



HAL
open science

Self-stabilizing K -out-of- L exclusion on tree network

Ajoy Datta, Stéphane Devismes, Florian Horn, Lawrence Larmore

► **To cite this version:**

Ajoy Datta, Stéphane Devismes, Florian Horn, Lawrence Larmore. Self-stabilizing K -out-of- L exclusion on tree network. 2008. hal-00344193v2

HAL Id: hal-00344193

<https://hal.science/hal-00344193v2>

Preprint submitted on 5 Dec 2008 (v2), last revised 12 Feb 2009 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Self-Stabilizing k -out-of- ℓ Exclusion on Tree Networks

Ajoy K. Datta¹ Stéphane Devismes² Florian Horn³ Lawrence L. Larmore¹

¹School of Computer Science, University of Nevada Las Vegas, {datta,larmore}@cs.unlv.edu

²VERIMAG, Université Grenoble 1 Joseph Fourier, stephane.devismes@imag.fr

³LIAFA, Université Paris 7 Denis Diderot, horn@liafa.jussieu.fr

Abstract

In this paper, we address the problem of k -out-of- ℓ exclusion, a generalization of the mutual exclusion problem, in which there are ℓ units of a shared resource, and any process can request up to k units ($1 \leq k \leq \ell$). We propose the first deterministic self-stabilizing distributed k -out-of- ℓ exclusion protocol in message-passing systems for asynchronous oriented tree networks which assumes bounded local memory for each process.

Keywords: Fault-tolerance, self-stabilization, resource allocation, k -out-of- ℓ exclusion, oriented tree networks.

1 Introduction

The basic problem in resource allocation is the management of shared resources, such as printers or shared variables. The use of such resources by an agent affects their availability for the other users. In the aforementioned cases (*e.g.*, printers), at most one agent can access to the resource at any time, using a special section of code called a *critical section*. The associated protocols must guarantee the *mutual exclusion* property [13]: the critical section can be executed by at most one process at any time. The ℓ -exclusion property [6] is a generalization of mutual exclusion, where ℓ processes can execute the critical section simultaneously. Thus, in ℓ -exclusion, ℓ units of a same resource (*e.g.*, a pool of IP addresses) can be allocated. This problem can be generalized still further by considering heterogeneous requests, *e.g.*, bandwidth for audio or video streaming. The k -out-of- ℓ exclusion property [12] allows us to deal with such requests; requests may vary from 1 to k units of a given resource, where $1 \leq k \leq \ell$.

Contributions. In this paper, we propose a (deterministic) *self-stabilizing* distributed k -out-of- ℓ exclusion protocol for asynchronous oriented tree networks. A distributed protocol is self-stabilizing [5] if, after transient faults hit the system and place it in some arbitrary global state, the system recovers from this catastrophic situation without external (*e.g.* human) intervention in finite time. Our protocol is written in the message-passing model, and assumes bounded memory per process. To the best of our knowledge, there is no prior protocol of this type in the literature.

Obtaining a self-stabilizing solution for the k -out-of- ℓ exclusion problem in oriented trees is desirable, but also complex. Our main reason for dealing with oriented trees is that extension to general rooted networks is trivial; it consists of running the protocol concurrently with a spanning tree construction (for message passing systems), such as given in [1, 4]. In the other hand, the complexity of the solution comes from the fact that the problem is a generalization of mutual exclusion. This is exacerbated by the difficulty of obtaining self-stabilizing solutions in message-passing system (the more realistic model), as underlined by the impossibility result of Gouda and Multari [7].

Designing protocols for such problems on realistic systems often leads to obfuscated solutions. A direct consequence is then the difficulty of checking, or analyzing the solution. To circumvent this problem, we propose here a step-by-step approach. We start from a “naive” non-operating circulation of ℓ resource tokens. Incrementally, we then augment this solution with several other types of tokens until we obtain a correct non fault-tolerant solution. We then introduce an additional control mechanism that guarantees self-stabilization,

assuming unbounded local memory. Finally, we modify the protocol to accommodate bounded local memory. We validate our approach by showing correctness and analyzing waiting time, a crucial parameter in resource allocation.

Related Work. Two kinds of protocols are widely used in the literature to solve the k -out-of- ℓ exclusion problem: permission-based protocols, and ℓ token circulation. All non self-stabilizing solutions currently in the literature are permission-based. In a permission-based protocol, any process can access a resource after receiving permissions from all processes [12], or from the processes constituting its quorum [10, 11]. There exist two self-stabilizing solutions for k -out-of- ℓ exclusion on the oriented rooted ring [2, 3]. These solutions are based on circulation of ℓ tokens, where each token corresponds to a resource unit.

Outline. The remainder of the paper is organized as follows: in the next section, we define the model used in this paper. In Section 3, we present our self-stabilizing k -out-of- ℓ exclusion protocol. In Section 4, we provide the proof of correctness of our protocol, and we analyze its waiting time. Finally, we conclude in Section 5.

2 Preliminaries

Distributed Systems. We consider *asynchronous distributed systems* having a *finite* number of *processes*. By asynchronous, we mean that there is no bound on message delay, clock drift, or process execution rate. Every process can directly communicate with a subset of processes called *neighbors*. We denote by Δ_p the number of neighbors of a process p . We consider the message-passing model where communications between neighboring processes are carried out by *messages* exchanging through *bidirectional links* (*i.e.*, each link can be seen as two channels in the opposite directions). The neighbor relation defines a *network*. We assume that the topology of the network is that of an *oriented tree*. *Oriented* means that there is a distinguished process called *root* (denoted r in the following) and that every non-root process knows which neighbor is its *parent* in the tree, *i.e.*, the neighbor that is nearest to the root. We say that process q is a *child* of process p if and only if p is the parent of q .

A process is a sequential deterministic machine that uses a bounded local memory, a local algorithm, and input/output capabilities. Intuitively, a process executes its local algorithm by taking steps. In a step, a process executes two actions in sequence: (1) either it tries to receive a message from another process, sends a message to another process, or does nothing; and then (2) modifies some variables. Our local algorithms are structured as infinite loop that contains a finite number of actions. We denote by x_p the variable x in the code of process p .

We assume that the channels incident to a process are locally distinguished by a *number*. For sake of simplicity, every process p numbers its channels from 0 to $\Delta_p - 1$, and we simply denote by the *label* q the channel number of the process q in the code any other process. Without loss of generality, we also assume that any non-root process locally designates by 0 its channel to its parent. We assume that the channels are *reliable* and *FIFO*. *Reliable* means that no message can be lost (after the end of the transient faults). *FIFO* means that the messages are received in the order they are sent. We also assume that the channels initially contain a bounded number of arbitrary messages, we denote such a bound by C_{MAX} .¹

The messages are of the following form: $\langle \text{type}, \text{value} \rangle$. The *value* field is omitted if the message does not carry any value. The messages can contain more than one *value*.

A *distributed protocol* is a collection of n local algorithms, one for each process. We define the *state* of each process to be the state of its local memory and the contents of each of its incoming channels. The global state of the system, referred to as a *configuration*, is defined as the product of the states of processes. We denote by \mathcal{C} the set of all possible configuration. An *execution* of a distributed protocol \mathcal{P} in a system \mathcal{S} is an infinite sequence of configurations (of \mathcal{S}) $\gamma_0 \gamma_1 \dots \gamma_i \dots$ such that in any transition $\gamma_i \mapsto \gamma_{i+1}$ either a process take a step, or an external (*w.r.t.* the distributed protocol) application modifies an input variable. Any execution are assumed to be *asynchronous* but *fair*: Every process takes an infinite number of steps in the execution but the time between two steps of a process is unbounded.

