

Local Mutual Exclusion for Dynamic, Anonymous, Bounded Memory Message Passing Systems

Joshua J. Daymude  

Biodesign Center for Biocomputing, Security and Society, Arizona State University, Tempe, AZ

Andréa W. Richa  

School of Computing and Augmented Intelligence, Arizona State University, Tempe, AZ

Christian Scheideler  

Department of Computer Science, Paderborn University, Paderborn, Germany

Abstract

Mutual exclusion is a classical problem in distributed computing that provides *isolation* among concurrent action executions that may require access to the same shared resources. Inspired by algorithmic research on distributed systems of weakly capable entities whose connections change over time, we address the *local mutual exclusion* problem that tasks each node with acquiring exclusive locks for itself and the maximal subset of its “persistent” neighbors that remain connected to it over the time interval of the lock request. Using the established *time-varying graphs* model to capture adversarial topological changes, we propose and rigorously analyze a local mutual exclusion algorithm for nodes that are anonymous and communicate via asynchronous message passing. The algorithm satisfies *mutual exclusion* (non-intersecting lock sets) and *lockout freedom* (eventual success with probability 1) under both semi-synchronous and asynchronous concurrency. It requires $\mathcal{O}(\Delta)$ memory per node and messages of size $\Theta(1)$, where Δ is the maximum number of connections per node. We conclude by describing how our algorithm can implement the pairwise interactions assumed by *population protocols* and the concurrency control operations assumed by the *canonical amoebot model*, demonstrating its utility in both passively and actively dynamic distributed systems.

2012 ACM Subject Classification Theory of computation → Distributed algorithms; Theory of computation → Concurrency; Software and its engineering → Mutual exclusion

Keywords and phrases Mutual exclusion, dynamic networks, message passing, concurrency

Funding *Joshua J. Daymude*: NSF (CCF-1733680), U.S. ARO (MURI W911NF-19-1-0233), the Momenta Foundation’s Mistletoe Research Fellowship, and the ASU Biodesign Institute.

Andréa W. Richa: NSF (CCF-1733680, CCF-2106917) and U.S. ARO (MURI W911NF-19-1-0233).

Christian Scheideler: DFG Project SCHE 1592/6-1.

1 Introduction

Distributed computing research has grown increasingly concerned with characterizing the capabilities and limitations of systems composed of *dynamic entities* (or *nodes*). Recently, these studies have considered both biological collectives such as social insects [2, 14, 26], spiking neural networks [49], and DNA and molecular computers [13, 39, 51] as well as engineered systems such as overlay networks and the Internet of Things (IoT) [23], swarm and modular self-reconfigurable robotics [25, 28, 40, 53], and programmable matter [3, 17, 19, 37]. Entities in these systems often make decisions based only on their own knowledge (or “state”), locally-perceptible measures of their environment (e.g., pheromones, the number or density of nearby neighbors, etc.), and information communicated to them by their neighbors.

Compared to the static setting where acting nodes’ neighborhoods do not change, designing correct distributed algorithms in the dynamic setting is a challenging task. In this paper, we use the established *time-varying graphs* (TVGs) model [11, 12] to capture adversarial changes in network topology and consider weakly capable nodes that are *anonymous*, have *bounded*

memory, communicate via *asynchronous message passing*, and execute their algorithms *semi-synchronously* or *asynchronously*. The classical *mutual exclusion* problem [20] regulates how nodes enter their critical sections using locks, defined as a pair of operations **LOCK** and **UNLOCK**. Our *local mutual exclusion* problem—designed to enable nodes to locally coordinate their interactions in the dynamic, concurrent setting—defines **LOCK** as a node acquiring locks for itself and the maximal subset of its “persistent” neighbors that remain connected to it while the request is processed. A core challenge in designing such a **LOCK** operation in the dynamic setting lies in the nodes’ inability to know, when issuing lock requests, which neighbors will be persistent and which others will later be removed.

This locking mechanism greatly simplifies the design of local distributed algorithms in highly dynamic settings by providing *isolation* for concurrently executed actions. An algorithm’s actions can first be designed for the simpler sequential setting in which at most one node is active (potentially changing the system configuration) at a time. When considering the concurrent setting, each action is then treated as a critical section wrapped in a **LOCK**/**UNLOCK** pair; this ensures that no two simultaneously executing actions can involve overlapping neighborhoods. Our locking mechanism gracefully handles neighbor disconnections, ensuring that the locked and connected subset of an acting node’s neighborhood remains fixed throughout the execution of its action, just as it would be in the sequential setting. Thus, our locking mechanism restricts the algorithm designer’s concern from all possible complications arising from concurrent dynamics to just one: New connections may concurrently be established with a node while it is executing an action.

Our Contributions. We summarize our contributions as follows.

- A formalization of the *local mutual exclusion problem* in an extension of the time-varying graphs model that captures topological changes, asynchronous message passing, and semi-synchronous or asynchronous node activation (Section 2).
- A randomized algorithm implementing the **LOCK** and **UNLOCK** operations for local mutual exclusion that satisfies *mutual exclusion* (non-intersecting lock sets) and *lockout freedom* (eventual success with probability 1) under both semi-synchronous and asynchronous concurrency. This algorithm requires $\mathcal{O}(\Delta)$ memory per node and messages of size $\Theta(1)$, where Δ is the maximum number of connections per node (Sections 3–5).
- Applications of this local mutual exclusion algorithm to *population protocols* [3], establishing an underlying mechanism for guaranteeing pairwise interactions in a broader class of concurrent activation models, and the *canonical amoebot model* [17], implementing the model’s concurrency control operations (Section 6).

Related Work

Designing algorithms for concurrent computing environments is a challenging task requiring the careful control of simultaneously interacting processes and coordinated access to shared resources. Since its introduction by Dijkstra [20], the closely related *mutual exclusion* problem has received much attention from the research community. For shared memory systems, mutual exclusion can be conveniently solved by atomic operations like compare-and-swap, test-and-set, and fetch-and-add [29]. In contrast, our present focus is on asynchronous message passing. Classical approaches to mutual exclusion in asynchronous message passing systems often assume that nodes have unique identifiers and global coordination (see, e.g., the survey [47]) or make use of unbounded counters like Lamport clocks (e.g., [35, 43]), neither of which are appropriate for the anonymous, bounded memory nodes we consider here. The most relevant classical algorithm to our setting is the *arrow protocol* [18, 41] that requires

only constant memory per node to locally maintain a spanning tree rooted at the node with exclusive access to the shared resource; however, despite recent improvements [27, 32], it is not clear how to adapt this protocol to systems with dynamic topologies.

Our local variant of the mutual exclusion problem blurs the usual delineation between processes and the shared resources they’re accessing as nodes compete to gain exclusive access to their neighborhoods. Like the well-studied *k-mutual exclusion* [24] and *group mutual exclusion* [30, 31] variants, ours allows multiple nodes to be in their critical sections simultaneously; however, these variants allow multiple process to access the same shared resource(s) concurrently while ours requires that concurrently locked neighborhoods be non-intersecting. This constraint is similar to ensuring the active nodes form a *distance-3 independent set* from graph theory and is related to the more general (α, β) -*ruling sets* [5] recently solved under the LOCAL and CONGEST models [33, 44]; however, these distributed algorithms rely on static topologies, unique identifiers, and synchronous message delivery. The recent results on mutual exclusion for *fully anonymous* systems [42], like our nodes and their neighborhoods, assume that neither the processes nor the shared resources have unique identifiers. However, like the earlier classical results above and other recent models of weak finite automata [21, 22], these do not extend to dynamic network topologies.

Research on *mobile ad hoc networks* (MANETs) directly embraces node and edge dynamics, modeling wireless communication links that form and fail as nodes move in and out of each other’s transmission radii. Mutual exclusion has been exhaustively studied under MANET models [4, 6, 7, 15, 45, 50], and many of those ideas inspired recent work on mutual exclusion for intersection traffic control for autonomous vehicles [46, 52]. Mutual exclusion for MANETs is almost always solved using a token-based approach, sometimes combined with the imposition of a logical structure like a ring or tree. These approaches only apply to competitions for a single shared resource or critical section per token type; our nodes’ competitions over their local neighborhoods would need one token type per neighboring node which is not addressed by prior work. More relevant to our local variant of mutual exclusion are randomized backoff mechanisms for local contention resolution used by MANETs and wireless networks [8, 9, 10] to ensure no two nodes are broadcasting in overlapping neighborhoods; however, these rely on nodes’ chosen backoff delays to correspond to a consistent wall clock that is incompatible with our weaker model of concurrency. In any case, the standard MANET communication model of wireless broadcast with time-ordered, instantaneous receipt of messages is more powerful than our asynchronous message passing. Like MANETs, algorithms for *self-stabilizing overlay networks* (see [23] for a recent survey) similarly embrace node and edge dynamics, but often use more memory than our present algorithm and assume unique node identifiers.

