Self-Stabilizing Labeling and Ranking in Ordered Trees

Ajoy K. Datta[1]   Stéphane Devismes[2]   Lawrence L. Larmore[1]

Yvan Rivierre[2]

[1] School of Computer Science, University of Nevada Las Vegas

[2] VERIMAG UMR 5104, Université Joseph Fourier

Abstract

We propose two self-stabilizing algorithms for tree networks. The first one computes a special label, called *guide pair* of each process P in $O(h)$ rounds (h being the height of the tree) using $O(\delta_P \log n)$ space per process P, where δ_P is the degree of P and n the number of processes in the network. Guide pairs have numerous applications, including ordered traversal or navigation of the processes in the tree. Our second self-stabilizing algorithm, which uses the guide pairs computed by the first algorithm, solves the *ranking problem* in $O(n)$ rounds and has space complexity $O(b + \delta_P \log n)$ in each process P, where b is the number of bits needed to store a value. The first algorithm orders the tree processes according to their topological positions. The second algorithm orders (ranks) the processes according to the values stored in them.

Keywords: Self-stabilization, tree networks, tree labeling, ranking problem.

1 Introduction

Self-stabilization [Dij74, Dol00] is a versatile property, enabling an algorithm to withstand transient faults in a distributed system. A distributed algorithm is self-stabilizing if, after transient faults hit the system and place it in some arbitrary global state, the system recovers without external intervention in finite time.

An *ordered tree* $\mathcal{T}$ is a rooted tree, together with an order (called a left-to-right order) on the children of every node. In this paper, we give two self-stabilizing distributed algorithms for ordered trees. None of the two algorithms assumes knowledge of the size of the network n, or of a known upper bound of n, although, as it is usual in the literature, we assume that each process can store an integer in the range $1..n$, using $O(\log n)$ space. We choose the ordered tree topologies because results in such topologies can be easily extended to arbitrary rooted networks by composing our solutions with any existing self-stabilizing spanning tree construction algorithm (see [Dol00] for the literature). However, the meaning of "traversing" or "ranking" processes in a general network is not clear.

Our first algorithm, GUIDE, computes a *guide pair* for each process P, which we write as $P.guide = (P.pre\_ind, P.post\_ind)$, where $P.pre\_ind$ and $P.post\_ind$ are the rank of P in the *preorder* and *reverse postorder* traversal, respectively, of the ordered tree. Figure 1 shows an example of ordered tree labeled with guide pairs. The guide pairs provide a labeling scheme that can be used for various applications [FEP+06]. In this work, we use these labels to navigate in the tree $\mathcal{T}$. We can define a partial ordering on the guide pairs as follows: We say $(i, j) \leq (k, \ell)$ if $i \leq k$ and $j \leq \ell$. Then, A process Q is a member of the subtree $\mathcal{T}_P$ rooted at P if and only if $P.guide \leq Q.guide$. The guide pairs can be used to implement routing between any two processes of the tree. If the two nodes satisfy the above partial ordering, then the routing path simply follows the list of ancestors/descendants. Otherwise, the routing must be established via the nearest common ancestor.

Our second algorithm, RANK, uses GUIDE, hence shows another application of guide pairs. The input of our second algorithm is a value $P.weight$, of some ordered type, for each process P. RANK computes the *rank* of each process, which is defined to be the index of that process if all processes were sorted by their weights.

Figure 1: Guide pairs.

1.1 Contributions

GUIDE has time complexity $O(h)$ rounds, where h is the height of $\mathcal{T}$. The time complexity of RANK is $O(n)$ rounds. The space complexity of GUIDE in each process P is $O(\delta_P \log n)$, where δ_P is the degree of P. RANK, which uses GUIDE as a subroutine, has space complexity $O(b + \delta_P \log n)$ in each process P, where b is the number of bits needed to store a value. GUIDE and RANK are self-stabilizing. GUIDE is *silent*, that is, it eventually reaches a terminal configuration where all actions of all processes are disabled. RANK correctly computes the rank of every process within $O(n)$ rounds. Unless the weights change, the ranking do not change once the system stabilizes. However, the algorithm repeatedly computes them to detect possible change of weights. If the weights do not change, the repeated computation of RANK will be transparent to the application that uses the output of RANK.

1.2 Related Work

The notion of guide pairs appeared first in [FEP+06], but that solution is not self-stabilizing. To the best of our knowledge, there exist no self-stabilizing algorithms for computing the guide pairs.

The only self-stabilizing solution to the ranking problem was given in [BDN95]. This algorithm works in rooted trees. Like ours, that algorithm is not silent. Moreover, it assumes that each process has a unique identifier in the range 1..n. The algorithm stabilizes in $\Omega(n^2)$ rounds using $O(\log n)$ space per process. The ranking problem is related to the *sorting problem*. There exist numerous self-stabilizing solutions to sorting in a tree, *e.g.*, [HP01, HM01, BDV05]. However, all those previous problems are quite different than ours.

1.3 Roadmap

In the next section, we present the model we use throughout this paper. In Section 3, we present our self-stabilizing silent algorithm for computing guide pairs. In Section 4, we present our self-stabilizing algorithm for the ranking problem, which uses the guide pairs.

2 Preliminaries

Let $G = (V, E)$ be an undirected graph, where V is a set of nodes and E is a set of undirected edges linking nodes. Two nodes $P, Q \in V$ are said to be neighbors if $\{P, Q\} \in E$. The set of P's neighbors is denoted by $N(P)$. The degree of P *i.e.,* $|N(P)|$, is denoted by δ_P. $G = (V, E)$ is a *tree* if it is connected and acyclic. A

tree $\mathcal{T}$ can be *rooted* at some node, meaning that one of its nodes *Root* is distinguished as the *root* (all other nodes are anonymous). In a rooted tree $\mathcal{T}$, we denote by *P.par*, the parent of node P in $\mathcal{T}$: If $P = Root$, then *P.par* $= P$; otherwise *P.par* $= Q$, where Q is the neighbor of P that is the closest from the root (in this case, P is said to be a *child* of Q in $\mathcal{T}$). Let *Chldrn*$(P) = \{Q \in N(P) : Q.par = P\}$, the *children* of P in the tree $\mathcal{T}$. An *ordered tree* is a rooted tree $\mathcal{T}$, together with an (local) order (called a left-to-right order) on the children of every node. We denote by $\prec_P$ the local order relation among the children of node P. Let $P_1, P_2, \ldots P_m$ be the children of the root of $\mathcal{T}$ in the left-to-right order. We denote by $\mathcal{T}_i$ be the subtree rooted at any P_i. Finally, we denote by $Q \in \mathcal{T}_i$ the fact that node Q is a node of $\mathcal{T}_i$.

We model our network topology as an ordered tree $\mathcal{T} = (V, E)$, where V is a set of n nodes representing processes and E is a set of edges, each representing the ability of two processes to communicate directly. (We will use the terms "node" and "process" interchangeably.) We denote by $h(P)$ the height of process P in $\mathcal{T}$, *i.e.*, its distance to the root. We denote by h the height of $\mathcal{T}$, *i.e.*, $\max_{P \in V} h(P)$.