¹This assumption is required to obtain a deterministic self-stabilizing solution working with bounded process memories, see [7].

k-out-of- ℓ exclusion. In the k-out-of- ℓ exclusion problem, the existence of ℓ units of a shared resource is assumed. Any process can request at most k units of the shared resource (with $k \leq \ell$). We say that a distributed protocol satisfies the k-out-of- ℓ exclusion specification if it satisfies the following three properties:

- **Safety:** At any given time, each resource unit (*n.b.*, here any resource unit corresponds to a token) is used by at most one process, each process uses at most k resource units, and at most ℓ resource units are used.
- **Fairness:** If a process requests at most k resource units, then its request is eventually satisfied (*i.e.* it can eventually use the resource unit it requests using a special section of code called *critical section*).
- **Efficiency:** Informally, this means that as many requests as possible must be satisfied simultaneously.

The above mentioned notion of *efficiency* is difficult to define precisely. A convenient parameter was introduced in [3] to formally characterize efficiency: (k, ℓ) -*liveness*. In this definition, it is assumed that there is a subset I of processes such that every process in I executes its critical section forever (*i.e.* they holds some resource units forever). Let u_i the number of resource units held forever by each process $p \in I$. Let $\alpha = \sum_{p \in I} u_i$. Let R be the set of processes not in I that request some resource units. Let r_p be the number of resource units requested by $p \in R$. To satisfy the (k, ℓ) -*liveness*, a distributed k-out-of- ℓ exclusion protocol must guarantee that if $(\alpha < \ell) \wedge (|R| > 0) \wedge (\forall p \in R, r_p \leq \ell - \alpha)$, then at least one process in R eventually satisfies its request.

Waiting Time. The *waiting time* [14] is the maximal number of times that all processes can enter in the critical section before some process p , starting from the moment p requests the critical section.

Interface. In any k-out-of- ℓ exclusion protocol, the process needs to interact with the application that requests the resource units. To manage these interactions, we use the following interface at each process:

- **State** $\in \{\text{Req}, \text{In}, \text{Out}\}$. This variable is used to know the status of the application. **State** = *Req* means that the application is requesting some resource units. Our protocol switches **State** from *Req* to *In* when the application is allowed to access to the requested resource units. Finally, **State** switches from *In* to *Out* when the requested resource units are released into the system. Of course, the switching of **State** from *Req* to *In* and from *In* to *Out* is managed by the k-out-of- ℓ exclusion protocol itself while the switching from *Out* to *In* is managed by the application. Note that all other transitions (for instance, *In* to *Req*) are forbidden.
- **Need** $\in [0 \dots k]$. In this input variable, the application is assumed to assign the number of resource units it requests.
- **EnterCS()**: *function*. This function is called by the distributed protocol to execute the *critical section*. From this call, the application has control of the resource units until the application releases them (we assume that the critical section is always executed in finite, yet unbounded, time).
- **ReleaseCS()**: *Boolean*. This predicate is defined over the variables of the application. **ReleaseCS()** is true if and only if the application is not executing its critical section.

Self-Stabilization [5]. In the following, we define a *specification* to be a predicate over the set of all executions. A set of configurations $\mathcal{C}_1 \subseteq \mathcal{C}$ is an *attractor* for a set of configurations $\mathcal{C}_2 \subseteq \mathcal{C}$ if for any $\gamma \in \mathcal{C}_2$ and any execution whose initial configuration is γ , the execution contains a configuration of \mathcal{C}_1 .

Definition 1 An protocol \mathcal{P} is self-stabilizing for the specification SP in a system \mathcal{S} if there exists a non-empty subset of configurations \mathcal{L} such that:

- Any execution of \mathcal{P} in \mathcal{S} starting from a configuration of \mathcal{L} satisfies SP (Closure Property).
- \mathcal{L} is an attractor for \mathcal{C} (Convergence Property).

Algorithm 1 k-out-of- ℓ exclusion on oriented trees, code for the root r

```

1: variables:
2:    $C, myC \in [0 \dots 2(n-1)(C_{MAX} + 1)]; Succ \in [0 \dots \Delta_r - 1]$ 
3:    $RSet$ : multiset of at most  $k$  values taken in  $[0 \dots \Delta_r - 1]$ 
4:    $Need \in [0 \dots k]; State \in \{Req, In, Out\}$ 
5:    $Prio \in \{\perp, 0, \dots, \Delta_r - 1\}$ 
6:    $R, Reset$ : Booleans;  $SToken, PT \in [0 \dots \ell + 1]$ 
7:    $SPush, SPrio, PPr \in [0 \dots 2]$ 
8: repeat forever
9:   for all  $q \in [0 \dots \Delta_r - 1]$  do
10:    if  $(receive(ResT) \text{ from } q) \wedge \neg Reset$  then
11:      if  $(State = Req) \wedge (|RSet| < Need)$  then
12:         $RSet \leftarrow RSet \cup \{q\}$ 
13:      else
14:        if  $q = \Delta_r - 1$  then
15:           $SToken \leftarrow \min(SToken + 1, \ell + 1)$ 
16:        end if
17:         $send(ResT) \text{ to } q + 1$ 
18:      end if
19:    end if
20:    if  $(receive(PushT) \text{ from } q) \wedge \neg Reset$  then
21:      if  $(Prio \neq \perp) \wedge (State \neq Req \vee |RSet| < Need) \wedge$ 
22:         $(State \neq In)$  then
23:        for all  $i \in RSet$  do
24:          if  $i = \Delta_r - 1$  then
25:             $SToken \leftarrow \min(SToken + 1, \ell + 1)$ 
26:          end if
27:           $send(ResT) \text{ to } i + 1$ 
28:        end for
29:         $RSet \leftarrow \emptyset$ 
30:      end if
31:      if  $q = \Delta_r - 1$  then
32:         $SPush \leftarrow \min(SPush + 1, 2)$ 
33:      end if
34:       $send(PushT) \text{ to } q + 1$ 
35:    end if
36:    if  $(receive(PrioT) \text{ from } q) \wedge \neg Reset$  then
37:      if  $Prio = \perp$  then
38:         $Prio \leftarrow q$ 
39:      else
40:         $send(PrioT) \text{ to } q + 1$ 
41:      end if
42:    end if
43:    if  $(receive(ctrl, C, R, PT, PPr) \text{ from } q)$  then
44:      if  $(q = Succ) \wedge (myC = C)$  then
45:         $Succ \leftarrow Succ + 1$ 
46:        if  $Succ = 0$  then
47:           $myC \leftarrow myC + 1$ 
48:           $Reset \leftarrow (PT + SToken > \ell) \vee$ 
49:             $(PPr + SPrio > 1) \vee (SPush > 1)$ 
50:          if  $Reset$  then
51:             $RSet \leftarrow \emptyset$ 
52:             $Prio \leftarrow \perp$ 
53:          end if
54:        end if
55:        if  $PPr + SPrio < 1$  then
56:           $send(PrioT) \text{ to } 0$ 
57:        end if
58:        while  $PT + SToken < \ell$  do
59:           $send(ResT) \text{ to } 0$ 
60:           $SToken \leftarrow \min(SToken + 1, \ell + 1)$ 
61:        end while
62:        if  $SPush < 1$  then
63:           $send(PushT) \text{ to } 0$ 
64:        end if
65:        end if
66:         $PT \leftarrow \min(PT + |RSet|_q, \ell + 1)$ 
67:        if  $Prio = q$  then
68:           $PPr \leftarrow \min(PPr + 1, 2)$ 
69:        end if
70:         $send(ctrl, myC, Reset, PT, PPr) \text{ to } Succ$ 
71:         $RestartTimer()$ 
72:      end if
73:    end if
74:  end for
75:  if  $(State = Req) \wedge (|RSet| \geq Need)$  then
76:     $State \leftarrow In$ 
77:     $EnterCS()$ 
78:  end if
79:  if  $(State = In) \wedge ReleaseCS()$  then
80:    for all  $i \in RSet$  do
81:      if  $i = \Delta_r - 1$  then
82:         $SToken \leftarrow \min(SToken + 1, \ell + 1)$ 
83:      end if
84:       $send(ResT) \text{ to } i + 1$ 
85:    end for
86:     $RSet \leftarrow \emptyset$ 
87:     $State \leftarrow Out$ 
88:  end if
89:  if  $(Prio \neq \perp) \wedge (State \neq Req \vee |RSet| \geq Need)$  then
90:    if  $Prio = \Delta_r - 1$  then
91:       $SPrio \leftarrow \min(SPrio + 1, 2)$ 
92:    end if
93:     $send(PrioT) \text{ to } Prio + 1$ 
94:     $Prio \leftarrow \perp$ 
95:  end if
96:  if  $TimeOut()$  then
97:     $send(ctrl, myC, Reset, 0, 0) \text{ to } Succ$ 
98:     $RestartTimer()$ 
99:  end if
100: end repeat

