For *fetch-and-increment* primitives, the input parameter α is extraneous. However, this is not the case for other primitives. Consider a *fetch-and-store* primitive on a variable with $2N + 1$ distinct values ($2N$ pairs $(p, 0)$ and $(p, 1)$, where p is a process, and an additional initial value \perp). By defining $\alpha_p[j] = (p, j \bmod 2)$, it is easily shown that *fetch-and-store* has infinite rank. (Informally, each process may write the information “this is an (even/odd)-indexed invocation by process p ” each time.) It also follows that an unbounded *fetch-and-store* primitive has infinite rank.

Finally, *test-and-set* has rank two: only the first *test-and-set* invocation on a variable initially *false* returns its initial value. *compare-and-swap* also has rank two.

3 A Constant-Time Generic Algorithm

In this section, we present an $O(1)$ mutual exclusion algorithm that uses a generic *fetch-and- ϕ* primitive, which is assumed to have rank at least $2N$. Two variants of the algorithm are presented, one for CC machines and one for DSM machines. These two architectural paradigms have been considered extensively in work on local-spin algorithms. Both are illustrated in Figure 1. In a DSM machine, each processor has its own memory module that can be accessed without accessing the global interconnection network. On such a machine, a shared variable can be made locally accessible by storing it in a local memory module. In a CC machine, each processor has a private cache, and some hardware protocol is used to enforce cache consistency (*i.e.*, to ensure that all copies of the same variable in different local caches are consistent). On such a machine, a shared variable becomes locally accessible by migrating to a local cache line. In this paper, we consider a DSM machine with caches that are kept coherent to be a CC machine.⁴ We also assume that there is a unique process executing the algorithm on each processor and that these processes do not migrate.

In local-spin algorithms for DSM machines, each process must have its own dedicated spin variables (which must be stored in its local memory module). In contrast, in algorithms for CC machines, processes may *share* spin variables, because each process can read a different cached copy. Because of this flexibility, algorithms

⁴Although virtually every modern multiprocessor is cache-coherent, non-cache-coherent DSM systems are still used in embedded applications, where cheaper computing technology often must be used due to cost limitations. Thus, the DSM model is of relevance for reasons other than historical interest.

for CC machines tend to be a bit simpler than those for DSM machines. This is why we present separate algorithms. Our CC algorithm, denoted G-CC, is presented first, and then its DSM counterpart, denoted G-DSM, is obtained by means of a fairly simple transformation. The two algorithms are shown in Figures 2 and 5. In both algorithms, “**await** B ,” where B is a boolean expression, is used as a shorthand for the busy-waiting loop “**while** $\neg B$ **do** /* null */ **od**.”

As noted earlier, we assess time complexity using the *RMR (remote-memory-references) time complexity measure*. As its name suggests, only remote memory references that cause an interconnect traversal are counted under this measure. We will assess the RMR time complexity of an algorithm by counting the total number of remote memory references required by one process to enter and then exit its critical section once. An algorithm may have different RMR time complexities under the CC and DSM models because the notion of a remote memory reference differs under these two models. In the CC model, we assume that, once a spin variable has been cached, it remains cached until it is either updated or invalidated as a result of being modified by another process on a different processor. (Effectively, we are assuming an idealized cache of infinite size: a cached variable may be updated or invalidated but it is never replaced by another variable because of associativity or capacity limitations.)

3.1 ALGORITHM G-CC: A Generic Algorithm for CC Machines

In this section, we present our generic algorithm for CC machines. Our intent here is to explain the various mechanisms used in the algorithm in an intuitive way. In so doing, we state and explain several formal properties of the algorithm. However, these explanations cannot be taken as a formal proof of correctness. For that, we refer the reader to Appendix A. At the end of this section, we illustrate some of the algorithm’s properties by considering a fairly comprehensive example execution.

When trying to implement a mutual exclusion algorithm using a generic *fetch-and- ϕ* primitive — of which only its rank r is known — the primary problem that arises is the following.

If the primitive is invoked more than r times to access a variable, then it may not be able to provide enough information for processes to order themselves. Therefore, *the algorithm must provide a means of resetting such a variable before it is accessed r times.*

Because we are using a primitive of rank $2N$ in ALGORITHM G-CC, we need to reset a variable accessed by the primitive before it is accessed $2N$ times. We do this by using two “waiting queues,” indexed 0 and 1. Associated with each queue j is a “tail pointer,” $Tail[j]$. In its entry section, a process enqueues itself onto one of these two queues by using the *fetch-and- ϕ* primitive to update its tail pointer, and waits on its predecessor, if necessary. At any time, one of the queues is designated as the “current” queue, which is indicated by the shared variable *CurrentQueue*. The other queue is called the “old” queue. The algorithm switches between the two queues over time in a way that ensures that each tail pointer is reset before being accessed $2N$ times. We now describe the reset mechanism in detail.

When a process p begins its entry section, it determines which queue is the current queue by reading the variable *CurrentQueue* (statement 3 of Figure 2), and then enqueues itself onto that queue using the *fetch-and- ϕ* primitive (statement 5). If p is not at the head of its queue ($p.prev \neq \perp$),⁵ then it waits until its predecessor in the queue updates the spin variable *Signal*[$p.idx$][$p.prev$] (statement 6), which p then resets (statement 7). Note that the local variable $p.prev$ (if not \perp) indicates a spin variable written by p ’s predecessor in the queue. The spin variable that a potential successor to p may wait on is indicated, from p ’s perspective, by the local variable $p.self$.

As explained below, it is possible for a process q to read *CurrentQueue* before another process updates *CurrentQueue* to switch to the other queue. Such a process q will then enqueue itself onto the old queue. Thus, *both* queues may possibly hold waiting processes. To arbitrate between processes in the two queues, an extra two-process mutual exclusion algorithm is used. A process competes in this two-process algorithm (after reaching the head of its waiting queue) using the routines **Entry**₂ and **Exit**₂, with the index of its queue as a “process identifier” (statements 8 and 12), as illustrated in Figure 3(a). Note that this extra two-process algorithm can be implemented from reads and writes in $O(1)$ time [13].

⁵We use $s.p$ to denote the statement with label s of process p , and $p.v$ to represent p ’s private variable v .

shared variables <i>CurrentQueue</i> : 0..1; <i>Tail</i> : array [0..1] of <i>Vartype</i> initially \perp ; <i>Position</i> : array [0..1] of 0..2 <i>N</i> - 1 initially 0; <i>Signal</i> : array [0..1][<i>Vartype</i>] of boolean initially <i>false</i> ; <i>Active</i> : array [0.. <i>N</i> - 1] of boolean initially <i>false</i> ; <i>QueueIdx</i> : array [0.. <i>N</i> - 1] of (\perp , 0..1)	private variables <i>idx</i> : 0..1; <i>counter</i> : integer ; <i>prev</i> , <i>self</i> , <i>tail</i> : <i>Vartype</i> ; <i>pos</i> : 0..2 <i>N</i> - 1
--	---


```

process p :: /* 0 ≤ p < N */
while true do
0:  Noncritical Section;
1:  QueueIdx[p] :=  $\perp$ ;
2:  Active[p] := true;
3:  idx := CurrentQueue;
4:  QueueIdx[p] := idx;
5:  prev := fetch-and-φ(Tail[idx],  $\alpha_p$ [counter]);
   self :=  $\phi$ (prev,  $\alpha_p$ [counter]);
   counter := counter + 1;
   if prev ≠  $\perp$  then
6:     await Signal[idx][prev];
7:     Signal[idx][prev] := false
   fi;
8:  Entry2(idx);
9:  Critical Section;
10: pos := Position[idx];
11: Position[idx] := pos + 1;
12: Exit2(idx);
13: if (pos < N) ∧ (pos ≠ p) then
14:     await ¬Active[pos] ∨
15:     (QueueIdx[pos] = idx)
   elseif pos = N then
16:     tail := Tail[1 - idx];
17:     Signal[1 - idx][tail] := false;
18:     Tail[1 - idx] :=  $\perp$ ;
19:     Position[1 - idx] := 0;
20:     CurrentQueue := 1 - idx
   fi;
21: Signal[idx][self] := true;
22: Active[p] := false
od

```

Figure 2: ALGORITHM G-CC: Generic *fetch-and-φ*-based mutual exclusion algorithm for CC machines.

As explained above, some process must reset the current queue before it is accessed $2N$ times. To facilitate this, each queue j has an associated shared variable $Position[j]$. This variable indicates the relative position of the current head of the queue, starting from 0. For example, in Figure 3(a), the head of queue 0 is at position 2, and hence $Position[0]$ equals 2. A process in queue j updates $Position[j]$ while still effectively in its critical section (statements 10 and 11). Thus, $Position[j]$ cannot be concurrently updated by different processes.

A process exchanges the role of the two queues in its exit section if it is at position N in the current queue (statements 16–20). (These statements will be explained in detail shortly.) Figures 3(b) and 3(c) show the state of the two queues before and after such an exchange. In order to exchange the queues, we must ensure the following property.

Property 1 If a process executes its critical section after having acquired position N of the current queue, then no process is in the old queue.

(A process is considered to be “in” the old queue if it read the index of that queue from *CurrentQueue*. In particular, that process may not yet have updated the queue’s tail pointer.) Given Property 1, a process at position N may safely reset the old queue and exchange the queues. (More formally, the fact that such an exchange is safe follows from invariant (I13) in Appendix A.) Property 1 is a direct consequence of the following property.

Property 2 If a process executes its critical section after having acquired position pos of the current queue, and if $pos > q$, then, by the time process p reaches statement 21, process q is not in the old queue.

To maintain Property 2, each process p has two associated variables, $Active[p]$ and $QueueIdx[p]$, which indicate (respectively) whether process p is active, and if so, which queue it is executing in (statements 1, 2, 4, and 22). If a process p executes at position $q < N$ in the current queue, then in its exit section, p waits

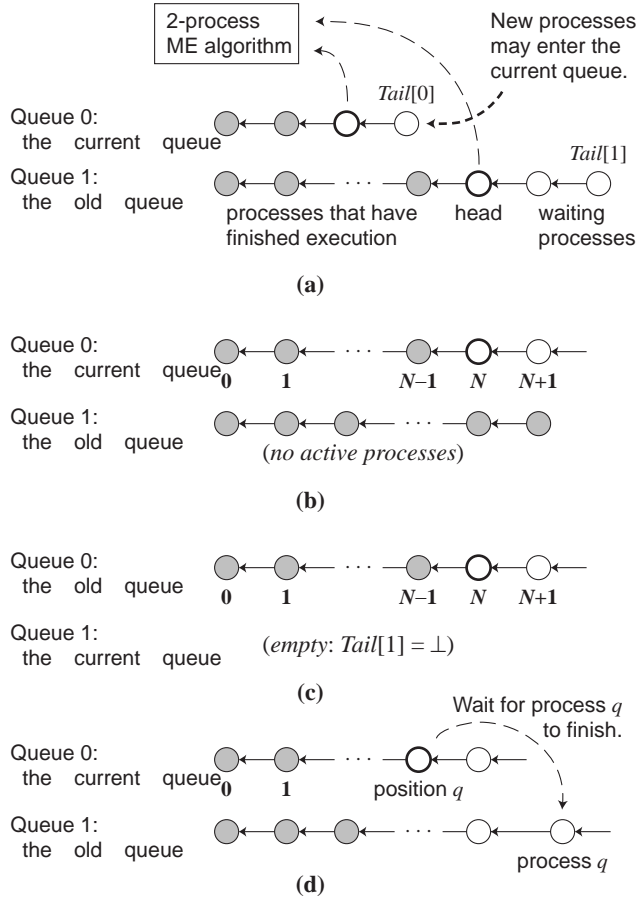


Figure 3: The structure of ALGORITHM G-CC. **(a)** The overall structure. This figure shows a possible state of execution when the current queue is queue 0. (The “finished” processes may be duplicated, because a process may execute its critical section multiple times.) **(b)** A state just before $CurrentQueue$ is updated. **(c)** A state just after $CurrentQueue$ is updated. **(d)** A process (in its exit section) in the current queue waiting for another in the old queue.

until either q finishes its exit-section execution (*i.e.*, $Active[p] = false$; statement 14) or enters the current queue (statement 15). p thus ensures that process q does not execute in the old queue, and then signals a possible successor (*i.e.*, a process at position $q + 1$ in the current queue) that it is now at the head of the current queue (statement 21). This situation is depicted in Figure 3(d).

Although p waits for q , starvation-freedom is guaranteed, because q is in the old queue, and hence makes progress independently of the current queue. Only the current queue is stalled until q finishes execution. (The fact that p may have to wait for a significant duration in its exit section may be a cause for concern. However, such waiting can be eliminated, if process p instructs process q to signal p 's successor *after* q finishes its critical section. Thus, p may finish execution without waiting for q . For simplicity, this handshake has not been added to ALGORITHM G-CC.)

We now explain statements 16–20, which are executed in order to exchange the role of the two queues. Without loss of generality, suppose that a process p executes these statements with $p.idx = 0$. (See Figure 3(b) and Figure 3(c).) Variables $Tail[1]$ and $Position[1]$ are initialized by statements 18 and 19, respectively. In addition, we must ensure that each entry of the $Signal[1][. . .]$ array is reset to *false*. Note that, if a process q in queue 1 establishes $Signal[1][x] = true$ (where $x = q.self$) by executing statement 21, then its successor r (which spins on $Signal[1][x]$ at statement 6) resets $Signal[1][x]$ by executing statement 7. Thus, by the time p executes statements 16–20, Property 1 ensures that every entry of $Signal[1][. . .]$ is reset to *false*, *except for the one set by the last process in queue 1*. (Clearly, the last process does not have a successor, so this entry is not

reset.)

Property 1 again ensures that this last entry of queue 1 is indicated by $Tail[1]$. Therefore, statements 16 and 17 properly reset this entry and thereby complete the reinitialization of $Signal[1][\dots]$. Finally, statement 20 exchanges the two queues.

We still must show that using a *fetch-and-φ* primitive of rank $2N$ is sufficient. Suppose that process p acquires position N of queue 0 when it is the current queue. We claim that at most $N - 1$ processes may be enqueued onto queue 0 after p and before the queues are exchanged again. For a process q to enqueue itself onto queue 0 after p , it must have read the value of $CurrentQueue$ before it was updated by p . For q to enqueue itself a *second* time onto queue 0, it must read $CurrentQueue = 0$ again, *after* $CurrentQueue = 1$ was established by p . This implies that the two queues have been exchanged again. (We remind the reader that, by the explanation above, the queues will not be exchanged again until there are no processes in queue 0.) Thus, after p establishes that queue 1 is current, and while queue 0 continues to be the old queue, at most $N - 1$ processes may be enqueued (after p) onto queue 0. Thus, we have the following property.

Property 3 Each process’s position is between 0 and $2N - 1$ (inclusive).

(Formally, this property follows from invariants (I15) and (I22) in Appendix A.) It follows that a rank of $2N$ is sufficient.

Example: To better illustrate the various mechanisms used in the algorithm, we will construct an example execution, which is illustrated in Figure 4. In this example, each $Tail$ variable is assumed to range over $0, \dots, 2N - 1$, with its initial value defined to be 0. (That is, the type $Vartype$ is taken to be $0..2N - 1$, and in the declaration of $Tail$ and in statements 5 and 18, the symbol \perp is taken to be 0.) Furthermore, we assume that the *fetch-and-φ* primitive being used is a $2N$ -bounded *fetch-and-increment* primitive.

The execution is constructed as follows. Starting from the initial state, processes 0 and 1 alternately enter and exit their critical sections while $CurrentQueue = 0$ continues to hold until process 0 obtains $0.prev = N$ (statement 5) and establishes $0.pos = N$ (statement 10). Assume at this point in the execution that process 0 is about to execute statement 13 and all other processes are in their noncritical sections.

Before continuing, note that, in the execution described so far, had one of processes 0 and 1 waited on the other at statement 6, then when its waiting ceased, it would have immediately reset the correspond $Signal$ variable to *false* in statement 7. Thus, at this point in the execution, all $Signal$ variables in queue 0 are *false*. The situation so far is illustrated in Figure 4(a).

To continue, suppose that process 1 executes its entry section until it busy waits at statement 6, then process 2 does the same, followed in the same way (in order) by processes 3 through $N - 1$. Then, prior to waiting at statement 6, each process k , where $1 \leq k < N$, obtains $k.prev = N + k$. That is, the spin variable that process k ’s predecessor in queue 0 updates is at position $N + k$. (Note that, for our particular choice of *fetch-and-φ* primitive, we have $p.pos = p.prev$ for each process p , but in general these two values may have no relation with each other.)

Now, suppose that process 0 executes statements 16–20 to switch the queues. This establishes the following:

- $0.tail = 0$ (statement 16; recall that $\perp = 0$ is the initial value of $Tail[1]$);
- $Signal[1][0] = false$ (statement 17; this actually has no impact because all $Signal$ variables are initially *false*);
- $Tail[1] = 0$ (statement 18; again, \perp is defined to be 0);
- $Position[1] = 0$ (statement 19); and
- $CurrentQueue = 1$ (statement 20).

Next, suppose that process 0 establishes $Signal[0][N + 1] = true$ (statement 21; note that $0.prev = N$ and $0.self = N + 1$) and $Active[0] = false$ (statement 22) and returns to its noncritical section. Before continuing with the execution, note that, at this point, because $Signal[0][N + 1] = true$ holds, process 1 is enabled to transit from statement 6 to statement 7. When it later executes statement 7, it will re-establish $Signal[0][N + 1] = false$.

