

Self-Stabilizing Depth-First Token Circulation In Arbitrary Rooted Networks *

Ajoy K. Datta,¹ Colette Johnen,² Franck Petit,³
Vincent Villain³

¹ Department of Computer Science, University of Nevada, Las Vegas

² L.R.I./C.N.R.S., Université de Paris-Sud, France

³ LaRIA, Université de Picardie Jules Verne, France

Abstract: We present a deterministic distributed depth-first token passing protocol on a rooted network. This protocol uses neither the processor identifiers nor the size of the network, but assumes the existence of a distinguished processor, called the root of the network. The protocol is self-stabilizing, meaning that starting from an arbitrary state (in response to an arbitrary perturbation modifying the memory state), it is guaranteed to reach a state with no more than one token in the network. Our protocol implements a *1-fair* token circulation scheme, i.e., in every round, every processor obtains the token once. The proposed protocol has extremely small state requirement—only $3(\Delta + 1)$ states per processor, i.e., $O(\log \Delta)$ bits per processor, where Δ is the degree of the network.

The protocol can be used to implement a fair distributed mutual exclusion in any rooted network. This protocol can also be used to construct a DFS spanning tree.

Keywords: Distributed mutual exclusion, self-stabilization, spanning tree, token passing.

Correspondance:

Franck PETIT

Email. petit@laria.u-picardie.fr

LaRIA, Université de Picardie Jules Verne

5, rue du Moulin Neuf F - 80039 AMIENS Cedex 1 FRANCE

Tel. +33-322-827-874

Fax. +33-322-827-654

*A preliminary version of this work was presented at SIROCCO '98 [6].

1. Introduction

Robustness is one of the most important requirements of modern distributed systems. Various types of faults are likely to occur at various parts of the system. These systems go through the transient states because they are exposed to constant change of their environment. The concept of self-stabilization [7] is the most general technique to design a system to tolerate arbitrary transient faults. A self-stabilizing system, regardless of the initial states of the processors and initial messages in the links, is guaranteed to converge to the intended behavior in finite time.

The depth-first token circulation problem is to implement a token circulating from one processor to the next in the depth-first order such that every processor gets the token at least once in every round (defined more formally later). In this paper, the token is initiated by the root of the network.

Related Work. Dijkstra introduced the property of self-stabilization in distributed systems by applying it to algorithms for mutual exclusion on a ring [7]. Several deterministic self-stabilizing token passing algorithms for different topologies have been proposed in the literature: [1, 4, 7, 11, 9] for a ring; [2, 10, 21] for a linear array of processors, and [16, 17] for a tree network. Huang and Chen [13] presented a token circulation protocol for a connected network in non-deterministic depth-first-search order, and Dolev, Israeli, and Moran [8] gave a mutual exclusion protocol on a tree network under the model whose actions only allow read/write atomicity.

One of the important performance issues of self-stabilizing algorithms is the memory requirement per processor. The memory requirement of a processor depends on the total number of states of the processor. The solution in [13] requires $O(n\Delta)$ states per processor, where n is the number of processors. The algorithm in [8] constructs a spanning tree and implements a token circulation scheme on the constructed spanning tree. The spanning tree protocol uses a *distance* variable (which needs n states) and the token circulation algorithm maintains a descendant pointer (which needs Δ states). So, the algorithm [8] requires at least $n\Delta$ states or $\log(n\Delta)$ bits.

A state-efficient token passing protocol on general network is presented in [15]. In this protocol, a processor p_i needs to maintain $3(\Delta_i + 1)$ states ($\lceil \log(3(\Delta_i + 1)) \rceil$), where Δ_i is the degree of p_i . Subsequently, this result was improved by Petit and Villain [18] to $2(\Delta_i + 1)$ states for a processor p_i . Both protocols do not use the *distance* variable. But, in these algorithms, a processor needs the knowledge of the state of the neighbors of its neighbors. Since the algorithms assume the atomic execution of the actions, this requirement makes the atomic step bigger—in one atomic step, a processor reads the state of its neighbors, the state of the neighbors of its neighbors, and finally changes its own state. This drawback has been removed in [14]. In this protocol, a processor only reads the state of its neighbors in an atomic step. Thus, this algorithm has a smaller atomicity than that in [15, 18]. The state requirement of this protocol is $12(\Delta_i + 1)$ states for a processor p_i . Petit and Villain [20] and [19] adapted the result of [15] and [18], respectively, in the message passing model.

Contributions. In this paper, we present a self-stabilizing depth-first token circulation scheme on a general network with a distinguished root, called Algorithm \mathcal{TC} . Algorithm \mathcal{TC} has all the desirable features of the algorithm in [14]. In addition, we reduced the state requirement for a processor p to $3(\Delta_p + 1)$ states (only $2(\Delta_p + 1)$ states for the root). Also, our algorithm is simpler

(less number of actions) than that in [14]. Algorithm \mathcal{TC} implements a *1-fair* circulation of token, i.e., while each processor p is waiting for the token, every processor q ($q \neq p$) gets the token at most Δ_q times. Hence, Algorithm \mathcal{TC} can also be used to implement a fair distributed mutual exclusion among the processors.

Outline of the Paper. We present a self-stabilizing depth-first token circulation scheme on a general network with a distinguished root. The token passing problem is formally defined in Section 2.2. The rest of the paper is organized as follows: In Section 2, we describe the distributed systems and the model in which our token circulation scheme is written, and give a formal statement of the token passing problem solved in this paper. In Section 3, we present the token passing protocol, and in the following section (Section 4), we give the proof of correctness of the protocol. The state complexity of the protocol is given in Section 5. Finally, we make concluding remarks in Section 6.

2. Preliminaries

In this section, we define the distributed systems and programs considered in this paper, and state what it means for a protocol to be self-stabilizing. We then present the statement of the token passing problem and its properties.

2.1. Self-Stabilizing System

System. A *distributed system* is an undirected connected graph, $S = (V, E)$, where V is a set of processors ($|V| = n$) and E is the set of bidirectional communication link. We consider networks which are *asynchronous* and *rooted*, i.e., all processors, except the root are *anonymous*. We denote the root processor by r . The numbers, $1..n$, are used to identify the processors to present our ideas here, but no processor, except the root (identified by r), has any identity. A communication link (p, q) exists iff p and q are neighbors. Each processor p maintains its set of neighbors, denoted as N_p . We assume that N_p is a constant and is maintained by an underlying protocol.

Programs. Each processor executes the same program except the root r . The program consists of a set of *shared variables* (henceforth referred to as variables) and a finite set of actions. A processor can only write to its own variables and can only read its own variables and variables owned by the neighboring processors. So, the variables of p can be accessed by p and its neighbors.

Each action is uniquely identified by a label and is of the following form:

$$< label > :: < guard > \longrightarrow < statement >$$

The guard of an action in the program of p is a boolean expression involving the variables of p and its neighbors. The statement of an action of p updates zero or more variables of p . An action can be executed only if its guard evaluates to true. We assume that the actions are atomically executed: the evaluation of a guard and the execution of the corresponding statement of an action, if executed, are done in one atomic step. The atomic execution of an action of p is called a *step* of p .

The *state* of a processor is defined by the values of its variables. The *state* of a system is a product of the states of all processors ($\in V$). In the sequel, we refer to the state of a processor and system as a (*local*) *state* and *configuration*, respectively. Let a distributed protocol \mathcal{P} be a collection of binary transition relations denoted by \mapsto , on \mathcal{C} , the set of all possible configurations of the system. A *computation* of a protocol \mathcal{P} is a *maximal* sequence of configurations $e = (\gamma_0, \gamma_1, \dots, \gamma_i, \gamma_{i+1}, \dots)$, such that for $i \geq 0, \gamma_i \mapsto \gamma_{i+1}$ (a single *computation step*) if γ_{i+1} exists, or γ_i is a terminal configuration. *Maximality* means that the sequence is either infinite, or it is finite and no action of \mathcal{P} is enabled in the final configuration. All computations considered in this paper are assumed to be maximal. During a computation step, one or more processors execute a step and a processor may take at most one step. This execution model is known as the *distributed daemon* [3]. We use the notation $Enable(A, p, \gamma)$ to indicate that the guard of the action A is true at processor p in the configuration γ . A processor p is said to be *enabled* at γ ($\gamma \in \mathcal{C}$) if there exists an action A such that $Enable(A, p, \gamma)$. We assume a *weakly fair* daemon, meaning that if a processor p is continuously *enabled*, then p will be eventually chosen by the daemon to execute an action.