```

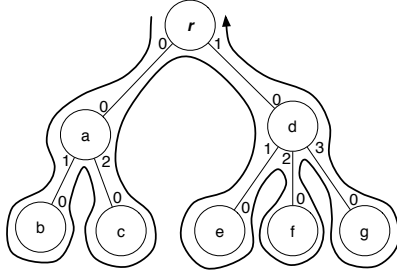


Figure 1: Depth-first token circulation on oriented trees.

3 Protocol

In this section we present our self-stabilizing k -out-of- ℓ exclusion protocol for oriented trees (Algorithms 1 and 2). Our solution uses circulation of several types of tokens. To clearly understand the goal of these tokens, we adopt a step-by-step approach: We start from a “naive” non-operating circulation of ℓ resource tokens. Incrementally, we then augment this solution with several other types of tokens until we obtain a correct non fault-tolerant solution. We then define an additional control mechanism that guarantees self-stabilization assuming unbounded local memory of processes. Finally, we modify our protocol to work with bounded memory.

A non fault-tolerant protocol. The basic principle of our protocol is to use ℓ circulating *resource tokens* (the `ResT` messages) following depth-first search (DFS) order: when a process p receives a token from channel number i , it keeps the token if necessary and then retransmits it in the channel number $i + 1$ (modulo Δ_p). Figure 1 shows a possible path followed by depth-first token circulation into an oriented tree (remember that any non-root process locally numbers the channel to its parent by 0). As explained Section 2, the requests are managed by the variables `State` and `Need`. Each process also uses the multiset² variable `RSet` to collect the tokens, any collected token is said “reserved”. While `State = Req` and $|\text{RSet}| < \text{Need}$, a process collects all the tokens it receives: it stores in `RSet` the channel numbers from which it receives the tokens. When `State = Req` and $|\text{RSet}| \geq \text{Need}$, it enters into the critical section: `State` is set to `In` and the function `EnterCS()` is called. Once the critical section is done (*i.e.*, when `State = In` and the predicate `ReleaseCS()` holds) `State` is set to `Out`, and every token in `RSet` is retransmitted following DFS order³, and `RSet` is set to \emptyset . Note also that when a process receives a token it does not need, it immediately retransmits it to its successor, in DFS order.

Unfortunately, such a simple protocol does not always guarantee liveness. Figure 2 show a case where liveness is not maintained. In this example, there are five resources tokens (*i.e.*, $\ell = 5$) and each process can request up to three tokens (*i.e.*, $k = 3$). In the configuration shown on the left side of the figure, processes a , b , c , and d request more tokens than they will receive. This configuration will lead to the deadlock configuration shown on the right side of the figure: processes a , b , c , and d reserve all the tokens they receive and never release them because their requests are never satisfied.

We correct this deadlock by adding a particular token called the *pusher* (the message `PushT`). It permanently circulates in the tree (following DFS order) and prevents a process that is not in the critical section from holding resource tokens forever: when a process receives the pusher, it releases all its reserved tokens, unless if it is either in its critical section (`State = In`) or it has the guarantee to enter it soon (`State = Req` and $|\text{RSet}| \geq \text{Need}$). In either case, it retransmits the pusher to its successor in the DFS order.

The pusher protects the system from deadlock. However, it can cause *livelock*; an example is shown in Figure 3, for 2-out-of-3 exclusion in a tree of three processes. In configuration (i), every process is a requester: r and b request one resource token and a requests two resource tokens. Also, every process has a resource token in one of its incoming channels, and holds no token. Finally, the pusher is in the channel from

²*N.b.* a multiset can contain several identical items.

³Thanks to the channel numbers stored in `RSet`, the process knows in which channel it must send each token to follow DFS order.

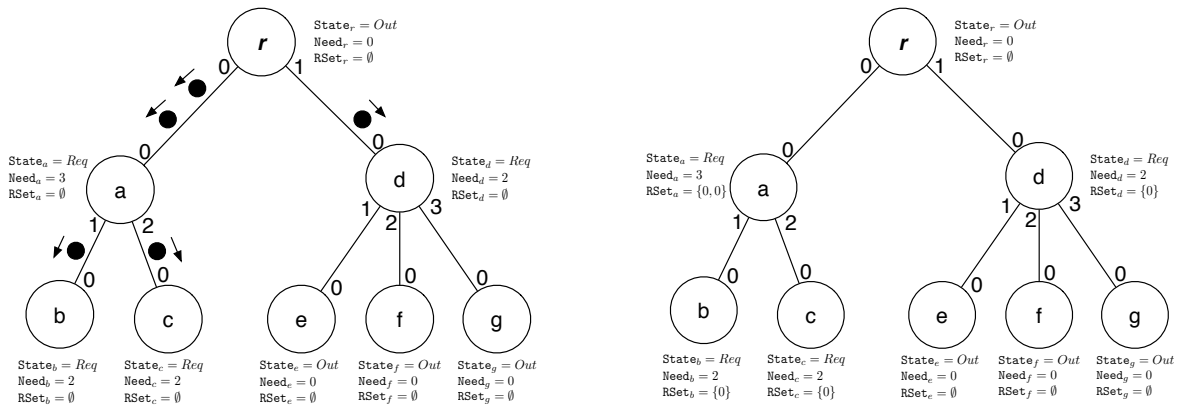


Figure 2: Possible deadlock.