Finally, we briefly highlight related models of *dynamic networks*, i.e., those whose structural properties change over time. Our model is closely related to the *time-varying graphs* (TVGs) model [11, 12] which unifies prior models of dynamic networks by capturing graph structural evolution over time through adversarial dynamics. We join recent work on message passing algorithms for TVGs that address the challenge of rapidly changing network topology [1]. The local nature of our mutual exclusion problem enables us to weaken the assumptions considered by prior works in this area. For example, we allow messages to have arbitrary but finite delays akin to asynchronous message passing, we assume weaker “semi-synchronous” and asynchronous models of concurrency, and we trade globally unique node identifiers for local port labels. In Section 6, we demonstrate how the rapid dynamics modeled by TVGs combined with these weak assumptions on node capabilities facilitate application of our mutual exclusion algorithm to both systems with *passive* dynamics [3, 48], in which nodes have no control over topological changes, and those with *active* dynamics [17, 36, 38].

in which nodes control the connections they establish and sever (e.g., via movements).

2 Preliminaries

2.1 Computational Model

We consider a distributed system composed of a fixed set of nodes V . Each node is assumed to be *anonymous*, lacking a unique identifier, and has a local memory storing its *state*. Nodes communicate with each other via message passing over a communication graph whose topology changes over time. We model this topology using a *time-varying graph* $\mathcal{G} = (V, E, T, \rho)$ where V is the set of nodes, E is a (static) set of undirected pairwise edges between nodes, $T = \{0, \dots, t_{\max}\}$ for some (possibly infinite) $t_{\max} \in \mathbb{N}$ is called the *lifetime* of \mathcal{G} , and $\rho : E \times T \rightarrow \{0, 1\}$ is the *presence* function indicating whether or not a given edge exists at a given time. A *snapshot* of \mathcal{G} at time $t \in T$ is the undirected graph $G_t = (V, \{e \in E : \rho(e, t) = 1\})$ and the *neighborhood* of a node $u \in V$ at time $t \in T$ is the set $N_t(u) = \{v \in V : \rho(\{u, v\}, t) = 1\}$. For $i \geq 0$, the i -th *round* lasts from time i to the instant just before time $i + 1$; thus, the communication graph in round i is G_i .

We assume that an adversary controls the presence function ρ and that E is the complete set of edges on nodes V ; i.e., we do not limit which edges the adversary can introduce. The only constraint we place on the adversary's topological changes is $\forall t \in T, u \in V, |N_t(u)| \leq \Delta$, where $\Delta > 0$ is the fixed number of *ports* per node. When the adversary establishes a new connection between nodes u and v , it must assign the endpoints of edge $\{u, v\}$ to open ports on u and v (and cannot do so if either node has no open ports). Node u locally identifies $\{u, v\}$ using its corresponding *port label* $\ell \in \{1, \dots, \Delta\}$ and v does likewise. For convenience of notation, we use $\ell_u(v)$ to refer to the label of the port on node u that is assigned to the edge $\{u, v\}$; this mapping of port labels to nodes is not available to the nodes. Edge endpoints remain in their assigned ports (and thus labels remain fixed) until disconnection, but nodes u and v may label $\{u, v\}$ differently and their labels are not known to each other a priori. Each node has a *disconnection detector* that adds the label of any port whose connection is severed to a set $D \subseteq \{1, \dots, \Delta\}$. A node's disconnection detector provides it a snapshot of D whenever it starts an action execution (see below) and then resets D to \emptyset .¹

Nodes communicate via message passing. A node u sends a message m to a neighbor v by calling $\text{SEND}(m, \ell_u(v))$. Message m remains in transit until either v receives and processes m at a later round chosen by the adversary, or u and v are disconnected and m is lost. Multiple messages in transit from u to v may be received by v in a different order than they were sent. A node always knows from which port it received a given message.

All nodes execute the same distributed algorithm \mathcal{A} , which is a set of *actions* each of the form $\langle \text{label} \rangle : \langle \text{guard} \rangle \rightarrow \langle \text{operations} \rangle$. An action's *label* specifies its name. Its *guard* is a Boolean predicate determining whether a node u can execute it based on the state of u and any message in transit that u may receive. An action is *enabled* for a node u if its guard is TRUE for u ; a node u is *enabled* if it has at least one enabled action. An action's *operations* specify what a node does when executing the action, structured as

1. Receiving at most one message chosen by the adversary,
2. A finite amount of internal computation and state updates, and
3. At most one call to $\text{SEND}(m, \ell)$ per port label ℓ .

¹ Without this assumption, the adversary could disconnect an edge assigned to port ℓ of node u and then immediately connect a different edge to ℓ , causing an indistinguishability issue for node u .

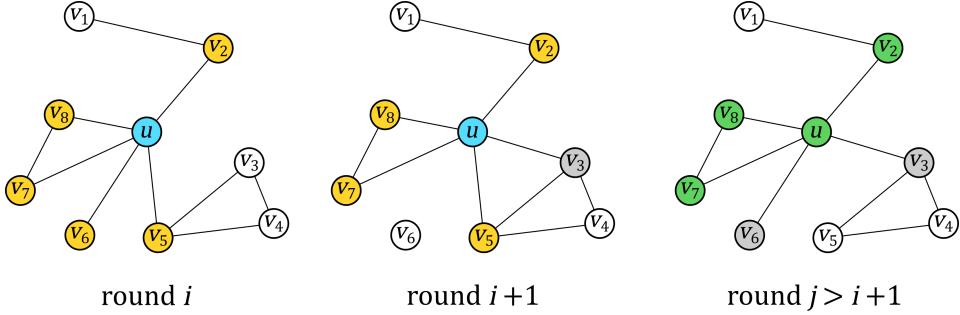


Figure 1 A node u (blue) and its persistent neighborhood (yellow) over time. Node u succeeds in locking itself and its persistent neighbors (green) in round j . Neighbors that are disconnected from u before round j (such as v_6 at time $i+1$) as well as new neighbors that connect to u from time $i+1$ to j (grey) are not persistent.

Each node executes its own instance of \mathcal{A} independently, sequentially (executing at most one action at a time), and reliably (meaning we do not consider crash or Byzantine faults). We assume an adversary controls the timing of node activations and action executions. When the adversary activates a node, it also chooses exactly one of the node's enabled actions for the node to execute; we note that this choice must be compatible with any message the adversary chooses to deliver to the node. In this work, we primarily focus on *semi-synchronous* activations, which we interpret in the time-varying graph context to mean that in each round, the adversary activates any (possibly empty) subset of enabled nodes concurrently and the activated nodes execute their specified actions within that round. In Section 5, we additionally consider *asynchronous* activations in which action executions may span arbitrary finite time intervals. We only constrain the adversary by *weak fairness*, meaning it must activate nodes such that any continuously enabled action is eventually executed and any message in transit on a continuously existent edge is eventually processed.

2.2 Local Mutual Exclusion

In the classical problem of *mutual exclusion*, nodes enter their critical sections using locks, defined as a pair of operations `LOCK` and `UNLOCK` (or “acquire” and “release”). A node issues a lock request by calling `LOCK`; once acquired, it is assumed that a node eventually releases these locks by calling `UNLOCK`. Our *local mutual exclusion* variant is concerned with nodes acquiring exclusive access to themselves and their immediate neighbors, though in the present context of dynamic networks, these neighborhoods may change over time.

Formally, each node u stores a variable $\text{lock} \in \{\perp, 0, \dots, \Delta\}$ that is equal to \perp if u is unlocked, 0 if u has locked itself, and $\ell_u(v) \in \{1, \dots, \Delta\}$ if u is locked by v . The *lock set* of a node u in round i is $\mathcal{L}_i(u) = \{v \in N_i(u) : \text{lock}(v) = \ell_u(v)\}$ which additionally includes u itself if $\text{lock}(u) = 0$. Suppose that in round i , a node u calls `LOCK` to issue a lock request of its current closed neighborhood $N_i[u] = \{u\} \cup N_i(u)$. This lock request *succeeds* at some later round $j > i$ if round j is the first in which $\mathcal{L}_j(u) = \{u\} \cup \{v \in N_i(u) : \forall t \in [i, j], \{u, v\} \in G_t\}$; i.e., j is the earliest round in which u obtains locks for itself and every *persistent* neighbor that remained connected to u in rounds i through j (see Figure 1). Our goal is to design an algorithm \mathcal{A} implementing `LOCK` and `UNLOCK` that satisfies the following properties:

- *Mutual Exclusion.* For all rounds $i \in T$ and all pairs of nodes $u, v \in V$, $\mathcal{L}_i(u) \cap \mathcal{L}_i(v) = \emptyset$.
- *Lockout Freedom.* Every issued lock request eventually succeeds with probability 1.