2.1 Computational Model

We consider the locally shared memory model, introduced by Dijkstra [Dij74]. In this model, communications are carried out by locally shared variables. Each process has the finite set of shared variables (henceforth, referred to as variables) whose domains are finite. A process P can read its own variables and that of its neighbors, but can write only to its own variables. We assume that every process P can read the local names of its neighbors, so that if $Q \in N(P)$, P can tell, for example, whether *Q.par* $= P$. Each process writes its variables according to its (local) *program*. A *distributed algorithm* is a collection of n *programs*, each one operating on a single process. The *program* of each process is a finite set of actions $\langle label \rangle :: \langle guard \rangle \mapsto \langle statement \rangle$. *Labels* are only used to identify actions in the discussion. The *guard* of an action in the program of a process P is a Boolean expression involving the variables of P and its neighbors. The *statement* of an action of P updates one or more variables of P. An action can be executed only if it is *enabled*, *i.e.*, its guard evaluates to *true*. A process is said to be *enabled* if at least one of its actions is enabled.

Let $\mathcal{A}$ be a distributed algorithm operating of a network of topology G. The values of $\mathcal{A}$'s variables at some process P define $\mathcal{A}$'s *(local) state* of P in G. A configuration of $\mathcal{A}$ in G is an instance of $\mathcal{A}$'s states of all processes in G. In the following, if there is no ambiguity, configurations of $\mathcal{A}$ in G will be simply denoted by *configurations*.

Let $\mapsto$ be the binary relation over configurations of $\mathcal{A}$ in G such that $\gamma \mapsto \gamma'$ if and only if it is possible for the network of topology G to change from configuration γ to configuration γ' in one step of $\mathcal{A}$. An *execution* of $\mathcal{A}$ is a maximal sequence of configurations $\varrho = \gamma_0 \gamma_1 \ldots \gamma_i \ldots$ such that $\gamma_{i-1} \mapsto \gamma_i$ for all $i > 0$. The term "maximal" means that the execution is either infinite, or ends at a *terminal* configuration in which no action of any process is enabled. Each step $\gamma_i \mapsto \gamma_{i+1}$ consists of one or more *enabled* processes executing an action. The evaluations of all guards and executions of all statements of those actions are presumed to take place in one atomic step; this model is called *composite atomicity* [Dol00].

We assume that each step from a configuration to another is driven by a *scheduler*, also called a *daemon*. If one or more processes are enabled, the scheduler selects at least one of these enabled processes to execute an action. We assume that the scheduler is *weakly fair*, meaning that, every continuously enabled process P is selected by the scheduler within finite time.

We say that a process P is *neutralized* in the step $\gamma_i \mapsto \gamma_{i+1}$ if P is enabled in γ_i and not enabled in γ_{i+1}, but does not execute any action between these two configurations. The neutralization of a process represents the following situation: at least one neighbor of P changes its state between γ_i and γ_{i+1}, and this change effectively makes the guard of all actions of P false.

We use the notion of *round*. The first *round* of an execution ϱ, noted ϱ', is the minimal prefix of ϱ in which every process that is enabled in the initial configuration either executes an action or becomes neutralized. Let ϱ'' be the suffix of ϱ starting from the last configuration of ϱ'. The second round of ϱ is the first round of ϱ'', the third round of ϱ is the second round of ϱ'', and so forth.

2.2 Self-stabilization and Silence

In the following, we define a *specification* as a set of executions. We said that an execution ϱ satisfies the specification SP if $\varrho \in SP$.

A distributed algorithm $\mathcal{A}$ is *self-stabilizing with respect to* the specification SP in a network of topology G if and only if there exists a set of configurations $\mathcal{C}$ such that:

1. Every execution of $\mathcal{A}$ in a network of topology G starting from a configuration in $\mathcal{C}$ satisfies SP (*closure*).

2. Every execution of $\mathcal{A}$ in a network of topology G eventually reaches a configuration in $\mathcal{C}$ (*convergence*).

All configurations of $\mathcal{C}$ are said to be legitimate, all other configurations are said to be illegitimate.

We say that an algorithm is *silent* [DGS96] if each of its executions is finite. In other words, starting from an arbitrary configuration, the network will eventually reach a configuration where no process is enabled.

2.3 Composition

To simplify the design of our algorithms, we use a variant of the well-known *collateral composition* [Tel00]. Roughly speaking, when we collaterally compose two algorithms $\mathcal{A}$ and $\mathcal{B}$, $\mathcal{A}$ and $\mathcal{B}$ run concurrently and $\mathcal{B}$ uses the outputs of $\mathcal{A}$ in its executions. In the variant we use, we modify the code of $\mathcal{B}$ so that a process executes an action of $\mathcal{B}$ only when it has no enabled action in $\mathcal{A}$.

Let $\mathcal{A}$ and $\mathcal{B}$ be two algorithms such that no variable written by $\mathcal{B}$ appears in $\mathcal{A}$. The *hierarchical collateral composition* [DDH+11] of $\mathcal{A}$ and $\mathcal{B}$, noted $\mathcal{B} \circ \mathcal{A}$, is the algorithm defined as follows:

1. $\mathcal{B} \circ \mathcal{A}$ contains all variables of $\mathcal{A}$ and $\mathcal{B}$.

2. $\mathcal{B} \circ \mathcal{A}$ contains all actions of $\mathcal{A}$.

3. For every action "$L_i :: G_i \mapsto S_i$" of $\mathcal{B}$, $\mathcal{B} \circ \mathcal{A}$ contains the action "$L_i :: \neg D \wedge G_i \mapsto S_i$" where D is the disjunction of all guards of actions in $\mathcal{A}$.

The following sufficient condition is given in [DDH+11] to show the correctness of the composite algorithm:

Theorem 1 *The composite algorithm $\mathcal{B} \circ \mathcal{A}$ self-stabilizes to specification SP in a network of topology G assuming a weakly fair scheduler if the following conditions hold: (i) in a network of topology G, Algorithm $\mathcal{A}$ is a silent algorithm under a weakly fair scheduler; (ii) in a network of topology G, Algorithm $\mathcal{B}$ stabilizes to SP under a weakly fair daemon, starting from any configuration where no action of $\mathcal{A}$ is enabled.*

3 Computing Guide Pairs

3.1 Guide Pairs

Given an ordered tree $\mathcal{T}$, the guide pair of a node P in $\mathcal{T}$ is the pair of integers i and j such that i and j are, respectively, the rank of P in the *preorder* and *reverse postorder* traversal of $\mathcal{T}$. Below, we define these notions. Recall that we denote by $P_1, P_2, \ldots P_m$ the children of the root of $\mathcal{T}$ in the left-to-right order, and we denote by $\mathcal{T}_i$ be the subtree rooted at any P_i. The *preorder traversal* of $\mathcal{T}$ is defined, recursively, as follows:

1. Visit the root of $\mathcal{T}$.

2. For each i from 1 to m in increasing order, visit the nodes of $\mathcal{T}_i$ in *preorder*.

Postorder traversal $\mathcal{T}$ is similarly defined:

1. For each i from 1 to m in increasing order, visit the nodes of $\mathcal{T}_i$ in *postorder*.

2. Visit the root of $\mathcal{T}$.

Preorder traversal is top-down, while postorder traversal is bottom-up. However, we can also traverse $\mathcal{T}$ in *reverse postorder*, which is top-down, as follows.