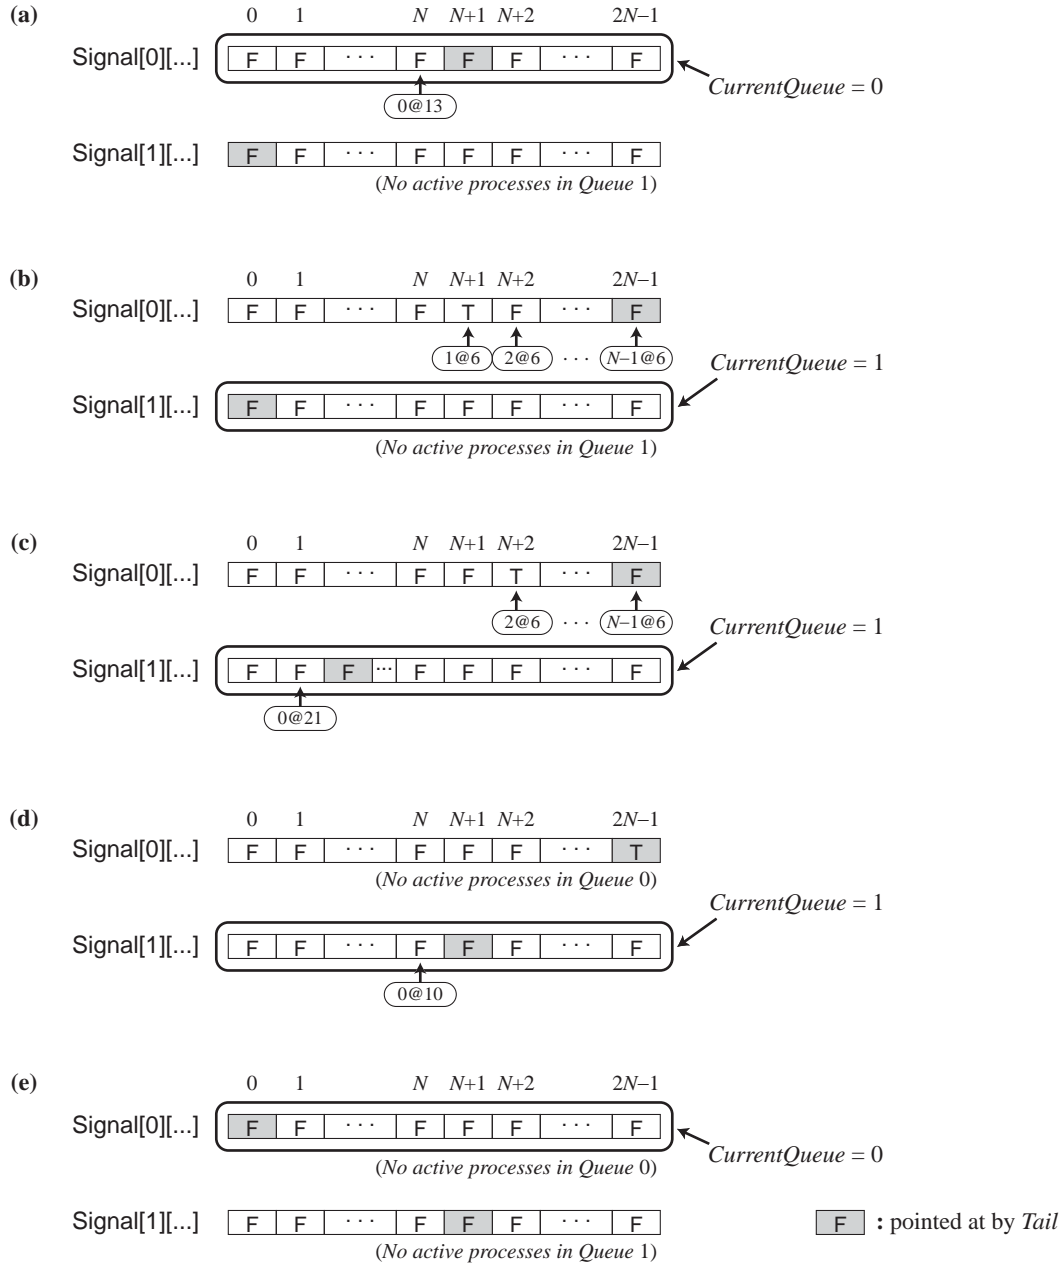


Figure 4: Example execution of of ALGORITHM G-CC. The notation $p@s$ means that process p is enabled to execute statement s . In each inset, the value of each active process p 's $prev$ variable is indicated; this value indicates a spin variable that p 's predecessor updates and that p may potentially block on. For example, in inset (a), $0.prev = N$.

In a similar way, when processes 2 through $N - 1$ later execute statement 7, each of $Signal[0][N + 2]$ through $Signal[0][2N - 1]$ will be reset to *false*. However, process $N - 1$ later will establish $Signal[0][2N - 1] = true$ when it executes statement 21. (A $2N$ -bounded *fetch-and-increment* operation on a variable with the value $2N - 1$ leaves the variable's value unchanged, so $(N - 1).prev = 2N - 1$ and $(N - 1).self = 2N - 1$.) We must make sure that, as the execution progresses, $Signal[0][2N - 1] = false$ is re-established before queue 0 becomes the current queue again. Note that, at this point in the execution, $Tail[0] = 2N - 1$ holds. (This fact will be used later.) The situation so far is illustrated in Figure 4(b).

To continue the execution, note that, to the point described so far, process 0 is in its noncritical section, and

each other process is in queue 0, with process 1 stalled at statement 6 and processes 2 through $N - 1$ blocked there. Now, assume that process 0 enters its entry section once again. Then, it will read $CurrentQueue = 1$ (statement 3) and obtain $0.prev = 0$ (statement 5). Assuming that each of the remaining processes remains at statement 6 (so none acquires the two-process lock), process 0 may progress to statement 10, where it establishes $0.pos = 0$, and then to statement 13. Because $0.pos = 0$ holds, process 0 next transits to statement 21. After executing statements 21 and 22, suppose that process 0 returns to its noncritical section.

Now, suppose that process 0 repeats the above actions while the other processes remain at statement 6. Then, the scenario is similar, but process 0 now obtains $0.prev = 1$ (statement 5) and then establishes $0.pos = 1$ (statement 10). Thus, when process 0 reaches statement 13, it will transit to statement 14, where it will be blocked until process 1 is either no longer active or has accessed queue 1. Suppose that process 1 completes its entry, critical, and exit sections, and returns to its noncritical section, after which, process 0 (because it is no longer blocked) transits to statement 21, as illustrated in Figure 4(c). At this point in the execution, process 1 is no longer in queue 0 (see Property 2). Suppose now that process 0 completes its exit section and also returns to its noncritical section.

Suppose that we repeat this same scenario until process 0 establishes $0.pos = N$ by executing statement 10. Then, by this point in the execution, no processes can be in queue 0 (see Property 1). This is illustrated in Figure 4(d). Thus, when process 0 switches the queues again by subsequently executing statements 16–20, it is safe to switch the two queues. Further, note that, when process 0 performs statement 16, it will obtain $0.tail = Tail[0] = 2N - 1$. (This is the most recent value assigned to $Tail[0]$ in the execution.) Thus, when it performs statement 17, it will establish $Signal[0][2N - 1] = false$, which as noted earlier, must be done before queue 0 becomes the current queue again. The execution at this point is depicted in Figure 4(e). \square

The busy-waiting loops at statements 6, 14, and 15 in Figure 2 are read-only loops in which variables are read that may be updated by a unique process. On a CC machine, each such loop incurs $O(1)$ RMR time complexity. It follows that ALGORITHM G-CC has $O(1)$ RMR time complexity on CC machines.

As noted earlier, a detailed correctness proof of ALGORITHM G-CC is given in Appendix A. From the discussion so far, we have the following lemma.

Lemma 1 *If the underlying fetch-and- ϕ primitive has rank at least $2N$, then ALGORITHM G-CC is a correct, starvation-free mutual exclusion algorithm with $O(1)$ RMR time complexity in CC machines.* \square

3.2 ALGORITHM G-DSM: A Generic Algorithm for DSM Machines

We now explain how to convert ALGORITHM G-CC into ALGORITHM G-DSM, which is illustrated in Figure 5. The key idea of this conversion is a simple transformation of each busy-waiting loop, which we examine here in isolation. In ALGORITHM G-CC, all busy waiting is by means of statements of the form “**await** B ,” where B is some boolean condition. Moreover, if a process p is waiting for condition B to hold, then there is a unique process that can establish B , and once B is established, it remains true, until p ’s “**await** B ” statement terminates.

In ALGORITHM G-DSM, each statement of the form “**await** B ” has been replaced by the code fragment on the left below (see statements 10–17 and 25–33 in Figure 5), and each statement of the form “ $B := true$ ” by the code fragment on the right (see statements 4–8, 40–44, and 45–49).

<pre> a: Entry₂(\mathcal{J}, 0); b: $flag := B$; c: $Waiter[\mathcal{J}] := \text{if } flag \text{ then } \perp \text{ else } p$; d: $Spin[p] := false$; e: Exit₂(\mathcal{J}, 0); f: if $\neg flag$ then g: await $Spin[p]$; h: $Waiter[\mathcal{J}] := \perp$ fi </pre>	<pre> i: Entry₂(\mathcal{J}, 1); j: $B := true$; k: $next := Waiter[\mathcal{J}]$; l: Exit₂($\mathcal{J}$, 1); m: if $next \neq \perp$ then $Spin[next] := true$ fi </pre>
---	---

The variable $Waiter[\mathcal{J}]$ is assumed to be initially \perp , and $Spin[p]$ is a spin variable used exclusively by process

/* all variable declarations are as defined in Figure 2 except as noted here */

shared variables

Waiter1: **array**[0..*N* - 1] **of** (\perp , 0..*N* - 1);
Waiter2: **array**[0..1][*Vartype*] **of** (\perp , 0..*N* - 1);
Spin: **array**[0..*N* - 1] **of** **boolean** **initially** *false*

private variables

next: (\perp , 0..*N* - 1);
flag: **boolean**;
q: 0..*N* - 1

process *p* :: /* 0 ≤ *p* < *N* */

while *true* **do**

0: Noncritical Section;
1: *QueueIdx*[*p*] := \perp ;
2: *Active*[*p*] := *true*;
3: *idx* := *CurrentQueue*;
4: **Entry**₂(*p*, 1);
5: *QueueIdx*[*p*] := *idx*;
6: *q* := *Waiter1*[*p*];
7: **Exit**₂(*p*, 1);
8: **if** *q* ≠ \perp **then** *Spin*[*q*] := *true* **fi**;
9: *prev* := *fetch-and-φ*(*Tail*[*idx*], α_p [*counter*]);
self := ϕ (*prev*, α_p [*counter*]);
counter := *counter* + 1;
if *prev* ≠ \perp **then**
10: **Entry**₂((*idx*, *prev*), 0);
11: *flag* := *Signal*[*idx*][*prev*];
12: *Waiter2*[*idx*][*prev*] :=
if *flag* **then** \perp **else** *p*;
13: *Spin*[*p*] := *false*;
14: **Exit**₂((*idx*, *prev*), 0);
15: **if** \neg *flag* **then**
16: **await** *Spin*[*p*];
17: *Waiter2*[*idx*][*prev*] := \perp
fi;
18: *Signal*[*idx*][*prev*] := *false*
fi;
19: **Entry**₂(*idx*)
20: Critical Section;
21: *pos* := *Position*[*idx*];
22: *Position*[*idx*] := *pos* + 1;
23: **Exit**₂(*idx*);

24: **if** (*pos* < *N*) ∧ (*pos* ≠ *p*) **then**
q := *pos*;
25: **Entry**₂(*q*, 0);
26: *flag* := \neg *Active*[*q*] ∨
27: (*QueueIdx*[*q*] = *idx*);
28: *Waiter1*[*q*] :=
if *flag* **then** \perp **else** *p*;
29: *Spin*[*p*] := *false*;
30: **Exit**₂(*q*, 0);
31: **if** \neg *flag* **then**
32: **await** *Spin*[*p*];
33: *Waiter1*[*q*] := \perp
fi
elseif *pos* = *N* **then**
34: *temp* := *Tail*[1 - *idx*];
35: *Signal*[1 - *idx*][*temp*] := *false*;
36: *Tail*[1 - *idx*] := \perp ;
37: *Position*[1 - *idx*] := 0;
38: *CurrentQueue* := 1 - *idx*
fi;
39: **if** *pos* < 2*N* - 1 **then**
40: **Entry**₂((*idx*, *self*), 1);
41: *Signal*[*idx*][*self*] := *true*;
42: *next* := *Waiter2*[*idx*][*self*];
43: **Exit**₂((*idx*, *self*), 1);
44: **if** *next* ≠ \perp **then** *Spin*[*next*] := *true* **fi**
fi;
45: **Entry**₂(*p*, 1);
46: *Active*[*p*] := *false*;
47: *next* := *Waiter1*[*p*];
48: **Exit**₂(*p*, 1);
49: **if** *next* ≠ \perp **then** *Spin*[*next*] := *true* **fi**
od

Figure 5: ALGORITHM G-DSM: Generic *fetch-and-φ*-based mutual exclusion algorithm for DSM machines. Statements different from Figure 2 are shown with **boldface** line numbers.

p (and, hence, it can be stored in memory local to p). Entry_2 and Exit_2 represent an instance of a two-process mutual exclusion algorithm, indexed by \mathcal{J} , *i.e.*, \mathcal{J} is used to identify a particular instance of the two-process mutual exclusion algorithm. To see that this transformation is correct, assume that a process p executes lines a–h while another process q executes lines i–m. Since lines b–d and j–k execute within a critical section, lines b–d precede lines j–k, or vice versa. If b–d precede j–k, and if $B = \text{false}$ holds before the execution of b–d, then p assigns $\text{Waiter}[\mathcal{J}] := p$ at line c, and initializes its spin variable at line d. Process q subsequently reads $\text{Waiter}[\mathcal{J}] = p$ at line k, and establishes $\text{Spin}[p] = \text{true}$ at line m, which ensures that p is not blocked. On the other hand, if lines j–k precede lines b–d, then process q reads $\text{Waiter}[\mathcal{J}] = \perp$ (the initial value) at line k, and does not update any spin variable at line m. Since process p executes line b after q executes line j, p preserves $\text{Waiter}[\mathcal{J}] = \perp$, and does not execute lines g and h.

Note that for the above correctness argument to be valid, it is crucial that, at any time, at most one process executes within lines a–h and within lines i–m. Otherwise, the correctness of the underlying two-process mutual exclusion algorithm cannot be guaranteed. In order to satisfy this property, we have to apply a minor modification to ALGORITHM G-CC before applying the transformation shown above. Namely, we change statement 21 of ALGORITHM G-CC as follows. (See also statement 39 in Figure 5.)

```

21: if  $pos < 2N - 1$  then
     $\text{Signal}[idx][self] := \text{true}$ 
fi;

```

We hereafter denote this modified algorithm as ALGORITHM G-CC'. We now informally argue that this modification does not affect the algorithm's correctness. Consider a process p that executes statement 21 while at position $2N - 1$. Let $i = p.idx$ and $x = p.self$. By Property 3, p cannot possibly have a successor. Moreover, a process may read $\text{Signal}[i][x]$ (at statement 6) only if its predecessor has written (or will eventually write) $\text{Signal}[i][x] := \text{true}$ by executing statement 21. Hence, in ALGORITHM G-CC, the value of $\text{Signal}[i][x]$ written by p at statement 21 is never read by any process until the queues are exchanged again, at which point the process that performs the exchange overwrites $\text{Signal}[i][x]$ by executing statement 17.

From the preceding discussion, it follows that ALGORITHM G-CC' is a correct, starvation-free mutual exclusion algorithm, to which we can then apply the transformation shown above and obtain ALGORITHM G-DSM. Moreover, all the invariants stated in Appendix A also remain valid for ALGORITHM G-CC', except for those that directly refer to the *Signal* array. (We will explain shortly why we cannot apply the above transformation to ALGORITHM G-CC directly; see the reasoning for statements 40–43 below.) A formal version of the argument given here regarding ALGORITHM G-CC' is given in Appendix B.

We now argue that the various two-process mutual exclusion algorithms added to Figure 5 are safe by considering each pair of Entry/Exit calls in this figure in turn. By *safe*, we mean that, if some process is executing within a code fragment that begins with $\text{Entry}_2(\mathcal{J}, b)$ and ends with $\text{Exit}_2(\mathcal{J}, b)$, where \mathcal{J} identifies this particular two-process mutual exclusion instance and b is 0 or 1, then no other process may be concurrently executing within a code fragment that begins with $\text{Entry}_2(\mathcal{J}, b)$ and ends with $\text{Exit}_2(\mathcal{J}, b)$. (The arguments below are also presented more formally in Appendix B.)

Statements 4–7 and statements 45–48: The Entry/Exit calls here are clearly safe because of the usage of the identifier ' p ', which uniquely corresponds to process p .

Statements 10–14: Because the *fetch-and- ϕ* primitive used in ALGORITHM G-CC' has rank at least $2N$, by Property 3, each process that executes statements 5–7 of ALGORITHM G-CC' within a particular queue must have a distinct value of *prev*. This implies that the Entry/Exit pair at statements 10–15 in Figure 5 is safe.

Statements 25–30: Since the queues are switched after a process at position N executes its critical section, if a process executes statements 14 and 15 of ALGORITHM G-CC', then at that time, no process with the same value of *pos* may be executing within the other queue. Since the algorithm clearly ensures that each process executing within the same queue has a distinct value of *pos*, this implies that the Entry/Exit pair at statements 25–30 in Figure 5 is safe.