The set of computations of a protocol \mathcal{P} in system S starting with a particular configuration $\alpha \in \mathcal{C}$ is denoted by \mathcal{E}_α . The set of all possible computations of \mathcal{P} in system S is denoted as \mathcal{E} . A configuration β is *reachable* from α , denoted as $\alpha \rightsquigarrow \beta$, if there exists a computation $e = (\gamma_0, \gamma_1, \dots, \gamma_i, \gamma_{i+1}, \dots) \in \mathcal{E}_\alpha$ such that $\beta = \gamma_i$ ($i \geq 0$).

Predicates. Let \mathcal{X} be a set. $x \vdash P$ means that an element $x \in \mathcal{X}$ satisfies the predicate P defined on the set \mathcal{X} . A predicate is non-empty if there exists at least one element that satisfies the predicate. We define a special predicate *true* as follows: *for any* $x \in \mathcal{X}$, $x \vdash \text{true}$.

Self-Stabilization. We use the following term, *attractor* in the definition of self-stabilization.

Definition 2.1 (Attractor). Let X and Y be two predicates of a protocol \mathcal{P} defined on \mathcal{C} of system \mathcal{S} . Y is an attractor for X if and only if the following condition is true:

$$\forall \alpha \vdash X : \forall e \in \mathcal{E}_\alpha : e = (\gamma_0, \gamma_1, \dots) :: \exists i \geq 0, \forall j \geq i, \gamma_j \vdash Y. \text{ We denote this relation as } X \triangleright Y.$$

Definition 2.2 (Self-stabilization). The protocol \mathcal{P} is self-stabilizing for the specification $\mathcal{SP}_\mathcal{P}$ on \mathcal{E} if and only if there exists a predicate $\mathcal{L}_\mathcal{P}$ (called the legitimacy predicate) defined on \mathcal{C} such that the following conditions hold:

1. $\forall \alpha \vdash \mathcal{L}_\mathcal{P} : \forall e \in \mathcal{E}_\alpha :: e \vdash \mathcal{SP}_\mathcal{P}$ (correctness).
2. $\text{true} \triangleright \mathcal{L}_\mathcal{P}$ (closure and convergence).

2.2. Specification of the Depth-First Token Passing Protocol

Definition 2.3 (Token Circulation Round). We define a computation in the protocol \mathcal{TC} starting from a configuration δ_{c0} to another configuration δ_{c1} as a token circulation round (in the sequel referred to as *round*) if the following conditions are true:

- (i) Exactly one processor holds a token in any configuration.
- (ii) r holds a token in both δ_{c0} and δ_{c1} .

- (iii) The token is passed among the processors in the depth-first search order.
- (iv) There is at least one configuration in between δ_{c0} and δ_{c1} ¹.

We are now ready to define the specification, $\mathcal{SP}_{\mathcal{TC}}$ of the protocol \mathcal{TC} . We consider a computation e of \mathcal{TC} to satisfy $\mathcal{SP}_{\mathcal{TC}}$ if the following conditions are true:

- (SP1) In every configuration in e , at most one processor has a token.
- (SP2) The token is passed among the processors in the depth-first search order.
- (SP3) Starting from δ_{c0} or δ_{c1} , the token circulation rounds are repeated.

We also require the protocol \mathcal{TC} to be *self-stabilizing*.

Condition (SP1) specifies the *Single Token* property, whereas Conditions (SP2) and (SP3) specify the *Fairness* property.

3. Depth-First Token Passing Algorithm

In this section, we propose the self-stabilizing depth-first token circulation algorithm. We first present the data structure used by the processors. Then we present the formal algorithm. Next, we define some terms to be used later in the paper. We then explain the process of token circulation, followed by the method of error correction. In particular, we do not use the distance variable used in [13] to destroy the cycles. We use a method similar to the one introduced in [15] to remove the cycles in the network.

3.1. Data Structures and Algorithm \mathcal{TC}

To distinguish each token round, each processor p uses a variable C_p , called the *round color*, which contains a value $\in \{0, 1\}$ when the system is stabilized. A third color E , called the *Error* color, is used by processors, except the root, during the stabilization. The descendant relationship is indicated by the variable D_p ($D_p \in N_p \cup \{\perp\}$). To choose its descendant, each processor p locally distinguishes each neighbor by some ordering, denoted as \succ_p .

The self-stabilizing depth-first token circulation is shown in Algorithm 3.1. To make the algorithm readable, we present it in three parts: the *macros*, *predicates*, and *actions*. The macros are not variables and they are dynamically evaluated. Anc_p denotes the set of ancestors of p , i.e., $Anc_p = \{q \in N_p \mid D_q = p\}$. UV_p is the set of neighbors not visited by the token. $Search_p$ chooses the next neighbor from UV_p . In the following, $\#Anc_p$ denotes the current number of ancestors of p . If $\#Anc_p = 1$, then the only ancestor of p is denoted as a_p . The predicates are used to describe the guards of the actions in Algorithm 3.1. Actions $TC1$ and $TC2$ implement the *token circulation*, i.e., the correct behavior of the system. The token circulates in the network according to Definition 2.3. Actions $EC1$, $EC2$, and $EC3$ implement the *error correction* of the system, i.e., they are used to bring the system from an illegitimate configuration to a legitimate one. All these predicates will be explained in detail in Section 3.2.

When the system stabilizes, the system must contain only one token which circulates in the DFS order. In such a configuration, a processor can make a move only if it holds the token. Holding the token means either $Forward(p)$ or $Backtrack(p)$ is true. Formally:

¹We assume that the network has at least one processor other than r .

Algorithm 3.1 (\mathcal{TC}) Self-Stabilizing Depth-First Token Circulation in Rooted Network.

Macro

$$\begin{aligned}
Anc_p &= \{q \in N_p : D_q = p\} \\
UV_p &= \left\{ q \in N_p : \left(\begin{array}{l} (q \succ_p D_p) \wedge (C_q \neq C_p) \wedge (D_q \neq p) \wedge (q \neq r) \\ \wedge ((C_q \neq E) \vee (D_q \neq \perp)) \end{array} \right) \right\} \\
Search_p &= \begin{cases} \min_{\succ_p}(UV_p) & \text{if } (UV_p \neq \emptyset) \\ \perp & \text{otherwise} \end{cases}
\end{aligned}$$

Predicates

$$\begin{aligned}
Forward(p) &\equiv (D_p = \perp) \wedge ((p = r) \vee ((\#Anc_p = 1) \wedge (C_p \neq E) \wedge (C_{a_p} = (C_p + 1) \bmod 2))) \\
Backtrack(p) &\equiv (D_p \neq \perp) \wedge (D_p \neq r) \wedge (D_{D_p} = \perp) \wedge (C_{D_p} = C_p) \wedge (C_p \neq E) \\
&\quad \wedge ((p = r) \vee (\#Anc_p = 1)) \\
Break(p) &\equiv (p \neq r) \wedge (D_p \neq \perp) \\
&\quad \wedge \left(\begin{array}{l} (D_p = r) \\ \vee ((C_p = E) \wedge ((D_{D_p} = \perp) \vee ((\#Anc_p > 1) \wedge (C_{D_p} = E)))) \\ \vee ((D_{D_p} = \perp) \wedge (\#Anc_p = 0) \wedge (C_{D_p} \neq (C_p + 1) \bmod 2)) \end{array} \right) \\
EDetect(p) &\equiv (p \neq r) \wedge (C_p \neq E) \\
&\quad \wedge \left(\begin{array}{l} ((D_p \neq \perp) \wedge (C_{D_p} = E) \wedge (\#Anc_p = 1)) \\ \vee ((D_p \neq r) \wedge (\#Anc_p > 1)) \end{array} \right) \\
EEnd(p) &\equiv (p \neq r) \wedge (C_p = E) \wedge (D_p = \perp) \wedge ((\#Anc_p = 0) \vee (Anc_p = \{r\}))
\end{aligned}$$

Actions

$$\begin{aligned}
TC1 &:: Forward(p) \longrightarrow C_p := (C_p + 1) \bmod 2; D_p := Search_p; \\
TC2 &:: Backtrack(p) \longrightarrow D_p := Search_p; \\
EC1 &:: Break(p) \longrightarrow D_p := \perp; \\
EC2 &:: EDetect(p) \longrightarrow C_p := E; \\
EC3 &:: EEnd(p) \longrightarrow \text{if } (\#Anc_p = 0) \text{ then } C_p := 0; \text{ else } C_p := C_r;
\end{aligned}$$

$$Token(p) \equiv Forward(p) \vee Backtrack(p)$$

3.2. Informal Explanation of Algorithm \mathcal{TC}

The proposed algorithm has two major tasks: (i) to circulate the token in the network in a deterministic depth-first order and (ii) to handle the abnormal situations (illegal configurations) due to the unpredictable initial configurations and transient errors. The tasks (i) and (ii) are explained with examples in Paragraphs **Token Circulation** and **Error Correction**, respectively.