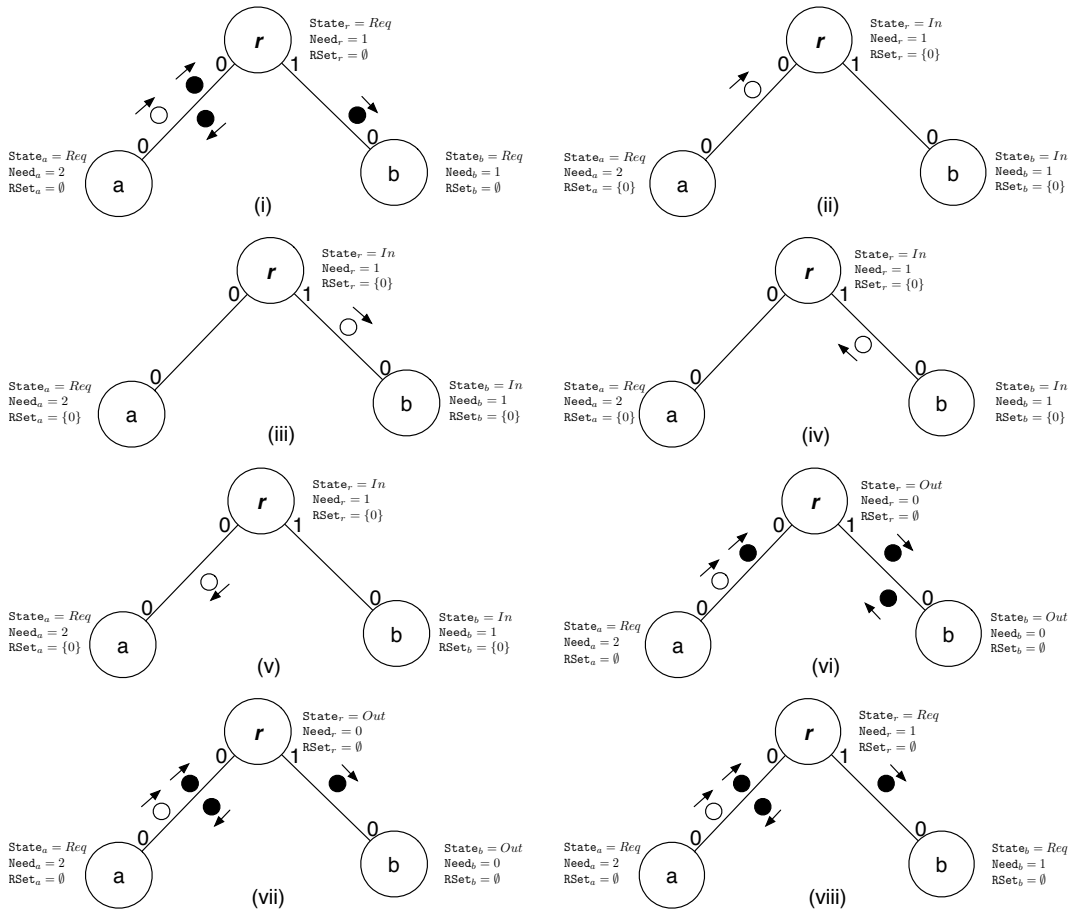


Figure 3: Possible livelock.

Algorithm 2 k-out-of- ℓ exclusion on oriented trees, code for the other process p

```

1: variables:
2:    $C, myC \in [0 \dots 2(n-1)(C_{MAX} + 1)]$ ;  $Succ \in [0 \dots \Delta_p - 1]$ 
3:   RSet: multiset of at most  $k$  values taken in  $[0 \dots \Delta_p - 1]$ 
4:   Need  $\in [0 \dots k]$ ; State  $\in \{Req, In, Out\}$ 
5:   Prio  $\in \{\perp, 0, \dots, \Delta_p - 1\}$ 
6:    $R, Ok$ : Booleans;  $PT \in [0 \dots \ell + 1]$ ;  $PPr \in [0 \dots 2]$ 
7: repeat forever
8:   for all  $q \in [0 \dots \Delta_p - 1]$  do
9:     if (receive(ResT) from  $q$ ) then
10:      if (State = Req)  $\wedge$  ( $|RSet| < Need$ ) then
11:        RSet  $\leftarrow$  RSet  $\cup \{q\}$ 
12:      else
13:        send(ResT) to  $q + 1$ 
14:      end if
15:    end if
16:    if (receive(PushT) from  $q$ ) then
17:      if (Prio  $\neq \perp$ )  $\wedge$  (State  $\neq Req \vee |RSet| < Need$ )  $\wedge$ 
18:        (State  $\neq In$ ) then
19:        for all  $i \in RSet$  do
20:          send(ResT) to  $i + 1$ 
21:        end for
22:        RSet  $\leftarrow \emptyset$ 
23:      end if
24:    end if
25:    if (receive(PrioT) from  $q$ ) then
26:      if Prio =  $\perp$  then
27:        Prio  $\leftarrow q$ 
28:      else
29:        send(PrioT) to  $q + 1$ 
30:      end if
31:    end if
32:    if (receive(ctr1, C, R, PT, PPr) from  $q$ ) then
33:      Ok  $\leftarrow false$ 
34:      if ( $q = Succ$ )  $\wedge$  ( $myC = C$ )  $\wedge$  ( $Succ \neq 0$ ) then
35:        Succ  $\leftarrow Succ + 1$ 
36:        Ok  $\leftarrow true$ 
37:        if  $R$  then
38:          RSet  $\leftarrow \emptyset$ 
39:          Prio  $\leftarrow \perp$ 
40:        end if
41:      end if
42:      if ( $q = 0$ ) then
43:        Ok  $\leftarrow true$ 
44:        if  $myC \neq C$  then
45:          Succ  $\leftarrow \min(1, \Delta_p - 1)$ 
46:          if  $R$  then
47:            RSet  $\leftarrow \emptyset$ 
48:            Prio  $\leftarrow \perp$ 
49:          end if
50:        end if
51:        myC  $\leftarrow C$ 
52:      end if
53:      if Ok then
54:        PT  $\leftarrow \min(PT + |RSet|_q, \ell + 1)$ 
55:        if Prio =  $q$  then
56:          PPr  $\leftarrow \min(PPr + 1, 2)$ 
57:        end if
58:        send(ctr1, myC, R, PT, PPr) to Succ
59:      end if
60:    end if
61:  end for
62:  if (State = Req)  $\wedge$  ( $|RSet| \geq Need$ ) then
63:    State  $\leftarrow In$ 
64:    EnterCS()
65:  end if
66:  if (State = In)  $\wedge$  ReleaseCS() then
67:    for all  $i \in RSet$  do
68:      send(ResT) to  $i + 1$ 
69:    end for
70:    RSet  $\leftarrow \emptyset$ 
71:    State  $\leftarrow Out$ 
72:  end if
73:  if (Prio  $\neq \perp$ )  $\wedge$  (State  $\neq Req \vee |RSet| \geq Need$ ) then
74:    send(PrioT) to Prio + 1
75:    Prio  $\leftarrow \perp$ 
76:  end if
77: end repeat

```

a to r behind a resource token. So, every process will collect the incoming resource token and the system will reach the configuration (ii) where r and b execute their critical section while a is still waiting for a resource token and the pusher is reaches r . When r receives the pusher, it directly retransmits it to b , while keeping its resource token, as shown in configuration (iii) . Similarly, b receives the pusher while executing its critical section, and then b retransmits it to r , as shown in configuration (iv) , and then r retransmits the pusher to a (configuration (v)). Assume now that a receives the pusher while r and b leave their critical sections. We obtain configuration (vi) : a must release its resource tokens because of the pusher. In configuration (vii) , r and b directly retransmit the resource token they receive because they are not requesters. Finally, r and b again become requesters for one resource token in configuration $(viii)$, identical to configuration (i) . We can repeat this cycle indefinitely, and process a never satisfies its request.