Statements 40–43: Because the *fetch-and-φ* primitive has rank at least $2N$, each process at queue i (for some i) whose position is between 0 and $2N - 2$ (inclusive) writes a distinct value to $Tail[i]$ at statement 6 of ALGORITHM G-CC', and hence, has a distinct value of *self*. This implies that the **Entry/Exit** pair at statements 40–43 in Figure 5 is safe.

Note that, if process p is at position $2N - 1$ of queue i , then p is allowed to write any value to $Tail[i]$, by the definition of rank. (By Property 3, p cannot have a successor.) Hence, in ALGORITHM G-CC, p 's execution of statement 21 may write the same element of *Signal* as some other process residing at an earlier position. In particular, in ALGORITHM G-CC, if some other process q has acquired position h ($< 2N - 1$) of queue i , then both statements 21. p and 21. q may write the same element of *Signal*.

This does not pose a problem for ALGORITHM G-CC, because the execution of statement 21 is atomic, and hence p may execute statements 8–21 only after all of its predecessors have finished executing statement 21. However, this *would* be a problem if we were to implement statement 21 with statements i–m as in ALGORITHM G-DSM. This is why we introduce an additional “**if**” clause in ALGORITHM G-CC'.⁶

The above transformation can also be applied within other algorithms, as long as a safety proof like that above can be established. For example, this transformation can be applied to convert the algorithm of Graunke and Thakkar [7] to a variant that locally spins on DSM machines.⁷ Given the correctness of the above transformation, we have the following.

Lemma 2 *If the underlying fetch-and-φ primitive has rank at least $2N$, then ALGORITHM G-DSM is a correct, starvation-free mutual exclusion algorithm with $O(1)$ RMR time complexity in DSM machines.* □

If we have a *fetch-and-φ* primitive with rank r ($4 \leq r < 2N$), then we can arrange instances of ALGORITHM G-DSM in an arbitration tree, where each process is statically assigned a leaf node and each non-leaf node consists of an $\lfloor r/2 \rfloor$ -process mutual exclusion algorithm, implemented using ALGORITHM G-DSM. Because this arbitration tree is of $\Theta(\log_r N)$ height, we have the following theorem. (Note that for $r = 2$ or 3, a $\Theta(\log_r N)$ algorithm is possible without even using the *fetch-and-φ* primitive [13].)

Theorem 1 *Using any fetch-and-φ primitive of rank $r \geq 2$, starvation-free mutual exclusion can be implemented with $\Theta(\max(1, \log_r N))$ RMR time complexity on either CC or DSM machines.* □

It is possible to combine the arbitration-tree algorithm just described with an adaptive mutual exclusion algorithm presented previously by us [10]. The adaptive algorithm utilizes two trees, called *renaming* and *overflow trees*, respectively, as illustrated in Figure 6. A “name” is associated with each node within the renaming tree. A process begins execution by attempting to acquire one of these names. It does so by entering the renaming tree at its root and descending. As a process descends this tree, it may either stop, move left, or move right at a node. If a process stops at a node, then it begins moving upwards within the renaming tree — in this phase of its execution, the renaming tree is just an arbitration tree. It is possible for a process to fail to acquire a name within the renaming tree, in which case it enters the overflow tree at a designated leaf node. A process can fail to acquire a name only if the point contention⁸ it experiences exceeds (asymptotically) the renaming tree’s height. The overflow tree is simply a second arbitration tree. An extra two-process mutual exclusion algorithm is used to arbitrate between the “winning” processes coming from these two trees. As shown in [10], if a process experiences point contention k , then it will acquire a name in the renaming tree in $O(k)$ time (provided the renaming tree is of sufficient height). In this case, its overall RMR time complexity for entering and exiting its critical section is $O(k)$. By defining the renaming tree’s height to be $\Theta(\log N)$, the

⁶In the conference version of this paper [3], we presented a DSM algorithm as a direct transformation of ALGORITHM G-CC. We later discovered that an additional “**if**” clause is necessary, as explained here.

⁷In [3], we stated that the above transformation could also be applied to the algorithm of T. Anderson [5]. However, Hyonho Lee of the University of Toronto later found that this assertion is not correct because in that algorithm it is possible for more than two processes to invoke one of the added two-process mutual exclusion algorithms at the same time. Nonetheless, a similar transformation can be applied to that algorithm where the availability of the *fetch-and-increment* primitive is exploited to avoid the **Entry** and **Exit** calls.

⁸The *point contention* experienced by a process p is the maximum number of processes that are *simultaneously* active (*i.e.*, outside of their noncritical sections) over a computation that starts when p becomes active and ends when it once again becomes inactive [1].

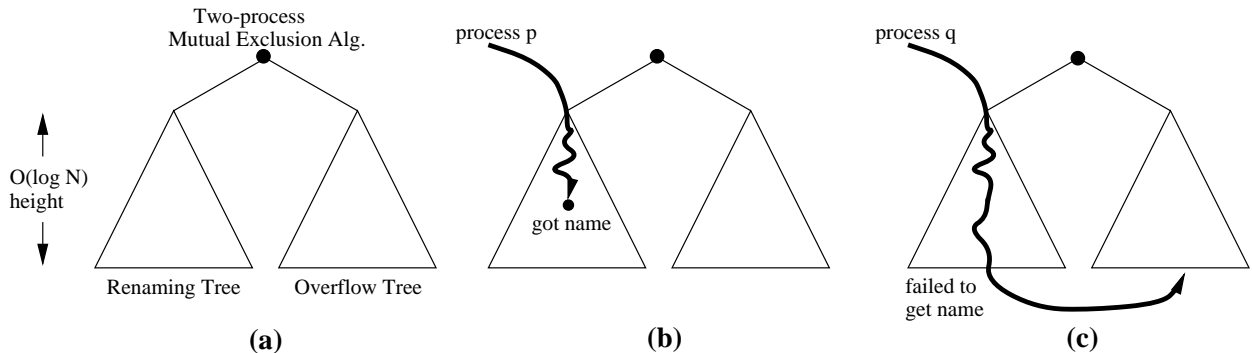


Figure 6: (a) Renaming tree and overflow tree. (b) Process p gets a name in the renaming tree. (c) Process q fails to get a name and must compete within the overflow tree.

overall RMR time complexity becomes $O(\min(k, \log N))$. The renaming tree is implemented using only atomic reads and writes. If we change its height to $\Theta(\max(1, \log_r N))$ and use the *fetch-and- ϕ* -based arbitration-tree algorithm described above to implement the overflow tree, then the overall RMR time complexity becomes $O(\min(k, \max(1, \log_r N)))$. Thus, we have the following theorem.

Theorem 2 *Using any fetch-and- ϕ primitive of rank $r \geq 2$, starvation-free adaptive mutual exclusion can be implemented with $O(\min(k, \max(1, \log_r N)))$ RMR time complexity on either CC or DSM machines, where k is point contention. \square*

In presenting our algorithms, we have assumed that process identifiers range over $0, \dots, N - 1$. With one exception, this assumption is simply a matter of convenience. That one exception is the code sequence in lines 13–15 of ALGORITHM G-CC in Figure 2, where a process p with $p.pos = k$, where $0 \leq k < N$, waits until process k is either inactive or is accessing the same queue as p . This code sequence, which is part of the reset mechanism, creates a linkage between the range of process identifiers and the range of queue positions. This linkage can be eliminated by introducing a third queue, which each process uses only for its *first* critical-section execution. Because each process uses this queue only once, it does not need to be reset. If a process p obtains $p.pos = k$ in this queue, then it uses the process identifier k in its subsequent critical-section executions, which are implemented using the two queues considered earlier. Because three queues are being used now, a three-process mutual exclusion algorithm must be used instead of a two-process algorithm in order to arbitrate among the “winning” processes from these queues.

4 Concluding Remarks

We have presented a ranking of *fetch-and- ϕ* primitives based on the time complexity with which such primitives can be used to implement mutual exclusion. We have also shown that any *fetch-and- ϕ* primitive of rank r can be used to implement a $\Theta(\max(1, \log_r N))$ mutual exclusion algorithm, on either DSM or CC machines. $\Theta(\max(1, \log_r N))$ is clearly optimal for $r = \Omega(N)$. However, as remarked earlier, it follows from work appearing in the second author’s Ph.D. dissertation [9] that $\Theta(\max(1, \log_r N))$ is *not* optimal for certain “self-resettable” primitives of rank at least three. Devising asymptotically optimally algorithms for primitives of arbitrary ranks remains as future work. It is important to note that, in designing these algorithms, our main goal was to achieve certain asymptotic time complexities. In particular, we have *not* concerned ourselves with designing algorithms that can be practically applied. Indeed, it is difficult to design practical algorithms when assuming so little of the *fetch-and- ϕ* primitives being used. It is likely that by exploiting the semantics of a particular primitive, our algorithms could be optimized considerably.

We believe that the notion of rank defined in this paper may be a suitable way of characterizing the “power” of primitives from the standpoint of blocking synchronization, much like the notion of a *consensus number*, which is used in Herlihy’s wait-free hierarchy [8], reflects the “power” of primitives from the standpoint of

nonblocking synchronization. Interestingly, primitives like *compare-and-swap* that are considered to be powerful according to Herlihy's hierarchy are weak from a blocking synchronization standpoint (since they are subject to our $\Omega(\log N / \log \log N)$ lower bound [2]). Also, primitives like *fetch-and-increment* and *fetch-and-store* that are considered to be powerful from a blocking synchronization standpoint are considered quite weak according to Herlihy's hierarchy. (They have consensus number two.) This difference arises because in nonblocking algorithms, the need to reach consensus is fundamental (as shown by Herlihy), while in blocking algorithms, the need to order competing processes is important.

References

- [1] Y. Afek, H. Attiya, A. Fouren, G. Stupp, and D. Touitou. Long-lived renaming made adaptive. In *Proceedings of the 18th Annual ACM Symposium on Principles of Distributed Computing*, pages 91–103. ACM, May 1999.
- [2] J. Anderson and Y.-J. Kim. An improved lower bound for the time complexity of mutual exclusion. *Distributed Computing*, 15(4):221–253, December 2003.
- [3] J. Anderson and Y.-J. Kim. Local-spin mutual exclusion using fetch-and- ϕ primitives. In *Proceedings of the 23rd IEEE International Conference on Distributed Computing Systems*, pages 538–547. IEEE, May 2003.
- [4] J. Anderson and J.-H. Yang. Time/contention tradeoffs for multiprocessor synchronization. *Information and Computation*, 124(1):68–84, January 1996.
- [5] T. Anderson. The performance of spin lock alternatives for shared-memory multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 1(1):6–16, January 1990.
- [6] T. Craig. Queuing spin lock algorithms to support timing predictability. In *Proceedings of the 14th IEEE Real-time Systems Symposium*, pages 148–156. IEEE, December 1993.
- [7] G. Graunke and S. Thakkar. Synchronization algorithms for shared-memory multiprocessors. *IEEE Computer*, 23:60–69, June 1990.
- [8] M. Herlihy. Wait-free synchronization. *ACM Transactions on Programming Languages and Systems*, 13(1):124–149, 1991.
- [9] Y.-J. Kim. *Time Complexity Bounds for Shared-memory Mutual Exclusion*. Ph.D. thesis, University of North Carolina, Chapel Hill, NC, 2003.
- [10] Y.-J. Kim and J. Anderson. Adaptive mutual exclusion with local spinning. *Distributed Computing*, to appear.
- [11] P. Magnusson, A. Landin, and E. Hagersten. Queue locks on cache coherent multiprocessors. In *Proceedings of the 8th International Symposium on Parallel Processing*, pages 165–171. IEEE, April 1994.
- [12] J. Mellor-Crummey and M. Scott. Algorithms for scalable synchronization on shared-memory multiprocessors. *ACM Transactions on Computer Systems*, 9(1):21–65, February 1991.
- [13] J.-H. Yang and J. Anderson. A fast, scalable mutual exclusion algorithm. *Distributed Computing*, 9(1):51–60, August 1995.

Appendix A: Correctness Proof for ALGORITHM G-CC

In this appendix, we formally prove that ALGORITHM G-CC, presented in Section 3.1, satisfies the Exclusion and Starvation-Freedom properties. We use invariants and *leads-to* and *unless* properties in our proof. To define such notions, some additional definitions are required. We define a *state* of a program to be an assignment of values to the variables of the program, including process program counters. An *event* is a particular execution of a statement. A *history* of a program is a sequence $t_1, e_1, t_2, e_2, \dots$, where each t_j is a state, each e_j is an event, t_1 is an initial state, and state t_{j+1} can be reached from state t_j via the execution of the statement corresponding to e_j . A history is *fair* if and only if each continuously enabled statement is eventually executed. A non-**await** statement is enabled if the program counter of the process in which it appears equals the label of that statement. (We further comment on statement labels below.) To ease the reasoning a bit, we assume that each **await** statement induces a single state transition, even though it actually represents a busy-waiting loop. That is, we assume that a statement of the form “**await** B ” is enabled if the above condition for non-**await** statements holds, and in addition, the boolean condition B is true; when executed, such a statement causes a single transition to the next executable statement. An assertion is an *invariant* for a program if and only if it is true at each state of every history of that program. Letting A and B denote assertions, A *leads-to* B is true for a program if and only if the following holds for any *fair* history of the program: if A holds at some state in the history, then B holds at either the same state or some later state. A *unless* B is true for a program if and only if, for any history of the program, if $A \wedge \neg B$ holds at some state in the history, then $A \vee B$ holds at the next state. Informally, A is not falsified until B is established.

The following notational conventions are used in the remainder of this appendix.

Notational Conventions: We use $s.p$ to denote the statement with label s of process p , and $p.v$ to represent p 's private variable v (see Footnote 5). Let S be a subset of the statement labels in process p . Then, $p@S$ holds if and only if the program counter for process p equals some value in S . (Note that if s is a statement label, then $p@\{s\}$ means that statement s of process p is *enabled*, i.e., p has not yet executed s .)

We assume that each labeled sequence of statement(s) is atomic. For example, consider statement 15 of Figure 2. When executed by process p , this statement (atomically) updates $p.prev$, $p.self$, and $p.counter$, and establishes $p@\{6\}$ if the value assigned to $p.prev$ differs from \perp , and establishes $p@\{8\}$, otherwise. We number statements in this manner to reduce the number of cases that must be considered in the correctness proof. It can be seen that each labeled statement accesses at most one non-auxiliary (see below) shared variable, and does so via a single read, write, or *fetch-and- ϕ* operation. \square

Our proof also makes use of a number of auxiliary variables. In Figure 7, ALGORITHM G-CC is shown with these added auxiliary variables, which are accessed only by statements 5 and 18. We begin with a description of these variables.

Private auxiliary variable $p.position$ keeps track of p 's position in the queue; as shown by invariant (I49) below, when p is in its exit section, $p.position$ equals the non-auxiliary variable $p.pos$.

For $i = 0$ or 1 , $HistLen[i]$, $Hist[i][0..\infty]$, and $Param[i][0..\infty]$ keep the history of queue i since the last time it was initialized. (As explained in Section 3.1, each queue is accessed at most $2N$ times before it is re-initialized. Thus, only finite prefixes of $Hist[i][..]$ and $Param[i][..]$ are actually used.)

Example. Assume that queue i is initially empty. Then, initially we have the following.

$$\begin{aligned} Tail[i] &= \perp; \\ HistLen[i] &= 0; \\ Hist[i] &= (\perp, \perp, \perp, \dots); \\ Param[i] &= (\perp, \perp, \perp, \dots). \end{aligned}$$

Now suppose that processes p , q , r , and p execute statement 5 (in that order), and that the private variables $p.counter$, $q.counter$, and $r.counter$ initially equal 0, 3, and 5, respectively.