1. Visit the root of $\mathcal{T}$.

2. For i from m to 1 in decreasing order, visit the nodes of $\mathcal{T}_i$ in *reverse postorder*.

If a node P is the i^{th} node of $\mathcal{T}$ visited in a preorder traversal of $\mathcal{T}$, we say that the *preorder rank* of P is i. If a node P is the j^{th} node of $\mathcal{T}$ visited in a reverse postorder traversal of $\mathcal{T}$, we say that the *reverse postorder rank* of P is j. Write $pre\_ind(P)$ and $post\_ind(P)$ for the *preorder* rank and *reverse postorder* rank of P, respectively. We define the *guide pair* of P to be the ordered pair $guide(P) = (pre\_ind(P), post\_ind(P))$. Figure 1 shows an ordered tree where each process is labeled with its guide pair.

If (i, j) and (k, ℓ) are guide pairs, we write $(i, j) \leq (k, \ell)$ if $i \leq k$ and $j \leq \ell$. Thus, the set of guide pairs is partially ordered by $\leq$.

Remark 2 *[Property 2 in [FEP$^+$06]] If P and Q are nodes of an ordered tree $\mathcal{T}$, then $guide(P) \leq guide(Q)$ if and only if P is an ancestor of Q.*

3.2   Algorithm GUIDE

Algorithm GUIDE is a hierarchical collateral composition of two algorithms: GUIDE = CGP $\circ$ COUNT, where both COUNT and CGP (for *Compute Guide Pairs*) use *P.par* as input in the program of every process P. Note that *P.par* either designates the actual parent link of P or is computed by a distributed spanning tree algorithm with which GUIDE must be composed using the hierarchical collateral composition.

3.2.1   Algorithm COUNT.

COUNT acts as a bottom-up wave that computes the number of processes in each subtree. In COUNT, each process P has only one variable: *P.subcount*. Moreover, each process P can compute the following function: $Subcount(P) = 1 + \sum_{Q \in Chldrn(P)} Q.subcount$. Thus, the program of P consists of the following action:

$$\texttt{SetCnt} \; :: \; P.subcount \neq Subcount(P) \; \mapsto \; P.subcount \leftarrow Subcount(P)$$

Lemma 3 COUNT *is self-stabilizing and silent, converges within $h + 1$ rounds from an arbitrary initial configuration to a legitimate configuration where $P.subcount = |\{Q \in \mathcal{T}_P\}|$ for all processes P, and works under the weakly fair scheduler.*

Proof: For any process P, define $D(P)$ to be the maximum distance from P to any leaf of $\mathcal{T}_P$. Then $D(Root) = h$, the height of $\mathcal{T}$.

Claim: For any $d \geq 0$ and any process P such that $D(P) = d$, $P.subcount = Subcount(P) = |\{Q \in \mathcal{T}_P\}|$ if at least $d + 1$ rounds have elapsed.

We prove the previous claim by induction on d. If $d = 0$, then P is a leaf, and $Subcount(P) = 1$. In one round, P executes Action SetCnt if necessary.

If $d > 0$, then, by the inductive hypothesis, all processes in $\mathcal{T}_P$, except possibly P itself, have calculated their correct values of *subcount*, if at least d rounds have elapsed. In one more round, $P.subcount$ will have the correct value.

With $d = h$, we can deduce that under a weakly fair scheduler, COUNT converges within $h + 1$ rounds to a legitimate configuration from an arbitrary initial configuration. Finally, as in a legitimate configuration, all values of *subcount* are correct, no process can execute Action SetCnt. Thus, COUNT is silent. $\square$

3.2.2   Algorithm CGP.

Using the values of *subcount* computed by COUNT, each process P evaluates in CGP for each of its children Q the number of processes before Q in the *preorder* and *reverse postorder* traversal of the tree $\mathcal{T}$, respectively (using Actions SetChldPrePred and SetChldPostPred, respectively). Then, reading these values from its parent, each process, except the root, can compute its guide pair (using Actions SetPreInd and SetPostInd). The guide pair of the root is $(1, 1)$ (see Actions SetPreInd and SetPostInd for the root).

Variables of CGP.   In CGP, each process maintains several variables. First, the following array variable enables each non-root process to know its index in the local left-to-right order of its parent:

1. $P.chld[i] \in N(P) \cup \{\bot\}$, for all $1 \leq i \leq \delta_P$. This array is maintained by Action `SetChld`. For all $1 \leq i \leq |Chldrn(P)|$, $P.chld[i]$ is set to the i^{th} child in P's local ordering of $N(P)$, while for all $|Chldrn(P)| < i \leq \delta_P$, $P.chld[i]$ is set to $\bot$.[*]

Then, CGP uses the following additional variables:

2. $P.pre\_ind$, $P.post\_ind$, integers. In stabilized state, they contain the *preorder* and *reverse postorder* ranks of P, respectively. Thus, we will write $P.guide = (P.pre\_ind, P.post\_ind)$, the guide pair of P.

3. $P.chld\_pre\_pred[i]$, $P.chld\_post\_pred[i]$, integer, defined for all $1 \leq i \leq |\delta_P|$:

   - For all $1 \leq i \leq |Chldrn(P)|$, $P.chld\_pre\_pred[i]$ is set to the number of predecessors of the i^{th} child of P (that is, $P.chld[i]$) in the *preorder* traversal of $\mathcal{T}$; and $P.chld\_post\_pred[i]$ is set to the number of predecessors of the i^{th} child of P in the *reverse postorder* traversal of $\mathcal{T}$.

   - For all $|Chldrn(P)| < i \leq \delta_P$, $P.chld\_pre\_pred[i]$ and $P.chld\_post\_pred[i]$ are set to 0.

   Hence, each process P computes its guide pair to be

   $$(P.par.chld\_pre\_pred[j] + 1, P.par.chld\_post\_pred[j] + 1)$$

   where P is the j^{th} child of its parent in left-to-right order.

Functions of CGP.   Based on the previous variables, each process P can compute the following functions:

- $my\_order(P)$. If P is not the root and there exists i, $1 \leq i \leq \delta_{P.par}$, such that $P.par.chld[i] = P$, then $my\_order(P)$ returns i. If the values of $P.par.chld$ did not stabilize, $my\_order(P)$ returns 1.

  Once the system has stabilized, $my\_order(P)$ returns the index of the non-root process P in the local left-to-right order of its parent.

- $Chld\_index(Q) = |\{Q' \in Chldrn(P) : Q' \prec_P Q\}| + 1$. It returns the index of the child Q of process P in the local left-to-right order of P.

- $Eval\_chld(i)$ returns the local name of the i^{th} child of P. That is, if $\exists Q \in Chldrn(P)$ such that $Chld\_index(Q) = i$, then $Eval\_chld(i)$ returns Q; otherwise, $Eval\_chld(i)$ returns $\bot$.