To solve this problem, we add a *priority token* (message PrioT) whose goal is to cancel the effect of the pusher. A process which receives the priority token retransmits it immediately to its successor (in DFS order), unless it has an unsatisfied request. In this case, the process holds the priority token (the variable Prio is set from \perp to the channel number from which the process receives the priority token) until its request is satisfied: the token will be released when the process enters its critical section. A process that holds the priority token does not have to release its reserved resource tokens when it receives the pusher: it only retransmits the pusher to its DFS order successor.

Using these three types of tokens, we obtain a simple non self-stabilizing k-out-of- ℓ exclusion protocol. Below, we present the module that turns this protocol into a self-stabilizing one.

A controller for self-stabilization. After a finite period of transient faults, some tokens may have disappeared or may be duplicated. To resume correct behavior, we need an additional self-stabilizing mechanism that regulates the number of tokens in the network: to achieve that, we use a *self-stabilizing*

controller. Such a mechanism is similar to the one introduced in [8] for a self-stabilizing ℓ -exclusion protocol on a ring.

Basically, the controller is a special token (message `ctrl`) that is used as a flag to count the other tokens: at the end of each full depth-first circulation of the controller into the tree, the root learns the number of tokens of each type (resource, pusher, priority) and adjusts this number if it is incorrect.

Of course, the controller can be effected by transient faults, too. So, we use the *counter flushing* technique of Varghese [15] to design the controller as a self-stabilizing *depth-first token circulation* (DFTC) in the tree. We will see in the next paragraph (page 9) how to design such self-stabilizing token circulation.

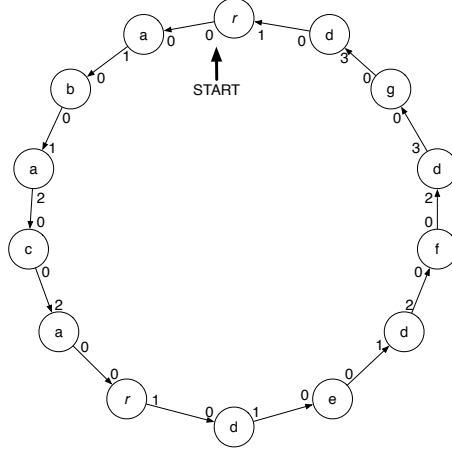


Figure 4: Virtual ring.

We now explain how we count the tokens using the controller. To that end, consider the problem of counting the resource tokens. The counting of the other types of tokens derives from this method. Any token circulation follows a virtual ring of the network: *e.g.*, Figure 4 shows the virtual ring followed by the depth-first token circulation of Figure 1. The controller uses this virtual structure to count the (resource) tokens. We split the count of the resource tokens into two subcounts:

- *The “passed” tokens.* When a process holds some resource tokens that came from channel i and receives the controller from the channel i , it retransmits the controller through channel $i + 1$ while keeping the resource tokens: in this case, we say that the controller *passes* these tokens in the virtual ring. Indeed, these tokens were ahead the controller (in the virtual ring) before the process received the controller, and are behind afterwards. The field *PT* of the controller message is used to compute the number of the passed resource tokens.
- *The tokens that are never passed by the controller.* These tokens are counted in the variable *SToken* maintained at the root process. At the beginning of any circulation of the controller, the variable *SToken* is reset to 0. Then, until the end of the circulation of the controller, each time a resource token starts a new circulation (*i.e.* the token leaves the root from channel 0), *SToken* is incremented.

Hence, when the controller terminates its circulation, the number of resource tokens in the network is equal to $PT + \text{SToken}$. Three cases are then possible:

- *The number of tokens is correct.* in this case, the system is stabilized.
- *There are too few tokens.* In this case, the root creates the number of necessary tokens at the end of the traversal; the system is then stabilized.
- *There are too many tokens.* In this case, we reset the network. To that goal, we mark the controller token with a special flag (the field *R* in the message `ctrl`). The root sends the marked controller, erases its reserved tokens as well as all the tokens it receives until the termination of the controller’s traversal. Upon receiving the controller, every other process erases its reserved tokens. When the

controller finishes its traversal, there is no token in the network, and then the root creates exactly ℓ resource tokens; the system is stabilized.

By extending this technique to the other types of tokens, we obtain a self-stabilizing controller.

Self-stabilizing DFTC. Using the counter flushing technique, we design a self-stabilizing *DFTC* to implement the controller. The principle of counter flushing is the following: after transient faults, the token message can be lost. Hence, the root must use a timeout mechanism to retransmit the token in case of deadlock. The timeout is managed using the function `RestartTimer()` (that allows it to reinitialize the timeout) and the predicate `TimeOut()`, which holds when a specified time interval is exceeded.⁴

Due to the use of the timeout, we must now deal with duplicated messages. Also, arbitrary messages may exist in the network after faults (however they are assumed to be bounded). To distinguish the duplicates from the valid token and to flush the system of corrupted messages, every process maintains a counter variable `myC` that takes values in $[0 \dots 2(n-1)(C_{\text{MAX}}+1)]$ and marks each message with that value. Every process also maintains a pointer `Succ` to know to which process it must send the token. The effects of the reception of a token message differs for the root and the other processes:

- The root considers a token message as valid when the message comes from `Succ` and is marked with a value c such that `myC` = c . Otherwise, it simply ignores the message. In case of a valid message, the root increments `Succ` (modulo Δ_r) and retransmits the token with the flag value `myC` to `Succ` so that the valid token follows DFS order. If `Succ` = 0, this means that the token just finished its previous circulation. As a consequence, the root increments `myC` (modulo $2(n-1)(C_{\text{MAX}}+1)$) before retransmitting the token.
- A non-root process p considers a message as valid in two cases: (1) When it receives a token message from its parent (channel 0) marked with a value c such that `myC` $\neq c$ or (2) when it receives a token message from `Succ` and the message is marked with a value c such that `myC` = c . In case (1), p sets `myC` to c and `Succ` to $\min(1, \Delta_p - 1)$ (*n.b.* in case of a leaf process `Succ` is set to 0) before retransmitting the token message marked with `myC` to `Succ`. In case (2), p increments `Succ` (modulo Δ_p) and then sends the token marked with `myC` to `Succ` so that the valid token follows DFS order. In all other cases, p considers the message to be invalid. In the case of an invalid message coming from channel 0 with `myC` = c , p does not consider the message in the computations but retransmits it to prevent deadlocks. In all other cases, p simply ignores the message.