If the underlying *fetch-and- ϕ* primitive is *fetch-and-store* with $2N+1$ distinct values (as defined in Section 2),

shared variables

CurrentQueue: 0..1;
Tail: **array**[0..1] of *Vartype* **initially** \perp ;
Position: **array**[0..1] of 0..2*N* - 1 **initially** 0;
Signal: **array**[0..1][*Vartype*] of **boolean** **initially** *false*;
Active: **array**[0..*N* - 1] of **boolean** **initially** *false*;
QueueIdx: **array**[0..*N* - 1] of (\perp , 0..1)

type

ParamType = **record**
 proc: 0..*N* - 1;
 counter: **integer**
end

shared auxiliary variables

HistLen: **array**[0..1] of 0.. ∞ **initially** 0;
Hist: **array**[0..1][0.. ∞] of *Vartype* **initially** \perp ;
Param: **array**[0..1][0.. ∞] of **ParamType** **initially** \perp

process *p* :: /* 0 ≤ *p* < *N* */

while true do

<p>0: Noncritical Section;</p> <p>1: <i>QueueIdx</i>[<i>p</i>] := \perp;</p> <p>2: <i>Active</i>[<i>p</i>] := <i>true</i>;</p> <p>3: <i>idx</i> := <i>CurrentQueue</i>;</p> <p>4: <i>QueueIdx</i>[<i>p</i>] := <i>idx</i>;</p> <p>5: /* atomically execute lines 5a–5h */</p> <p>5a: <i>prev</i> := <i>Tail</i>[<i>idx</i>];</p> <p>5b: <i>self</i> := ϕ(<i>prev</i>, α_p[<i>counter</i>]);</p> <p>5c: <i>Tail</i>[<i>idx</i>] := <i>self</i>;</p> <p>5d: <i>position</i> := <i>HistLen</i>[<i>idx</i>];</p> <p>5e: <i>Param</i>[<i>idx</i>][<i>position</i>] := (<i>p</i>, <i>counter</i>);</p> <p>5f: <i>Hist</i>[<i>idx</i>][<i>position</i> + 1] := <i>self</i>;</p> <p>5g: <i>HistLen</i>[<i>idx</i>] := <i>position</i> + 1;</p> <p>5h: <i>counter</i> := <i>counter</i> + 1;</p> <p>if <i>prev</i> \neq \perp then</p> <p>6: await <i>Signal</i>[<i>idx</i>][<i>prev</i>];</p> <p>7: <i>Signal</i>[<i>idx</i>][<i>prev</i>] := <i>false</i></p> <p> fi;</p> <p>8: Entry₂(<i>idx</i>);</p> <p>9: Critical Section;</p>	<p>10: <i>pos</i> := <i>Position</i>[<i>idx</i>];</p> <p>11: <i>Position</i>[<i>idx</i>] := <i>pos</i> + 1;</p> <p>12: Exit₂(<i>idx</i>);</p> <p>13: if (<i>pos</i> < <i>N</i>) \wedge (<i>pos</i> \neq <i>p</i>) then</p> <p>14: await \neg<i>Active</i>[<i>pos</i>] \vee</p> <p>15: (<i>QueueIdx</i>[<i>pos</i>] = <i>idx</i>)</p> <p> elseif <i>pos</i> = <i>N</i> then</p> <p>16: <i>tail</i> := <i>Tail</i>[1 - <i>idx</i>];</p> <p>17: <i>Signal</i>[1 - <i>idx</i>][<i>tail</i>] := <i>false</i>;</p> <p>18: /* atomically execute lines 18a–18d */</p> <p>18a: <i>Tail</i>[1 - <i>idx</i>] := \perp;</p> <p>18b: <i>HistLen</i>[1 - <i>idx</i>] := 0;</p> <p>18c: forall <i>j</i> do <i>Param</i>[1 - <i>idx</i>][<i>j</i>] := \perp od;</p> <p>18d: forall <i>j</i> do <i>Hist</i>[1 - <i>idx</i>][<i>j</i>] := \perp od;</p> <p>19: <i>Position</i>[1 - <i>idx</i>] := 0;</p> <p>20: <i>CurrentQueue</i> := 1 - <i>idx</i></p> <p> fi;</p> <p>21: <i>Signal</i>[<i>idx</i>][<i>self</i>] := <i>true</i>;</p> <p>22: <i>Active</i>[<i>p</i>] := <i>false</i></p> <p> od</p>
--	--

private variables

idx: 0..1;
counter: **integer**;
prev, *self*, *tail*: *Vartype*;
pos: 0..2*N* - 1

private auxiliary variable

position: 0.. ∞

Figure 7: ALGORITHM G-CC with auxiliary variables added.

then we have the following after these four *fetch-and-φ* invocations take place.

$$\begin{aligned}
Tail[i] &= (p, 1); \\
HistLen[i] &= 4; \\
Hist[i] &= (\perp, (p, 0), (q, 1), (r, 1), (p, 1), \perp, \perp, \dots); \\
Param[i] &= ((p, 0), (q, 3), (r, 5), (p, 1), \perp, \perp, \dots).
\end{aligned}$$

On the other hand, if the underlying *fetch-and-φ* primitive is *fetch-and-increment* (where we define $\perp = 0$), then $Tail[i]$ and $Hist[i]$ have the following final value, while the other two variables hold the same values as above.

$$\begin{aligned}
Tail[i] &= 4; \\
Hist[i] &= (\perp (= 0), 1, 2, 3, 4, \perp, \perp, \dots).
\end{aligned}$$

List of Invariants

We will establish the Exclusion property by proving that the conjunction of a number of assertions is an invariant. This proves that each of these assertions individually is an invariant. These invariants are listed below. Unless stated otherwise, we assume the following: i ranges over 0 and 1; h, j, k , and l range over $0..\infty$; x and y range over $Vartype$; p, q , and r range over $0..N - 1$.

$$\text{invariant (Exclusion)} \quad |\{p :: p@\{9..12\}\}| \leq 1 \quad (I1)$$

$$\text{invariant} \quad |\{p :: p@\{7..21\} \wedge p.idx = i\}| \leq 1 \quad (I2)$$

$$\text{invariant} \quad p@\{6..22\} \wedge p.idx = i \wedge p.position = h \Rightarrow \\ HistLen[i] > h \wedge p.prev = Hist[i][h] \wedge p.self = Hist[i][h + 1] \wedge Param[i][h].proc = p \quad (I3)$$

$$\text{invariant} \quad 0 < HistLen[i] < 2N \Rightarrow Tail[i] \neq \perp \quad (I4)$$

$$\text{invariant} \quad p@\{8..21\} \wedge p.idx = i \Rightarrow (\forall x :: Signal[i][x] = false) \quad (I5)$$

$$\text{invariant} \quad Signal[i][x] = true \Rightarrow x = Hist[i][Position[i]] \quad (I6)$$

$$\text{invariant} \quad Signal[i][x] = Signal[i][y] = true \Rightarrow x = y \quad (I7)$$

$$\text{invariant} \quad p@\{7\} \wedge p.idx = i \Rightarrow Signal[i][p.prev] = true \quad (I8)$$

$$\text{invariant} \quad |\{p :: p@\{6, 7\} \wedge p.idx = i \wedge p.prev = x\}| \leq 1 \quad (I9)$$

$$\text{invariant} \quad Position[i] = q + 1 \Rightarrow \\ q@\{0..3\} \vee q.idx = i \vee (\exists p :: p@\{12..20\} \wedge p.idx = i) \quad (I10)$$

$$\text{invariant} \quad q + 1 < Position[i] \leq N \Rightarrow q@\{0..3\} \vee q.idx = i \quad (I11)$$

$$\text{invariant} \quad p@\{11..13\} \wedge p.idx = i \wedge p.pos = N \Rightarrow \\ CurrentQueue = i \wedge (\forall q :: q@\{0..3\} \vee q.idx = i) \quad (I12)$$

$$\text{invariant} \quad p@\{16..20\} \wedge p.idx = i \Rightarrow \\ CurrentQueue = i \wedge (\forall q :: q@\{0..3\} \vee q.idx = i) \quad (I13)$$

$$\text{invariant} \quad p@\{11\} \wedge p.idx = i \Rightarrow Position[i] = p.pos \quad (I14)$$

$$\text{invariant} \quad p@\{12..21\} \wedge p.idx = i \Rightarrow Position[i] = p.pos + 1 \quad (I15)$$

$$\text{invariant} \quad HistLen[i] > h \wedge Param[i][h] = (p, c) \Rightarrow \\ Hist[i][h + 1] = \phi(Hist[i][h], \alpha_p[c]) \quad (I16)$$

$$\text{invariant} \quad Tail[i] = Hist[i][HistLen[i]] \quad (I17)$$

$$\text{invariant} \quad Hist[i][0] = \perp \quad (I18)$$

$$\text{invariant} \quad 0 \leq j < k < HistLen[i] \wedge \\ Param[i][j] = (p, c_1) \wedge Param[i][k] = (p, c_2) \wedge \\ (\forall l : j < l < k :: Param[i][l].proc \neq p) \Rightarrow$$

$$c_2 = c_1 + 1 \tag{I19}$$

invariant $0 \leq j < \text{HistLen}[i] \wedge$
 $\text{Param}[i][j] = (p, c) \wedge$
 $(\forall k : j < k < \text{HistLen}[i] :: \text{Param}[i][k].\text{proc} \neq p) \Rightarrow$
 $(p.\text{counter} = c + 1 \wedge p@\{4..22\} \wedge p.\text{idx} = i) \vee$
 $(p.\text{counter} = c + 1 \wedge p@\{0..3\}) \vee$
 $(p@\{0..3\} \wedge \text{CurrentQueue} = 1 - i) \vee$
 $(p@\{4..22\} \wedge p.\text{idx} = \text{CurrentQueue} = 1 - i)$ (I20)

invariant $0 \leq \text{HistLen}[i] \leq 2N$ (I21)

invariant $0 \leq \text{Position}[i] \leq 2N$ (I22)

invariant $\text{HistLen}[i] = 0 \Rightarrow (\forall x :: \text{Signal}[i][x] = \text{false})$ (I23)

invariant $p@\{18\} \wedge p.\text{idx} = i \Rightarrow (\forall x :: \text{Signal}[1 - i][x] = \text{false})$ (I24)

invariant $p@\{17\} \wedge p.\text{idx} = i \Rightarrow p.\text{tail} = \text{Tail}[1 - i]$ (I25)

invariant $p@\{7..11\} \wedge p.\text{idx} = i \Rightarrow \text{Position}[i] = p.\text{position}$ (I26)

invariant $p@\{12..21\} \wedge p.\text{idx} = i \Rightarrow \text{Position}[i] = p.\text{position} + 1$ (I27)

invariant $[\text{HistLen}[i] - \text{Position}[i] = |\{p :: p@\{6..11\} \wedge p.\text{idx} = i\}|] \vee$
 $(\exists p :: p@\{19\} \wedge p.\text{idx} = 1 - i)$ (I28)

invariant $|\{p :: p@\{6..22\} \wedge p.\text{idx} = i \wedge p.\text{position} = h\}| \leq 1$ (I29)

invariant $\text{CurrentQueue} = i \Rightarrow$
 $|\{p :: p@\{4, 5\} \wedge p.\text{idx} = 1 - i\}| \leq 2N - \text{HistLen}[1 - i]$ (I30)

invariant $\text{Position}[i] = N + 1 \Rightarrow$
 $\text{CurrentQueue} = 1 - i \vee$
 $(\exists p :: p@\{12..20\} \wedge p.\text{idx} = i \wedge p.\text{pos} = N)$ (I31)

invariant $\text{Position}[i] > N + 1 \Rightarrow \text{CurrentQueue} = 1 - i$ (I32)

invariant $\text{Active}[p] = p@\{3..22\}$ (I33)

invariant $p@\{5..22\} \wedge p.\text{idx} = i \Rightarrow \text{QueueIdx}[p] = i$ (I34)

invariant $p@\{2..4\} \Rightarrow \text{QueueIdx}[p] = \perp$ (I35)

invariant $\text{Position}[i] \leq N \Rightarrow$
 $\text{CurrentQueue} = i \vee$
 $(\exists p :: p@\{20\} \wedge p.\text{idx} = 1 - i) \vee$
 $(\text{Position}[i] = 0 \wedge (\forall q :: q@\{0..3\} \vee q.\text{idx} = 1 - i))$ (I36)

invariant $1 \leq \text{Position}[i] \leq N \Rightarrow \text{CurrentQueue} = i$ (I37)

invariant $p@\{4, 5\} \wedge p.\text{idx} = i \Rightarrow \text{HistLen}[i] < 2N$ (I38)

invariant $p@\{19, 20\} \wedge p.\text{idx} = i \Rightarrow \text{HistLen}[1 - i] = 0$ (I39)

invariant $p@\{20\} \wedge p.\text{idx} = i \Rightarrow \text{Position}[1 - i] = 0$ (I40)

invariant $\text{Position}[i] \leq h < \text{HistLen}[i] \Rightarrow$
 $(\exists p :: p@\{6..11\} \wedge p.\text{idx} = i \wedge p.\text{position} = h)$ (I41)

invariant $\text{Position}[i] = h > 0 \Rightarrow$
 $\text{Signal}[i][\text{Hist}[i][h]] = \text{true} \vee$
 $(\exists p :: p@\{12..21\} \wedge p.\text{idx} = i \wedge p.\text{position} = h - 1) \vee$
 $(\exists q :: q@\{8..11\} \wedge q.\text{idx} = i \wedge q.\text{position} = h)$ (I42)

invariant $0 < j < k \leq \text{HistLen}[i] \wedge k < 2N \wedge$
 $\text{Param}[i][j - 1].\text{proc} \neq \text{Param}[i][k - 1].\text{proc} \Rightarrow$
 $\text{Hist}[i][j] \neq \text{Hist}[i][k]$ (I43)

invariant $0 < j \leq \text{HistLen}[i] \wedge j < 2N \Rightarrow \text{Hist}[i][j] \neq \perp$ (I44)

invariant $0 < j < k \leq \text{HistLen}[i] \wedge k < 2N \wedge$
 $\text{Param}[i][j - 1] = (p, c) \wedge \text{Param}[i][k - 1] = (p, c + 1) \Rightarrow$

$$\text{Hist}[i][j] \neq \text{Hist}[i][k] \quad (\text{I45})$$

$$\text{invariant } \text{Position}[i] \leq h < \text{HistLen}[i] \wedge \text{Param}[i][h].\text{proc} = p \Rightarrow p@ \{6..11\} \wedge p.\text{idx} = i \wedge p.\text{position} = h \quad (\text{I46})$$

$$\text{invariant } p@ \{6\} \wedge p.\text{idx} = i \Rightarrow \text{Position}[i] \leq p.\text{position} \quad (\text{I47})$$

$$\text{invariant } \text{Position}[i] \leq j < k < \text{HistLen}[i] \Rightarrow \text{Hist}[i][j] \neq \text{Hist}[i][k] \quad (\text{I48})$$

$$\text{invariant } p@ \{11..22\} \Rightarrow p.\text{pos} = p.\text{position} \quad (\text{I49})$$

Proof of the Exclusion Property

We now prove that each of (I1)–(I48) is an invariant. For each invariant I , we prove that for any pair of consecutive states t and u , if all invariants hold at t , then I holds at u . (It is easy to see that each invariant is initially true, so we leave this part of the proof to the reader.) If I is an implication (which is the case for most of our invariants), then it suffices to check only those program statements that may establish the antecedent of I , or that may falsify the consequent if executed while the antecedent holds.

$$\text{invariant (Exclusion)} \quad |\{p :: p@ \{9..12\}\}| \leq 1 \quad (\text{I1})$$

Proof: Since the Entry_2 and Exit_2 routines (statements 8 and 12) are assumed to be correct, (I1) follows easily from (I2). \square

$$\text{invariant } |\{p :: p@ \{7..21\} \wedge p.\text{idx} = i\}| \leq 1 \quad (\text{I2})$$

Proof: The only statements that may potentially falsify (I2) are $5.p$ and $6.p$. Statement $5.p$ may falsify (I2) by establishing $p@ \{7..21\} \wedge p.\text{idx} = i$ only if executed when

$$p.\text{idx} = i \wedge \text{Tail}[i] = \perp \wedge q@ \{7..21\} \wedge q.\text{idx} = i$$

holds, where q is any arbitrary process, different from p . (In this case, process p transits from statement 5 to statement 8.) By (I4) and (I38), $p@ \{5\} \wedge p.\text{idx} = i \wedge \text{Tail}[i] = \perp$ implies $\text{HistLen}[i] = 0$. However, by applying (I3) with ' $p' \leftarrow q$ ', we have $\text{HistLen}[i] > q.\text{position} \geq 0$, a contradiction. Thus, statement $5.p$ cannot falsify (I2).

Statement $6.p$ may falsify (I2) by establishing $p@ \{7..21\} \wedge p.\text{idx} = i$ only if executed when

$$p.\text{idx} = i \wedge \text{Signal}[i][p.\text{prev}] = \text{true} \wedge q@ \{7..21\} \wedge q.\text{idx} = i,$$

holds, where q is any arbitrary process, different from p . By (I5), we have $q@ \{7\}$. Thus, by (I8), we have $\text{Signal}[i][q.\text{prev}] = \text{true}$, which in turn implies $p.\text{prev} = q.\text{prev}$ by (I7). However, by (I9), this implies $p = q$, a contradiction. Thus, statement $6.p$ cannot falsify (I2). \square

$$\text{invariant } p@ \{6..22\} \wedge p.\text{idx} = i \wedge p.\text{position} = h \Rightarrow \text{HistLen}[i] > h \wedge p.\text{prev} = \text{Hist}[i][h] \wedge p.\text{self} = \text{Hist}[i][h+1] \wedge \text{Param}[i][h].\text{proc} = p \quad (\text{I3})$$

Proof: The only statement that may establish the antecedent is $5.p$, which may do so only if executed when $p.\text{idx} = i \wedge \text{HistLen}[i] = h$ holds. In this case, $5.p$ establishes $\text{HistLen}[i] = h+1 \wedge p.\text{self} = \text{Hist}[i][h+1] \wedge \text{Param}[i][h].\text{proc} = p$. Moreover, by (I17), $\text{Tail}[i] = \text{Hist}[i][h]$ holds before $5.p$ is executed, and hence, $p.\text{prev} = \text{Hist}[i][h]$ holds afterward.

The only statements that may falsify the consequent are $5.q$ and $18.q$, where q is any arbitrary process. As shown above, statement $5.p$ cannot falsify (I3). For $q \neq p$, if statement $5.q$ is executed when $\text{HistLen}[i] > h$ holds, then it preserves $\text{HistLen}[i] > h$, and does not update $\text{Hist}[i][h]$, $\text{Hist}[i][h+1]$, or $\text{Param}[i][h]$. Thus, statement $5.q$ preserves (I3).

Statement $18.q$ may falsify the consequent only if executed when $q.\text{idx} = 1-i$ holds. However, by applying (I13) with ' $p' \leftarrow q$ and ' $i' \leftarrow 1-i$ ', this implies $p@ \{0..3\} \vee p.\text{idx} = 1-i$. Thus, in this case, the antecedent is false before and after the execution of $18.q$. \square

$$\mathbf{invariant} \quad 0 < \text{HistLen}[i] < 2N \Rightarrow \text{Tail}[i] \neq \perp \quad (\text{I4})$$

Proof: Invariant (I4) follows easily by applying (I44) with ‘ j ’ \leftarrow $\text{HistLen}[i]$, and using (I17). \square

$$\mathbf{invariant} \quad p@{\{8..21\}} \wedge p.idx = i \Rightarrow (\forall x :: \text{Signal}[i][x] = \text{false}) \quad (\text{I5})$$

Proof: The only statements that may establish the antecedent are $5.p$ and $7.p$. Statement $5.p$ may establish the antecedent only if executed when $p.idx = i \wedge \text{Tail}[i] = \perp$ holds. In this case, by (I4) and (I38), we have $\text{HistLen}[i] = 0$. Hence, by (I23), the consequent is true before and after the execution of $5.p$.