- $Eval\_chld\_pre\_pred(i)$. If $i = 1$, then $Eval\_chld\_pre\_pred(i)$ returns $P.pre\_ind$; else if $2 \leq i \leq |Chldrn(P)|$, then $Eval\_chld\_pre\_pred(i)$ returns $P.chld\_pre\_pred[i-1] + P.chld[i-1].subcount$; otherwise it returns 0.

  Once the system has stabilized, $Eval\_chld\_pre\_pred(i)$ returns the number of predecessors of the i^{th} child of P in the *preorder* traversal of $\mathcal{T}$.

- $Eval\_chld\_post\_pred(i)$. If $i = |Chldrn(P)|$, then $Eval\_chld\_post\_pred(i)$ returns $P.post\_ind$; else if $1 \leq i < |Chldrn(P)|$, then $Eval\_chld\_post\_pred(i)$ returns $P.chld\_post\_pred[i+1] + P.chld[i+1].subcount$; otherwise $Eval\_chld\_post\_pred(i)$ returns 0.

  Once the system has stabilized, $Eval\_chld\_post\_pred(i)$ returns the number of predecessors of the i^{th} child of P in the reverse postorder traversal of $\mathcal{T}$.

---

[*]Actually, cells from index $|Chldrn(P)| + 1$ to δ_P are useless. However, as the tree may be obtained by a spanning tree construction, we cannot know the number of children of P in advance, but this number is bounded by δ_P.

Actions of CGP. Actions of CGP are given below. To simplify the presentation, we assume priorities on actions, and list them below in the order from the highest to the lowest priority. If several actions are enabled simultaneously at a process, only the one of the highest priority can be executed. In other words, the actual guard of any action "$L :: G \mapsto S$" of process P is $\neg D \wedge G$, where D is the disjunction of the guards of all actions at P that appear before in the text.

For every process P:

SetChld	$:: \exists i \in [1..\delta_P],$	$\mapsto$	$\forall i \in [1..\delta_P],$
	$\quad P.chld[i] \neq Eval\_chld(i)$		$\quad P.chld[i] \leftarrow Eval\_chld(i)$
SetChldPrePred	$:: \exists i \in [1..\delta_P],$	$\mapsto$	$\forall i \in [1..\delta_P],$
	$\quad P.chld\_pre\_pred[i] \neq Eval\_chld\_pre\_pred(i)$		$\quad P.chld\_pre\_pred[i] \leftarrow Eval\_chld\_pre\_pred(i)$
SetChldPostPred	$:: \exists i \in [1..\delta_P],$	$\mapsto$	$\forall i \in [1..\delta_P],$
	$\quad P.chld\_post\_pred[i] \neq Eval\_chld\_post\_pred(i)$		$\quad P.chld\_post\_pred[i] \leftarrow Eval\_chld\_post\_pred(i)$

For the root process $Root$ only:

SetPreInd	$:: Root.pre\_ind \neq 1$	$\mapsto$	$Root.pre\_ind \leftarrow 1$
SetPostInd	$:: Root.post\_ind \neq 1$	$\mapsto$	$Root.post\_ind \leftarrow 1$

For every non-root process P only:

SetPreInd	$:: P.pre\_ind \neq 1 + P.par.chld\_pre\_pred[my\_order(P)]$	$\mapsto$	$P.pre\_ind \leftarrow 1 + P.par.chld\_pre\_pred[my\_order(P)]$
SetPostInd	$:: P.post\_ind \neq 1 + P.par.chld\_post\_pred[my\_order(P)]$	$\mapsto$	$P.post\_ind \leftarrow 1 + P.par.chld\_post\_pred[my\_order(P)]$

Overview of CGP. We now give an intuitive explanation of how CGP computes the values of $P.pre\_ind$ for all P. The values of $P.post\_ind$ are computed similarly.

Suppose that P is the i^{th} process visited in a *preorder* traversal of $\mathcal{T}$. Then i is the correct value of $P.pre\_ind$. CGP works by computing the number of predecessors of P, *i.e.*, the number of processes visited before P is visited. Let us call that number $Num\_Preorder\_Predecessors(P)$. It is the correct value of $P.pre\_ind - 1$.

$Num\_Preorder\_Predecessors(Root) = 0$; otherwise, $Num\_Preorder\_Predecessors(P)$ is computed by $P.par$ and stored in the variable $P.par.chld\_pre\_pred[j]$, where P is the j^{th} child of $P.par$ in left-to-right order. In order to compute these values for all its children, $P.par$ must have computed its own value of pre_ind, as well as the sizes of all of its subtrees. If $j = 1$, then $Num\_Preorder\_Predecessors(P) = P.par.pre\_ind$, since $P.par$ is the immediate predecessor of its leftmost child in the *preorder* visitation. Thus, $P.par.chld\_pre\_pred[1] \leftarrow P.par.pre\_ind$. $P.par.chld\_pre\_pred[2]$ is obtained by adding the size of the leftmost subtree of $P.par$ to $P.par.chld\_pre\_pred[1]$, since all members of that subtree are predecessors of the second child of $P.par$.

In general, the number of predecessors of P is equal to $P.par.pre\_ind$ plus the sum of the sizes of the leftmost $j - 1$ subtrees of $P.par$. The values of the array $P.par.chld\_post\_pred$ are computed from right to left, similarly. P then executes:

$$P.pre\_ind \quad \leftarrow \quad P.par.chld\_pre\_pred[j] + 1$$
$$P.post\_ind \quad \leftarrow \quad P.par.chld\_post\_pred[j] + 1$$

Theorem 4 GUIDE *is self-stabilizing and silent, computes the guide pairs of all processes in $O(h)$ rounds from an arbitrary initial configuration, and works under the weakly fair scheduler.*

Proof: According to Theorem 1 and Lemma 3, to show that GUIDE is self-stabilizing, it is sufficient to show that CGP stabilizes from any silent legitimate configuration of COUNT.

From such configurations, the values of *subcount*'s variables are correct. The variables of CGP are then computed in a top-down wave which takes $O(h)$ rounds. (We can prove this by induction on the height of processes in the tree, similar to the proof for COUNT.) Once a legitimate configuration is reached, no action is enabled.

Finally, the round convergence time of GUIDE is equal to the round convergence time of COUNT ($O(h)$ rounds) plus the number of rounds for CGP to reach a terminal configuration from any configuration where the values of *subcount* variables are correct ($O(h)$ rounds). $\qquad\square$

4 Rank Ordering

In this section, we give an algorithm, RANK, that uses guide pairs to solve the *ranking problem* on an ordered tree, $\mathcal{T}$. We are given a value *P.weight* for each process P in $\mathcal{T}$. (For convenience, we assume that the weights are integers.) The problem is to find the *rank* of each P. If $P_1, P_2, \ldots, P_n$ is the list of processes in $\mathcal{T}$ sorted by weight, then i is the rank of P_i. We allow ties to be broken arbitrarily, but deterministically.

Our algorithm RANK is a hierarchical collateral composition of two algorithms: RANK = CRK ∘ GUIDE. RANK computes the rank of each process P in $\mathcal{T}$, and sets the variable *P.rank* to that value. RANK is self-stabilizing, and requires $O(n)$ rounds and $O(b + \delta_P \log n)$ space for each process P.