Some Definitions. A path μ_p is a sequence (p_1, p_2, \dots, p_l) such that (i) $p = p_1$, (ii) $l \geq 2$, (iii) $\forall i \in [1, l[, D_{p_i} = p_{i+1}$ and (iv) $D_{p_l} = \perp$ or $D_{p_l} \in \{p_1, p_2, \dots, p_{l-1}\}$. $\forall i \in [1, l], p_i$ is said to belong to the path μ_p and is denoted as $p_i \in \mu_p$.

If $Anc_p = \emptyset$, then μ_p is called a *rooted path* (the path is rooted at p). A path μ_p rooted at $p \neq r$ is called an *illegal rooted path* and p is called an *illegal root*. A path μ_p rooted at r is called a *legal (rooted) path*.

The processor $p_l \in \mu_p$ is termed as a *leaf* if $D_{p_l} = \perp$. The leaf of a legal (respectively, illegal) rooted path is called *legal* (respectively, *illegal*) *leaf*. A leaf p' is termed as a *live* (respectively, *dead*) *leaf*, if $C_{p'} \neq E$ (respectively, $C_{p'} = E$).

The path μ_p is called a *cycle* if $D_{p_l} \in \{p_1, p_2, \dots, p_{l-1}\}$. A cycle μ_p is called a *strict cycle* if $\forall i \in [2, l]$, p_{i-1} is the only ancestor of p_i ($a_p = p_{i-1}$) and p_l is the only ancestor of p_1 ($a_1 = a_l$). If there exists at least one rooted path μ_q such that all processors in a cycle μ_p belong to μ_q , i.e., $\exists p_i \in \mu_p$ such that $i \in [2, l]$ and $\sharp Anc_{p_i} > 1$, then the rooted path μ_p is called a *rooted cycle*.

A rooted path with a live leaf is termed as a *live rooted path*. All other rooted paths are called *dead rooted paths*.

Every processor p such that $Anc_p = \emptyset$ and $D_p = \perp$ is called *path-free*, meaning p does not belong to any path.




Figure 3.1: A Possible Configuration.

These definitions are illustrated in Figure 3.1. Processors c, f, g , and j are illegal roots. d, e, i, h form a strict cycle. f and c are roots of a rooted cycle. k is a dead leaf. l and m are path-free.

Token Circulation. The root r initiates the token circulation round. The token then traverses all processors during a token circulation round (Definition 2.3).

We use $\delta_{c0} \in \mathcal{C}$ to denote a configuration where every processor in the system is path-free and has the color 0. Similarly, δ_{c1} denotes the configuration in which every processor is path-free and has the color 1. That is,

$$\delta_{c0} \equiv \forall p \in V :: Anc_p = 0 \wedge D_p = \perp \wedge C_p = 0$$

$$\delta_{c1} \equiv \forall p \in V :: Anc_p = 0 \wedge D_p = \perp \wedge C_p = 1.$$

Both δ_{c0} and δ_{c1} are among the possible configurations from where the algorithm behaves correctly, i.e., starting from δ_{c0} (respectively, from δ_{c1}), Algorithm \mathcal{TC} circulates the token (represented by the predicate $Token()$) in the depth-first search order to reach δ_{c1} (respectively, δ_{c0}). This is called one token circulation round (*crownd*). From δ_{c1} (respectively, δ_{c0}), the system reaches δ_{c0} (respectively, δ_{c1}) again in the same manner; the *crownd* is said 0 colored (1 colored). The token circulation consists of successive *crownds*, alternately colored with 0 and 1. After stabilization, the system repeats the *crownds* forever. The *crownd* is implemented by Actions $TC1$ and $TC2$. Every suffix of the computation starting from δ_{c0} or δ_{c1} is a legitimate configuration.

Consider the example in Figure 3.2. Step (i) corresponds to the configuration δ_{c0} . In this configuration, $Forward(r)$ is true and the only process *enabled* is r and the only action enabled at r is $TC1$. The root changes its color ($C_r := (C_r + 1) \bmod 2$) and builds μ_r by choosing a descendant




Figure 3.2: Depth-First Search Token Circulation.

(*Search* predicate). The root chooses the processor b as the descendant. This is shown in Step (ii). Similarly, b changes its color and chooses a descendant (Step (iii)). This process of extending the path continues until c executes Action $TC1$. c does not have any neighbor to choose from. So, c executes $D_c := \perp$ (*Search*). This indicates to its ancestor d that the token has traversed all processors reachable from c in the DFS tree (Step (v)). Now, $Backtrack(d)$ becomes true and d can execute $TC2$. Since d has no more unvisited neighbors, D_d becomes equal to \perp (Step (vi)). Actions $TC1$ and $TC2$ are repeated until all processors are visited by the token (Steps (vii) to (x)). Step (x) corresponds to δ_{c1} . Now, r changes its color to 0 and starts a new round with this color.

Error Correction. We now consider the transient failures. An illegitimate configuration is shown in Figure 3.1.

Actions $EC1$, $EC2$ and $EC3$ are used to bring the system into a legitimate configuration. Illegal configurations are locally detected by the predicates $Break$, $EDetect$, and $EEnd$. We split the predicates $Break$ and $EDetect$ into simpler predicates in Figure 3.3 to help explain them better.

$$\begin{aligned}
BrkA(p) &\equiv (D_p = r) \\
BrkB(p) &\equiv (C_p = E) \wedge (\#Anc_p > 1) \wedge (C_{D_p} = E) \\
BrkC(p) &\equiv (C_p = E) \wedge (D_{D_p} = \perp) \\
BrkD(p) &\equiv (D_{D_p} = \perp) \wedge (\#Anc_p = 0) \wedge (C_{D_p} \neq (C_p + 1) \bmod 2) \\
Break(p) &\equiv (p \neq r) \wedge (D_p \neq \perp) \wedge (BrkA(p) \vee BrkB(p) \vee BrkC(p) \vee BrkD(p)) \\
EDetectA(p) &\equiv (D_p \neq \perp) \wedge (C_{D_p} = E) \wedge (\#Anc_p = 1) \\
EDetectB(p) &\equiv (D_p \neq r) \wedge (\#Anc_p > 1) \\
EDetect(p) &\equiv (p \neq r) \wedge (C_p \neq E) \wedge (EDetectA(p) \vee EDetectB(p))
\end{aligned}$$

Figure 3.3: Predicates $Break$ and $EDetect$.

First consider the illegal configuration in which r has ancestors. For every ancestor p of r , p

satisfies $BrkA$ ($D_p = r$) and hence, $Break(p)$. Upon executing Action $EC1$, p eventually destroys the descendant pointer to r and since, r cannot be chosen as a descendant in the algorithm (see macro UV_p), r eventually will not be on any illegal paths.

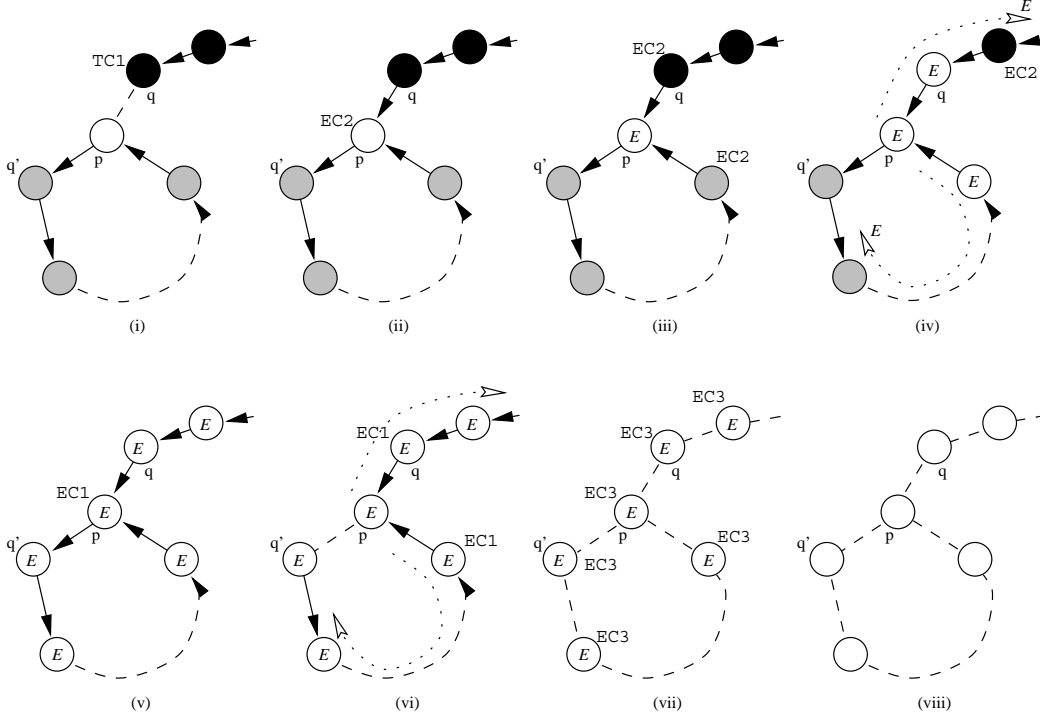


Figure 3.4: Cycles Destruction.