Statement $7.p$ may establish the antecedent only if executed when $p@{\{7\}} \wedge p.idx = i$ holds. In this case, by (I7) and (I8), $\text{Signal}[i][p.prev]$ is the only entry among $\text{Signal}[i][..]$ that is *true*. Thus, statement $7.p$ establishes the consequent.

The only statement that may falsify the consequent is $21.q$, where q is any arbitrary process. Statement $21.q$ may potentially falsify (I5) only if executed when $p@{\{8..21\}} \wedge p.idx = i \wedge q@{\{21\}} \wedge q.idx = i$ holds, which implies $p = q$ by (I2). Thus, $21.q$ falsifies the antecedent in this case. \square

$$\mathbf{invariant} \quad \text{Signal}[i][x] = \text{true} \Rightarrow x = \text{Hist}[i][\text{Position}[i]] \quad (\text{I6})$$

Proof: The only statement that may establish the antecedent is $21.p$, which may do so only if $p.idx = i \wedge p.self = x$ holds. In this case, before $21.p$ is executed, $\text{Position}[i] = p.position + 1$ and $p.self = \text{Hist}[i][p.position + 1]$ hold, by (I27) and (I3), respectively. Thus, the consequent is true before and after the execution of $21.p$.

The only statements that may falsify the consequent are $5.p$ and $18.p$ (which may update $\text{Hist}[i][\text{Position}[i]]$) and $11.p$ (which may update $\text{Position}[i]$), where p is any arbitrary process. Statement $5.p$ may update $\text{Hist}[i][\text{Position}[i]]$ only if executed when

$$p@{\{5\}} \wedge p.idx = i \quad (1)$$

holds. In this case, $5.p$ updates only one entry of $\text{Hist}[i][..]$, namely, $\text{Hist}[i][h + 1]$, where h is the value of $\text{HistLen}[i]$ before its execution. Moreover, by (I28), either $h \geq \text{Position}[i]$ holds, or there exists a process q satisfying $q@{\{19\}} \wedge q.idx = 1 - i$. In the former case, statement $5.p$ does not update $\text{Hist}[i][\text{Position}[i]]$, and hence preserves (I5). In the latter case, by (I13), we have $p@{\{0..3\}} \vee p.idx = 1 - i$, which contradicts (1). It follows that the latter case in fact cannot arise.

Statement $18.p$ may update $\text{Hist}[i][\text{Position}[i]]$ only if executed when $p.idx = 1 - i$ holds, in which case, by (I24), the antecedent of (I6) is false before and after its execution. Similarly, statement $11.p$ may update $\text{Position}[i]$ only if executed when $p.idx = i$ holds, in which case, by (I5), the antecedent is false before and after its execution. \square

$$\mathbf{invariant} \quad \text{Signal}[i][x] = \text{Signal}[i][y] = \text{true} \Rightarrow x = y \quad (\text{I7})$$

Proof: This invariant follows trivially from (I6). \square

$$\mathbf{invariant} \quad p@{\{7\}} \wedge p.idx = i \Rightarrow \text{Signal}[i][p.prev] = \text{true} \quad (\text{I8})$$

Proof: The only statement that may establish the antecedent is $6.p$, which may do so only if the consequent holds.

The only statements that may falsify the consequent are $5.p$ (which assigns $p.prev$) and $7.q$ and $17.q$, where q is any arbitrary process. After the execution of $5.p$, $p@{\{7\}}$ is false. Statement $7.q$ may potentially falsify (I8) only if executed when $p@{\{7\}} \wedge p.idx = i \wedge q@{\{7\}} \wedge q.idx = i$ holds. In this case, by (I2), we have $p = q$, and hence $7.q$ falsifies the antecedent. Statement $17.q$ may falsify the consequent only if executed when $q.idx = 1 - i$ holds, which implies that $p@{\{0..3\}} \vee p.idx = 1 - i$ holds, by (I13). Thus, the antecedent is false before and after the execution of $17.q$. \square

$$\mathbf{invariant} \quad |\{p :: p@{\{6, 7\}} \wedge p.idx = i \wedge p.prev = x\}| \leq 1 \quad (\text{I9})$$

Proof: Assume the following.

$$p@{6, 7} \wedge q@{6, 7} \wedge p.idx = q.idx = i \wedge p.prev = q.prev = x.$$

Our proof obligation is to show $p = q$. Define $j = p.position$ and $k = q.position$. By (I3), we have $(p.prev = \text{Hist}[i][j] = x) \wedge (j < \text{HistLen}[i])$. Similarly, we also have $(q.prev = \text{Hist}[i][k] = x) \wedge (k < \text{HistLen}[i])$. Moreover, by (I26) and (I47), we have $j \geq \text{Position}[i]$ and $k \geq \text{Position}[i]$.

Combining these assertions, by (I48), we have $j = k$. By (I29), this implies $p = q$. \square

invariant $\text{Position}[i] = q + 1 \Rightarrow$

$$q@{0..3} \vee q.idx = i \vee (\exists p :: p@{12..20} \wedge p.idx = i) \quad (\text{I10})$$

Proof: The only statements that may establish the antecedent are 11. r and 19. r , where r is any arbitrary process. Statement 11. r may establish the antecedent only if executed when $r.idx = i$ holds, in which case it establishes the last disjunct of the consequent. If 19. r updates $\text{Position}[i]$, then it establishes $\text{Position}[i] = 0$, and hence cannot establish the antecedent. (Recall that q is assumed to range over $0..N - 1$.)

The only statement that may falsify $q@{0..3} \vee q.idx = i$ is 3. q , which may do so only if executed when $\text{CurrentQueue} = 1 - i$ holds, which in turn implies that $\text{Position}[i] = 0 \vee \text{Position}[i] > N$ holds, by (I37). Thus, since $0 \leq q < N$, the antecedent is false before and after 3. q is executed.

The only statements that may falsify $p@{12..20} \wedge p.idx = i$ are 13. p , 14. p , 15. p , and 20. p . Assume that the antecedent holds before the execution of each of these statements. By (I15), we have $p.pos = q < N$, and hence statement 13. p establishes $p@{14}$, and statement 20. p cannot be executed (since $p.pos = N$ holds when it is executed).

Statement 14. p may falsify $p@{12..20}$ only if executed when $\text{Active}[q] = \text{false}$ holds. By (I33), this implies that $q@{0..2}$ holds, and hence the consequent of (I10) holds after 14. p is executed. Statement 15. p may falsify $p@{12..20} \wedge p.idx = i$ only if executed when $\text{QueueIdx}[q] = i$ holds. By (I34) and (I35), this implies $q@{0, 1} \vee q.idx = i$, and hence the consequent of (I10) holds after its execution. \square

invariant $q + 1 < \text{Position}[i] \leq N \Rightarrow q@{0..3} \vee q.idx = i$ (I11)

Proof: The only statements that may establish the antecedent are 11. r and 19. r , where r is any arbitrary process. Statement 11. r updates $\text{Position}[i]$ only if executed when $r.idx = i$ holds, in which case, by (I14), it increments $\text{Position}[i]$ by one. Therefore, statement 11. r may establish the antecedent only if executed when $\text{Position}[i] = q + 1$ holds. In this case, by (I10), either the consequent holds, or there exists a process p satisfying $p@{12..20} \wedge p.idx = i$, before the execution of 11. r . However, since we have $r@{11} \wedge r.idx = i$, the latter is precluded by (I2). Therefore, the consequent is true before 11. r is executed, and hence it is also true afterward.

If statement 19. r updates $\text{Position}[i]$, then it establishes $\text{Position}[i] = 0$, and hence cannot establish the antecedent.

The only statement that may falsify the consequent is 3. q , which may do so only if $\text{CurrentQueue} \neq i$ holds. In this case, by (I37), the antecedent is false before and after the execution of 3. q . \square

invariant $p@{11..13} \wedge p.idx = i \wedge p.pos = N \Rightarrow$

$$\text{CurrentQueue} = i \wedge (\forall q :: q@{0..3} \vee q.idx = i) \quad (\text{I12})$$

Proof: The only statement that may establish the antecedent is 10. p , which may do so only if executed when

$$p@{10} \wedge p.idx = i \wedge \text{Position}[i] = N \quad (2)$$

holds. By (I37), this implies that $\text{CurrentQueue} = i$ holds before and after 10. p is executed.

In order to prove that $(\forall q :: q@{0..3} \vee q.idx = i)$ holds after the execution of 10. p , we consider two cases, depending on the value of q . If $0 \leq q < N - 1$, then by (I11), $q@{0..3} \vee q.idx = i$ holds before and after the execution of 10. p . On the other hand, if $q = N - 1$, then since $\text{Position}[i] = N$, by (I10), either

$q@\{0..3\} \vee q.idx = i$ holds, or there exists a process r (different from p) satisfying $r@\{12..20\} \wedge r.idx = i$. However, the latter is precluded by (2) and (I2).

The only statement that may falsify $CurrentQueue = i$ is $20.r$ (where r is any arbitrary process), which may do so only if executed when $r.idx = i$ holds. Taken together with the antecedent, this implies that $r = p$ holds, by (I2). Thus, statement $20.r$ falsifies the antecedent in this case.

The only statement that may falsify $q@\{0..3\} \vee q.idx = i$ is $3.q$. However, if $3.q$ is executed when the consequent is true, then $3.q$ establishes $q.idx = i$, and hence preserves the consequent. \square

$$\begin{aligned} \text{invariant } p@\{16..20\} \wedge p.idx = i &\Rightarrow \\ CurrentQueue = i \wedge (\forall q :: q@\{0..3\} \vee q.idx = i) &\end{aligned} \quad (\text{I13})$$

Proof: The only statement that may establish the antecedent is $13.p$, which may do so only if executed when $p.idx = i \wedge p.pos = N$ holds. In this case, by (I12), the consequent holds before and after the execution of statement $13.p$.

The only statements that may falsify the consequent are $3.q$ and $20.q$ (where q is any arbitrary process). However, each preserves the consequent as shown in the proof of (I12). \square

$$\text{invariant } p@\{11\} \wedge p.idx = i \Rightarrow Position[i] = p.pos \quad (\text{I14})$$

$$\text{invariant } p@\{12..21\} \wedge p.idx = i \Rightarrow Position[i] = p.pos + 1 \quad (\text{I15})$$

Proof: The only statement that may establish the antecedent of (I14) (respectively, (I15)) is $10.p$ (respectively, $11.p$), which clearly establishes the corresponding consequent.

The only other statements that may falsify either consequent are $11.q$ and $19.q$, where q is any arbitrary process. Statement $11.q$ may falsify either consequent only if executed when $q.idx = i$ holds. Taken together with either antecedent, this implies that $q = p$ holds, by (I2). Thus, statement $11.q$ falsifies the antecedent of (I14), and establishes the antecedent and consequent of (I15).

Statement $19.q$ may falsify either consequent only if executed when $q.idx = 1 - i \wedge q.pos = N$ holds. By (I13), this implies that $p@\{0..3\} \vee p.idx = 1 - i$ holds. Hence, the antecedents of (I14) and (I15) are false before and after the execution of $19.q$. \square

$$\begin{aligned} \text{invariant } HistLen[i] > h \wedge Param[i][h] = (p, c) &\Rightarrow \\ Hist[i][h + 1] = \phi(Hist[i][h], \alpha_p[c]) &\end{aligned} \quad (\text{I16})$$

Proof: The only statement that may establish the antecedent is $5.q$, where q is any arbitrary process. (Note that statement $18.q$ assigns $HistLen[1 - q.idx] := 0$, and hence cannot establish the antecedent.)

Statement $5.q$ may establish the antecedent only if executed when $HistLen[i] = h \wedge q.idx = i$ holds. In this case, by (I17), $5.q$ establishes

$$Hist[i][h + 1] = q.self = \phi(Hist[i][h], \alpha_q[c']) \quad \text{and} \quad Param[i][h] = (q, c'),$$

where c' is the value of $q.counter$ before the execution of $5.q$. Thus, the antecedent is established only if $q = p$ and $c = c'$, in which case the consequent easily follows.

The only statements that may falsify the consequent are $5.q$ and $18.q$, where q is any arbitrary process. Statement $5.q$ may falsify the consequent only if executed when $q.idx = i \wedge (HistLen[i] = h - 1 \vee HistLen[i] = h)$ holds. If $HistLen[i] = h - 1$ holds before its execution, then $HistLen[i] = h$ holds after its execution, and hence the antecedent is false. On the other hand, if $HistLen[i] = h$ holds before its execution, then statement $5.q$ preserves (I16) as shown above.

Statement $18.q$ may falsify the consequent only if executed when $q.idx = 1 - i$ holds, in which case it establishes $HistLen[i] = 0$, and hence the antecedent is false after its execution. \square

$$\text{invariant } Tail[i] = Hist[i][HistLen[i]] \quad (\text{I17})$$

$$\text{invariant } Hist[i][0] = \perp \quad (\text{I18})$$

Proof: These invariants follow trivially from inspecting the code. In particular, lines 5c, 5f and 5g, as well as lines 18a and 18d, ensure that (I17) is maintained. Also, since $\text{HistLen}[i]$ is always nonnegative (by (I21)), line 5f cannot update $\text{Hist}[i][0]$, and hence (I18) is maintained. \square

$$\begin{aligned}
\text{invariant } & 0 \leq j < k < \text{HistLen}[i] \wedge \\
& \text{Param}[i][j] = (p, c_1) \wedge \text{Param}[i][k] = (p, c_2) \wedge \\
& (\forall l : j < l < k :: \text{Param}[i][l].\text{proc} \neq p) \Rightarrow \\
& c_2 = c_1 + 1
\end{aligned} \tag{I19}$$

Proof: The only statement that may establish the antecedent is 5.q, where q is any arbitrary process. (Note that statement 18.q assigns $\text{HistLen}[1-1.\text{idx}] := 0$ and does not update any entry of Param .) Since 5.q increments $\text{HistLen}[q.\text{idx}]$ by one, it may establish the antecedent only if executed when $q.\text{idx} = i \wedge \text{HistLen}[i] = k$ holds. Note that 5.q establishes $\text{Param}[i][k] = (q, c)$ in this case, where c is the value of $q.\text{counter}$ before the execution of 5.q. Thus, 5.q may establish the antecedent only if executed when the following holds.

$$\begin{aligned}
& q = p \wedge c_2 = q.\text{counter} \wedge q@5 \wedge q.\text{idx} = i \wedge \\
& 0 \leq j < k = \text{HistLen}[i] \wedge \\
& \text{Param}[i][j] = (p, c_1) \wedge \\
& (\forall l : j < l < k :: \text{Param}[i][l].\text{proc} \neq p).
\end{aligned}$$

Thus, by applying (I20) with ' $c \leftarrow c_1$ ', the first disjunct of the consequent of (I20) follows, and hence we have $c_2 = q.\text{counter} = c_1 + 1$. \square

$$\begin{aligned}
\text{invariant } & 0 \leq j < \text{HistLen}[i] \wedge \\
& \text{Param}[i][j] = (p, c) \wedge \\
& (\forall k : j < k < \text{HistLen}[i] :: \text{Param}[i][k].\text{proc} \neq p) \Rightarrow \\
& (p.\text{counter} = c + 1 \wedge p@4..22 \wedge p.\text{idx} = i) \vee \tag{A} \\
& (p.\text{counter} = c + 1 \wedge p@0..3) \vee \tag{B} \\
& (p@0..3 \wedge \text{CurrentQueue} = 1 - i) \vee \tag{C} \\
& (p@4..22 \wedge p.\text{idx} = \text{CurrentQueue} = 1 - i) \tag{D}
\end{aligned} \tag{I20}$$

Proof: The only statement that may establish the antecedent is 5.q, where q is any arbitrary process. (As with (I19), 18.q need not be considered here.) Since 5.q increments $\text{HistLen}[q.\text{idx}]$ by one, it may establish the antecedent only if executed when $q.\text{idx} = i \wedge \text{HistLen}[i] = j$ holds. Note that 5.q establishes $\text{Param}[i][k] = (q, c')$, where c' is the value of $q.\text{counter}$ before the execution of 5.q. Thus, the antecedent may be established only if $q = p \wedge c' = c$ holds. It follows that statement 5.q establishes disjunct \mathcal{A} in this case.

Disjunct \mathcal{A} may be falsified only by statements 5.p (which may update $p.\text{counter}$) and 22.p (which may falsify $p@4..22$). If statement 5.p is executed while the antecedent holds, then it establishes $\text{Param}[i][\text{HistLen}[i] - 1] = (p, c + 1)$, and hence falsifies the last conjunct of the antecedent. Statement 22.p establishes disjunct \mathcal{B} if executed when disjunct \mathcal{A} holds.

Disjunct \mathcal{B} may be falsified only by statement 3.p, which establishes either disjunct \mathcal{A} or disjunct \mathcal{D} , depending on the value of CurrentQueue .