■ **Table 1** The notation, domain, initialization, and description of the local variables used in the algorithm for local mutual exclusion by a node u .

Var.	Domain	Init.	Description
<code>lock</code>	$\{\perp, 0, \dots, \Delta\}$	\perp	\perp if u is unlocked, 0 if u has locked itself, and $\ell_u(v)$ if u is locked by v
<code>state</code>	$\{\perp, \text{PREPARE}, \text{COMPETE}, \text{WIN}, \text{LOCKED}, \text{UNLOCK}\}$	\perp	The lock state of node u
<code>phase</code>	$\{\perp, \text{PREPARE}, \text{COMPETE}\}$	\perp	The algorithm phase node u is in
L	$\subseteq N[u]$	\emptyset	Ports (nodes) u intends to lock
R	$\subseteq N[u]$	\emptyset	Ports via which u has received <code>ready()</code> , <code>ack-lock()</code> , or <code>ack-unlock()</code> responses
W	$\subseteq N[u] \times \{\text{TRUE}, \text{FALSE}\}$	\emptyset	Port-outcome pairs of <code>win()</code> messages u has received
H	$\subseteq N[u]$	\emptyset	Ports (nodes) on hold for the competition to lock u
A	$\subseteq N[u]$	\emptyset	Ports (nodes) of applicants that can join the competition to lock u
C	$\subseteq N[u]$	\emptyset	Ports (nodes) of candidates competing to lock u
P	$\subseteq C(u) \times \{0, \dots, K - 1\}$	\emptyset	Port-priority pairs of the candidates

Following the long tradition of mutual exclusion problem definitions, our local mutual exclusion problem is defined in terms of mutual exclusion and fairness properties. However, because each `lock` variable points to at most one node per time, it is impossible for two nodes' lock sets to intersect, trivially satisfying the mutual exclusion property. Nevertheless, satisfying lockout freedom remains challenging, especially in highly dynamic settings. When issuing lock requests, nodes do not know which of their connections will remain stable and which will disconnect by the time their coordination is complete. Thus, our problem variant captures what it means for nodes to lock their maximal persistent neighborhoods despite unpredictable and rapid topological changes.

3 Algorithm for Local Mutual Exclusion

Our randomized algorithm for the local mutual exclusion problem specifies actions for the execution of `LOCK` and `UNLOCK` operations satisfying mutual exclusion and lockout freedom. An execution of the `LOCK` operation by a node u is organized into two phases: a *preparation phase* (Algorithm 1) in which u determines and notifies the nodes $L(u)$ it intends to lock, and a *competition phase* (Algorithm 2) in which u attempts to lock all nodes in $L(u)$, contending with any other nodes v for which $L(v) \cap L(u) \neq \emptyset$. An execution of the `UNLOCK` operation (Algorithm 3) by node u is straightforward, simply notifying all nodes in $L(u)$ that their locks are released. All local variables used in our algorithm are listed in Table 1 as they appear in the pseudocode. In a slight abuse of notation, we use $N[u]$ and the subsets thereof to represent both the nodes in the closed neighborhood of u and the port labels of u they are connected to. For clarity of presentation, the algorithm pseudocode allows for a node to send messages to itself (via "port 0") just as it sends messages to its neighbors, though in reality these self-messages would be implemented with in-memory variable updates.

We refer to nodes that call `LOCK/UNLOCK` as *initiators* and the nodes that are being locked or unlocked as *participants*; it is possible for a node to be an initiator and participant simultaneously. Initiators progress through a series of *lock states* associated with the `state` variable; participants advance through the algorithm's *phases* as indicated by the `phase`

Algorithm 1 The LOCK Operation: Preparation Phase for Node u

```

1: INITLOCK: On LOCK being called →                               ▷ Initiator initiates a lock request.
2:   if state = ⊥ then                                         ▷ Only one locking operation at a time.
3:     CLEANUP().
4:     Set state ← PREPARE and  $L \leftarrow N[u]$ .
5:     for all  $\ell \in L$  do SEND(prepare(),  $\ell$ ).
6: RECEIVEPREPARE: On receiving prepare() via port  $\ell$  →
7:   CLEANUP().
8:   if phase = COMPETE then set  $H \leftarrow H \cup \{\ell\}$ .      ▷ Put  $\ell$  on hold if already competing.
9:   else
10:    Set  $A \leftarrow A \cup \{\ell\}$  and phase ← PREPARE.          ▷ Add  $\ell$  as an applicant otherwise.
11:    SEND(ready(),  $\ell$ ).
12: RECEIVEREADY: On receiving ready() via port  $\ell$  →
13:   CLEANUP().
14:   Set  $R \leftarrow R \cup \{\ell\}$ .
15: CHECKSTART: (state = PREPARE)  $\wedge$  ( $R = L$ ) →             ▷ All ready() messages received.
16:   CLEANUP().
17:   Set state ← COMPETE,  $R \leftarrow \emptyset$ , and  $W \leftarrow \emptyset$ .
18:   Choose priority  $p \in \{0, \dots, K-1\}$  uniformly at random.
19:   for all  $\ell \in L$  do SEND(request-lock( $p$ ),  $\ell$ ).
20: CLEANUP: (phase ≠ ⊥)  $\vee$  (state = UNLOCK) →
21:   CLEANUP().
22: function CLEANUP()                                         ▷ Helper function for processing disconnections  $D$ .
23:   for all  $\ell \in D$  do
24:     if lock =  $\ell$  then lock ← ⊥.
25:     Remove  $\ell$  from all sets:  $L \leftarrow L \setminus \{\ell\}$ ,  $R \leftarrow R \setminus \{\ell\}$ ,  $W \leftarrow W \setminus \{(\ell, \cdot)\}$ ,  $H \leftarrow H \setminus \{\ell\}$ ,
26:      $A \leftarrow A \setminus \{\ell\}$ ,  $C \leftarrow C \setminus \{\ell\}$ , and  $P \leftarrow P \setminus \{(\ell, \cdot)\}$ .
27:   if  $C = \emptyset$  then
28:     for all  $\ell \in H$  do SEND(ready(),  $\ell$ ).
29:   Set  $A \leftarrow A \cup H$  and  $H \leftarrow \emptyset$ .                      ▷ All nodes on hold become applicants.
30:   if  $A \neq \emptyset$  then set phase ← PREPARE.
31:   else set phase ← ⊥.

```

variable. We first describe the algorithm from an initiator's perspective and then describe the complementary participants' actions. A special CLEANUP helper function ensures that the nodes adapt to any disconnections affecting their variables that may have occurred since they last acted, so we omit the handling of these disconnections in the following description.

When an initiator u calls LOCK, it advances to the PREPARE state, sets $L(u)$ to all nodes in its closed neighborhood $N[u]$, and then sends `prepare()` messages to all nodes of $L(u)$. Once it has received `ready()` responses from all nodes of $L(u)$, it advances to the COMPETE state and joins the competitions for each node in $L(u)$ by sending `request-lock(p)` messages to all nodes of $L(u)$, where p is a priority chosen uniformly at random from $\{0, \dots, K-1\}$ for a fixed $K = \Theta(1)$. It then waits for the outcomes of these competitions. If it receives at least one `win(FALSE)` message, it lost this competition and must compete again. Otherwise, if all responses are `win(TRUE)`, it advances to the WIN state and sends `set-lock()` messages to all nodes of $L(u)$. Once it has received `ack-lock()` responses from all nodes of $L(u)$, it advances to the LOCKED state indicating $L(u)$ now represents the lock set $\mathcal{L}(u)$.