We illustrate the strategy to destroy strict cycle with the Figure 3.4. The basic idea is to detect the cycle by a processor which does not belong to the cycle (that idea was presented for the first time in [15]). The coloring scheme is designed such that during a *round* 1 colored. (resp 0 colored), only processors that have the 1 should have a child (resp. color 0). Therefore, if in *round* of color 1 (resp. color 0), a processor having a child and not having the 1 color is faulty and can be detected as faulty by the legal leaf: the leaf will choose it as child. In the configuration in Step (i), p belongs to a cycle. The grey processors in the figure can have any color. Assume that all processors of the cycle are 0 colored. During a 1-colored *round*, a processor of the cycle (p) is eventually chosen as a descendant by the legal leaf (as it is explained above). Let q be the legal leaf. (this is shown in Step (ii)). After being chosen as a child, p detects that it has more than one ancestor ($EDetectB(p)$ is true) and executes $EC2$ to become a E -colored processor (Step (iii)). The key point of our strategy is that the color E is propagated from a descendant to its ancestors. In our example, for each ancestor q of p , $EDetectA(q)$ is true. The ancestors of the E colored processors execute $EC2$ to propagate the color E along all the paths q belongs to (Steps (iii) and (iv)).

Since p belongs to a cycle, its descendant (q' in the figure) is eventually E -colored (Step (v)). p then satisfies $BrkB$, executes $EC1$, and detaches q' to break the cycle (Step (vi)). Now, the cycle

is replaced by two rooted paths that end at a dead leaf. Next, for every ancestor $q \neq r$ of p , either $BrkD(q)$ or $BrkC(q)$ becomes true depending on if q is an illegal root or not a root, respectively. Processors in a dead rooted path will eventually be E -colored by executing $EC2$. The ancestor of the dead leaf satisfies either $BrkD$ or $BrkC$, and eventually executes Action $EC1$ (Step (vii)). Finally, the E -colored, path-free processors are made 0-colored by Action $EC3$ (Step (viii)).

It is easy to observe that the rooted cycles can be destroyed using the same mechanism as above (Step (ii) and the following steps).

Finally, the protocol must destroy all illegal rooted paths. If a rooted path μ_p is a rooted cycle, it is destroyed using the cycle destruction mechanism described above. Otherwise, it has a dead (E -colored) or a live (not E -colored) leaf. In the first case, μ_p is self-destroyed as above (Step (vi) and the following steps). In the second case, μ_p is self-destroyed by just allowing the token to circulate. The macro $Search_p$ uses the local ordering among the neighbors of a processor to avoid repeating a path during a token circulation round. Thus, a token circulation round cannot loop, and hence, eventually terminates. Once the current token circulation round completes, the processors which were in the illegal path become path-free. Moreover, the illegal root cannot initiate a new token circulation, i.e., cannot create a new path because only the root can initiate a token circulation round.

4. Correctness of the Token Passing Protocol \mathcal{TC}

We define the legitimacy predicate, $\mathcal{L}_{\mathcal{TC}}$ as follows: A configuration γ satisfies $\mathcal{L}_{\mathcal{TC}}$ if it is reachable from δ_{c0} , i.e.,

$$\mathcal{L}_{\mathcal{TC}} \equiv \delta_{c0} \rightsquigarrow \gamma.$$

Note that we could also define $\mathcal{L}_{\mathcal{TC}}$ in terms of δ_{c1} since $\delta_{c0} \rightsquigarrow \delta_{c1}$ by Actions TC1 and TC2.

We apply the convergence stair method [12] to prove the closure and convergence of our protocol. We exhibit a finite sequence of state predicates $\mathcal{A}_0, \mathcal{A}_1, \dots, \mathcal{A}_m$, of Protocol \mathcal{TC} such that the following conditions hold:

- (i) $\mathcal{A}_0 \equiv \text{true}$ (meaning any arbitrary state)
- (ii) $(\mathcal{A}_m \equiv \mathcal{L}_{\mathcal{TC}}) \vee (\mathcal{A}_m \vdash \mathcal{L}_{\mathcal{TC}})$
- (iii) $\forall j : 0 \leq j < m :: \mathcal{A}_j \triangleright \mathcal{A}_{j+1}$

The proof outline is as follows:

In Section 4.1, we show that eventually no processor has the root as a descendant. Then, we prove that a locked processor (which never executes any action) cannot be the root, and either it does not have any descendant, or it belongs to a strict cycle (Section 4.2). This last result trivial leads to the proof of liveness of the algorithm and is also used to prove that all illegal live rooted paths are eventually destroyed (Section 4.3). Once no illegal live rooted path exists, the system contains only one token. In Section 4.4, we show that since the root changes its color infinitely often, the legal path will be eventually colored with the color of the root. Then in Section 4.5, we prove that all cycles are eventually detected and destroyed. Finally, in Section 4.6, we prove that the system reaches a configuration which satisfies $\mathcal{L}_{\mathcal{TC}}$, and Protocol \mathcal{TC} is self-stabilizing.

4.1. Root Without An Ancestor

In this section, we show that the system trivially reaches a configuration in which r does not have any ancestor.

We define $\mathcal{A}_1 \equiv (\forall p \in V \setminus \{r\} : D_p \neq r)$.

Theorem 4.1. $\mathcal{A}_0 \triangleright \mathcal{A}_1$.

Proof: \mathcal{A}_1 is closed: The root r can not be chosen as a descendant by a process $p \neq r$ (see in macro UV_p). Hence, $\#Anc_r$ cannot increase.

Every computation leads to \mathcal{A}_1 : $(\forall p \in V, \forall \alpha \in \mathcal{C} :: D_p = r) \Rightarrow \text{Enable}(EC1, p, \alpha)$. p executes $EC1$ in the configuration α or $\text{Enable}(EC1, p, \beta)$ where $\alpha \mapsto \beta$. By fairness, $\exists \beta : \alpha \rightsquigarrow \beta$ such that p executes $EC1$. Hence, $\#Anc_r$ decreases. Since $\#Anc_r$ cannot increase, $\exists \gamma \in \mathcal{C} : \alpha \rightsquigarrow \gamma :: Anc_r = \emptyset$. \square

4.2. Properties of Locked Processors

We need the following term throughout this section:

A processor p is said to be *Locked* in a configuration α , if in all configurations reachable from α , C_p and D_p remain constant. Formally:

$$\text{Locked}(p, \alpha) \equiv (\forall \beta : \alpha \rightsquigarrow \beta :: (C_{p_\alpha} = C_{p_\beta}) \wedge (D_{p_\alpha} = D_{p_\beta})), \text{ where } V_{p_\gamma} \text{ denotes the value of } V_p \text{ in the configuration } \gamma$$

Since the daemon is weakly fair, $\text{Locked}(p, \alpha)$ implies that p is not continuously *enabled* in all configurations reachable from α and that p never executes an action.

We first prove that if a processor q is the descendant of a *Locked* processor p , then q is eventually *Locked* (Lemma 4.2). Then we establish that if p is not the root r , then q has a descendant and will maintain the descendant forever (Lemma 4.3).

Lemma 4.2. $\forall p, q \in V, \forall \alpha \vdash \mathcal{A}_1 : (\text{Locked}(p, \alpha) \wedge (D_p = q)) \Rightarrow (\exists \beta : \alpha \rightsquigarrow \beta :: \text{Locked}(q, \beta))$.