Disjunct \mathcal{C} may be falsified only by statements 3.p and 20.q, where q is any arbitrary process. Statement 3.p establishes disjunct \mathcal{D} if executed when disjunct \mathcal{C} holds. Statement 20.q may falsify disjunct \mathcal{C} only if executed when $q.\text{idx} = 1 - i$. In this case, by (I39), $\text{HistLen}[i] = 0$ holds before and after the execution of 20.q, and hence the antecedent of (I20) is false before and after its execution.

Disjunct \mathcal{D} may be falsified only by statements 22.p and 20.q, where q is any arbitrary process. Statement 22.p establishes disjunct \mathcal{C} if executed when disjunct \mathcal{D} holds. As shown above, the antecedent is false after the execution of 20.q. \square

$$\text{invariant } 0 \leq \text{HistLen}[i] \leq 2N \tag{I21}$$

Proof: The only statement that may potentially falsify (I21) is $5.p$, where p is any arbitrary process. Since $5.p$ increments $\text{HistLen}[p.idx]$ by one, it may falsify (I21) only if executed when $p.idx = i \wedge \text{HistLen}[i] = 2N$ holds. However, this is precluded by (I38). \square

invariant $0 \leq \text{Position}[i] \leq 2N$ (I22)

Proof: The only statement that may potentially falsify (I22) is $11.p$ (where p is any arbitrary process), which may do so only if executed when $p.idx = i$. In this case, by (I14) and (I26), statement $11.p$ establishes $\text{Position}[i] = p.pos + 1 = p.\text{position} + 1$. Moreover, by (I3) and (I21), $p.\text{position} < \text{HistLen}[i] \leq 2N$ holds before $11.p$ is executed. Thus, statement $11.p$ preserves (I22). \square

invariant $\text{HistLen}[i] = 0 \Rightarrow (\forall x :: \text{Signal}[i][x] = \text{false})$ (I23)

Proof: The only statement that may establish the antecedent is $18.p$, where p is any arbitrary process. Statement $18.p$ may establish the antecedent only if executed when $p@\{18\} \wedge p.idx = 1 - i$ holds. In this case, by (I24), the consequent holds before and after $18.p$ is executed.

The only statement that may falsify the consequent is $21.p$, where p is any arbitrary process. Statement $21.p$ may falsify the consequent only if executed when $p.idx = i$ holds. In this case, by (I3), $\text{HistLen}[i] > p.\text{position}$ holds before and after $21.p$ is executed. Thus, the antecedent is false before and after the execution of $21.p$. \square

invariant $p@\{18\} \wedge p.idx = i \Rightarrow (\forall x :: \text{Signal}[1 - i][x] = \text{false})$ (I24)

Proof: The only statement that may establish the antecedent is $17.p$, which may do so only if $p.idx = i$ holds. Assume that $\text{Signal}[1 - i][x] = \text{true}$ holds for some x before the execution of $17.p$. It suffices to show $x = p.tail$.

By (I6), we have

$$x = \text{Hist}[1 - i][\text{Position}[1 - i]]. \quad (3)$$

Also, by (I13), we have

$$(\forall q :: q@\{0..3\} \vee q.idx = i). \quad (4)$$

Moreover, by (I2), $p@\{17\} \wedge p.idx = i$ implies

$$\neg(\exists r :: r@\{19\} \wedge r.idx = i). \quad (5)$$

Combining (4) and (5), and applying (I28) with ' $i' \leftarrow 1 - i$, we have $\text{HistLen}[1 - i] = \text{Position}[1 - i]$. Hence, by (3) and (I17), we have

$$x = \text{Hist}[1 - i][\text{HistLen}[1 - i]] = \text{Tail}[1 - i].$$

Thus, by (I25), $x = p.tail$ holds.

The only statement that may falsify the consequent is $21.q$, where q is any arbitrary process. Statement $21.q$ may falsify the consequent only if executed when $q.idx = 1 - i$ holds. However, when the antecedent holds, $q@\{21\} \wedge q.idx = 1 - i$ is false, by (I13). \square

invariant $p@\{17\} \wedge p.idx = i \Rightarrow p.tail = \text{Tail}[1 - i]$ (I25)

Proof: The only statement that may establish the antecedent is $16.p$, which may do so only if $p.idx = i$ holds. In this case, $16.p$ establishes the consequent.

The only statements that may falsify the consequent are $5.q$ and $18.q$, where q is any arbitrary process. Statement $5.q$ may falsify the consequent only if executed when $q.idx = 1 - i$ holds, which implies that the antecedent is false, by (I13). Similarly, statement $18.q$ may falsify the consequent only if executed when $q.idx = i$ holds, which implies that the antecedent is false, by (I2). \square

invariant $p@\{7..11\} \wedge p.idx = i \Rightarrow \text{Position}[i] = p.\text{position}$ (I26)

Proof: The only statements that may establish the antecedent are 5.*p* and 6.*p*. Statement 5.*p* may establish the antecedent only if executed when

$$p@\{5\} \wedge Tail[i] = \perp \wedge p.idx = i \quad (6)$$

holds. In this case, by (I4) and (I38), we have $HistLen[i] = 0$. Thus, statement 5.*p* establishes $p.position = 0$. By (I28), $HistLen[i] = 0$ also implies that either $Position[i] = 0$ or $(\exists q :: q@\{19\} \wedge q.idx = 1 - i)$ holds. In the former case, the consequent is established. In the latter case, by applying (I13) with '*p*' \leftarrow *q* and '*i*' \leftarrow $1 - i$, we have $p@\{0..3\} \vee p.idx = 1 - i$, which contradicts (6). It follows that the latter case in fact cannot arise.

Statement 6.*p* may establish the antecedent only if executed when

$$p@\{6\} \wedge p.idx = i \wedge Signal[i][p.prev] = true \quad (7)$$

holds. By (I6), this implies $p.prev = Hist[i][Position[i]]$. Let $k = p.position$. By (I3) (with '*h*' \leftarrow *k*) and (I47), we have $p.prev = Hist[i][k]$ and $Position[i] \leq k < HistLen[i]$. If $k > Position[i]$, then by applying (I48) with '*j*' \leftarrow $Position[i]$, we have $Hist[i][Position[i]] \neq Hist[i][k]$, a contradiction. It follows that $Position[i] = k = p.position$ holds before the execution of 6.*p*. Thus, it also holds after its execution.

The only statements that may falsify the consequent are 11.*q* and 19.*q*, where *q* is any arbitrary process. Statement 11.*q* may falsify the consequent (when the antecedent holds) only if executed when $q.idx = i$ holds. Taken together with the antecedent, this implies that $q = p$ holds, by (I2). Thus, statement 11.*q* falsifies the antecedent.

Statement 19.*q* may falsify the consequent only if executed when $q.idx = 1 - i$ holds. By applying (I13) with '*p*' \leftarrow *q* and '*i*' \leftarrow $1 - i$, this implies $p@\{0..3\} \vee p.idx = 1 - i$. Hence, the antecedent is false before and after the execution of 19.*q*. \square

invariant $p@\{12..21\} \wedge p.idx = i \Rightarrow Position[i] = p.position + 1 \quad (I27)$

Proof: The only statement that may establish the antecedent is 11.*p*, which may do so only if $p.idx = i$ holds. In this case, by (I14) and (I26), 11.*p* establishes the consequent.

The only statements that may falsify the consequent (while the antecedent holds) are 11.*q* and 19.*q*, where *q* is any arbitrary process. Statement 11.*q* may falsify the consequent only if executed when $q@\{11\} \wedge q.idx = i$ holds. In this case, if $q = p$, then statement 11.*p* preserves (I27) as shown above. If $q \neq p$, then by (I2), the antecedent is false before and after the execution of 11.*q*.

The proof that 19.*q* preserves (I27) is similar to that given in the proof of (I26). \square

invariant $\left[HistLen[i] - Position[i] = \left| \{p :: p@\{6..11\} \wedge p.idx = i\} \right| \right] \vee (\exists p :: p@\{19\} \wedge p.idx = 1 - i) \quad (I28)$

Proof: Define

$$X = \left| \{p :: p@\{6..11\} \wedge p.idx = i\} \right|.$$

The only statements that may potentially falsify (I28) are 5.*q* (which may modify $HistLen[i]$ and X), 11.*q* (which may modify $Position[i]$ and X), 18.*q* (which may modify $HistLen[i]$), and 19.*q* (which may modify $Position[i]$ and also falsify the second disjunct), where *q* is any arbitrary process.

Statements 5.*q* and 11.*q* may modify $HistLen[i]$, $Position[i]$, or X only if executed when $q.idx = i$ holds. In this case, 5.*q* increments both $HistLen[i]$ and X by one, and hence preserves (I28). Similarly, by (I14), 11.*q* increments $Position[i]$ and decrements X by one, and hence preserves (I28).

Statement 18.*q* may update $HistLen[i]$ only if executed when $p.idx = 1 - i$, in which case it establishes the second disjunct. Statement 19.*q* may falsify the second disjunct only if executed when $q.idx = 1 - i$. In this case, by (I13) and (I39) (with '*p*' \leftarrow *q* and '*i*' \leftarrow $1 - i$), we have $X = 0$ and $HistLen[i] = 0$, respectively. Since 19.*q* establishes $Position[i] = 0$, it establishes the first disjunct of (I28). \square

invariant $\left| \{p :: p@\{6..22\} \wedge p.idx = i \wedge p.position = h\} \right| \leq 1 \quad (I29)$

Proof: Assume that there exists a process p satisfying $p@{6..22} \wedge p.idx = i \wedge p.position = h$. By (I3), this implies $\text{HistLen}[i] > h$.

The only statement that may potentially falsify (I29) is $5.q$, where q is any arbitrary process different from p . Statement $5.q$ may falsify (I29) only if executed when $q.idx = i$ holds. However, since $\text{HistLen}[i] > h$, statement $5.q$ establishes $q.position > h$, and hence cannot increase the left-hand side of (I29). \square

invariant $CurrentQueue = i \Rightarrow$

$$|\{p :: p@{4, 5} \wedge p.idx = 1 - i\}| \leq 2N - \text{HistLen}[1 - i] \quad (\text{I30})$$

Proof: The only statement that may establish the antecedent is $20.q$, where q is any arbitrary process. Let X denote the value of

$$|\{p :: p@{6..11} \wedge p.idx = 1 - i\}|$$

prior to the execution of $20.q$. Statement $20.q$ may establish the antecedent only if executed when

$$q@{20} \wedge q.idx = 1 - i \quad (8)$$

holds, which also implies the following.

$$|\{p :: p@{6..11, 20} \wedge p.idx = 1 - i\}| \geq X + 1$$

By (8), and by applying (I13) with ' $p' \leftarrow q$ ' and ' $i' \leftarrow 1 - i$ ', $\neg(\exists r :: r@{19} \wedge r.idx = i)$ holds, and hence, by (I28), we have

$$\text{HistLen}[1 - i] - \text{Position}[1 - i] = X.$$

Also, since $20.q$ may be executed only if $q.pos = N$ holds, by applying (I15) with ' $p' \leftarrow q$ ' and ' $i' \leftarrow 1 - i$ ', we have

$$\text{Position}[1 - i] = N + 1.$$

Combining these assertions, we have the following.

$$\begin{aligned} |\{p :: p@{4, 5} \wedge p.idx = 1 - i\}| &\leq N - X - 1 \\ &= N - (\text{HistLen}[1 - i] - \text{Position}[1 - i]) - 1 \\ &= N - (\text{HistLen}[1 - i] - N - 1) - 1 \\ &= 2N - \text{HistLen}[1 - i]. \end{aligned}$$

Thus, the consequent holds before the execution of $20.q$, and hence it also holds after its execution.

The only statements that may falsify the consequent are $3.q$ (which may increment X) and $5.q$ (which may increment $\text{HistLen}[1 - i]$), where q is any arbitrary process. If statement $3.q$ is executed while the antecedent holds, then it establishes $q.idx = i$, and hence cannot increment X . Statement $5.q$ may increment $\text{HistLen}[1 - i]$ (by one) only if executed when $q.idx = 1 - i$ holds, in which case it also decrements X by one, and hence preserves the consequent. \square

invariant $Position[i] = N + 1 \Rightarrow$

$$\begin{aligned} &CurrentQueue = 1 - i \vee \\ &(\exists p :: p@{12..20} \wedge p.idx = i \wedge p.pos = N) \end{aligned} \quad (\text{I31})$$

Proof: The only statement that may establish the antecedent is $11.q$, where q is any arbitrary process. However, if $11.q$ establishes the antecedent, then by (I14), it also establishes the second disjunct of the consequent.

The only statements that may falsify the consequent are $20.q$ (which may update $CurrentQueue$) and $13.p$, $14.p$, $15.p$, and $20.p$ (which may falsify $p@{12..20} \wedge p.idx = i \wedge p.pos = N$), where q is any arbitrary process. Statement $20.q$ may falsify the consequent only if executed when $q.idx = 1 - i$ holds. In this case, by (I40), $Position[i] = 0$ holds before and after the execution of $20.q$. Thus, the antecedent is false before and after $20.q$ is executed.

Since $p.pos = N$, statement $13.p$ establishes $p@{16}$, and statements $14.p$ and $15.p$ cannot be executed. If statement $20.p$ is executed when $p.idx = i$ holds, then it establishes $CurrentQueue = 1 - i$, and hence preserves the consequent. \square

$$\text{invariant } Position[i] > N + 1 \Rightarrow CurrentQueue = 1 - i \quad (I32)$$

Proof: The only statement that may establish the antecedent is 11. p , where p is any arbitrary process. By (I14), 11. p may establish the antecedent only if executed when $p@\{11\} \wedge p.idx = i \wedge Position[i] = N + 1$ holds. In this case, by (I2) and (I31), we have the consequent.

The only statement that may falsify the consequent is 20. q , where q is any arbitrary process. As shown in the proof of (I31), if 20. q falsifies the consequent, then the antecedent is false after its execution. \square

$$\text{invariant } Active[p] = p@\{3..22\} \quad (I33)$$

$$\text{invariant } p@\{5..22\} \wedge p.idx = i \Rightarrow QueueIdx[p] = i \quad (I34)$$

$$\text{invariant } p@\{2..4\} \Rightarrow QueueIdx[p] = \perp \quad (I35)$$

Proof: These invariants follow trivially from inspecting ALGORITHM G-CC. \square

$$\begin{aligned} \text{invariant } Position[i] \leq N \Rightarrow & \\ & CurrentQueue = i \vee \quad \mathcal{A} \\ & (\exists p :: p@\{20\} \wedge p.idx = 1 - i) \vee \quad \mathcal{B} \\ & (Position[i] = 0 \wedge (\forall q :: q@\{0..3\} \vee q.idx = 1 - i)) \quad \mathcal{C} \end{aligned} \quad (I36)$$

Proof: The only statements that may establish the antecedent are 11. r and 19. r , where r is any arbitrary process. However, if statement 11. r updates $Position[i]$, then by (I14), it increments $Position[i]$ by one. It follows that, although statement 11. r may preserve the antecedent, it cannot establish it. Statement 19. r may establish the antecedent only if executed when $r.idx = 1 - i$ holds, in which case it establishes disjunct \mathcal{B} .

The only statement that may falsify disjunct \mathcal{A} is 20. r , where r is any arbitrary process. Statement 20. r may falsify disjunct \mathcal{A} only if executed when $r.idx = i \wedge r.pos = N$ holds, which implies that $Position[i] = N + 1$ holds, by (I15). Thus, in this case, the antecedent is false before and after the execution of 20. r .

The only statement that may falsify disjunct \mathcal{B} is 20. p , which establishes disjunct \mathcal{A} .

The only statements that may falsify disjunct \mathcal{C} are 3. q and 11. q , where q is any arbitrary process. Statement 3. q may falsify disjunct \mathcal{C} only if executed when $CurrentQueue = i$ holds, in which case disjunct \mathcal{A} holds before and after its execution. Statement 11. q may falsify disjunct \mathcal{C} only if executed when $q@\{11\} \wedge q.idx = i$ holds, which is precluded when disjunct \mathcal{C} holds. \square

$$\text{invariant } 1 \leq Position[i] \leq N \Rightarrow CurrentQueue = i \quad (I37)$$

Proof: By (I36), the antecedent implies one of the following.

$$\begin{aligned} \mathcal{A} : & CurrentQueue = i, \\ \mathcal{B} : & (\exists p :: p@\{20\} \wedge p.idx = 1 - i), \quad \text{or} \\ \mathcal{C} : & Position[i] = 0 \wedge (\forall q :: q@\{0..3\} \vee q.idx = 1 - i). \end{aligned}$$

By (I40), \mathcal{B} implies $Position[i] = 0$. Also, \mathcal{C} clearly implies $Position[i] = 0$. Thus, both are precluded by the antecedent. It follows that \mathcal{A} is true. \square

$$\text{invariant } p@\{4, 5\} \wedge p.idx = i \Rightarrow HistLen[i] < 2N \quad (I38)$$

Proof: For the sake of contradiction, assume

$$p@\{4, 5\} \wedge p.idx = i \wedge HistLen[i] \geq 2N. \quad (9)$$

By applying (I30) with ' $i \leftarrow 1 - i$ ', we have $CurrentQueue \neq 1 - i$, i.e.,

$$CurrentQueue = i. \quad (10)$$

Thus, by (I32), we have

$$Position[i] \leq N + 1. \quad (11)$$

Also, (9) implies

$$|\{q :: q@{6..11} \wedge q.idx = i\}| \leq N - 1.$$

Hence, by (I28), we have

$$HistLen[i] - Position[i] \leq N - 1 \quad \vee \quad (\exists r :: r@{19} \wedge r.idx = 1 - i).$$

However, if there exists a process r satisfying $r@{19} \wedge r.idx = 1 - i$, then by (I13) (with ‘ p ’ $\leftarrow r$ and ‘ i ’ $\leftarrow 1 - i$), we have $p@{0..3} \vee p.idx = 1 - i$, which contradicts (9). Therefore, we have $HistLen[i] - Position[i] \leq N - 1$.