4.1 Overview of CRK

4.1.1 Flow of Packages.

The key part of the algorithm CRK is the *flow of packages*. Each package is an ordered pair $x = (x.value, x.guide)$, where *x.value* is its *value* and *x.guide* is its *guide pair*. We identify a package with its *guide pair*. Moreover, for every two packages, x and y, we have $x \geq y$ (resp. $x > y$) if and only if $x.value \geq y.value$ (resp. $x.value > y.value$).

Each package has a *home process* (the node from which the package is originally issued), although its location can be at any process in the chain between its home and the root. The guide pair of a package is the same as the guide pair of its home process, and its value is either the weight of its home process or the rank that CRK will assign to its home process.

Each process P initiates its flow of packages by creating an *up-package* whose value is *P.weight*. This up-package then moves to the root by successive copying. The flow of packages is organized so that packages with smaller weights reach the root before packages with larger weights, in a manner similar to the standard technique for maintaining min-heap order in a tree.

After the root copies an up-package from a child, it creates a *down-package* with the same home process as the up-package, but whose value is a number (a rank) in the range $1..n$. The root maintains a counter so that the first down-package it creates has value 1, the second value 2, and so forth. Each down-package then moves back to its home process by copying. When its home process copies a down-package, it assigns, or re-assigns, its rank to be the value of that package.

The purpose (in fact even the name) of the guide pair is now obvious. It is used to guide the down-package to its home process.

Since the root copies up-packages in weight order, it creates down-packages in that same order. The i^{th} down-package created by the root will carry rank i and will use the same guide pair as the i^{th} up-package copied by the root. Its home process will then be the process whose weight is the i^{th} smallest in $\mathcal{T}$.

When the root detects that it has created all down-packages, it initiates a broadcast wave which resets the variables of CRK (except the rank and weight variables) and starts a new epoch.

4.1.2 Redundant Packages.

In our model of computation, if a variable of a process P is copied by a neighbor Q, it also remains at P. In the algorithm CRK, each process P can be home to at most one package, but we cannot avoid the existence of multiple copies of that package (up and/or down). We handle that problem by defining a package variable currently held by a process (not necessarily its home process, rather any process on the chain from its home to the root) as being either *active* or *redundant*. A redundant package can freely be overwritten, but not an active package.

If x is an up-package currently held by some process Q which is not the root, then x is redundant if x has already been copied by *Q.par*. If x is an up-package currently held by the root, then x is redundant if the root has already created a down-package with the same guide pair as x. Any other up-package is active.

If x is a down-package held by some process Q which is not its home process, then Q is redundant if it has been copied by some child of Q. (The child that copies x must be the process whose subtree contains the home process of x.) If x is a down-package held by its home process P, then x is redundant if *P.rank* is equal to the value of x. This indicates that P has already copied its rank from x, or that *P.rank* was correct before x arrived. Any other down-package is active.

4.1.3 Status Waves.

As it is typical for distributed algorithms which are self-stabilizing, but not silent, CRK endlessly repeats the calculation of the ranks of the processes in $\mathcal{T}$. We call one (complete) pass through this cycle of computations an *epoch*. At the end of each epoch, the variables of CRK at all processes, other than the variables for weight and rank, are reset for the next epoch. If an epoch has a clean start, it will calculate the correct rank for each process. Subsequent epochs will simply recalculate the same value, and *P.rank* will never change again.

On the other hand, in case of an arbitrary initial configuration, it is possible for incorrect values of rank to be calculated, but eventually a configuration will be reached when the next epoch will get a clean start.

This system is controlled by the *status* variables of the processes. At the beginning of an epoch, a broadcast wave starting from the root changes the status of every process from either 0 or 4 to 1, and all variables of CRK except rank and weight are set to their initial values. When this wave reaches the leaves of $\mathcal{T}$, a convergecast wave changes the status of all processes to 2. All computation of the ranking algorithm, as discussed above, takes place while processes have status 2. After the root has created the last down-package, it initiates a broadcast wave where the status of all processes changes to 3. The return convergecast wave then changes the status of all processes to 4, and when this wave reaches the root, the new epoch begins.

Status zero is used for error correction. If any process detects that the current epoch is erroneous, it changes its status to 0. Status 0 spreads down the tree, as well as up the tree unless it meets a process whose status is 1. If *Root.status* becomes 0 (and all its children have status 0 or 4), then *Root* initiates a status 1 broadcast wave starting a new epoch. However, this may cause an endless cycle of 0 and 1 wave, going up and down the tree, respectively. We solve this problem by adding a special rule for the non-root processes. If *P.status* = 0 and *P.par.status* = 1, the status 0 wave cannot move up; instead, the status 0 wave moves down followed by status 1 wave.

4.2 Formal Definition of CRK

4.2.1 Variables of CRK.

Let P be any process. *P.par*, *P.guide*, and *P.weight* are inputs of CRK. Then, the output of CRK is *P.rank*, an integer. To compute this output, P maintains the following additional variables:

1. *P.up_pkg* and *P.down_pkg* are respectively of package type (that is, a guide pair and an integer) or $\perp$ (undefined).

   If *P.up_pkg* (resp. *P.down_pkg*) is defined, then its home process is some $Q \in \mathcal{T}_P$.

2. *P.started*, Boolean.

   This variable indicates whether P has already generated its up-package during this epoch. (*P.up_pkg* may or may not still contain that up-package.)

3. *P.up_done*, Boolean.

   It indicates whether all processes in $\mathcal{T}_P$ have created their own up-package in the current epoch and whether $\mathcal{T}_P$ contains no active up-package. (Active up-packages whose home processes are in $\mathcal{T}_P$ could exist at processes above P.)

4. *P.status* $\in [0..4]$.

   Status variables are used to control the order of computation and to correct errors.

Finally, *Root* contains the following additional variable:

5. *Root.counter* $\in \mathbb{N}$

   This incrementing integer variable assigns the rank to packages. It is initialized to be 0 every time a new epoch begins.

4.2.2 Predicates of CRK.

The predicate *Clean_State(P)* below indicates if P is in a good initial or "clean" state.

$$Clean\_State(P) \quad \equiv \quad P.up\_pkg = \bot \ \wedge \ P.down\_pkg = \bot \ \wedge \neg P.started \ \wedge \neg P.up\_done$$