Proof: We will prove this by contradiction. We assume the contrary, i.e., $\exists p, q \in V, \exists \alpha \vdash \mathcal{A}_1 : (\text{Locked}(p, \alpha) \wedge (D_p = q)) \wedge (\forall \beta : \alpha \rightsquigarrow \beta :: \neg \text{Locked}(q, \beta))$. Thus, in all computations starting from α , q executes an action infinity often.

As p is locked, $\forall \beta : \alpha \rightsquigarrow \beta :: p \in Anc_q$.

1. Assume that $\exists \beta : \alpha \rightsquigarrow \beta :: C_q = E$. Since $C_q = E$, only the guards of Actions $EC1$ and $EC3$ can be true. Thus, q can execute only $EC1$ or $EC3$ at β or β' , where $\beta \rightsquigarrow \beta'$.
 - a. Assume that $D_q = \perp$ at β . Then, q can execute only $EC3$ in β . Anc_q does not increase while q does not execute $EC3$ because no neighbor of q can select q (see macro UV_p). Since q is not locked, $\exists \beta' : \beta \rightsquigarrow \beta' :: \text{Enable}(EC3, q, \beta')$. Since $p \in Anc_q$, $\text{Enable}(EC3, q, \beta')$ implies $[p = r \text{ and } Anc_q = \{r\}]$ in β' . Execution of $EC3$ makes $[C_p := C_r \text{ and } \text{Enable}(TC2, r, \beta')]$ ($\text{Backtrack}(r)$ is true). In all β'' such that $\beta' \rightsquigarrow \beta''$, $\text{Enable}(TC2, p, \beta'')$ since p does not execute an action. By fairness, p eventually runs $TC2$ which contradicts the assumption, $\text{Locked}(p, \alpha)$.

- b. Assume that $D_q \neq \perp$ in β . Then, q can execute $EC1$ only in β . Since q is not locked, $\exists \beta' : \beta \rightsquigarrow \beta'$ such that q eventually executes $EC1$ in β' . After the execution of $EC1$ by q , $D_q = \perp$. Thus, we arrive at the assumed state of Case 1a.
2. Assume that $\forall \beta : \alpha \rightsquigarrow \beta :: C_q \neq E$. So, $\forall \beta : \alpha \rightsquigarrow \beta$, $Enable(EC3, q, \beta)$ is not satisfied. Similarly, $Enable(EC1, q, \beta)$ is not satisfied ($C_q \neq E$ and $\sharp Anc_q > 0$ because $p \in Anc_q$). Furthermore, q cannot execute $EC2$ because otherwise, $\exists \beta : \alpha \rightsquigarrow \beta$ such that $C_q = E$, which contradicts the assumption. Therefore, $\forall \beta : \alpha \rightsquigarrow \beta$, q cannot execute $EC1$, $EC2$, and $EC3$.
 - a. Assume that q executes $TC2$ infinitely often. Since D_q strictly increases with respect to \succ_q (see macro $Search_p$), $\exists \beta : \alpha \rightsquigarrow \beta :: D_q = \perp$. In this case, $\forall \beta' : \beta \rightsquigarrow \beta'$, $Enable(TC2, p, \beta')$, $Enable(EC1, p, \beta')$, or $Enable(EC2, p, \beta')$ depending on $\sharp Anc_p$ and C_p . By fairness, p eventually executes $TC2$, $EC1$, or $EC2$, which contradicts the assumption, $Locked(p, \alpha)$.
 - b. Assume that q executes $TC2$ a finite number of times only. So, $\exists \beta : \alpha \rightsquigarrow \beta$ after which q executes only $TC1$. But, after the execution of $TC1$, $C_q = C_p$. If $D_q = \perp$ then, $Enable(TC2, p, \beta')$, $Enable(EC1, p, \beta')$, or $Enable(EC2, p, \beta')$; by fairness, p eventually executes an action which contradicts the assumption, $Locked(p, \alpha)$. If $D_q \neq \perp$, then q can only execute $TC2$, which contradicts the hypotheses.

□

Lemma 4.3. $\forall p, q \in V, \forall \alpha \vdash \mathcal{A}_1 :$

$(Locked(p, \alpha) \wedge (D_p = q) \wedge (p \neq r)) \Rightarrow (\exists \beta : \alpha \rightsquigarrow \beta : \forall \beta' : \beta \rightsquigarrow \beta' :: D_q \neq \perp)$

Proof: By Lemma 4.2, $\exists \beta : \alpha \rightsquigarrow \beta :: Locked(q, \beta)$. So, $\forall \beta' : \beta \rightsquigarrow \beta'$, D_q remains unchanged. Assume $D_q = \perp$ in β . We will consider the following cases to arrive at the contradiction.

1. Assume that $C_p = E$. Then, irrespective of the color of q , $\forall \beta' : \beta \rightsquigarrow \beta' :: Enable(EC1, p, \beta')$ ($C_p = E$ and $D_{D_p} = \perp$). By fairness, p eventually executes $EC1$. But, that is not possible since p is locked.
2. Assume that $C_p \neq E$.
 1. Assume that $C_q = E$. Then $\forall \beta' : \beta \rightsquigarrow \beta'$, $Enable(EC1, p, \beta')$ or $Enable(EC2, p, \beta')$ (depending on the value of $\sharp Anc_p$). By fairness, p eventually executes either $EC1$ or $EC2$, which is not possible since p is locked.
 2. Assume that $C_q = C_p \neq E$. Then, $\forall \beta' : \beta \rightsquigarrow \beta'$, $Enable(TC2, p, \beta')$, $Enable(EC1, p, \beta')$, or $Enable(EC2, p, \beta')$ (if $\sharp Anc_p = 1, 0$, or > 1 , respectively). By fairness, p eventually executes $TC2$, $EC1$, or $EC2$, which is not possible since p is locked.
3. Assume that $C_q = (C_p + 1) \bmod 2$. Then either $[\exists \beta' : \beta \rightsquigarrow \beta' :: \sharp Anc_q > 1]$, or $[\forall \beta' : \beta \rightsquigarrow \beta' :: Anc_q = \{p\}]$. If $\sharp Anc_q = 1$, then $Enable(TC1, q, \beta')$. If $\sharp Anc_q > 1$, $Enable(EC2, q, \beta')$. By fairness, q eventually executes either $TC1$ or $EC2$, both of which are not possible since q is locked.

□

Now, we will prove the main results of this section. We first establish that a locked processor cannot be in a rooted cycle (Lemma 4.4). We prove this by contradiction. If p belongs to a rooted cycle, then a processor q in that cycle has at least two ancestors. The processor q will eventually get the color E and the color E will be propagated along the cycle in the direction from the descendants towards the ancestors. So, q 's descendant will also be E -colored. Then, q will break the cycle. By induction, every processor in the rooted cycle will eventually detach its descendant. Thus, p is not *Locked*.

Then we show that a locked processor ($\neq r$) has no descendant or it belongs to a strict cycle (Theorem 4.5). Finally, we show that the root cannot be locked (Theorem 4.6), which implies the liveness of the algorithm.

Lemma 4.4. $\forall p \in V, \forall \alpha \vdash \mathcal{A}_1 : \text{Locked}(p, \alpha) \Rightarrow p \text{ does not belong to a rooted cycle.}$

Proof: Let $\mu_q = (p_1, p_2, \dots, p_{l-1}, p_l)$ be a rooted cycle where $p \in \mu_q$ and $D_{p_l} = p_i$, $i \in [2, l[$ ($\#Anc_{p_i} > 1$). We will prove this lemma by contradiction by assuming the contrary, i.e., there is a processor p such that $\text{Locked}(p, \alpha)$ is true and p belongs to the rooted cycle μ_q .