Note that the only common solution to $HistLen[i] \geq 2N$ (given in (9)), (11), and $HistLen[i] - Position[i] \leq N - 1$ is $HistLen[i] = 2N$ and

$$Position[i] = N + 1.$$

By (10) and (I31), this implies that $(\exists r :: r@{12..20} \wedge r.idx = i)$ holds. From this and (9), we have

$$|\{q :: q@{6..11} \wedge q.idx = i\}| \leq N - 2.$$

Hence, by (I28), we have

$$HistLen[i] - Position[i] \leq N - 2 \quad \vee \quad (\exists r :: r@{19} \wedge r.idx = 1 - i).$$

The second disjunct is precluded by (I13) as shown above, and hence we have $HistLen[i] - Position[i] \leq N - 2$. However, this cannot hold simultaneously with (9) and (11). Thus, we have reached a contradiction. \square

invariant $p@{19, 20} \wedge p.idx = i \Rightarrow HistLen[1 - i] = 0$ (I39)

Proof: The antecedent may be established only by statement 18. p , which may do so only if $p.idx = i$ holds. In this case, 18. p also establishes the consequent.

The only statement that may falsify the consequent is 5. q , where q is any arbitrary process. Statement 5. q may potentially falsify (I39) only if executed when

$$p@{19, 20} \wedge p.idx = i \wedge q@{5} \wedge q.idx = 1 - i$$

holds. However, this contradicts (I13). \square

invariant $p@{20} \wedge p.idx = i \Rightarrow Position[1 - i] = 0$ (I40)

Proof: The antecedent may be established only by statement 19. p , which may do so only if $p.idx = i$ holds. In this case, 19. p also establishes the consequent.

The only statement that may falsify the consequent is 11. q , where q is any arbitrary process. Statement 11. q may potentially falsify (I40) only if executed when $p@{20} \wedge p.idx = i \wedge q@{11} \wedge q.idx = 1 - i$ holds. However, this contradicts (I13). \square

invariant $Position[i] \leq h < HistLen[i] \Rightarrow$
 $(\exists p :: p@{6..11} \wedge p.idx = i \wedge p.position = h)$ (I41)

Proof: The only statements that may establish the antecedent are 5. q and 18. q (which may modify $HistLen[i]$) and 11. q and 19. q (which may modify $Position[i]$), where q is any arbitrary process. Statement 5. q may establish $HistLen[i] > h$ only if executed when $q.idx = i \wedge HistLen[i] = h$ holds, in which case it establishes the consequent.

If statement 18. q modifies $HistLen[i]$, then it establishes $HistLen[i] = 0$, and hence falsifies the antecedent.

Statement 11. q may modify $Position[i]$ only if executed when $q.idx = i$ holds. In this case, by (I14), it increments $Position[i]$ by one. Hence, although 11. q may preserve the antecedent, it cannot establish it.

Statement 19. q may establish $Position[i] \leq h$ only if executed when $q.idx = 1 - i$ holds, in which case, by (I39), it establishes $Position[i] = HistLen[i] = 0$. Thus, the antecedent is false after its execution.

The only statement that may falsify the consequent is 11. p , which may do so only if executed when $p.idx = i \wedge p.position = h$ holds (which comes from the consequent itself, which is assumed to hold before being falsified). In this case, by (I14) and (I26), statement 11. p establishes $Position[i] = h + 1$, and hence falsifies the antecedent. \square

$$\begin{aligned}
\text{invariant } Position[i] = h > 0 &\Rightarrow \\
Signal[i][Hist[i][h]] = true &\vee & \mathcal{A} \\
(\exists p :: p@\{12..21\} \wedge p.idx = i \wedge p.position = h - 1) &\vee & \mathcal{B} \\
(\exists q :: q@\{8..11\} \wedge q.idx = i \wedge q.position = h) & & \mathcal{C} \\
& & \text{(I42)}
\end{aligned}$$

Proof: The only statement that may establish the antecedent is 11. p , where p is any arbitrary process. By (I14) and (I26), statement 11. p may establish the antecedent only if executed when $p.idx = i \wedge p.position = h - 1$ holds, in which case it establishes disjunct \mathcal{B} .

The only statements that may falsify disjunct \mathcal{A} are 5. p , 7. p , and 18. p , where p is any arbitrary process. Statement 5. p may falsify disjunct \mathcal{A} only if executed when

$$p@\{5\} \wedge p.idx = i \wedge HistLen[i] = h - 1 \quad \text{(I2)}$$

holds. Combining this with the antecedent, and using (I28), this implies $(\exists q :: q@\{19\} \wedge q.idx = 1 - i)$. Hence, by applying (I13) with ' $p \leftarrow q$ ' and ' $i \leftarrow 1 - i$ ', we have $p@\{0..3\} \vee p.idx = 1 - i$, which contradicts (I2). It follows that statement 5. p cannot falsify disjunct \mathcal{A} while the antecedent holds.

Statement 7. p may falsify disjunct \mathcal{A} only if executed when $p.idx = i$ holds, in which case, by (I26), we have $p.position = Position[i] = h$. Hence, statement 7. p establishes disjunct \mathcal{C} in this case.

Statement 18. p may falsify disjunct \mathcal{A} only if executed when $p.idx = 1 - i$ holds. In this case, by (I24), disjunct \mathcal{A} is already false before 18. p is executed.

The only statement that may falsify disjunct \mathcal{B} is 21. p (where p is any arbitrary process), which may do so only if executed when $p.idx = i \wedge p.position = h - 1$ holds. In this case, by (I3), $p.self = Hist[i][h]$ holds before its execution. Thus, statement 21. p establishes disjunct \mathcal{A} .

The only statement that may falsify disjunct \mathcal{C} is 11. p (where p is any arbitrary process), which may do so only if executed when $p.idx = i \wedge p.position = h$ holds. In this case, by (I14) and (I26), 11. p establishes $Position[i] = h + 1$, and hence falsifies the antecedent. \square

$$\begin{aligned}
\text{invariant } 0 < j < k \leq HistLen[i] \wedge k < 2N \wedge \\
Param[i][j - 1].proc \neq Param[i][k - 1].proc &\Rightarrow \\
Hist[i][j] \neq Hist[i][k] & \text{(I43)}
\end{aligned}$$

$$\text{invariant } 0 < j \leq HistLen[i] \wedge j < 2N \Rightarrow Hist[i][j] \neq \perp \quad \text{(I44)}$$

$$\begin{aligned}
\text{invariant } 0 < j < k \leq HistLen[i] \wedge k < 2N \wedge \\
Param[i][j - 1] = (p, c) \wedge Param[i][k - 1] = (p, c + 1) &\Rightarrow \\
Hist[i][j] \neq Hist[i][k] & \text{(I45)}
\end{aligned}$$

Proof: Invariants (I43)–(I45) follow easily from invariants (I16), (I17), (I18), and (I19), together with the assumption that the underlying *fetch-and- ϕ* primitive has rank at least $2N$. In particular, (I43) states that any two invocations (among the first $2N - 1$) by different processes write different values to $Tail[i]$; (I44) states that each of the first $2N - 1$ invocations writes to $Tail[i]$ a value different from \perp ; (I45) states that any two successive invocations (among the first $2N - 1$) by the same process write different values to $Tail[i]$. \square

$$\begin{aligned}
\text{invariant } Position[i] \leq h < HistLen[i] \wedge Param[i][h].proc = p &\Rightarrow \\
p@\{6..11\} \wedge p.idx = i \wedge p.position = h & \text{(I46)}
\end{aligned}$$

Proof: The only statements that may establish the antecedent are 5. q and 18. q (which may update $HistLen[i]$ and $Param[i][h].proc$) and 11. q and 19. q (which may update $Position[i]$), where q is any arbitrary process.

Statement 5. q may establish $h < \text{HistLen}[i] \wedge \text{Param}[i][h].\text{proc} = p$ only if executed when $\text{HistLen}[i] = h \wedge p = q \wedge q.\text{idx} = i$ holds, in which case it establishes the antecedent.

Statement 18. q may update $\text{HistLen}[i]$ or $\text{Param}[i][h]$ only if $q.\text{idx} = 1 - i$, in which case it establishes $\text{HistLen}[i] = 0$. Thus, the antecedent is false after its execution.

Statement 11. q may update $\text{Position}[i]$ only if executed when $q.\text{idx} = i$ holds. In this case, by (I14), it increments $\text{Position}[i]$ by one. Hence, although 11. q may preserve the antecedent, it cannot establish it.

Statement 19. q may establish $\text{Position}[i] \leq h$ only if executed when $q.\text{idx} = 1 - i$ holds, in which case, by (I39), it establishes $\text{Position}[i] = \text{HistLen}[i] = 0$. Thus, the antecedent is false after its execution.

The only statement that may falsify the consequent is 11. p , which may do so only if executed when $p.\text{idx} = i \wedge p.\text{position} = h$ holds. In this case, by (I14) and (I26), statement 11. p establishes $\text{Position}[i] = h + 1$, and hence the antecedent is false after its execution. \square

$$\text{invariant } p@6 \wedge p.\text{idx} = i \Rightarrow \text{Position}[i] \leq p.\text{position} \quad (\text{I47})$$

Proof: The only statement that may establish the antecedent is 5. p , which may do so only if $p.\text{idx} = i$. Let h be the value of $\text{HistLen}[i]$ before the execution of 5. p . By (I28), we have either $\text{Position}[i] \leq h$ or $(\exists q :: q@19 \wedge q.\text{idx} = 1 - i)$. However, by applying (I13) with ' $i \leftarrow 1 - i$ ', the latter implies $p@0..3 \vee p.\text{idx} = 1 - i$, which implies that the antecedent is false. On the other hand, if $\text{Position}[i] \leq h$ holds before the execution of 5. p , then $\text{Position}[i] \leq h = p.\text{position}$ is established.

The only statement that may falsify the consequent (while the antecedent holds) is 11. q (where q is any arbitrary process), which may do so only if $q.\text{idx} = i$. In this case, by (I14), 11. q increments $\text{Position}[i]$ by one, and hence it may falsify the consequent only if executed when $\text{Position}[i] = p.\text{position}$ holds. By applying (I26) with ' $p \leftarrow q$ ', we also have $\text{Position}[i] = q.\text{position}$. Combining these assertions with the antecedent, and using (I29), we have $p = q$. However, in this case, the antecedent is false after the execution of 11. q . \square

$$\text{invariant } \text{Position}[i] \leq j < k < \text{HistLen}[i] \Rightarrow \text{Hist}[i][j] \neq \text{Hist}[i][k] \quad (\text{I48})$$

Proof: The only statements that may establish the antecedent are 5. p and 18. p (which may update $\text{HistLen}[i]$, $\text{Hist}[i][j]$, or $\text{Hist}[i][k]$) and 11. p and 19. p (which may update $\text{Position}[i]$), where p is any arbitrary process.

Statement 5. p may establish the antecedent only if executed when $p.\text{idx} = i \wedge \text{Position}[i] \leq j < k = \text{HistLen}[i]$ holds. In this case, by (I38),

$$\text{Position}[i] \leq j < k = \text{HistLen}[i] < 2N \quad (13)$$

holds before its execution. We consider two cases.

- If $j = 0$, then by (I18), we have $\text{Hist}[i][j] = \perp$. Also, by applying (I44) with ' $j \leftarrow k$ ', and using (13), we have $\text{Hist}[i][k] \neq \perp$. Hence, the consequent of (I48) holds before and after the execution of 5. p .
- If $j > 0$, then let (q, c_1) denote the value of $\text{Param}[i][j - 1]$ and (r, c_2) denote the value of $\text{Param}[i][k - 1]$ prior to the execution of 5. p . If $q \neq r$, then by (13) and (I43), the consequent holds before and after the execution of 5. p .

Therefore, assume that $q = r$. Let q_l denote the value of $\text{Param}[i][l - 1].\text{proc}$ prior to the execution of 5. p , for each $0 < l \leq k$. Then, we have $q = q_j = q_k$.

For each l satisfying $\text{Position}[i] < l \leq k$, by applying (I46) with ' $p \leftarrow q_l$ ' and ' $h \leftarrow l - 1$ ', and using (13), we have $q_l.\text{position} = l - 1$ prior to the execution of 5. p . In particular, we have $q.\text{position} = q_k.\text{position} = k - 1$, and

$$(\forall l : \text{Position}[i] < l < k :: q_l \neq q). \quad (14)$$

Since $q_j = q$, this implies that $\text{Position}[i] < j < k$ is false. Thus, by (13), we have $j = \text{Position}[i]$. Combining this with (14), we also have

$$(\forall l : j < l < k :: \text{Param}[i][l - 1].\text{proc} \neq q)$$

prior to the execution of 5.p.

Therefore, by applying (I19) with ‘ $p \leftarrow q$ ’, ‘ $j \leftarrow j - 1$ ’, and ‘ $k \leftarrow k - 1$ ’, and using (13) and the assertions above, we have $c_2 = c_1 + 1$. Combining this with (13), and using (I45), it follows that the consequent holds both before and after the execution of 5.p.

If statement 18.p updates $\text{HistLen}[i]$, then it establishes $\text{HistLen}[i] = 0$, and hence the antecedent is false after its execution.

If statement 11.p updates $\text{Position}[i]$, then by (I14), it increments $\text{Position}[i]$ by one. Hence, although 11.p may preserve the antecedent, it cannot establish it.

Statement 19.p may establish $\text{Position}[i] \leq j$ only if executed when $q.\text{idx} = 1 - i$ holds, in which case, by (I39), it establishes $\text{Position}[i] = \text{HistLen}[i] = 0$. Thus, the antecedent is false after its execution.

The only statement that may falsify the consequent is 5.p (which may update either $\text{Hist}[i][j]$ or $\text{Hist}[i][k]$), where p is any arbitrary process. Statement 5.p may update $\text{Hist}[i][j]$ only if executed when $p.\text{idx} = i \wedge \text{HistLen}[i] = j - 1$ holds, in which case it establishes $\text{HistLen}[i] = j$. Thus, in this case, the antecedent is false after 5.p is executed. Similar reasoning applies to $\text{Hist}[i][k]$. \square

invariant $p@\{11..22\} \Rightarrow p.\text{pos} = p.\text{position}$ (I49)

Proof: The only statement that may establish the antecedent is 10.p, which also establishes the consequent by (I26). The consequent cannot be falsified while the antecedent holds. \square

Proof of Starvation-Freedom

To establish the Starvation-Freedom property, we begin with proving the following *unless* and *leads-to* properties. Informally, (U1) states that if a process p is waiting at statement 6, and if the busy-waiting condition is established, then it holds continuously until p exits the busy-waiting loop. (L1) (respectively, (L2)) is used to prove that, if p has entered the current queue (respectively, the old queue), and waits at statement 6, then the busy-waiting condition is eventually established. Throughout this section, we assume that the Entry_2 and Exit_2 routines are starvation-free.

$p@\{6\} \wedge p.\text{idx} = i \wedge \text{Signal}[i][p.\text{prev}] = \text{true} \text{ unless } p@\{7\}$ (U1)

Proof: The only statement that may falsify $p@\{6\} \wedge p.\text{idx} = i$ is 6.p, which establishes $p@\{7\}$.

$\text{Signal}[i][p.\text{prev}] = \text{true}$ may be falsified only if some process $q (\neq p)$ executes statement 7.q when $q.\text{idx} = i \wedge q.\text{prev} = p.\text{prev}$ holds. However, if the left-hand side of (U1) is true, then this is precluded by (I9). \square

$\text{CurrentQueue} = i \wedge p@\{6\} \wedge p.\text{idx} = i \wedge p.\text{position} = h \wedge \text{Position}[i] = k$
leads-to $(\text{Signal}[i][p.\text{prev}] = \text{true} \vee \text{Position}[i] = k + 1)$ (L1)

$\text{CurrentQueue} = 1 - i \wedge p@\{6\} \wedge p.\text{idx} = i \wedge p.\text{position} = h \wedge \text{Position}[i] = k$
leads-to $(\text{Signal}[i][p.\text{prev}] = \text{true} \vee \text{Position}[i] = k + 1)$ (L2)

Proof: We prove (L1) and (L2) by induction on h . Since their proofs are nearly identical, we simply say the “left-hand side” when the argument applies to both (L1) and (L2).

First, assume $h = 0$. The assertion $p@\{6\} \wedge p.\text{idx} = i \wedge p.\text{position} = h$ may be established only if statement 5.p is executed when $\text{HistLen}[i] = 0$ holds. However, by (I17) and (I18), this implies that $\text{Tail}[i] = \text{Hist}[i][0] = \perp$ holds. Thus, statement 5.p establishes $p@\{8\}$, and cannot establish the left-hand side. It follows that (L1) and (L2) hold vacuously for $h = 0$.

Now assume that $h > 0$, and that (L1) and (L2) hold for smaller values of h . Consider a state t that satisfies the left-hand side.

By (I13), the left-hand side implies $\neg(\exists q :: q@\{19\} \wedge q.\text{idx} = 1 - i)$. Thus, by (I28), and using $p@\{6\} \wedge p.\text{idx} = i$, we have $k < \text{HistLen}[i]$. By applying (I41) with ‘ $h \leftarrow k$ ’, this in turn implies that a process q exists such that

$$q@\{6..11\} \wedge q.\text{idx} = i \wedge q.\text{position} = \text{Position}[i] = k. \quad (15)$$

We consider two cases, depending on the value of k .