Using this method, the root increments its counter `myC` infinitely often and, due to the size of the `myC`'s domain, the `myC` variable of the root eventually takes a value that does not exist anywhere else in the system (because the number of possible values initially in the system is bounded by $2(n-1)(C_{\text{MAX}}+1)$). In this case, the token marked with the new value will traverse the whole network and will be considered to be a valid token by every process. Until the end of that traversal, the root will ignore all other token messages. Hence, at the end of the traversal, the system will be stabilized.

Dealing with bounded memory. Due to the use of reset, the root does not need to know the exact number of tokens at the end of the controller's traversals. Actually, the root must only know if the number of tokens is too high, or the number of tokens it needs to add if the number is too low. Hence, the counting variables can be bounded by $\ell + 1$ for the resource tokens and by 2 for the other types of token. The fact that a variable is assigned to its maximal value will mean that there are too many tokens in the network and so a reset must be started. Otherwise, the value of the counting variable will state whether there is a deficient number of tokens, and in that case, how many must be added. For any assignment to one of these bounded variables, the value is set to the minimum between its new computed value and the maximal value of its domain.

4 Correctness and Waiting Time

In this section, we first prove that our protocol is a self-stabilizing k -out-of- ℓ exclusion protocol. We then analyze its waiting time.

⁴We assume that this time interval is sufficiently large to prevent congestion.

Correctness. We split the proof into three steps. (1) Recall that the controller part of our protocol is a self-stabilizing DFS token circulation. (2) We show that once the controller is stabilized to DFS token circulation, the system eventually stabilizes to the expected number of different tokens. (3) We show that once the system contains the expected tokens, the system stabilizes to the k -out-of- ℓ exclusion specification.

In our protocol, when a process receives a `ctrl` message, either it considers the message as valid or not. The process takes account of the messages for computations only when they are valid. Assume that a process p receives a `ctrl` message marked with the flag value c from channel q . Process p considers this message as valid if and only if $(q = \text{Succ}_p \wedge c = \text{myC}_p) \vee (p \neq r \wedge (q = 0 \wedge c \neq \text{myC}_p))$.

In the following, we call any `ctrl` message a *control token*. Each time a process receives a valid `ctrl` message, it makes some local computations, and then sends another `ctrl` message. In the case of a non-root process, the sent message is marked with the same flag as the received message: we consider it to be the same control token. In the case of the root, the sent message is marked either with the same value or with a new one. In the former case, we consider it to be the same control token, while in the latter case, we consider the received control token to have terminated its traversal, and the transmitted control token to be new.

To implement the control part, we use the counter flushing techniques introduced by Varghese in [15]. Hence, from [15], we can deduce the following lemma:

Lemma 1 *Starting from any configuration, the system converges to a configuration at which:*

1. *There exists at most one valid control token in the network.*
2. *The root regularly creates a new valid control token.*
3. *Any valid control token visits all processes following in DFS order.*

Remark 1 *In our protocol, only the valid control tokens are considered in the computations. Hence, from now on, we only consider the valid control tokens and we simply refer to them as control tokens.*

We now show that starting from any configuration, the system eventually contains the expected number of each type of token.

Note that each resource token is either in a link (in this case, the token is said to be *free*) or it is stored in the `RSet` of a process (in this case, the token is said to be *reserved*). Hence, at any time the number of resource tokens in the network is equal to the sum of the size of the `RSet` multisets plus the number of *free* resource tokens.

Similarly, note that at any time the number of priority tokens is equal to the number of processes satisfying $\text{Prio} \neq \perp$ plus the number of *free* priority tokens.

Finally, as a process cannot store any pusher token, the number of pusher tokens is equal to the number of free pusher tokens.

Lemma 2 *Let γ be the first configuration after the control part is stabilized. If, after γ , the root creates a control token having a reset field R equal to *true*, then the system contains no resource, priority, and pusher token at the end of the traversal of the control token.*

Proof. Consider any control token created by the root after configuration γ . Assume that the reset field R of the control token is set to *true*. Then, the `Reset` variable of the root is also *true* (see Line 73 in Algorithm 1). `Resetr` remains *true* until the control token terminates its traversal. Hence, during the traversal, any token (except the control token) that is received by the root is ignored by the root and so disappears from the network (see Lines 10, 20, and 35 in Algorithm 1). Also, during its traversal, each process erases all tokens (except the control token) it holds when visited by the control token (see Line 48 in Algorithm 1, and Lines 37, and 46 in Algorithm 2). Hence, every resource, priority, or pusher token is either erased at a process when the process is visited by the control token or is pushed to the root and then disappears: at the end of the traversal of the control token, the system contains no resource, priority, or pusher tokens. \square

Lemma 3 *Let γ be the first configuration after the control part is stabilized. When a control token created by the root after γ terminates its traversal, we have:*

- *If $PT + \text{SToken}_r > \ell$, then there are more than ℓ resource tokens in the network.*

- If $PT + \text{SToken}_r \leq \ell$, then there are exactly $PT + \text{SToken}_r$ resource tokens in the network.

Proof. Consider any control token created by the root after configuration γ . Let us study the following two cases:

- *The reset field R of the control token is true.* By Lemma 2, there is no resource token in the network when the control token terminates its circulation. So, to prove the lemma in this case, we must show that $PT + \text{SToken}_r = 0$ at the end of the circulation.
First, SToken_r is reset to 0 (Line 63) before the control token starts its circulation (Line 73). Also, Reset_r is *true* when the control token starts its circulation (see Line 73 in Algorithm 1). Thus, until termination of the circulation, r ignores any resource tokens it receives (see Lines 10, 20, and 35) and so SToken_r is still equal to 0 at the end of the control token circulation.
Consider now the PT field of the control token. Before the start of the control token circulation, r executes the following action: RSet is set to \emptyset (Line 49 in Algorithm 1), PT is set to 0 (Line 66 in Algorithm 1), and, as a consequence, PT is set to $\min(0, \ell + 1)$ (Line 69 in Algorithm 1). So, at the start of the control token circulation, the control token is sent with its PT field equal to 0. Since the reset field R of the control token is equal to *true*, each time the control token arrives at a process, the process resets its RSet variable to \emptyset (see Lines 49 in Algorithm 1, Lines 38, and 47 in Algorithm 2) before setting PT to $\min(PT + |\text{RSet}|_q, \ell + 1)$ (see Line 69 in Algorithm 1 and Line 54 in Algorithm 2) and then retransmitting the token. Hence, PT remains equal to 0 until the end of the circulation. When the control token terminates its circulation, $PT + \text{SToken}_r = 0$, and we are done.
- *The reset field R of the control token is false.* In this case, we can remark that no resource token is erased during the circulation of the control token, because both R and Reset_r are *false*.
(*) We now show that any resource token is counted at most once during the circulation of the control token. Due to the FIFO quality of the links and the fact that when the control token is received by a process, the process receives no other message before retransmitting the control token, we have the following property: A resource token is passed by the control token at most once during a circulation. So, during the circulation, either the resource token is counted into the PT field of the control token when the resource token is passed by the control token (see Line 69 in Algorithm 1 and Line 54 in Algorithm 2) or it is counted at the root when it terminates a loop of the virtual ring (Line 15). Hence, any resource token is counted at most once.
(**) Finally we show, by contradiction, that any resource token is counted at least once during the circulation of the control token. Assume that a resource token is not counted during that circulation. Then, the resource token is never passed by the control token. The links are FIFO, and when the control token is received by a process, the process receives no other message before retransmitting the control token. So, the resource token is always ahead the control token in the virtual ring. As a consequence, the resource token is eventually counted at the root when it terminates a loop of the virtual ring (Line 15), contradiction.
From (*), we know that if $PT + \text{SToken}_r > \ell$ at the end of the control token circulation, then there are more than ℓ resource tokens in the network. From (*) and (**), we know that if $PT + \text{SToken}_r \leq \ell$ at the end of the control token circulation, then there are exactly $PT + \text{SToken}_r$ resource tokens in the network, and we are done.