1. Assume that $p = p_i$ and $C_{p_i} \neq E$. Then $EDetect(p_i)$ is true. Thus, $Enable(EC2, p_i, \alpha)$ and $\forall \beta : \alpha \rightsquigarrow \beta$, $Enable(EC2, p_i, \beta)$ remains true because no action allows the ancestor of p_i to detach it because $D_p \neq \perp$. By fairness, p_i eventually executes $EC2$ which contradicts the assumption (p is *Locked*).
2. Assume that $p = p_i$ and $C_p = E$. Since p is *Locked* (according to our assumption), by Lemma 4.2, all processors $p_j, j \in [i, l]$, also are *Locked*.
 - a. Assume that $\exists j \in]i, l] :: C_{p_j} \neq E, C_{p_{j+1}} = E$. Then $EDetect(p_j)$ is true. Thus, $Enable(EC2, p_j, \alpha)$, and $\forall \beta : \alpha \rightsquigarrow \beta$, $Enable(EC2, p_j, \beta)$ remains true while p_j does not execute $EC2$, which contradicts the assumption, p_j is *Locked*.
 - b. Assume that $\forall j \in [i, l] :: C_{p_j} = E$. Then $Break(p_i)$ is true. Thus, $Enable(EC1, p_i, \alpha)$ is true and remains true while p_i does not execute $EC1$, which contradicts our assumption, p_i is *Locked*.
3. From Cases 1, 2a, and 2b, $p \neq p_i$, i.e., p_i cannot be *Locked*. Assume that $p = p_j, j \in]i, l]$. Then by Lemma 4.2, all processors $p_k, k \in [i, l]$, are also *Locked*, which contradicts the fact that p_i is not *Locked*. Thus, all processors $p_j, j \in [i, l]$, are not *Locked*.
4. Assume that $p = p_j, j \in [1, i-1]$. Then by Lemma 4.2, all processors $p_k, k \in [j+1, l]$, are also *Locked*, which is not possible according to Case 3.

□

Theorem 4.5. $\forall p \in V, \forall \alpha \vdash \mathcal{A}_1 :$
 $(\text{Locked}(p, \alpha) \wedge (p \neq r)) \Rightarrow ((D_p = \perp) \vee (p \text{ belongs to a strict cycle})).$

Proof: Assume that $D_p = q$. By Lemmas 4.2 and 4.3, $\exists \beta : \alpha \rightsquigarrow \beta :: \text{Locked}(q, \beta)$ and $D_q \neq \perp$. By induction, the descendant of q will also be eventually locked, and so on. Since the graph S is finite, p belongs to a cycle. By Lemma 4.4, p cannot belong to a rooted cycle. Thus, p belongs to a strict cycle. \square

Theorem 4.6. $\forall p \in V, \forall \alpha \vdash \mathcal{A}_1 : \text{Locked}(p, \alpha) \Rightarrow (p \neq r)$

Proof: We will prove by contradiction. Assume that $\exists \alpha \vdash \mathcal{A}_1 : \text{Locked}(p, \alpha) \wedge p = r$.

1. Assume that r has no descendant. Then, $\forall \beta : \alpha \rightsquigarrow \beta :: \text{Enable}(\text{TC1}, r, \alpha)$. Thus, by fairness, r is not *Locked*.
2. Assume that r has a descendant, q . Then by Lemma 4.2, q is also *Locked*. So, by Theorem 4.5, q is either inside a strict cycle or $D_q = \perp$.
 - a. $D_q = \perp$.
 - (i) Assume that $\sharp \text{Anc}_q = 1$ ($\text{Anc}_q = \{r\}$). If $C_q = E$, then $\forall \beta : \alpha \rightsquigarrow \beta :: \text{Enable}(\text{EC3}, q, \beta)$. So, by fairness, q is not *Locked*. If $C_q \neq E$, then $\forall \beta : \alpha \rightsquigarrow \beta :: \text{Enable}(\text{TC1}, q, \beta) \vee \text{Enable}(\text{TC2}, r, \beta)$. Thus, by fairness, either q or r is not *Locked*.
 - (ii) Assume that $\sharp \text{Anc}_q > 1$, and $\exists \beta : \alpha \rightsquigarrow \beta :: \sharp \text{Anc}_q = 1$. Thus, $\text{Anc}_q = \{r\}$ and by Case 2a(i), this is not possible. So, $\forall \beta : \alpha \rightsquigarrow \beta :: \sharp \text{Anc}_q > 1$. If $C_q \neq E$, then $\forall \beta : \alpha \rightsquigarrow \beta :: \text{Enable}(\text{EC2}, q, \beta)$ and q is not *Locked*. If $C_q = E$, then the ancestors of q ($\neq r$) can only execute Actions *EC1* or *EC2* until q remains their descendent. These ancestors of q can execute *EC2* at most once (to get the color E). After the execution of *EC2*, the ancestors can only execute *EC1*. Because $C_q = E$ and $D_q = \perp$, it cannot get a new ancestor. Thus, after repeated execution of *EC1*, eventually, q will have no ancestors except r . This contradicts the assumption, $\sharp \text{Anc}_q > 1$.
 - b. q has a descendant and is inside a strict cycle. Since $\forall \alpha \vdash \mathcal{A}_1$, r has no ancestor, q must belong to a rooted cycle (Theorem 4.1), which contradicts the assumption.

\square

Corollary 4.7 (Liveness). *In any configuration $\vdash \mathcal{A}_1$, at least one processor is enabled.*

4.3. Destruction of Live Illegal Rooted Paths

In this section, we show that all live illegal rooted paths are destroyed.

Lemma 4.8. *A processor cannot stays forever in an illegal and live rooted path.*

Proof: Assume the contrary, i.e., a processor p stays forever in an illegal rooted path. By Theorem 4.5, p is not *Locked*. Till p is in an illegal, live rooted path, p can perform one time the action *TC1*, δ_p times the action *TC2* and one time the action *TE2*. If p stays in an illegal live rooted path p is eventually *Locked*. \square

Let us denote the number of live illegal leaves by *LIL*. In the next Lemma (Lemma 4.9), we prove that Algorithm \mathcal{TC} cannot create a new live illegal leaf.

Lemma 4.9. $\forall \alpha \vdash \mathcal{A}_1, \forall \beta$ such that $\alpha \mapsto \beta$, the value of *LIL* in β is less than or equal to the value of *LIL* at α .

Proof: Assume the contrary, i.e., *LIL* in β is greater than *LIL* in α . Then one of the following is true: (1) A dead illegal leaf becomes a live illegal leaf, (2) A path is broken creating a live illegal leaf, and (3) A processor, other than the root, becomes the root of an illegal rooted path.

1. For any p such that $C_p = E$, only *EC3* changes C_p . If p executes *EC3*, then one of the following two conditions must be true: (i) $Anc_p = \emptyset$ and p is not a leaf, and (ii) $Anc_p = \{r\}$ and p is not a leaf of an illegal path. Both (i) and (ii) contradict our assumption.
2. In order to break a path $\mu_p = (p_1, p_2, \dots, p_l)$ so that a live leaf is created, $\exists p_i \in \mu_p$ such that p_i executes an action in α and p_i becomes a live leaf in β . Since, $D_{p_i} \neq \perp$, p_i can execute only *TC2*, *EC1*, and *EC2* in α .

If p_i executes *EC2*, then C_{p_i} becomes equal to E . So, p_i is not a live leaf.

If $Enable(EC1, p_i, \alpha)$, then $C_{p_i} = E$ in α because $\sharp Anc_{p_i} > 0$. Since the execution of *EC1* does not change the color, p_i cannot become a live leaf.

If $Enable(TC2, p_i, \alpha)$, then $D_{p_{i+1}} = \perp$ in α . Thus, after the execution of *TC2*, p_i becomes a live leaf, but p_{i+1} is no more a leaf.

3. A processor, $p \neq r$, without an ancestor, cannot select a new descendant because both *Forward*(p) and *Backward*(p) are disabled at p .

We proved the contradiction in all three cases. □

We define $\mathcal{A}_2 \equiv \mathcal{A}_1 \wedge (LIL = 0)$.

Theorem 4.10. $\mathcal{A}_1 \triangleright \mathcal{A}_2$.

Proof: Assume that the value of *LIL* eventually stays constant to a value $l > 0$. *LIL* cannot increase (lemma 4.9). An illegal live leaf belongs to at least one illegal live rooted path. According to the lemma 4.8, no processor stay forever in an illegal live path. Thus these paths infinitely often lose processors by the action *TC2* or *TE1* and get processors by the action *TC1* or *TC2*. To ensure that *LIL* never decreases, when one of the illegal live paths (by the action *TC1* or *TC2*) get a new processor either (1) this processor belongs to legal live path or (2) this processor has no child, no parent, and does not have the E color.

After the legal path is reduced to r , all processors in the legal live path have only one parent and does not have the E color until (1) the leaf of the legal live path q joins an illegal live rooted path by the action *TC1* or *TC2*, or (2) an illegal leaf chooses a processor t of the legal live path. Notice that in the second case, all processors between t and the legal leaf does not have the E color. In both cases, the legal live leaf is also an illegal live leaf. Therefore, no illegal live leaf chooses as child a processor of the legal live path (otherwise *LIL* would decrease). Thus, when the legal path is live, there is at most one processor t of the legal path that has two parents.