The four following predicates are used for error detection:

$$
\begin{aligned}
Is\_Consistent(P, g) \quad &\equiv \quad g = P.guide \ \vee \ \exists Q \in Chldrn(P), g \ \geq Q.guide \\[4pt]
Guide\_Error(P) \quad &\equiv \quad (P.up\_pkg \neq \bot \ \wedge \neg Is\_Consistent(P, P.up\_pkg.guide)) \ \vee \\
&\qquad (P.down\_pkg \neq \bot \ \wedge \neg Is\_Consistent(P, P.down\_pkg.guide)) \\[4pt]
Status\_Error(P) \quad &\equiv \quad (P.status \in \{1, 3\} \ \wedge \ P.par.status \neq P.status) \ \vee \\
&\qquad (P.status \in \{2, 4\} \ \wedge \ \exists Q \in Chldrn(P), Q.status \neq P.status) \ \vee \\
&\qquad (P.status \neq 0 \ \wedge \ P.par.status = 0) \ \vee \\
&\qquad (P.status \notin \{0, 1\} \ \wedge \ \exists Q \in Chldrn(P), Q.status = 0) \\[4pt]
Error(P) \quad &\equiv \quad Status\_Error(P) \ \vee \\
&\qquad (\neg Clean\_State(P) \ \wedge \ P.status = 1) \ \vee \\
&\qquad (Guide\_Error(P) \ \wedge \ P.status = 2) \ \vee \\
&\qquad (P.up\_done \ \wedge \neg P.started \ \wedge \ P.status = 2) \ \vee \\
&\qquad (P.up\_done \ \wedge \ P.status = 2 \ \wedge \ \exists Q \in Chldrn(P), \neg Q.up\_done)
\end{aligned}
$$

We say that a guide pair g is *consistent with* P if the predicate *Is_Consistent(P, g)* is true. If *Is_Consistent(P, g)* is false, g is the guide pair of no process in the subtree of P. *Guide_Error(P) = true* means that P holds a package whose home is not in the subtree of P. The predicate *Status_Error(P)* indicates whether P detects that its status is inconsistent with those of its neighbors. Status errors are always the result of arbitrary initializations; eventually, *Status_Error(P)* will become false and will remain false forever for all P. Finally, the predicate *Error(P)* detects error in the context of the current wave.

The four following predicates are used for flow control:

$$
\begin{aligned}
Up\_Redundant(P) \quad &\equiv \quad (P \neq Root \ \wedge \ P.up\_pkg \neq \bot \ \wedge \ P.par.up\_pkg \neq \bot \ \wedge \ P.par.up\_pkg \geq P.up\_pkg) \ \vee \\
&\qquad (P = Root \ \wedge \ P.up\_pkg \neq \bot \ \wedge \ P.down\_pkg \neq \bot \ \wedge \ P.down\_pkg.guide = P.up\_pkg.guide) \\[4pt]
Down\_Ready(P) \quad &\equiv \quad P.down\_pkg = \bot \ \vee \ (P.down\_pkg \neq \bot \ \wedge \\
&\qquad (P.down\_pkg.guide \neq P.guide \ \wedge \ \exists Q \in Chldrn(P), Q.down\_pkg = P.down\_pkg) \ \vee \\
&\qquad (P.down\_pkg.guide = P.guide \ \wedge \ P.rank = P.down\_pkg.value)) \\[4pt]
Can\_Start(P) \quad &\equiv \quad \neg P.started \ \wedge \ (P.up\_pkg = \bot \ \vee \ Up\_Redundant(P)) \ \wedge \forall Q \in Chldrn(P), \\
&\qquad (\neg Up\_Redundant(Q) \ \vee \ Q.up\_done) \ \wedge \ (Q.up\_pkg > (P.weight, P.guide) \ \vee \ Q.up\_done) \\[4pt]
Can\_Copy\_Up(P, Q) \quad &\equiv \quad Q \in Chldrn(P) \ \wedge \ (Q.up\_pkg \neq \bot \ \wedge \neg Up\_Redundant(Q)) \ \wedge \\
&\qquad (P.up\_pkg = \bot \ \vee \ Up\_Redundant(P)) \ \wedge \\
&\qquad (P.started \ \vee \ (P.weight, P.guide) > Q.up\_pkg) \ \wedge \forall R \in Chldrn(P), \\
&\qquad R.up\_done \ \vee \ (\neg Up\_Redundant(R) \ \wedge \ (R.up\_pkg \geq Q.up\_pkg \ \vee \ R.up\_done)
\end{aligned}
$$

$P.up\_pkg$ is redundant if *Up_Redundant(P)* is true. *Down_Ready(P)* states whether $P.down\_pkg$ is redundant or undefined, and thus P can create or copy a new down-package. *Can_Start(P)* decides whether P can create its own package, that is, if P can set $P.up\_pkg$ to $(P.weight, P.guide)$. *Can_Copy_Up(P)* indicates whether P can copy $Q.up\_pkg$ to $P.up\_pkg$. We note that P can evaluate *Up_Redundant(Q)* for any $Q \in Chldrn(P)$.

Predicate *Up_Done(P)* below decides whether all processes in $\mathcal{T}_P$ have created their own up-package in the current epoch and whether $\mathcal{T}_P$ contains no active up-package. The evaluation of *Up_Done(P)* gives the correct value for $P.up\_done$.

$$Up\_Done(P) \quad \equiv \quad P.started = true \ \wedge \ Up\_Redundant(P) \ \wedge \forall Q \in Chldrn(P), Q.up\_done$$

4.2.3 Actions of CRK.

Actions of CRK are given below. To simplify the design, we assume that the actions of CRK use the same priorities as those of CGP.

For the root process $Root$ **only:**

Err	$::$ $Error(Root)$	$\mapsto$	$Root.status \leftarrow 0$
NewEpoch	$::$ $Root.status \in \{0,4\} \wedge$ $\quad \forall Q \in Chldrn(Root),$ $\quad\quad Q.status \in \{0,4\}$	$\mapsto$	$Root.status \leftarrow 1;\ counter \leftarrow 0$ $Root.up\_pkg \leftarrow \bot;\ Root.down\_pkg \leftarrow \bot$ $Root.started \leftarrow false;\ Root.up\_done \leftarrow false$
ConvCast	$::$ $Root.status = 1 \wedge$ $\quad \forall Q \in Chldrn(Root), Q.status = 2$	$\mapsto$	$Root.status \leftarrow 2$
CreateUpPkg	$::$ $Root.status = 2 \wedge Can\_Start(Root)$	$\mapsto$	$Root.up\_pkg.value \leftarrow Root.weight$ $Root.up\_pkg.guide \leftarrow Root.guide$ $Root.started \leftarrow true$
CopyUpPkg	$::$ $Root.status = 2 \wedge$ $\quad \exists Q \in Chldrn(Root),$ $\quad\quad Can\_Copy\_Up(Root, Q)$	$\mapsto$	$Root.up\_pkg \leftarrow Q.up\_pkg,$ $\quad Q = min_{\prec_{Root}}\{R \in Chldrn(Root),$ $\quad\quad Can\_Copy\_Up(Root, R)\}$
EndUpPkg	$::$ $Root.started \wedge Up\_Redundant(Root) \wedge$ $\quad \forall Q \in Chldrn(Root), Q.up\_done$	$\mapsto$	$Root.up\_done \leftarrow true$
CreateDownPkg	$::$ $Down\_Ready(Root) \wedge$ $\quad Root.up\_pkg \neq \bot \wedge$ $\quad\quad \neg Up\_Redundant(Root)$	$\mapsto$	$counter \leftarrow counter + 1$ $Root.down\_pkg.value \leftarrow counter$ $Root.down\_pkg.guide \leftarrow Root.up\_pkg.guide$
SetRank	$::$ $Root.down\_pkg \neq \bot \wedge$ $\quad Root.down\_pkg.guide = Root.guide \wedge$ $\quad\quad Root.down\_pkg.value \neq Root.rank$	$\mapsto$	$Root.rank \leftarrow Root.down\_pkg.value$
BroadCast	$::$ $Root.status = 2 \wedge$ $\quad Root.up\_done \wedge Down\_Ready(Root)$	$\mapsto$	$Root.status \leftarrow 3$
EndEpoch	$::$ $Root.status = 3 \wedge$ $\quad \forall Q \in Chldrn(Root), Q.status = 4$	$\mapsto$	$Root.status \leftarrow 4$