□

Following similar reasoning, we obtain the following two lemmas:

Lemma 4 *Let γ be the first configuration after the control part is stabilized. When a control token created by the root after γ terminates its traversal, we have:*

- If $\text{SPrio}_r + PPr > 1$, there is more than one priority token in the network.
- If $\text{SPrio}_r + PPr \leq 1$, there are exactly $\text{SPrio}_r + PPr$ priority tokens in the network.

Lemma 5 *Let γ be the first configuration after the control part is stabilized. When a control token created by the root after γ terminates its traversal, we have:*

- If $SPush_r > 1$, then there is more than one pusher token in the network.
- If $SPush_r \leq 1$, then there are exactly $SPush_r$ pusher tokens in the network.

Lemma 6 *Starting from any configuration, the system eventually reaches a configuration from which there always exist exactly ℓ resource tokens.*

Proof. Let γ be the first configuration after the control part is stabilized. Consider any control token created by the root after γ . Let us study the two following cases:

- $PT + SToken_r \leq \ell$ at the end of the control token traversal. Then, $Reset_r$ is set to *false* (Line 47 in Algorithm 1) and, as a consequence, the reset field of the next control token will be equal to *false* (Line 73 of Algorithm 1). Hence, no resource token will be erased during the next circulation of a control token. If $PT + SToken_r < \ell$, then exactly $\ell - (PT + SToken_r)$ are created (see Lines 55 to 58 in Algorithm 1). Hence, the number of resource tokens will be exactly equal to ℓ at the beginning of the next control token circulation. By Lemma 3, $PT + SToken_r$ will be equal to ℓ at the end of the next control token circulation, no resource token will be added. Any later circulation of the control token cannot change the number of resource tokens. Hence, the system will contain ℓ resource tokens forever.
- $PT + SToken_r > \ell$ at the end of the control token traversal. Then, $Reset_r$ is set to *true* (Line 47 in Algorithm 1) and, as a consequence, the reset field of the next control token will be equal to *true* (Line 73 of Algorithm 1). By Lemmas 2 and 3, reducing to the previous case when the circulation of the next control token terminates, and we are done.

□

Following similar reasoning, we can deduce from Lemmas 2, 4, and 5, the following two lemmas:

Lemma 7 *Starting from any configuration, the system eventually reaches a configuration from which there always exists one priority token.*

Lemma 8 *Starting from any configuration, the system eventually reaches a configuration from which there always exists one pusher token.*

We now show that once the system contains the correct number of each type of token, the system stabilizes to the k -out-of- ℓ exclusion specification.

Lemma 9 *Starting from any configuration, every process receives a pusher token infinitely many times.*

Proof. By Lemmas 6, 7, and 8, starting any configuration, the system eventually reaches a configuration γ from which there are exactly ℓ resource tokens, one priority token, and one pusher token in the network. From γ , the system is then never again reset. So, from γ , the unique pusher token of the system always follows DFS order. Each time a process receives the pusher token, it retransmits it in finite time. Hence, every processes receives it infinitely often and the lemma holds. □

Lemma 10 *Starting from any configuration, every process receives a priority token infinitely many times.*

Proof. By Lemmas 6, 7, and 8, starting any configuration, the system eventually reaches a configuration γ from which there are ℓ resource tokens, one priority token, and one pusher token in the network. From γ , the system is never again reset. So from γ , the unique priority token of the system always follow the DFS order.

By way of contradiction, assume that, from γ , a process eventually stops receiving the priority token. Since the priority token circulate following the DFS order and it traverses any link in a finite time, we can deduce that some other process p eventually holds it forever. In this case, p is a requester and its request is never satisfied. Now, by Lemma 9 any other process receives the pusher token infinitely often. So, each other process retransmits the resource tokens it holds within finite time, because it eventually either satisfies its request, executes its critical sections, and then release its tokens or does satisfy its request, but receives

the pusher, and then releases its resource tokens. Similarly to the pusher token, the resource tokens always follows the DFS order from γ . Hence, p receives resource tokens infinitely many times, and, as it holds the priority token it never releases it even if it receives the pusher token. As $k \leq \ell$, the request of p is eventually satisfied, contradiction. \square

Lemma 11 *Starting from any configuration, every process receives resource tokens infinitely many times.*

Proof. By Lemmas 6, 7, and 8, starting any configuration, the system eventually reaches a configuration γ from which there is ℓ resource tokens, one priority token, and one pusher token in the network. From γ , the system is then never again reset. So, from γ , the resource tokens of the system always follow the DFS order.

Assume, by way of contradiction, that some process only receives resource tokens finitely many times. This implies that every resource token is eventually held forever by some process. Consider one process that holds at least one resource token forever. By Lemma 9, that process cannot hold the priority token forever. When it releases the priority token, either its request is satisfied, it executes the critical section within finite time, and then releases its resource tokens, or it is not a requester and thus must release its resource token. Either case is a contradiction, and we are done. \square

Lemma 12 *Starting from any configuration, the fairness property of the k -out-of- ℓ exclusion specification is eventually satisfied.*

Proof. Assume that there a request on some process p that is never satisfied.

By Lemmas 6, 7, and 8, starting any configuration, the system eventually reaches a configuration γ from which there are ℓ resource tokens, one priority token, and one pusher token into the network. From γ , the system is then never again reset. Hence, from γ , if p holds the priority token, it releases it only if its request is satisfied. By Lemma 11, p eventually receives the priority token. Again by Lemma 11, p eventually releases the priority token, and so its request must have been satisfied, contradiction. \square

Lemma 13 *Starting from any configuration, the safety property of the k -out-of- ℓ exclusion specification is eventually satisfied.*

Proof. First, each resource unit corresponds to a resource token. Hence each resource unit can be used by at most one process at any time.

By Lemma 6, there are eventually exactly ℓ resource tokens in the network. Hence, eventually, exactly ℓ resource unit are available in the system.

Finally, any process p that initially holds some resource tokens eventually releases them because either is not a requester or it eventually satisfies its request by Lemma 12. Hence, eventually p sets \mathbf{RSet} to \emptyset and then $|\mathbf{RSet}| \leq \mathbf{Need}$ forever because each time p receives a resource token while $|\mathbf{RSet}| \geq \mathbf{Need}$, it directly retransmits it (see Lines 10 to 19 in Algorithm 1 and Lines 9 to 15 in Algorithm 2). Now, \mathbf{Need} is always less or equal to k . Hence, every process eventually only uses at most k resource tokens (units) simultaneously. \square

Lemma 14 *Starting from any configuration, the efficiency property of the k -out-of- ℓ exclusion specification is eventually satisfied.*

Proof. We will use the definition of efficiency given in [3].

To show this lemma, we now prove that starting from any configuration, (k, ℓ) -liveness is eventually satisfied.