An infinitely of processors that have joined an illegal live path by *TC1* or *TC2* action will quit this path by the *TE1* action. Otherwise, the processor in the illegal live rooted path would

eventually perform only the TC1, TC2 and TE2 actions, and the root of this path will become locked (see proof of the lemma 4.8); but that is impossible (lemma 4.8).

Let p be a processor that joins the illegal live rooted path by the TC1 or TC2 actions and then quits the illegal live rooted path by the TE1 action (after that the legal path has been reduced to r). p and its descendants will not have the E color and will have one parent until (1) the live leaf (t) of the legal path chooses one of them (p') as its descendants or until (2) the illegal live leaf of p 's path chooses a processor q' of the legal live path as child. In both cases, all processors of the illegal rooted path between p and the leaf has one parent except p' (resp. q') and all processors of the illegal rooted path between p' (resp. q') and the illegal leaf do not have the E color. Then p and p' (resp. q') stays in the illegal live rooted path until p' (resp. q') becomes a dead leaf. p will not quit the illegal live rooted path by the TE1 action. \square

Corollary 4.11. *In any configuration $\vdash \mathcal{A}_2$, if there exists a live leaf, then it must be the legal leaf.*

4.4. Color Consistency

In this section, we show that eventually, either the system contains no live leaf, or every processor in the legal path (except the leaf) has the same color as r has. In such a configuration, the legal path cannot create a new cycle. We first show (by using Theorem 4.6) that r changes its color infinitely often. So, r starts a new *cround* with a new color infinitely often. If the legal path, μ_r , does not meet any illegal path, then it remains color consistent (all processors, except the leaf, have the same color). Otherwise, when μ_r meets an illegal path, its leaf becomes dead and it remains color consistent.

Lemma 4.12. *The root r changes its color infinitely often.*

Proof: By Theorem 4.6, r executes an action infinitely often. r can execute only TC1 and TC2. If r executes TC1 infinitely often, then r changes its color infinitely often and hence, the lemma is proven. Assume that r does not execute TC1 infinitely often. This implies that r executes TC2 infinitely often (by Theorem 4.6). Then, by the definition of $Search_r$, eventually $D_r = \perp$ must be true. This will enable r to execute TC1, which contradicts our assumption. \square

We define a predicate *ColorConsistent* in a configuration γ such that it is true if any of the following conditions is true: (CC1) $D_r = \perp$. (CC2) The leaf of the legal path is a live leaf and all processors on the legal path, except the leaf, are r -colored (with the same color as r). (CC3) The legal path does not have a leaf, i.e., the path is a rooted cycle, or has a dead leaf.

We define $\mathcal{A}_3 \equiv \mathcal{A}_2 \wedge ColorConsistent$.

Theorem 4.13. $\mathcal{A}_2 \triangleright \mathcal{A}_3$.

Proof: \mathcal{A}_3 is closed:

1. Assume that $D_r = \perp$. By Lemma 4.12, r changes its color infinitely often. r chooses a descendant by executing $TC1$. If r chooses a descendant which belongs to a cycle or to an illegal rooted path with a dead leaf, then *ColorConsistent* remains true (*CC3*). If r selects a path-free descendant p , then p becomes the new live leaf of the legal path, μ_r , and thus, *ColorConsistent* is preserved (*CC2*).
2. Assume that $D_r \neq \perp$. The only processor which can choose a descendant by executing $TC1$, is the live leaf of the legal path. Assume that p is the live leaf. If p chooses a path-free processor as the descendant (executing Action $TC1$ or $TC2$), then all processors except the leaf, are r -colored. Thus, *ColorConsistent* remains true (*CC2*). If p selects a processor q that is in a path, as the descendant, then the legal path ends in a cycle or a dead leaf (Theorem 4.10 and Corollary 4.11). Thus, *ColorConsistent* is preserved (*CC3*).

Every computation starting from a configuration satisfying \mathcal{A}_2 leads to a configuration in \mathcal{A}_3 : The proof follows from Lemma 4.12. \square

4.5. Cycle Destruction

In this section, we prove that all cycles are eventually destroyed. The process of destruction is as follows: All strict cycles are merged with the legal path and thus, become rooted cycles. Then by the repeated application of *EC1* and *EC2*, the rooted cycles will be destroyed.

We borrow the following term from [5] to simplify our presentation: The *first DFS tree* of the graph G is defined as the DFS spanning tree rooted at r , created by traversing the graph in the DFS manner, and visiting the adjacent edges of every processor in the order induced by \succ_p . We defined the macro *Search_p* such that Algorithm \mathcal{TC} circulates the token in the first DFS tree.

Lemma 4.14. *Starting from any configuration $\vdash \mathcal{A}_3$, all processors which do not belong either to the legal path or to any strict cycles, will be eventually path-free.*

Proof:

An illegal path that is not a strict cycle or a live rooted path is a dead rooted path or a rooted cycle. The only actions that a processor in illegal rooted path can perform is one time the action *TE2* and one time the action *TE1*. As this processor is not *Locked*, it will eventually not be anymore in an illegal rooted path. \square

Lemma 4.15. *Starting from any configuration $\vdash \mathcal{A}_3$, every processor which is path-free and E -colored, will be eventually path-free and 0-colored.*

Proof: By fairness, all E -colored and path-free processors eventually execute *EC3* because none of its neighbors can choose it as a descendant (in the macro UV_p , q cannot be chosen if $C_q = E$ and $D_q = \perp$). \square

Lemma 4.16. *Starting from any configuration $\vdash \mathcal{A}_3$, every strict cycle will be eventually transformed into a rooted cycle.*

Proof: By definition, in any configuration of $\vdash \mathcal{A}_3$, there exists no illegal rooted path ending by a live leaf.

So, our responsibility is to show that eventually a node on every strict cycle in the system will be selected as a descendant by the leaf of the legal path. Assume the contrary, i.e., there exists one strict cycle which will never be reached by the legal path.

So, there exists $\alpha \vdash \mathcal{A}_3$ such that all processors between r and the strict cycle on the first DFS-tree are path-free (by Lemma 4.14), or they belong to the legal path (r -colored in \mathcal{A}_3). By Lemma 4.15, $\exists \alpha' : \alpha \rightsquigarrow \alpha'$ such that every processor between r and the strict cycle is 0- or 1-colored. Also, by successive *crowds*, $\exists \alpha'' : \alpha' \rightsquigarrow \alpha''$ such that every processor between r and the strict cycle has the same color k (0 or 1).

Let q be the first processor in the strict cycle that is on the first DFS tree. Let $p \in N_q$ be the ancestor of q in the first DFS tree. Since no strict cycle is reachable by the legal path (by assumption), $C_q = k$. Otherwise, p will eventually select q as a descendant, which will contradict our assumption. But, in the next *crowd*, p will choose q as a descendant because C_p will be equal to $(k + 1) \bmod 2$. Thus, we arrive at the contradiction. \square

Lemma 4.17. *Starting from any configuration $\vdash \mathcal{A}_3$, every cycle is destroyed.*

Proof: By Lemma 4.16, every strict cycle is eventually transformed into a rooted cycle. All rooted cycles are eventually destroyed by the repeated application of *EC1* and *EC2*. \square

Let NC denote the number of cycles in the system.

We define $\mathcal{A}_4 \equiv \mathcal{A}_3 \wedge (NC = 0)$.

Theorem 4.18. $\mathcal{A}_3 \triangleright \mathcal{A}_4$.

Proof: \mathcal{A}_4 is closed: By the definition of *Search*, *Forward*, and Action *TC1*, the leaf of the legal path chooses a descendant of a different color. So, no new cycle can be created in \mathcal{A}_3 . Hence, NC cannot increase.

Every computation starting from a configuration in \mathcal{A}_3 leads to a state in \mathcal{A}_4 : Follows from Lemma 4.17. \square

4.6. Legitimacy Predicate

In this section, we prove that the legitimacy predicate $\mathcal{L}_{\mathcal{TC}}$ eventually holds. Then we show that Protocol \mathcal{TC} is self-stabilizing.

Lemma 4.19 (Single Token Property). $\forall \alpha \vdash \mathcal{A}_4 : \exists ! p :: \text{Token}(p)$.

Proof: Follows from Corollary 4.11. \square

We define the following for a configuration $\gamma \vdash \mathcal{A}_4$:

$\mathcal{A}_5 \equiv \mathcal{A}_4 \wedge \mathcal{L}_{\mathcal{TC}}$.

Theorem 4.20. $\mathcal{A}_4 \triangleright \mathcal{A}_5$.