For every non-root process P **only:**

Err	$::$ $Error(P)$	$\mapsto$	$P.status \leftarrow 0$
NewEpoch	$::$ $P.par.status = 1 \wedge P.status \in \{0,4\} \wedge$ $\quad \forall Q \in Chldrn(P), Q.status \in \{0,4\}$	$\mapsto$	$P.status \leftarrow 1$ $P.up\_pkg \leftarrow \bot$ $P.down\_pkg \leftarrow \bot$ $P.started \leftarrow false$ $P.up\_done \leftarrow false$
ConvCast	$::$ $P.status = 1 \wedge \forall Q \in Chldrn(P), Q.status = 2$	$\mapsto$	$P.status \leftarrow 2$
CreateUpPkg	$::$ $P.status = 2 \wedge Can\_Start(P)$	$\mapsto$	$P.up\_pkg.value \leftarrow P.weight$ $P.up\_pkg.guide \leftarrow P.guide$ $P.started \leftarrow true$
CopyUpPkg	$::$ $P.status = 2 \wedge$ $\quad \exists Q \in Chldrn(P), Can\_Copy\_Up(P, Q)$	$\mapsto$	$P.up\_pkg \leftarrow Q.up\_pkg,$ $\quad Q = min_{\prec_P}\{R \in Chldrn(P),$ $\quad\quad Can\_Copy\_Up(P, R)\}$
EndUpPkg	$::$ $P.started \wedge Up\_Redundant(P) \wedge$ $\quad \forall Q \in Chldrn(P), Q.up\_done$	$\mapsto$	$P.up\_done \leftarrow true$
CopyDownPkg	$::$ $Down\_Ready(P) \wedge$ $\quad P.par.down\_pkg \neq \bot \wedge$ $\quad\quad P.par.down\_pkg \neq P.down\_pkg \wedge$ $\quad\quad\quad Is\_Consistent(P, P.par.down\_pkg)$	$\mapsto$	$P.down\_pkg \leftarrow P.par.down\_pkg$
SetRank	$::$ $P.down\_pkg \neq \bot \wedge$ $\quad P.down\_pkg.guide = P.guide \wedge$ $\quad\quad P.down\_pkg.value \neq P.rank$	$\mapsto$	$P.rank \leftarrow P.down\_pkg.value$
BroadCast	$::$ $P.par.status = 3 \wedge P.status = 2 \wedge$ $\quad \forall Q \in Chldrn(P), Q.status = 2 \wedge Down\_Ready(P)$	$\mapsto$	$P.status \leftarrow 3$
EndEpoch	$::$ $P.status = 3 \wedge \forall Q \in Chldrn(P), Q.status = 4$	$\mapsto$	$P.status \leftarrow 4$

The actions above achieve three tasks. They are (1) error correction, (2) epochs, and (3) ranking computation (using the flow packages).

Error Correction. Action `Err` performs the error correction. If one process detects any inconsistency among its state and that of its neighbors, it initiates a reset of the network by changing its status to 0. This reset is contagious as previously explained.

Epochs. A new epoch starts by a reset initiated by Action `NewEpoch` at the root: If *Root.status* is either 0 or 4, and every child of *Root* has status 0 or 4, then *Root* broadcasts the status 1 and resets to a clean state.

When status 1 reaches the leaves, a convergecast wave starts and changes the status of all processes to 2 by Action `ConvCast`, so that actual ranks computation can begin.

When *Root* detects that there are no more up-packages in the tree, and it already sent every down-package, it initializes a broadcast of status 3 by Action `BroadCast`. Note that there could still be active down-packages below *Root*, but there could not be any active up-packages. Thus, *Root* is finished with its tasks for the current epoch. Non-root process *P* propagates the status 3 by Action `BroadCast` after sending all its down-packages. There could still be active down-packages below *P*, but no active up-packages. Since *P.par.status* = 3, *P* knows that its job for this epoch is done.

Once status 3 reaches the leaves, a convergecast of status 4 is initialized and propagated by Action `EndEpoch`. When *Root* changes to status 4, the current epoch is done, and *Root* initiates a new one.

Ranking computation. The computation of the ranking is bottom-up and starts when the convergecast of status 2 starts at the leaves. The flow of up-packages is organized using `CreateUpPkg` and `CopyUpPkg`, that is, a process either inserts its own package in the flow or copy some package coming from a child by ensuring that packages are moved up in ascending order of weight. Once a process *P* has detected that $\mathcal{T}_P$ has no active up-package, it sets *P.up_done* to true by Action `EndUpPkg`. *Root* initializes the broadcast of status 3 only after *Root.up_done* switches to true.

Upon receiving a new up-package (that is, *Root.up_pkg* is active), if *Root.down_pkg* is available (that is, it is either $\perp$ or redundant), *Root* is enabled to create a new down-package to send down to the home of its up-package by `CreateDownPkg`. If *counter* = i, then *Root.up_pkg* is the i^{th} up-package copied or created by *Root*, its weight is the i^{th} smallest weight in the network, and i will become the value of the down-package.

The new active down-package is propagated to its home process by successive copying using Action `CopyDownPkg`. When it reaches its home process *P*, the value field of that package contains the correct value of the rank of *P*, so *P* updates *P.rank* using Action `SetRank`, if necessary.

4.2.4 Correctness of CRK.

According to Theorem 1, to show the correctness of RANK, it is sufficient to show that variables of CRK stabilized to the expected result starting from any silent legitimate configuration of GUIDE. Let γ be such a configuration.

The first part of the proof deals with error correction.

Lemma 5 *After at most one round from γ, the following conditions hold forever for every process P:*

(a) *$\neg Clean\_State(P) \wedge (P.status = 1)$ is false.*

(b) *$P.up\_done \wedge \neg P.started \wedge (P.status = 2)$ is false.*

(c) *$(P.up\_done \wedge (P.status = 2)$, and there is some $Q \in Chldrn(P)$ such that $\neg Q.up\_done)$ is false.*

Proof: Let consider the three conditions separately:

(a) Assume $\neg Clean\_State(P) \wedge (P.status = 1)$ is false. Then, If *P.status* = 1, *P* cannot modify its other variables before changing its status. Moreover, every time *P.status* is reset to 1, the other variables are reset to a *clean* state (see Action `NewEpoch`). So, $\neg Clean\_State(P) \wedge (P.status = 1)$ remains false forever.

Assume $\neg Clean\_State(P) \wedge (P.status = 1)$ is true. Then, Action `Err` is enabled. Moreover, this condition only deals with local variables of *P*. So, Action `Err` is continuously enabled, and *P* executes *P.status* $\leftarrow$ 0 in at most one round. Then, $\neg Clean\_State(P) \wedge (P.status = 1)$ is false, and we obtain the previous case.