To see this, consider the configuration γ from which: (1) there are ℓ resource tokens, one priority token, and one pusher token; and (2) the safety properties of the k -out-of- ℓ exclusion are satisfied (such a configuration exists by Lemmas 6, 7, 8, and 13).

Assume that from γ the system reaches a configuration γ' after which there is a subset I of processes such that every process in I executes its critical section forever (in this case they hold some resource units forever). Let α be the total number of resource units hold forever by the processes in I .

Assume then that there are some processes not in I that request some resource units and each of these processes requests at most $\ell - \alpha$ resource units.

The priority token follows DFS order. Also, as every process in I executes the critical section forever, none of these processes keeps the priority token forever (see Lines 92 in Algorithm 1 and 73 in Algorithm 2). Finally, every non-requester directly retransmits the priority token when it receives it (see Line 92 in Algorithm 1 and Line 73 in Algorithm 2). Hence, there is a requesting process p which is not in I that eventually receives the priority token. From that point, p will release it only after its request is satisfied (see Line 92 in Algorithm 1 and Line 73 in Algorithm 2). As a consequence, p will keep every resource token it receives even if it receives the pusher token. Now, by checking the proof of Lemma 9, we can see that Lemma 9 still holds even if some processes execute the critical section forever. So, by Lemma 9 every processes that is not in $I \cup \{p\}$ receives the pusher token infinitely often and so cannot hold resource tokens forever. Finally, every process in I directly retransmits the resource tokens it receives when it is executing the critical section because they satisfy $|\text{RSet}| \geq \text{Need}$ by Lemma 13 (see Lines 10 to 19 in Algorithm 1 and Lines 9 to 15 in Algorithm 2). So, p eventually receives the resource tokens it needs to perform the critical section (remember that p requests at most $\ell - \alpha$ resource units) and we are done. \square

By Lemmas 13, 12, and 14, we obtain:

Theorem 1 *The distributed protocol proposed in Algorithms 1 and 2 is a k -out-of- ℓ exclusion protocol for tree networks.*

Waiting Time.

Theorem 2 *Once the protocol proposed in Algorithms 1 and 2 is stabilized, the waiting time is $\ell \times (2n - 3)^2$ in the worst case.*

Proof. We first show that the waiting time of a requesting process that holds the priority token is $\ell \times (2n - 3)$ in the worst case. Consider a process p that requests some resource units and holds the priority token. In the worst case, p appears only once in the virtual ring defined by the DFS order (if p is a leaf). Also in the worst case, the ℓ resource tokens may traverse the entire virtual ring before p receives the tokens it needs. The virtual ring can contain up to $(2n - 3)$ processes in addition to p . Any resource token may satisfy one request each time it traverses a process (in the worst case, the processes always requests one token). Hence, the ℓ resource tokens may satisfy up to $\ell \times (2n - 3)$ requests before p satisfies its request.

Using similar reasoning, we can see that a requesting process may wait until the priority token traverses the whole virtual ring (up to $2(n - 2)$ nodes) before it satisfy its request; during that time, up to $\ell \times (2n - 3)^2$ requests can be satisfied, and we are done. \square

5 Conclusion and Perspectives

In this paper, we propose the first (deterministic) self-stabilizing distributed k -out-of- ℓ exclusion protocol for asynchronous oriented tree networks. The proposed protocol uses a realistic model of computation, the message-passing model. The only restriction we make is to assume that the channels initially contain a bounded known number of arbitrary messages. We make this assumption to obtain a solution that uses bounded memory per process (see the results in [7]). However, if we assume unbounded process memory, our solution can be easily adapted to work without assumptions on channels (following the method presented in [9]).

The main interest in dealing with an oriented tree is that solutions on the oriented tree can be directly mapped to solutions for arbitrary rooted networks by composing the protocol with a spanning tree construction (e.g, [1, 4]).

There are several possible extensions of our work. On the theoretical side, one can investigate whether the waiting time of our solution ($\ell \times (2n - 3)^2$) can be improved. Possible extension to networks where processes are subject to other failure patterns, such as process crashes, remains open. On the practical side, our solution is designed in a realistic model and can be extended to arbitrary rooted networks. Hence, implementing our solution in a real network is a future challenge.

References

- [1] Y Afek and A Bremler. Self-stabilizing unidirectional network algorithms by power supply. *Chicago Journal of Theoretical Computer Science*, 1998:Article 3, 1998.
- [2] A K Datta, R Hadid, and V Villain. A new self-stabilizing k-out-of-l exclusion algorithm on rings. In Shing-Tsaan Huang and Ted Herman, editors, *Self-Stabilizing Systems*, volume 2704 of *Lecture Notes in Computer Science*, pages 113–128. Springer, 2003.
- [3] A K Datta, R Hadid, and V Villain. A self-stabilizing token-based k-out-of-l exclusion algorithm. *Concurrency and Computation: Practice and Experience*, 15(11-12):1069–1091, 2003.
- [4] Sylvie Delaët, Bertrand Ducourthial, and Sébastien Tixeuil. Self-stabilization with r-operators revisited. In Ted Herman and Sébastien Tixeuil, editors, *Self-Stabilizing Systems*, volume 3764 of *Lecture Notes in Computer Science*, pages 68–80. Springer, 2005.
- [5] EW Dijkstra. Self stabilizing systems in spite of distributed control. *Communications of the Association of the Computing Machinery*, 17:643–644, 1974.
- [6] M J Fischer, N A Lynch, J E Burns, and A Borodin. Distributed fifo allocation of identical resources using small shared space. *ACM Trans. Program. Lang. Syst.*, 11(1):90–114, 1989.
- [7] Mohamed G. Gouda and Nicholas J. Multari. Stabilizing communication protocols. *IEEE Trans. Computers*, 40(4):448–458, 1991.
- [8] Rachid Hadid and Vincent Villain. A new efficient tool for the design of self-stabilizing l-exclusion algorithms: The controller. In Ajoy Kumar Datta and Ted Herman, editors, *WSS*, volume 2194 of *Lecture Notes in Computer Science*, pages 136–151. Springer, 2001.
- [9] Shmuel Katz and Kenneth J. Perry. Self-stabilizing extensions for message-passing systems. *Distributed Computing*, 7(1):17–26, 1993.
- [10] Y Manabe, R Baldoni, M Raynal, and S Aoyagi. k-arbiter: A safe and general scheme for h-out of-k mutual exclusion. *Theor. Comput. Sci.*, 193(1-2):97–112, 1998.
- [11] Y Manabe and N Tajima. (k)-arbiter for h -out of- k mutual exclusion problem. In *ICDCS*, pages 216–223, 1999.
- [12] M Raynal. A distributed solution to the k-out of-m resources allocation problem. In F K H A Dehne, F Fiala, and W W Koczkodaj, editors, *ICCI*, volume 497 of *Lecture Notes in Computer Science*, pages 599–609. Springer, 1991.
- [13] Michel Raynal. *Algorithms for Mutual Exclusion*. MIT Press, 1986.
- [14] Michel Raynal. *Algorithmes du parallélisme, le problème de l'exclusion mutuelle*. Dunod informatique, 1990.
- [15] George Varghese. Self-stabilization by counter flushing. *SIAM J. Comput.*, 30(2):486–510, 2000.