Case 1: $k = 0$. If $k = 0$, then by (I18), and by applying (I3) with ‘ p ’ $\leftarrow q$ and ‘ h ’ $\leftarrow k$, we have $q.prev = \text{Hist}[i][0] = \perp$. Because $q@\{6, 7\} \Rightarrow q.prev \neq \perp$ is (trivially) an invariant, this implies that $q@\{8..11\}$ holds. Thus, q eventually executes statement 11 while $q.idx = i \wedge q.position = k$ holds. In this case, by (I14) and (I26), $11.q$ establishes $Position[i] = k + 1$, and hence the right-hand side of (L1)/(L2) is established.

Case 2: $k > 0$. Let x be the value of $\text{Hist}[i][k]$ at state t . By (I42), (15) implies that one of the following holds at state t , where r is some process.

$$\begin{aligned} \mathcal{A}: & \text{Signal}[i][x] = \text{true}, \\ \mathcal{B}: & r@\{12..21\} \wedge r.idx = i \wedge r.position = k - 1, \quad \text{or} \\ \mathcal{C}: & r@\{8..11\} \wedge r.idx = i \wedge r.position = k. \end{aligned}$$

Moreover, by (I3), $q.prev = x$ also holds at state t . We now prove that, in each of the three cases given by \mathcal{A} – \mathcal{C} , the right-hand side of (L1)/(L2) is eventually established. Toward this goal, we prove the following three claims.

Claim 1: If \mathcal{A} is true at state u , where u is either t or some later state, then the right-hand side of (L1)/(L2) is true at either t or some later state.

Claim 2: If \mathcal{B} is true at state t , then \mathcal{A} is true at either t or some later state u .

Claim 3: If \mathcal{C} is true at state t , then the right-hand side of (L1)/(L2) is true at either t or some later state.

Proof of Claim 1: First, if (15) is false at state u , then since it holds at state t , the execution of statement $11.q$ occurs between state t and state u . Note that $q.idx$ and $q.position$ do not change while $q@\{6..11\}$ holds. Hence, by (I14) and (I26), $q.position = q.pos = Position[i] = k$ holds before the execution of $11.q$. Therefore, $11.q$ establishes the right-hand side of (L1)/(L2).

On the other hand, assume that (15) is true at state u . Then, by (I3), $q.prev = x$ holds at state u . Moreover, by (U1), if $q@\{6\} \wedge q.prev = x$ holds at state u , then \mathcal{A} continues to hold until $q@\{7\}$ is established. It follows that q eventually executes statement $11.q$, which establishes the right-hand side of (L1)/(L2) as shown above. \square

Proof of Claim 2: Assume that \mathcal{B} holds at state t . In this case, by (I3), $r.self = \text{Hist}[i][k] = x$ holds at state t . Hence, if r eventually executes statement 21, then $21.r$ establishes \mathcal{A} .

Thus, it suffices to show that statement $21.r$ is eventually executed. Since $r@\{12..21\}$ holds at state t , it suffices to show the following.

If $r@\{14, 15\}$ holds at state u , where u is either t or some later state, then $r@\{21\}$ is eventually established.

Clearly, $r@\{14, 15\}$ implies $0 \leq r.pos < N$, which is true at state t (where \mathcal{B} holds) as well as state u . Thus, by (I15), $1 \leq Position[i] = r.pos + 1 \leq N$ holds at state t . Hence, by (I37), $CurrentQueue = i$ holds at state t . Thus, if the left-hand side of (L2) is true at state t , then \mathcal{B} is false at state t .

On the other hand, if the left-hand side of (L1) is true at state t , then by (L7), given later, $r@\{21\}$ is eventually established. (Note that the proof of (L2) does not depend on (L7). As explained shortly, this is necessary in order to avoid circular reasoning.) \square

Proof of Claim 3: Clearly, r eventually executes statement $11.r$ if \mathcal{C} holds. Thus, by (I14) and (I26), $11.r$ establishes $Position[i] = k + 1$, and hence the right-hand side of (L1)/(L2) is established. \square

Finally, from these three claims, (L1) and (L2) follow. \square

The reader may wonder why we have two separate properties (L1) and (L2), when they can be proved in essentially the same way. The reason is that the proof of (L7), given later, indirectly depends on (L2). Since the proof of (L1) depends on (L7), (L1) and (L2) must be kept separate to avoid circular reasoning.

The following properties are consequences of (U1), (L1), and (L2).

$$CurrentQueue = i \wedge p@{6} \wedge p.idx = i \text{ leads-to } p@{7} \quad (L3)$$

$$CurrentQueue = 1 - i \wedge p@{6} \wedge p.idx = i \text{ leads-to } p@{7} \quad (L4)$$

Proof: Since $Position[i]$ is bounded by (I22), by inductively applying (L1) and (L2), respectively, we have the following.

$$\begin{aligned} CurrentQueue = i \wedge p@{6} \wedge p.idx = i \text{ leads-to} \\ Signal[i][p.prev] = true; \end{aligned} \quad (16)$$

$$\begin{aligned} CurrentQueue = 1 - i \wedge p@{6} \wedge p.idx = i \text{ leads-to} \\ Signal[i][p.prev] = true. \end{aligned} \quad (17)$$

Assume that the left-hand side of (L3) holds at some state t . By (16), $Signal[i][p.prev] = true$ is eventually established at some later state u . If $p@{7}$ is established before state u , then (L3) holds. Otherwise, $p@{6}$ continues to hold from state t to u . Thus, $p.idx = i$ also continues to hold from state t to u . Therefore, the left-hand side of (U1) holds at state u . By (U1), $Signal[i][p.prev] = true$ is not falsified until $p@{7}$ is established, and hence p eventually establishes $p@{7}$ by executing statement 6.

The reasoning for (L4) is similar, except that (17) is used instead of (16). \square

Note that the proof of (L3) indirectly depends on (L7), while the proof of (L4) does not.

The following properties state that, if a process p is waiting for process q at statements 14 and 15, then the busy-waiting condition is eventually established. (Note that $q@{0}$ implies $Active[q] = false$ by (I33), and $q@{6} \wedge q.idx = i$ implies $QueueIdx[q] = i$ by (I34).)

$$\begin{aligned} p@{14, 15} \wedge p.idx = i \wedge p.pos = q \\ \text{leads-to } q@{0} \vee (q@{6} \wedge q.idx = i) \vee p@{21} \end{aligned} \quad (L5)$$

$$\begin{aligned} p@{14, 15} \wedge p.idx = i \wedge p.pos = q \wedge q@{1} \\ \text{leads-to } (q@{6} \wedge q.idx = i) \vee p@{21} \end{aligned} \quad (L6)$$

Proof: Assume that the left-hand side of either (L5) or (L6) holds at state t . By (I2), one of the following holds at t .

$$\begin{aligned} \mathcal{A} : q@{3..22} \wedge q.idx = 1 - i; \\ \mathcal{B} : q@{22} \wedge q.idx = i; \\ \mathcal{C} : q@{0}; \\ \mathcal{D} : q@{1..3}; \\ \mathcal{E} : q@{4, 5} \wedge q.idx = i; \\ \mathcal{F} : q@{6} \wedge q.idx = i. \end{aligned}$$

Also, by (I15), $Position[i] = q + 1$ holds at state t , and hence, by (I37), $CurrentQueue = i$ also holds at state t .

If $p@{21}$ is established at some future state, then (L5) and (L6) both hold. Thus, in the rest of the proof, we assume that $p@{14, 15} \wedge p.idx = i$ holds continuously at and after state t . We claim that $CurrentQueue = i$ also holds at all future states. Note that $CurrentQueue = i$ may be falsified only by statement 20. r (where r is any arbitrary process), which may do so only if executed when $r.idx = i$ holds. However, by (I2), this is precluded when $p@{14, 15} \wedge p.idx = i$ holds. Thus, we have the following.

$$\bullet p@{14, 15} \wedge p.idx = i \wedge CurrentQueue = i \text{ holds at } t \text{ and all later states.} \quad (18)$$

Note that the left-hand side of (L6) implies \mathcal{D} , and \mathcal{F} implies the right-hand side of (L6). Hence, in order to prove (L6), it suffices to prove the following.

- If $\mathcal{D} \vee \mathcal{E}$ holds at state u , where u is either t or some later state, then \mathcal{F} is established at some state after u . (19)

Also, since $\mathcal{C} \vee \mathcal{F}$ implies the right-hand side of (L5), in order to prove (L5), it suffices to prove the following claim in addition to (19).

- If $\mathcal{A} \vee \mathcal{B}$ holds at state u , where u is either t or some later state, then \mathcal{C} is established at some state after u . (20)

We prove (19) and (20) by considering each of \mathcal{A} , \mathcal{B} , \mathcal{D} , and \mathcal{E} .

- Assume that \mathcal{A} holds at state u . We claim that, in this case, \mathcal{B} is eventually established.

It suffices to show that the busy-waiting loops at statement 6 and statements 14 and 15 eventually terminate for q . By applying (L4) with ' p ' $\leftarrow q$ and ' i ' $\leftarrow 1 - i$, and using $CurrentQueue = i$ (given in (18)), it follows that the former loop eventually terminates. If $q@{14, 15}$, then by (I15), we have $Position[1 - i] = q.pos + 1$. Also, by (18) and (I37), we have $Position[1 - i] = 0 \vee Position[1 - i] > N$. Combining these two assertions, we have $q.pos = Position[1 - i] - 1 \geq N$, and hence $q@{14, 15}$ is false. It follows that q in fact does not execute the busy-waiting loop at statements 14 and 15 while \mathcal{A} holds.

- Assume that \mathcal{B} holds at state u . Clearly, \mathcal{C} is eventually established.
- Assume that \mathcal{D} holds at state u . Then, by (18), \mathcal{E} is eventually established.
- Assume that \mathcal{E} holds at state u . In this case, q eventually executes statement 5. By (I38), $HistLen[i] < 2N$ holds before the execution of 5. q . Moreover, by (18) and (I3), $HistLen[i] > p.position \geq 0$ also holds. (Note that $HistLen[i]$ is always nonnegative by (I21). Since $p.position$ is updated only by line 5d, it follows that $p.position$ is always nonnegative.) Combining these two assertions with (I4), we have $Tail[i] \neq \perp$. It follows that statement 5. q establishes \mathcal{F} .

From the reasoning above, assertions (20) and (19) follow. Therefore, we have (L5) and (L6). □

Note that the proofs of (L5) and (L6) do not depend on (L7).

The following property states that the busy-waiting loop at statements 14 and 15 eventually terminates.

$$p@{14, 15} \text{ leads-to } p@{21} \tag{L7}$$

Proof: For the sake of contradiction, assume that $p@{21}$ is never established. Let $i = p.idx$ and $q = p.pos$. By (L5), $q@{0} \vee (q@{6} \wedge q.idx = i)$ is eventually established.

First, assume that $q@{0}$ is established. If q remains in its noncritical section forever, then by (I33), $Active[q] = false$ holds forever. Thus, p eventually establishes $p@{21}$ by executing statement 14, a contradiction.

On the other hand, if q enters its entry section again, then it establishes $q@{1}$. In this case, by (L6), q eventually establishes $q@{6} \wedge q.idx = i$.

It follows that $q@{6} \wedge q.idx = i$ is eventually established. By (I5), and using $p@{14, 15}$, it follows that $q@{6}$ remains true forever. But then, by (I34), p eventually establishes $p@{21}$ by executing statement 15, a contradiction. □

Finally, by (L3), (L4), and (L7), it follows that each **await** statement in ALGORITHM G-CC eventually terminates. Thus, ALGORITHM G-CC is starvation-free.

Appendix B: Correctness Proof for ALGORITHM G-CC' and ALGORITHM G-DSM

In this appendix, we first prove that ALGORITHM G-CC' is equivalent to ALGORITHM G-CC, and also prove that the CC-to-DSM transformation introduced in Section 3.2 can be safely applied to ALGORITHM G-CC'. The following lemma proves the former claim.

Lemma 3 *In ALGORITHM G-CC, if some process p executes statement 21 while $p.pos \geq 2N - 1 \wedge p.idx = i \wedge p.self = x$ holds, and if some process p' later reads $Signal[i][x]$ by executing statement 6, then between these two events, some process p'' writes $Signal[i][x] := false$.*

Proof: Let s and t be the states right before the executions of statements 21. p and 6. p' , respectively. First, consider state s . By (I15) and (I22), we have the following.

- $Position[i] = 2N$ holds at s . (21)

If any process q satisfies $q@\{16..20\} \wedge q.idx = 1 - i$ at state s , then by applying (I13) with ' p ' $\leftarrow q$ and ' i ' $\leftarrow 1 - i$, every process r must satisfy $r@\{0..3\} \vee r.idx = 1 - i$. However, this contradicts $p@\{21\} \wedge p.idx = i$. Therefore, we have the following.

- $\neg(\exists q :: q@\{16..20\} \wedge q.idx = 1 - i)$ holds at s . (22)

Hence, by (I28), we have $HistLen[i] - Position[i] = |\{r :: r@\{6..11\} \wedge r.idx = i\}|$. Taken together with (I21) and (21), we have

$$HistLen[i] = 2N \quad \wedge \quad \neg(\exists r :: r@\{6..11\} \wedge r.idx = i).$$

By (I38), this also implies $\neg(\exists r :: r@\{4, 5\} \wedge r.idx = i)$. In particular, process p' satisfies $\neg(p'@\{4..6\} \wedge p'.idx = i)$ at s . Since $p'@\{6\} \wedge p'.idx = i$ holds at state t , p' must execute statement 3 while $CurrentQueue = i$ holds *between* states s and t .

However, by (21) and (I32), $CurrentQueue = 1 - i$ holds at state s . Hence, some process q must write $CurrentQueue := i$, by executing statement 20 while $q.idx = 1 - i$ holds, *after* state s but *before* the execution of 3. p' .

Clearly, q also executes 18. q (while $q.idx = 1 - i$ holds) at some earlier time. If 18. q is executed before state s , then $q@\{19, 20\} \wedge q.idx = 1 - i$ holds at s , which contradicts (22). Thus, we have the following sequence of events: 21. p , 18. q , 20. q , 3. p' , and 6. p' . By applying (I24) with ' p ' $\leftarrow q$ and ' i ' $\leftarrow 1 - i$, it follows that $Signal[i][x]$ equals *false* before the execution of 18. q . Since $Signal[i][x] = true$ holds right after the execution of 21. p , some process p'' must write $Signal[i][x] := false$, after the execution of 21. p but before the execution of 18. q (and hence, before the execution of 6. p').⁹ □

By Lemma 3, whenever ALGORITHM G-CC' causes a process to skip a write at statement 21, this behavior is guaranteed not to affect any process's behavior in the future. Moreover, as a result, all the invariants stated in Appendix A also remain valid for ALGORITHM G-CC', except for those that explicitly refer to the *Signal* array.

We now formally prove that the various two-process mutual exclusion algorithms added to ALGORITHM G-DSM (shown in Figure 5) are safe, by considering each pair of **Entry/Exit** calls. Statements 4–7 and 45–48 are clearly safe, as explained in Section 3.2.

Statements 10–14: By (I9),

$$|\{p :: p@\{6, 7\} \wedge p.idx = i \wedge p.prev = x\}| \leq 1$$

is an invariant of ALGORITHM G-CC', for any i and x . This implies that the **Entry/Exit** pair at statements 10–14 in Figure 5 is safe. □

⁹In fact, it can be easily shown that p'' must be q .

Statements 25–30: In order to show that the **Entry/Exit** pair at statements 25–30 in Figure 5 is safe, it suffices to show that

$$|\{p :: p@\{14, 15, 21\} \wedge p.pos = h\}| \leq 1 \quad (23)$$

is an invariant of ALGORITHM G-CC', for any h .

Consider a process p that satisfies $p@\{14, 15, 21\} \wedge p.pos = h$. By (I15) and (I27), $Position[p.idx] = h + 1 \wedge p.position = h$ holds as well. Since $h < N$ holds by assumption (see statement 13 in Figure 7), by (I36), we also have $p.idx = CurrentQueue$. Taken together with (I29), it follows that p is uniquely determined (i.e., (23) is true). \square

Statements 40–43: In order to show that the **Entry/Exit** pair at statements 40–43 in Figure 5 is safe, it suffices to show that

$$|\{p :: p@\{21, 22\} \wedge p.idx = i \wedge p.self = x \wedge p.pos < 2N - 1\}| \leq 1 \quad (24)$$

is an invariant of ALGORITHM G-CC', for any i and x .

For the sake of contradiction, assume that there exists two distinct processes p and q , satisfying

$$\begin{aligned} & p@\{21, 22\} \wedge q@\{21, 22\} \wedge \\ & p.idx = q.idx = i \wedge p.self = q.self = x \wedge \\ & p.pos < 2N - 1 \wedge q.pos < 2N - 1. \end{aligned}$$

Define $j = p.pos$ and $k = q.pos$. By (I49), we also have $j = p.position$ and $k = q.position$. Hence, by (I29), $j \neq k$ holds. Without loss of generality, assume $j < k$.

By applying (I3) to p and q separately, we have

$$\begin{aligned} & j < k < HistLen[i] \wedge \\ & x = p.self = Hist[i][j + 1] \wedge x = q.self = Hist[i][k + 1] \wedge \\ & Param[i][j].proc = p \wedge Param[i][k].proc = q. \end{aligned}$$

Therefore, by using $k + 1 < 2N$, and applying (I43) with ' j ' $\leftarrow j + 1$ and ' k ' $\leftarrow k + 1$, we have $Hist[i][j + 1] \neq Hist[i][k + 1]$, a contradiction. \square