A participant v is responsible for coordinating the competition among all initiators that want to lock v . To delineate successive competitions, v distinguishes among initiators that

■ **Algorithm 2** The LOCK Operation: Competition Phase for Node u

```

1: RECEIVEREQUEST: On receiving request-lock( $p$ ) via port  $\ell \rightarrow$ 
2:   CLEANUP().
3:   if  $\ell \in A$  then set  $A \leftarrow A \setminus \{\ell\}$  and  $C \leftarrow C \cup \{\ell\}$ .
4:   Set  $P \leftarrow P \cup \{(\ell, p)\}$  and phase  $\leftarrow$  COMPETE.                                 $\triangleright$  Close competition.
5: CHECKPRIORITIES: (phase = COMPETE)  $\wedge$  ( $|C| = |P|$ )  $\rightarrow$                                  $\triangleright$  All priorities received.
6:   CLEANUP().
7:   if lock =  $\perp$  and  $\exists(\ell, p) \in P$  with a unique highest  $p$  then
8:     SEND(win(TRUE),  $\ell$ ) and SEND(win(FALSE),  $\ell'$ ) for all  $\ell' \in C \setminus \{\ell\}$ .
9:   else SEND(win(FALSE),  $\ell$ ) for all  $\ell \in C$ .
10:  Reset  $P \leftarrow \emptyset$ .                                          $\triangleright$  Competition is over.
11: RECEIVEWIN: On receiving win( $b$ ) via port  $\ell \rightarrow$ 
12:   CLEANUP().
13:   Set  $W \leftarrow W \cup \{(\ell, b)\}$ .
14: CHECKWIN: (state = COMPETE)  $\wedge$  ( $|W| = |L|$ )  $\rightarrow$                                  $\triangleright$  All win( $b$ ) replies received.
15:   CLEANUP().
16:   if  $\exists(\cdot, \text{FALSE}) \in W$  then                                 $\triangleright$  Start new locking attempt.
17:     Choose priority  $p \in \{0, \dots, K - 1\}$  uniformly at random.
18:     for all  $\ell \in L$  do SEND(request-lock( $p$ ),  $\ell$ ).
19:   else                                          $\triangleright$  Succeeded in locking.
20:     Set state  $\leftarrow$  WIN and reset  $R \leftarrow \emptyset$ .
21:     for all  $\ell \in L$  do SEND(set-lock( $\cdot$ ),  $\ell$ ).
22:   Reset  $W \leftarrow \emptyset$ .
23: RECEIVESETLOCK: On receiving set-lock( $\cdot$ ) via port  $\ell \rightarrow$ 
24:   Set lock  $\leftarrow \ell$  and  $C \leftarrow C \setminus \{\ell\}$ .
25:   CLEANUP().
26:   SEND(ack-lock( $\cdot$ ),  $\ell$ ).
27: RECEIVEACKLOCK: On receiving ack-lock( $\cdot$ ) via port  $\ell \rightarrow$ 
28:   CLEANUP().
29:   Set  $R \leftarrow R \cup \{\ell\}$ .
30: CHECKDONE: (state = WIN)  $\wedge$  ( $R = L$ )  $\rightarrow$                                  $\triangleright$  All lock acknowledgements received.
31:   CLEANUP().
32:   Set state  $\leftarrow$  LOCKED and reset  $R = \emptyset$ .
33:   return  $L$ .                                          $\triangleright$  Locking complete.

```

are *candidates* in the current competition, *applicants* that may join the current competition, and those that are *on hold* for the next competition. When v receives a **prepare**(\cdot) message from an initiator u , it either puts u on hold if a competition is already underway or adds u as an applicant and replies **ready**(\cdot) otherwise. Participant v promotes its applicants to candidates when v receives their **request-lock**(p) messages. Once all such messages are received from the competition's candidates, v notifies the one with the unique highest priority of its success and all others of their failure (or, in the case of a tie, all candidates fail). A winning competitor is removed from the candidate set while all others remain to try again; once the candidate set is empty, v promotes all initiators that were on hold to applicants. Finally, when v receives a **set-lock**(\cdot) message, it sets its **lock** variable accordingly and acknowledges this with an **ack-lock**(\cdot) response.

Algorithm 3 The UNLOCK Operation for Node u

```

1: INITUNLOCK: On UNLOCK being called →                               ▷ Initiator initiates an unlock.
2:   if state = LOCKED then                                         ▷ Only one UNLOCK per successful LOCK.
3:     CLEANUP().
4:     Set state ← UNLOCK and reset  $R \leftarrow \emptyset$ .
5:     for all  $\ell \in L$  do SEND(release-lock(),  $\ell$ ).
6: RECEIVERELEASE: On receiving release-lock() via port  $\ell$  →
7:   CLEANUP().
8:   Set lock ←  $\perp$  and SEND(ack-unlock(),  $\ell$ ).
9: RECEIVEACKUNLOCK: On receiving ack-unlock() via port  $\ell$  →
10:  CLEANUP().
11:  Set  $R \leftarrow R \cup \{\ell\}$ .
12: CHECKUNLOCKED: (state = UNLOCK)  $\wedge$  ( $R = L$ ) → ▷ All unlock acknowledgements received.
13:  CLEANUP().
14:  Reset state ←  $\perp$  and  $R = \emptyset$ .                                     ▷ Unlocking complete.

```

4 Analysis

In this section, we prove the following theorem.

► **Theorem 1.** *If all nodes start with the initial values given by Table 1, the algorithm satisfies the mutual exclusion and lockout freedom properties under semi-synchronous concurrency, requires $\mathcal{O}(\Delta)$ memory per node and messages of size $\Theta(1)$, and has at most two messages in transit along any edge at any time.*

The algorithm in Section 3 is written with respect to local port labels; for ease of presentation, we use the corresponding nodes throughout this analysis and write $X_i(u)$ to denote the local variable X of node u at the start of round i . We begin with two straightforward lemmas demonstrating the eventual execution of enabled actions.

► **Lemma 2.** *Apart from CLEANUP, every enabled action will eventually be executed.*

Proof. Any enabled RECEIVE* action whose guard depends only on the receipt of some message must eventually be executed because it is assumed that every message in transit is eventually processed (unless the edge is disconnected, at which point the message is lost and the action is no longer enabled). Thus, it remains to consider the CHECK* actions.

Suppose CHECKSTART is enabled for a node u in some round i ; i.e., $\text{state}_i(u) = \text{PREPARE}$ and $R_i(u) = L_i(u)$. When **state** = PREPARE, only CHECKSTART can change the **state** variable or reset R to \emptyset . Any execution of the CLEANUP action does not change the **state** variable and maintains $R(u) = L(u)$ since it removes any disconnected neighbors from both sets. So CHECKSTART remains continuously enabled and thus must eventually be executed by the weakly fair adversary. An analogous argument also applies to CHECKPRIORITIES, CHECKWIN, CHECKDONE, and CHECKUNLOCKED. ◀

Lemma 2 shows that an enabled action will eventually be executed, but we also need to know that the actions become enabled in the first place. One potential obstacle is that CHECK* actions by a node u need to receive all responses from the nodes in $L(u)$ before becoming enabled. If some of nodes in $L(u)$ disconnect and their corresponding response messages are lost, the CHECK* action may be disabled indefinitely. This is one role of the CLEANUP action: removing disconnections from the algorithm's variables so other actions

stop waiting for neighbors that no longer exist. We call such an action *pre-enabled* if it is currently disabled but would become enabled after CLEANUP is executed.

► **Lemma 3.** *Every pre-enabled action eventually becomes enabled.*

Proof. Suppose that CHECKSTART is pre-enabled for node u . Then $\text{state}(u) = \text{PREPARE}$ and u must have sent a `prepare()` message to itself in its execution of INITLOCK. So RECEIVEPREPARE is enabled for u , and by Lemma 2 it is eventually executed, updating $\text{phase}(u) = \text{PREPARE}$. This enables CLEANUP for u , and it will remain enabled until executed because only CLEANUP itself can reset phase to \perp . Thus, CLEANUP must eventually be executed by the weakly fair adversary, enabling the pre-enabled CHECKSTART.

An analogous argument also applies to CHECKPRIORITIES, CHECKWIN, and CHECKDONE. For CHECKUNLOCKED, the condition $\text{state} = \text{UNLOCK}$ in the guard of CLEANUP ensures that CHECKUNLOCKED is eventually enabled. ◀

We continue our investigation of possible deadlocks resulting from actions remaining disabled by considering concurrent competitions. An initiator node u is *competing* if and only if $\text{state}(u) = \text{COMPETE}$, i.e., if u has executed CHECKSTART but has not yet received all `win()` messages needed to execute CHECKWIN. We model dependencies between competing initiators and participants at the start of round i as a directed bipartite graph $\mathcal{D}_i = (\mathcal{I}_i \cup \mathcal{P}_i, E_i)$ where $\mathcal{I}_i = \{u : \text{state}_i(u) = \text{COMPETE}\}$ is the set of competing initiators and $\mathcal{P}_i = \{u : \exists v \in \mathcal{I}_i \text{ s.t. } u \in L_i(v)\}$ is the set of participants. We note that some nodes belong to both partitions and consider their initiator and participant versions distinct. For nodes $u \in \mathcal{I}_i$ and $v \in \mathcal{P}_i \cap L_i(u)$ for which $u = v$ or $(u, v) \in G_i$ (i.e., the edge exists in round i), the directed edge $(u, v) \in E_i$ if and only if u has not yet sent a `request-lock()` message to v in response to the latest `win()` message from v ; analogously, $(v, u) \in E_i$ if and only if v has not yet sent a `win()` message to u in response to the latest `request-lock()` message from u .

► **Lemma 4.** *For all rounds i , \mathcal{D}_i is acyclic.*

Proof. Initially, no node has yet called LOCK and thus \mathcal{D}_0 is empty and trivially acyclic. So suppose that \mathcal{D}_j remains acyclic for all rounds $0 \leq j \leq i-1$ and consider the following events that may occur in round $i-1$ to form \mathcal{D}_i .