Proof: \mathcal{A}_5 is closed: Follows from the definition of δ_{c0} and Actions TC1 and TC2.

Every computation leads to \mathcal{A}_5 : By Lemma 4.14, $\exists \alpha' : \alpha \rightsquigarrow \alpha'$ such that all processors in the system are path-free or belong to the legal path. By Lemma 4.15, $\exists \alpha'' : \alpha' \rightsquigarrow \alpha''$ such that every path-free processor is 0- or 1-colored. Then by successive *crowds*, $\exists \beta : \alpha'' \rightsquigarrow \beta$ such that every processor is path-free and has the same color k (0 or 1). If $k = 0$ in β , then the lemma is proven ($\beta = \delta_{c0}$). If $k = 1$, then, in the next round, k becomes equal to 0. \square

Lemma 4.21 (Legitimacy Predicate). $\mathcal{A}_5 \vdash \mathcal{L}_{\mathcal{TC}}$.

Proof: Follows from the definition of \mathcal{A}_5 and $\mathcal{L}_{\mathcal{TC}}$. \square

Lemma 4.22 (Closure and Convergence). $\text{true} \triangleright \mathcal{L}_{\mathcal{TC}}$.

Proof: Follows from Theorems 4.1, 4.10, 4.13, 4.18, and 4.20 and Lemma 4.21. \square

Lemma 4.23 (Correctness). $\forall \alpha \vdash \mathcal{L}_{\mathcal{TC}} : \forall e \in \mathcal{E}_\alpha :: e \vdash \mathcal{SP}_{\mathcal{TC}}$.

Proof: In $\mathcal{L}_{\mathcal{TC}}$, only Actions TC1 and TC2 are executed. Thus, the token is passed among the processors in the depth-first search order during a *crowd*, and any computation is a repetition of *crowds*. In any configuration of $\mathcal{L}_{\mathcal{TC}}$, only one processor may execute an action. So, only one processor has a token in any configuration $\vdash \mathcal{L}_{\mathcal{TC}}$. \square

Theorem 4.24 (Self-Stabilizing). *Protocol \mathcal{TC} is self-stabilizing.*

Proof: Follows from Lemmas 4.22 and 4.23. \square

5. State Complexity

A processor p in Algorithm \mathcal{TC} uses two variables, D_p and C_p . The variable C_p , for a processor $p \neq r$, can have 3 different values (0, 1, and E), whereas C_r can have only 2 values (0 or 1). The variable D_p can have Δ_p ($|N_p|$) plus one (\perp) values. So, a processor, $p \neq r$, needs to maintain $3 \times (\Delta_p + 1)$ states and r needs $2 \times (\Delta_r + 1)$. Thus, the total number of configurations of the whole network is

$$2 (\Delta_r + 1) \times \prod_{p \in V, p \neq r} 3 (\Delta_p + 1)$$

It is worth mentioning here that all the previous papers computed the space complexity in terms of the number of *bits* only, not in terms of the *states*. We feel that the measurement in terms of the number of states is more accurate.

6. Concluding Remarks

We presented a self-stabilizing depth-first token circulation scheme on a general network with a distinguished root. Our algorithm can also be used to construct a DFS spanning tree simply by maintaining the ancestor pointers.

A solution to the problem of mutual exclusion in a network is to implement a token circulating from one processor to the next following some pattern. The token moves around the network. A processor having the token is granted access to the shared resource and can execute the code in the critical section.

Our solution to the depth-first token circulation problem can be used to solve the *mutual exclusion problem*. After stabilization, in Algorithm \mathcal{TC} , in each token circulation round, a processor p holds the token as many times as its degree Δ_p —once while satisfying $Forward(p)$ and $\Delta_p - 1$ times while $Backtrack(p)$ is true. Since the degree of the processors in the network is bounded, Algorithm \mathcal{TC} implements a *strictly fair* token circulation (and mutual exclusion). By *strict fairness*, we mean that, while a processor p is waiting for the token, any other processor q ($q \neq p$) can get the token at most a bounded number of times: here Δ_q times, where Δ_q is the degree of q .

It is also easy to implement the *1-fair* mutual exclusion, i.e., in each token round, all processors will enjoy the critical section access *exactly once*. In this case, a processor p can enter the critical section if and only if $Forward(p)$ is true.

The space requirement for a processor p is $3(\Delta_p + 1)$ states (only $2(\Delta_p + 1)$ states for the root). The question of the optimal state requirement for this problem is still open.

alpha

References

- [1] J Beauquier and O Debas. An optimal self-stabilizing algorithm for mutual exclusion on bidirectional non uniform rings. In *Proceedings of the Second Workshop on Self-Stabilizing Systems*, pages 17.1–17.13, 1995.
- [2] GM Brown, MG Gouda, and CL Wu. Token systems that self-stabilize. *IEEE Transactions on Computers*, 38:845–852, 1989.
- [3] JE Burns, MG Gouda, and RE Miller. On relaxing interleaving assumptions. In *Proceedings of the MCC Workshop on Self-Stabilizing Systems, MCC Technical Report No. STP-379-89*, 1989.
- [4] JE Burns and J Pachl. Uniform self-stabilizing rings. *ACM Transactions on Programming Languages and Systems*, 11:330–344, 1989.
- [5] Z Collin and S Dolev. Self-stabilizing depth-first search. *Information Processing Letters*, 49:297–301, 1994.
- [6] A.K. Datta, C. Johnen, F. Petit, and V. Villain. Self-stabilizing depth-first token circulation in arbitrary rooted networks. In *Structure, Information and Communication Complexity (SIROCCO98)*, 1998.

- [7] EW Dijkstra. Self stabilizing systems in spite of distributed control. *Communications of the Association of the Computing Machinery*, 17:643–644, 1974.
- [8] S Dolev, A Israeli, and S Moran. Self-stabilization of dynamic systems assuming only read/write atomicity. *Distributed Computing*, 7:3–16, 1993.
- [9] M Flatebo, AK Datta, and AA Schoone. Self-stabilizing multi-token rings. *Distributed Computing*, 8:133–142, 1994.
- [10] S Ghosh. An alternative solution to a problem on self-stabilization. *ACM Transactions on Programming Languages and Systems*, 15:735–742, 1993.
- [11] MG Gouda and FF Haddix. The stabilizing token ring in three bits. *Journal of Parallel and Distributed Computing*, 35:43–48, 1996.
- [12] MG Gouda and N Multari. Stabilizing communication protocols. *IEEE Transactions on Computers*, 40:448–458, 1991.
- [13] ST Huang and NS Chen. Self-stabilizing depth-first token circulation on networks. *Distributed Computing*, 7:61–66, 1993.
- [14] C Johnen, G Alari, J Beauquier, and AK Datta. Self-stabilizing depth-first token passing on rooted networks. In *WDAG97 Distributed Algorithms 11th International Workshop Proceedings, Springer-Verlag LNCS:1320*, pages 260–274, 1997.
- [15] C Johnen and J Beauquier. Space-efficient distributed self-stabilizing depth-first token circulation. In *Proceedings of the Second Workshop on Self-Stabilizing Systems*, pages 4.1–4.15, 1995.
- [16] HSM Kruijer. Self-stabilization (in spite of distributed control) in tree-structured systems. *Information Processing Letters*, 8:91–95, 1979.
- [17] F. Petit. Highly space-efficient self-stabilizing depth-first token circulation for trees. In *OPODIS97 Proceedings of the first international conference On Principles Of DISTRIBUTED Systems*, pages 221–235. Hermes, 1997.
- [18] F Petit and V Villain. Color optimal self-stabilizing depth-first token circulation. In *I-SPAN’97, Third International Symposium on Parallel Architectures, Algorithms and Networks Proceedings, IEEE Computer Society Press*. IEEE Computer Society Press, 1997. To appear.
- [19] F Petit and V Villain. Color optimal self-stabilizing depth-first token circulation protocol for asynchronous message-passing. In *PDCS-97 10th International Conference on Parallel and Distributed Computing Systems Proceedings*, pages 227–233. International Society for Computers and Their Applications, 1997.
- [20] F Petit and V Villain. A space-efficient and self-stabilizing depth-first token circulation protocol for asynchronous message-passing systems. In *Euro-par’97 Parallel Processing, Proceedings LNCS:1300*, pages 476–479. Springer-Verlag, 1997.

- [21] V. Villain. A new lower bound for self-stabilizing mutual exclusion algorithms. Technical Report RR97-17, Universite' de Picardie, Jules Vernes, LaRIA, 1997.