- A node u executes CHECKSTART. Then (v, u) is added to \mathcal{D}_i for each $v \in L_{i-1}(u)$ that u sends `request-lock()` messages to. But u is a sink, so \mathcal{D}_i remains acyclic.
- A node u executes CHECKWIN. If there exists $(\cdot, \text{FALSE}) \in W_{i-1}(u)$, then (u, v) is removed from \mathcal{D}_i and (v, u) is added to \mathcal{D}_i for each $v \in L_{i-1}(u)$ that u once again sends `request-lock()` messages to. As in the first case, this makes u a sink and \mathcal{D}_i remains acyclic. Otherwise, if all $(\cdot, b) \in W_{i-1}(u)$ have $b = \text{TRUE}$, u has won its competition and sets $\text{state}_i(u) = \text{WIN}$, meaning $u \notin \mathcal{D}_i$. So \mathcal{D}_i remains acyclic in this case as well.
- A node u executes CHECKPRIORITIES. Then (u, v) is removed from \mathcal{D}_i and (v, u) is added to \mathcal{D}_i for each $v \in C_{i-1}(u)$ that u sends `win()` messages to. For \mathcal{D}_i to be acyclic, it suffices to show it does not contain any outgoing edges from u ; i.e., there are no nodes w such that $u \in L_i(w)$, w has sent u a `request-lock()` message, but u has not yet sent a `win()` response to w . Such a node w could only have sent u a `request-lock()` message if it had previously received a `ready()` message from u , which in turn could only have been sent by u if u had included w as an applicant in $A(u)$. Thus, on receipt of the first `request-lock()` message from w , u would have promoted w to a candidate in $C(u)$, which is precisely the set that u responds to when executing CHECKPRIORITIES. So u has no outgoing edges in \mathcal{D}_i , as desired.

- An edge $\{u, v\}$ is disconnected in the TVG \mathcal{G} , for $u \in \mathcal{I}_{i-1}$ and $v \in \mathcal{P}_{i-1}$. This disconnection is processed by the CLEANUP helper function, removing v from $L(u)$ and thus any (u, v) edge from \mathcal{D}_i during the next execution of CHECKWIN by u ; an analogous statement holds for edges (v, u) in the next execution of CHECKPRIORITIES by v . As the removal of an edge cannot create a cycle, \mathcal{D}_i remains acyclic.

Therefore, \mathcal{D}_i remains acyclic in all cases, as claimed. \blacktriangleleft

► **Lemma 5.** *Every competing initiator eventually receives a `win()` response from its participants; likewise, every participant eventually receives a `request-lock()` response from its competing initiator(s).*

Proof. Suppose to the contrary that there exists a competing initiator u that waits indefinitely for a `win()` response from some participant v . Then the edge $\{u, v\}$ must never be disconnected in the TVG \mathcal{G} and the directed edge (v, u) must remain indefinitely in \mathcal{D} . By Lemmas 2 and 3, v can only be prohibited from sending the requisite `win()` message if CHECKPRIORITIES remains disabled for v indefinitely. This, in turn, is only possible if v waits indefinitely for a `request-lock()` response from some competing initiator $w \neq u$. This implies that $\{v, w\}$ is never disconnected in \mathcal{G} and the directed edge (w, v) remains indefinitely in \mathcal{D} . As before, Lemmas 2 and 3 can be applied iteratively to show that each node must be waiting on another. But since the set of nodes V is finite, some node must eventually be revisited, establishing a directed cycle in \mathcal{D} and contradicting Lemma 4. \blacktriangleleft

Lemma 5 directly implies the following corollary.

► **Corollary 6.** *Every competition trial of a competing initiator eventually completes.*

To demonstrate that our algorithm satisfies lockout freedom, it remains to show that every competing initiator u eventually wins a competition trial by receiving all `win(TRUE)` responses from $L(u)$. We first address the situation in which a competition trial of u is *open*, meaning none of the nodes $v \in L(u)$ are locked during the trial.

► **Lemma 7.** *If $K = \Theta(1)$, then an initiator that competes in an open competition trial infinitely often will eventually win a competition, with probability 1.*

Proof. Consider any competing initiator u and any open competition trial of u . By the start of its second competition trial, $u \in C(v)$ for all $v \in L(u)$, implying that $\text{phase}(v) = \text{COMPETE}$ and no other nodes will be added to $C(v) \cup A(v)$ while u is still competing for v . Since $|L(u) \setminus \{u\}| \leq \Delta$ and $|C(v) \cup A(v) \setminus \{u\}| \leq \Delta$ for each $v \in L(u) \setminus \{u\}$, node u can be competing against $c \leq \Delta^2$ other nodes. Every node chooses its priority uniformly at random from $\{0, \dots, K-1\}$, so it follows from symmetry that the probability u has the highest priority in a given trial is at least $1/\Delta^2$. In general,

$$\begin{aligned} \frac{\Pr[p(u) \text{ highest} \mid p(u) \text{ unique}]}{\Pr[p(u) \text{ highest}]} &= \frac{\sum_{p=0}^{K-1} \Pr[p(u) = p \wedge \forall v \neq u : p(v) \leq p(u) \mid p(u) \text{ unique}]}{\sum_{p=0}^{K-1} \Pr[p(u) = p \wedge \forall v \neq u : p(v) \leq p(u)]} \\ &= \frac{\sum_{p=0}^{K-1} \frac{1}{K} \left(\frac{p}{K-1}\right)^c}{\sum_{p=0}^{K-1} \frac{1}{K} \left(\frac{p+1}{K}\right)^c} \geq \frac{\sum_{p=0}^{K-1} p^c}{\sum_{p=1}^K p^c} \geq \frac{(K-1)^c}{2K^c} \geq \frac{(1-1/K)^{\Delta^2}}{2} \end{aligned}$$

Furthermore, the probability that u has a unique priority is $(1 - 1/K)^c \geq (1 - 1/K)^{\Delta^2}$. Thus, the probability that u has the unique highest priority in a given open trial is

$$\begin{aligned} \Pr [p(u) \text{ highest} \wedge p(u) \text{ unique}] &= \Pr [p(u) \text{ highest} \mid p(u) \text{ unique}] \cdot \Pr [p(u) \text{ unique}] \\ &\geq \frac{(1 - 1/K)^{\Delta^2}}{2\Delta^2} \cdot (1 - 1/K)^{\Delta^2} = \frac{(1 - 1/K)^{2\Delta^2}}{2\Delta^2} > 0. \end{aligned}$$

Since this probability is strictly positive, the probability that u never has the unique highest priority in an infinite sequence of open competition trials is

$$\lim_{n \rightarrow \infty} (1 - \Pr [p(u) \text{ highest} \wedge p(u) \text{ unique}])^n \leq \lim_{n \rightarrow \infty} \left(1 - \frac{(1 - 1/K)^{2\Delta^2}}{2\Delta^2}\right)^n = 0.$$

Therefore, with probability 1 there must eventually be an open competition trial in which u has the unique highest priority. Because the trial is open, all $v \in L(u)$ have $\text{lock}(v) = \perp$ and thus will send $\text{win}(\text{TRUE})$ responses to u .² ◀

We next show that a competing initiator competes in an open trial infinitely often. Recall from Section 2.2 that a LOCK operation by node u succeeds once u obtains locks for its persistent neighborhood, and once obtained, these locks are eventually released via UNLOCK.

► **Lemma 8.** *Every competing initiator eventually wins a competition trial with probability 1.*

Proof. Suppose to the contrary that a competing initiator u competes in an infinite number of competition trials. Only a finite number of these trials can be open, since u would eventually win one of an infinite number of open trials with probability 1 by Lemma 7. So an infinite number of trials of u must be closed; i.e., there are an infinite number of trials in which at least one $v \in L(u)$ has $\text{lock}(v) \neq \perp$. Since $|L(u) \setminus \{u\}| \leq \Delta$, there must be a node $v \in L(u)$ that is locked infinitely often. But by the start of its second competition trial, $u \in C(v)$ and no other nodes will be added to $C(v) \cup A(v)$ while u is still competing for v . Thus, only the nodes in $C(v) \cup A(v)$ and the node that had already locked v when u was added to $C(v)$ could possibly lock v . But whenever v sets its locks in RECEIVESETLOCK, it removes the locking node from $C(v)$. Moreover, any node that obtains locks must eventually release them, by supposition. So the set of nodes that could lock v is monotonically decreasing and thus nodes in $C(v) \cup A(v) \setminus \{u\}$ cannot lock v an infinite number of times, a contradiction. ◀

For an initiator u to benefit from eventual victory ensured by Lemma 8, it must become competing in the first place; i.e., it must advance to $\text{state}(u) = \text{COMPETE}$.

► **Lemma 9.** *Every initiator eventually becomes competing.*

Proof. Suppose to the contrary that an initiator u never becomes competing, i.e., it never executes CHECKSTART. By Lemmas 2 and 3, this is only possible if CHECKSTART remains disabled indefinitely. To be an initiator at all, u must have executed INITLOCK, set $\text{state}(u) = \text{PREPARE}$, and sent $\text{prepare}()$ messages to all nodes $v \in L(u)$. So u must be waiting for a $\text{ready}()$ response from at least one $v \in L(u)$ that remains connected to u indefinitely.

By Lemma 2, such a node v must eventually execute RECEIVEPREPARE. During this execution, it must be the case that $\text{phase}(v) = \text{COMPETE}$ and v adds u to $H(v)$; otherwise,

² This proof can be easily extended to show that if $K > \Delta^2$, u will win a competition within $\mathcal{O}(\Delta^2)$ open competition trials, in expectation. We chose to avoid this increase in message size requirements from $\Theta(1)$ to $\mathcal{O}(\log \Delta)$ since time complexity is not a focus of this work.

v would have added u to $A(v)$ and replied to u with a `ready()` message, a contradiction. Only the `CLEANUP` helper function can reset `phase(v)` to \perp , but it only does so when $C(v) \cup A(v) \cup H(v) = \emptyset$ which is not the case since $u \in H(v)$. So the `CLEANUP` action is continuously enabled for v and is eventually executed by the weakly fair adversary. During this execution, it must be the case that $C(v) \neq \emptyset$; otherwise, v would have sent `ready()` messages to all initiators on hold at v , including u , a contradiction. But for this situation to occur indefinitely, there must exist some competitor in the finite set $C(v) \cup A(v)$ that competes in an infinite number of trials, a contradiction of Lemma 8. \blacktriangleleft

Combining Corollary 6 with Lemmas 8 and 9 implies the following corollary.

► **Corollary 10.** *The local mutual exclusion algorithm satisfies lockout freedom.*

Recall from Section 2.2 that the mutual exclusion property is trivially satisfied by our construction of the lock sets. Thus, we conclude the proof of Theorem 1 with the following result regarding the algorithm’s memory and message size requirements.

► **Lemma 11.** *The algorithm requires $\mathcal{O}(\Delta)$ memory per node and messages of size $\Theta(1)$, and there are at most two messages in transit along any given edge at any time.*

Proof. Table 1 shows that `phase` and `state` can be stored in $\Theta(1)$ bits each and `lock` can be stored in $\log_2 \Delta$ bits. The remaining variables can be represented as linear registers of length Δ , where port ℓ is in the set variable X if and only if the ℓ -th bit of register X is TRUE. So the memory bound of $\mathcal{O}(\Delta)$ follows. Similarly, there are a constant number of message types, among which only `request-lock()` and `win()` carry additional data. A `win()` message carries one bit signaling whether a competition trial was won or lost. A `request-lock()` message carries a randomly chosen priority, which by Lemma 7 can be stored in $\Theta(1)$ bits.

To bound the number of messages in transit per edge per time, consider the execution of a `LOCK` operation by an initiator node u . The local mutual exclusion algorithm is structured around pairs of initiator messages and participant responses: `prepare()`/`ready()` messages in the preparation phase, `request-lock()`/`win()` messages in each competition trial, and `set-lock()`/`ack-lock()` messages once a node has won a trial. In each scenario, only one message per pair is in transit along $\{u, v\}$ per time for each $v \in L(u)$. Moreover, node u does not advance to the next phase and send any additional messages until all messages of the current phase are processed. An analogous argument applies to the `UNLOCK` operation with its `release-lock()`/`ack-unlock()` message pairs. Thus, there can be at most one message in transit per edge per time involved with any initiator’s `LOCK` or `UNLOCK` operation.

Furthermore, an initiator u can execute at most one `LOCK` or `UNLOCK` operation per time since u can only start a `LOCK` operation by executing `INITLOCK` if `state(u) = \perp` , implying it holds no locks; similarly, u can only start an `UNLOCK` operation by executing `INITUNLOCK` if `state(u) = LOCKED`, implying its previous `LOCK` operation has succeeded.

Thus, the lemma follows since there are at most two initiators u and v per edge $\{u, v\}$. \blacktriangleleft

5 Extending to Asynchronous Concurrency

Section 4 proved Theorem 1 under semi-synchronous concurrency in which (i) topological changes occur at discrete times in between rounds of action executions and (ii) the adversary chooses any non-empty subset of nodes to act in each round and those nodes’ action executions are guaranteed to end before the next round begins. In this section, we prove that Theorem 1 holds even in the more general asynchronous setting.

All assumptions from Section 2.1 about the time-varying graph \mathcal{G} , the nodes, their asynchronous message passing, and the structure of algorithms and their actions remain the same. However, in an *asynchronous schedule*, the adversary can schedule action executions over arbitrary finite time intervals, including those that are concurrent with topological changes and span multiple TVG rounds. In this setting, our prior assumptions about the disconnection detector now imply that any topological changes incident to u that are concurrent with one of its action executions are not observed or processed by u until its next action execution. We further assume for the asynchronous setting that any message sent by node u during one of its action executions starting at time t_1 is processed by a node v during some other action execution starting at time $t_2 > t_1$ if and only if the edge $\{u, v\} \in G_t$ for all $t \in [t_1, t_2]$. This implies that when an edge is disconnected, all messages in transit along that edge are immediately lost and no further messages can be sent or received by the corresponding ports until the corresponding action executions have finished.

► **Lemma 12.** *For any asynchronous schedule \mathcal{S} , there exists a semi-synchronous schedule \mathcal{S}' containing the same action executions as in \mathcal{S} that produces the same outcome for every action execution in \mathcal{S} .*

Proof. Consider any asynchronous schedule \mathcal{S} of the local mutual exclusion algorithm and let \mathcal{E} be the set of all action executions in \mathcal{S} . Analogous to Lamport [34], we define the *causal relation* \rightarrow on \mathcal{E} as the smallest relation satisfying the following three conditions:

- If $\alpha \in \mathcal{E}$ is an execution by node u and $\beta \in \mathcal{E}$ is the next execution by u , then $\alpha \rightarrow \beta$.
- If a message sent in $\alpha \in \mathcal{E}$ is processed in $\beta \in \mathcal{E}$, then $\alpha \rightarrow \beta$.
- If $\alpha \rightarrow \beta$ and $\beta \rightarrow \gamma$, then $\alpha \rightarrow \gamma$.

Since all causal relations are naturally forward in time, it follows that the graph represented by the causal relations on \mathcal{E} forms a DAG. Thus, the action executions of \mathcal{E} can be topologically sorted in some order $[\alpha_1, \alpha_2, \dots]$.

Now, consider the schedule $\hat{\mathcal{S}}$ containing the same action executions starting at the same times as those in \mathcal{S} , but (i) each action execution takes 0 time and (ii) any set of action executions starting at the same time as some edge changes is shifted before these edge changes without changing the order of the action executions. Then $\hat{\mathcal{S}}$ can be transformed into a semi-synchronous schedule \mathcal{S}' by adding filler time steps when no edges change so that each node executes at most one action per round and all action executions between two time steps start at the same time. Certainly, \mathcal{S}' is still a valid schedule since all causal relations remain forward in time and—by our assumption on asynchronous message processing—any message sent by action execution α that is processed by action execution β in \mathcal{S} can still be processed by β in \mathcal{S}' . Furthermore, since the causal relations haven't changed, the action executions in \mathcal{S}' can be sorted in the same order $[\alpha_1, \alpha_2, \dots]$ as for \mathcal{S} . Since any action execution can only change a node's state or send messages and, in both schedules, it only sees a snapshot of D at its start, it follows by induction on the ordering of the action executions that for any i , the outcome of α_i is identical in \mathcal{S} and \mathcal{S}' . ◀

As in the semi-synchronous setting, the mutual exclusion property is trivially satisfied in the asynchronous setting. But suppose to the contrary that there exists an asynchronous schedule in which at least one LOCK operation never succeeds. Lemma 12 shows that there must exist a semi-synchronous schedule in which at least one LOCK operation never succeeds, contradicting Theorem 1. So we have the following corollary.

► **Corollary 13.** *The local mutual exclusion algorithm also satisfies the mutual exclusion and lockout freedom properties under any asynchronous schedule.*

6 Applications

We next establish how our algorithm for local mutual exclusion can be used to implement key assumptions present in formal models of dynamic distributed systems. In particular, we focus on the assumptions of independent pairwise interactions in *population protocols* [3] and concurrency control operations in the *canonical amoebot model* of programmable matter [17].

Population Protocols. Inspired by passively mobile sensor networks, Angluin et al. [3] proposed the *population protocols* model. Each of agent in a population is assumed to have a finite state and a transition function defining how that state evolves as a result of a pairwise interaction with another agent. Agents cannot explicitly control their movements or who they interact with; i.e., they are passively dynamic. Instead, it is typically assumed that a sequential scheduler chooses one pair of agents to interact per time step. In reality, however, many agents within interacting distance might exist concurrently (see, e.g., [16]), requiring a mechanism to organize these agents into a matching of independent pairs.

This goal could be achieved directly using our algorithm for local mutual exclusion. Any agent u that wants to interact must first call `LOCK`. On success, u then chooses any locked neighbor to interact with, if it has one; if desired, one could even generalize the usual pairwise interactions to interactions among the full group of locked neighbors. Lockout freedom ensures u will eventually be allowed to make this choice, and mutual exclusion ensures this pairwise interaction is isolated from any others. After interacting, u then releases its locks with `UNLOCK`. If the expected number of competing agents is high, an alternative implementation of our algorithm could have u make its choice of interacting neighbor v first and then try to lock only u and v to avoid a lengthy competition. On success, u would then interact with v , isolated from any other interactions, and then unlock itself and v . In both implementations, it is possible that all neighbors may move out of interaction range, leaving u to lock only itself. In this situation, no interaction occurs and u simply unlocks itself.

Both implementations require $\mathcal{O}(\Delta)$ memory per agent and messages of size $\Theta(1)$. For many applications of population protocols where Δ is a fixed constant (e.g., proximity graphs or IoT), these requirements are reduced to $\Theta(1)$. Thus, our algorithm for local mutual exclusion could provide isolated pairwise interactions assumed by population protocols even in the presence of underlying network dynamics and asynchronous concurrency.

The Canonical Amoebot Model. The *amoebot model* abstracts active programmable matter as a collection of simple computational elements called *amoebots* that move and interact locally to collectively achieve tasks of coordination and movement. Each amoebot is typically assumed to be anonymous and have only constant-size memory, but can control its movements. The *canonical amoebot model* [17] is an updated formalization that addresses concurrency by partitioning amoebot functionality into a high-level application layer where algorithms call various operations and a low-level system layer where those operations are executed via asynchronous message passing. Two such operations are the concurrency control operations, `LOCK` and `UNLOCK`, which are used in a *concurrency control framework* to convert amoebot algorithms that terminate in the sequential setting and satisfy certain conventions into algorithms that exhibit equivalent behavior in the concurrent setting [17].

The canonical amoebot model treats the `LOCK` and `UNLOCK` operations as black boxes without giving an implementation. These operations facilitate amoebots gaining exclusive access to themselves and their neighbors, much like our mutual exclusion property, and are assumed to terminate (either successfully or in failure) in finite time. The asynchronous

extension of our local mutual exclusion algorithm presented in Section 5 could directly implement these operations, ensuring isolation of concurrent amoebot actions even as connections between amoebots change due to their movements. One interesting feature of such an implementation is that while the amoebot LOCK operation is allowed to fail—which must be taken into account by algorithm designers—our LOCK operation always succeeds due to lockout freedom, reducing complexity in algorithm design. Moreover, for the often-considered geometric space variant in which an (expanded) amoebot can have at most eight neighbors, our algorithm has $\Theta(1)$ amoebot memory and message size requirements.

7 Conclusion

We presented an algorithm for local mutual exclusion that enables weakly capable nodes to isolate concurrent actions involving their persistent neighborhoods despite dynamic network topology. Our algorithm ensures that nodes belong to at most one locked neighborhood at a time (mutual exclusion) and that every lock request eventually succeeds (lockout freedom). It requires $\mathcal{O}(\Delta)$ memory per node and messages of size $\Theta(1)$ —where Δ is the maximum number of connections per node—and is compatible with anonymous, message passing nodes that operate semi-synchronously or asynchronously. These weak requirements make our algorithm suitable for a wide range of application domains such as overlay networks, IoT, modular robots, and programmable matter. As two concrete examples, we demonstrated how our algorithm could implement the pairwise interactions assumed by population protocols [3] and the concurrency control operations assumed by the canonical amoebot model [17].

References

- 1 Karine Altisen, Stéphane Devismes, Anaïs Durand, Colette Johnen, and Franck Petit. Self-stabilizing Systems in Spite of High Dynamics. In *International Conference on Distributed Computing and Networking 2021*, pages 156–165, 2021. doi:[10.1145/3427796.3427838](https://doi.org/10.1145/3427796.3427838).
- 2 Marta Andrés Arroyo, Sarah Cannon, Joshua J. Daymude, Dana Randall, and Andréa W. Richa. A stochastic approach to shortcut bridging in programmable matter. *Natural Computing*, 17(4):723–741, 2018. doi:[10.1007/s11047-018-9714-x](https://doi.org/10.1007/s11047-018-9714-x).
- 3 Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, 2006. doi:[10.1007/s00446-005-0138-3](https://doi.org/10.1007/s00446-005-0138-3).
- 4 Hagit Attiya, Alex Kogan, and Jennifer L. Welch. Efficient and Robust Local Mutual Exclusion in Mobile Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 9(3):361–375, 2010. doi:[10.1109/TMC.2009.137](https://doi.org/10.1109/TMC.2009.137).
- 5 Baruch Awerbuch, Michael Luby, Andrew V. Goldberg, and Serge A. Plotkin. Network decomposition and locality in distributed computation. In *30th Annual Symposium on Foundations of Computer Science*, pages 364–369, 1989. doi:[10.1109/SFCS.1989.63504](https://doi.org/10.1109/SFCS.1989.63504).
- 6 Roberto Baldoni, Antonino Virgillito, and Roberto Petrassi. A distributed mutual exclusion algorithm for mobile ad-hoc networks. In *Proceedings ISCC 2002 Seventh International Symposium on Computers and Communications*, pages 539–544, 2002. doi:[10.1109/ISCC.2002.1021727](https://doi.org/10.1109/ISCC.2002.1021727).
- 7 Mahfoud Benchaïba, Abdelmadjid Bouabdallah, Nadjib Badache, and Mohamed Ahmed-Nacer. Distributed Mutual Exclusion Algorithms in Mobile Ad Hoc Networks: An Overview. *ACM SIGOPS Operating Systems Review*, 38(1):74–89, 2004. doi:[10.1145/974104.974111](https://doi.org/10.1145/974104.974111).
- 8 Michael A. Bender, Martin Farach-Colton, Simai He, Bradley C. Kuszmaul, and Charles E. Leiserson. Adversarial Contention Resolution for Simple Channels. In *Proceedings of the Seventeenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 325–332, 2005. doi:[10.1145/1073970.1074023](https://doi.org/10.1145/1073970.1074023).

- 9 Federico Cali, Marco Conti, and Enrico Gregori. IEEE 802.11 Protocol: Design and Performance Evaluation of an Adaptive Backoff Mechanism. *IEEE Journal on Selected Areas in Communications*, 18(9):1774–1786, 2000. doi:10.1109/49.872963.
- 10 John I. Capetanakis. Tree Algorithms for Packet Broadcast Channels. *IEEE Transactions on Information Theory*, 25(5):505–515, 1979. doi:10.1109/TIT.1979.1056093.
- 11 Arnaud Casteigts. A Journey through Dynamic Networks (with Excursions), 2018. HDR, available online at <https://hal.archives-ouvertes.fr/tel-01883384/>.
- 12 Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-Varying Graphs and Dynamic Networks. *International Journal of Parallel, Emergent and Distributed Systems*, 27(5):387–408, 2012. doi:10.1080/17445760.2012.668546.
- 13 Cameron Chalk, Austin Luchsinger, Eric Martinez, Robert Schweller, Andrew Winslow, and Tim Wylie. Freezing Simulates Non-freezing Tile Automata. In *DNA Computing and Molecular Programming*, volume 11145 of *Lecture Notes in Computer Science*, pages 155–172, 2018. doi:10.1007/978-3-030-00030-1_10.
- 14 Arjun Chandrasekhar, Deborah M. Gordon, and Saket Navlakha. A distributed algorithm to maintain and repair the trail networks of arboreal ants. *Scientific Reports*, 8(1):9297, 2018. doi:10.1038/s41598-018-27160-3.
- 15 Yu Chen and Jennifer L. Welch. Self-stabilizing dynamic mutual exclusion for mobile ad hoc networks. *Journal of Parallel and Distributed Computing*, 65(9):1072–1089, 2005. doi:10.1016/j.jpdc.2005.03.009.
- 16 Artur Czumaj and Andrzej Lingas. On Truly Parallel Time in Population Protocols. Available online at <https://arxiv.org/abs/2108.11613>, 2021.
- 17 Joshua J. Daymude, Andréa W. Richa, and Christian Scheideler. The Canonical Amoebot Model: Algorithms and Concurrency Control. In *35th International Symposium on Distributed Computing (DISC 2021)*, volume 209 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:19, 2021. doi:10.4230/LIPIcs.DISC.2021.20.
- 18 Michael J. Demmer and Maurice P. Herlihy. The arrow distributed directory protocol. In Shay Kutten, editor, *Distributed Computing*, volume 1499 of *Lecture Notes in Computer Science*, pages 119–133, 1998. doi:10.1007/BFb0056478.
- 19 Zahra Derakhshandeh, Shlomi Dolev, Robert Gmyr, Andréa W. Richa, Christian Scheideler, and Thim Strothmann. Amoebot - a new model for programmable matter. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 220–222, 2014. doi:10.1145/2612669.2612712.
- 20 E. W. Dijkstra. Solution of a problem in concurrent programming control. *Communications of the ACM*, 8(9):569, 1965. doi:10.1145/365559.365617.
- 21 Yuval Emek and Roger Wattenhofer. Stone age distributed computing. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, pages 137–146, Montréal, Québec, Canada, 2013. ACM. doi:10.1145/2484239.2484244.
- 22 Javier Esparza and Fabian Reiter. A Classification of Weak Asynchronous Models of Distributed Computing. In *31st International Conference on Concurrency Theory (CONCUR 2020)*, volume 171 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.CONCUR.2020.10.
- 23 Michael Feldmann, Christian Scheideler, and Stefan Schmid. Survey on Algorithms for Self-stabilizing Overlay Networks. *ACM Computing Surveys*, 53(4):74:1–74:24, 2020. doi:10.1145/3397190.
- 24 Michael J. Fischer, Nancy A. Lynch, James E. Burns, and Allan Borodin. Resource allocation with immunity to limited process failure. In *20th Annual Symposium on Foundations of Computer Science (SFCS 1979)*, pages 234–254, 1979. doi:10.1109/SFCS.1979.37.
- 25 Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro, editors. *Distributed Computing by Mobile Entities: Current Research in Moving and Computing*, volume 11340 of *Lecture Notes in Computer Science*. Springer International Publishing, Cham, 2019. doi:10.1007/978-3-030-11072-7.

26 Mohsen Ghaffari, Cameron Musco, Tsvetomira Radeva, and Nancy Lynch. Distributed House-Hunting in Ant Colonies. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, pages 57–66, 2015. doi:10.1145/2767386.2767426.

27 Abdolhamid Ghodselahi and Fabian Kuhn. Dynamic Analysis of the Arrow Distributed Directory Protocol in General Networks. In *31st International Symposium on Distributed Computing (DISC 2017)*, volume 91 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:16, 2017. doi:10.4230/LIPICS.DISC.2017.22.

28 Heiko Hamann. *Swarm Robotics: A Formal Approach*. Springer International Publishing, Cham, 2018. doi:10.1007/978-3-319-74528-2.

29 Maurice Herlihy. Wait-free synchronization. *ACM Transactions on Programming Languages and Systems*, 13(1):124–149, 1991. doi:10.1145/114005.102808.

30 Yuh-Jzer Joung. Asynchronous group mutual exclusion. *Distributed Computing*, 13(4):189–206, 2000. doi:10.1007/PL00008918.

31 Yuh-Jzer Joung. The congenial talking philosophers problem in computer networks. *Distributed Computing*, 15(3):155–175, 2002. doi:10.1007/s004460100069.

32 Pankaj Khanchandani and Roger Wattenhofer. The Arvy Distributed Directory Protocol. In *The 31st ACM Symposium on Parallelism in Algorithms and Architectures*, pages 225–235, 2019. doi:10.1145/3323165.3323181.

33 Fabian Kuhn, Yannic Maus, and Simon Weidner. Deterministic Distributed Ruling Sets of Line Graphs. In *Structural Information and Communication Complexity*, volume 11085 of *Lecture Notes in Computer Science*, pages 193–208, 2018. doi:10.1007/978-3-030-01325-7_19.

34 Leslie Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. *Communications of the ACM*, 21(7):558–565, 1978. doi:10.1145/359545.359563.

35 Mamoru Maekawa. A \sqrt{N} algorithm for mutual exclusion in decentralized systems. *ACM Transactions on Computer Systems*, 3(2):145–159, 1985. doi:10.1145/214438.214445.

36 Othon Michail, George Skretas, and Paul G. Spirakis. Distributed Computation and Re-configuration in Actively Dynamic Networks. *Distributed Computing*, 2021. doi:10.1007/s00446-021-00415-5.

37 Othon Michail and Paul G. Spirakis. Simple and efficient local codes for distributed stable network construction. *Distributed Computing*, 29(3):207–237, 2016. doi:10.1007/s00446-015-0257-4.

38 Othon Michail and Paul G. Spirakis. Connectivity preserving network transformers. *Theoretical Computer Science*, 671:36–55, 2017. doi:10.1016/j.tcs.2016.02.040.

39 Matthew J. Patitz. An introduction to tile-based self-assembly and a survey of recent results. *Natural Computing*, 13(2):195–224, 2014. doi:10.1007/s11047-013-9379-4.

40 Benoit Piranda and Julien Bourgeois. Designing a quasi-spherical module for a huge modular robot to create programmable matter. *Autonomous Robots*, 42:1619–1633, 2018. doi:10.1007/s10514-018-9710-0.

41 Kerry Raymond. A tree-based algorithm for distributed mutual exclusion. *ACM Transactions on Computer Systems*, 7(1):61–77, 1989. doi:10.1145/58564.59295.

42 Michel Raynal and Gadi Taubenfeld. Mutual exclusion in fully anonymous shared memory systems. *Information Processing Letters*, 158:105938, 2020. doi:10.1016/j.ipl.2020.105938.

43 Glenn Ricart and Ashok K. Agrawala. An optimal algorithm for mutual exclusion in computer networks. *Communications of the ACM*, 24(1):9–17, 1981. doi:10.1145/358527.358537.

44 Johannes Schneider, Michael Elkin, and Roger Wattenhofer. Symmetry breaking depending on the chromatic number or the neighborhood growth. *Theoretical Computer Science*, 509:40–50, 2013. doi:10.1016/j.tcs.2012.09.004.

45 Bharti Sharma, Ravinder Singh Bhatia, and Awadhesh Kumar Singh. A Token Based Protocol for Mutual Exclusion in Mobile Ad Hoc Networks. *Journal of Information Processing Systems*, 10(1):36–54, 2014. doi:10.3745/JIPS.2014.10.1.036.

- 46 Harisu Abdullahi Shehu, Md. Haidar Sharif, and Rabie A. Ramadan. Distributed Mutual Exclusion Algorithms for Intersection Traffic Problems. *IEEE Access*, 8:138277–138296, 2020. doi:10.1109/ACCESS.2020.3012573.
- 47 Mukesh Singhal. A Taxonomy of Distributed Mutual Exclusion. *Journal of Parallel and Distributed Computing*, 18(1):94–101, 1993. doi:10.1006/jpdc.1993.1048.
- 48 David Soloveichik, Matthew Cook, Erik Winfree, and Jehoshua Bruck. Computation with finite stochastic chemical reaction networks. *Natural Computing*, 7(4):615–633, 2008. doi:10.1007/s11047-008-9067-y.
- 49 Lili Su, Chia-Jung Chang, and Nancy Lynch. Spike-Based Winner-Take-All Computation: Fundamental Limits and Order-Optimal Circuits. *Neural Computation*, 31(12):2523–2561, 2019. doi:10.1162/neco_a_01242.
- 50 Jennifer E. Walter, Jennifer L. Welch, and Nitin H. Vaidya. A Mutual Exclusion Algorithm for Ad Hoc Mobile Networks. *Wireless Networks*, 7:585–600, 2001. doi:10.1023/A:1012363200403.
- 51 Damien Woods, Ho-Lin Chen, Scott Goodfriend, Nadine Dabby, Erik Winfree, and Peng Yin. Active self-assembly of algorithmic shapes and patterns in polylogarithmic time. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, pages 353–354, 2013. doi:10.1145/2422436.2422476.
- 52 Weigang Wu, Jiebin Zhang, Aoxue Luo, and Jiannong Cao. Distributed Mutual Exclusion Algorithms for Intersection Traffic Control. *IEEE Transactions on Parallel and Distributed Systems*, 26(1):65–74, 2015. doi:10.1109/TPDS.2013.2297097.
- 53 Mark Yim, Wei-Min Shen, Behnam Salemi, Daniela Rus, Mark Moll, Hod Lipson, Eric Klavins, and Gregory S. Chirikjian. Modular Self-Reconfigurable Robot Systems [Grand Challenges of Robotics]. *IEEE Robotics & Automation Magazine*, 14(1):43–52, 2007. doi:10.1109/MRA.2007.339623.