(b) Assume $P.up\_done \land \neg P.started \land (P.status = 2)$ is false. Then, P always sets $P.up\_done$ and $P.started$ to false together in Action `NewEpoch`. Moreover, P sets $P.up\_done$ to true only if $P.started$ holds (see Action `EndUpPkg`). So, $P.up\_done \land \neg P.started \land (P.status = 2)$ remains false forever.

Assume $\neg Clean\_State(P) \land (P.status = 1)$ is true. In this case, Action `Err` is enabled. Moreover, this condition only deals with local variables of P. So, Action `Err` is continuously enabled, and P executes $P.status \leftarrow 0$ in at most one round. Then, $P.up\_done \land \neg P.started \land (P.status = 2)$ is false, and we reach the previous case.

(c) Assume $(P.up\_done \land (P.status = 2)$, and there is some $Q \in Chldrn(P)$ such that $\neg Q.up\_done)$ is false. Then, P sets $P.up\_done$ to true only when every $Q \in Chldrn(P)$ satisfies $Q.up\_done = true$ (see Action `EndUpPkg`). Moreover, every $Q \in Chldrn(P)$ sets $Q.up\_done$ to false only when $P.up\_done = false$ (see Action `NewEpoch`). So, $(P.up\_done \land (P.status = 2)$, and there is some $Q \in Chldrn(P)$ such that $\neg Q.up\_done)$ remains false forever.

Assume $(P.up\_done \land (P.status = 2)$, and there is some $Q \in Chldrn(P)$ such that $\neg Q.up\_done)$ is true. Then, in one round, either every $Q \in Chldrn(P)$ satisfies $Q.up\_done = true$ or P executes Action `Err`. In both cases, $(P.up\_done \land (P.status = 2)$, and there is some $Q \in Chldrn(P)$ such that $\neg Q.up\_done)$ becomes false in at most one round, and we obtain the previous case.

$\square$

Lemma 6 *After at most one round from γ, for all process P, if $Status\_Error(P) = true$, then P satisfies one of the following two cases:*

- *$P \neq Root$ and $P.par.status = 0$.*

- *There is some $Q \in Chldrn(P)$ such that $Q.status = 0$.*

Proof: First, values 1 and 3 are propagated in the tree following a broadcast wave. Then, values 2 and 4 are propagated in the tree following a convergecast from the leaves. So, by definition of $Status\_Error(P)$, if P satisfies $Status\_Error(P) = false$ at some point, then $Status\_Error(P)$ will become true only after one neighbor of P switches its status to 0.

Finally, by Action `Err`, P cannot satisfy $Status\_Error(P) = true$ during one round without switching its status to 0. $\square$

Lemma 7 *If a process with status 0 holds an active package, this package remains blocked until it is removed or cleaned.*

Proof: If a process P has status 0, then no other process can copy its up or down packages because each of its neighbors has status 0, is its parent and has status 1, or has Action `Err` enabled — the action with the highest priority. Then, the next time P changes its status (Action `NewEpoch`), its state will become clean. $\square$

Lemma 8 *Within $O(n)$ rounds from γ, if process P contains an active package such that there is no process in its subtree which is the home process of that package, then $P.status = 0$.*

Proof:

Consider any configuration γ' after one round from γ. Consider an active package PK in γ' at any process P such that there is no process in the subtree of P that is the home process of that package.

(*) Assume that there is an ancestor of P with status 0, or a process in the subtree of P with status 0. Then, in at most h rounds, any process that holds PK as an active package has status 0 (Action `Err`), and by Lemma 7, PK cannot by copied anymore, so we are done.

Assume that no ancestor and no descendant of P have status 0. Consider the three following cases, according to the status of P:

(a) $P.status = 4$. Assume that there is an ancestor Q of P such that $Q.status = 1$. Then, by Lemma 6 and Definition of $Status\_Error$, all descendants of P have status 4, and for every ancestor R of P we have $R.status \in \{1, 4\}$ and $R.status = 1 \Rightarrow R = Root \vee R.par.status = 1$. So, in at most h rounds, the subtree of P has been reset to a clean state by Action `NewEpoch`, and we are done.

Assume that there is no ancestor Q of P such that $Q.status = 1$. Then, by Lemma 6 and Definition of $Status\_Error$, all descendants of P have status 4, and for every ancestor R of P we have $R.status \in \{3, 4\}$ and $R.status = 3 \Rightarrow R = Root \vee R.par.status = 3$. So, in at most h rounds, all ancestors of P have switched to status 4 by Action `EndEpoch`, and we reach the previous case.

(b) $P.status = 3$. If there is a process that has status 4, we reach the previous case, by Lemma 6 and Definition of $Status\_Error$.

Otherwise, every process of the tree has status 2 or 3, and if a process has status 3, then either it is the $Root$, or its parent also has status 3, by Lemma 6 and Definition of $Status\_Error$. In this case, PK can only be copied down in the tree (and only if it is a down package). Now, in $O(n)$ rounds, either (1) PK becomes an active package of a node Q that satisfies $Guide\_Error(Q) \wedge (Q.status = 2)$ (in the worst case Q is a leaf), its children cannot copy PK, after one additional round, Q has status 0, and PK cannot by copied anymore (by Lemma 7), so we are done; or (2) the broadcast wave of status 3 reaches the leaves of the tree, and in at most h additional rounds (the convergecast of status 4), we reach Case (a).

(c) $P.status = 2$. If there is a process Q such that $Q.status = 3$, then $Root.status = 3$ by Lemma 6 and Definition of $Status\_Error$, and we reach Case (b).

Otherwise, by Lemma 6, every process has status 1 or 2, and if a process has status 1, either it is the $Root$ or its parent has status 1. By Action `ConvCast`, in at most h rounds, all processes of $\mathcal{T}$ have status 2.

- If PK is an up-package, it can only be copied up the tree. Either P satisfies $Guide\_Error(P) \wedge (P.status = 2)$, its parent cannot copy PK, after one round P has status 0, and PK cannot by copied any more (by Lemma 7), so we are done; or in $O(n)$ rounds, PK becomes a down package at the $Root$, which has status 2, and is no longer an active up-package at any process.

- If PK is a down-package, it can only be copied down in the tree. After $O(n)$ rounds, the node Q where PK is active satisfies $Guide\_Error(Q) \wedge (Q.status = 2)$ (in the worst case Q is a leaf), its children cannot copy PK, after one additional round, Q has status 0, and PK cannot by copied anymore (by Lemma 7), so we are done.

(d) $P.status = 1$. By Lemma 6 and the definition of $Status\_Error$, two cases are possible:

- Every process has status 1 or 2, and if a process has status 1, then either it is the $Root$ or its parent has status 1. By Action `ConvCast`, in at most h rounds, $P.status = 2$ and we reach to Case (c).

- Every process has status 1 or 4 and if a process has status 1, then either it is the $Root$ or its parent has status 1. By Action `NewEpoch`, in at most h rounds, we reach the previous case.

$\square$

Lemma 9 *Within $O(n)$ rounds from γ, if process P contains a package, then there is a process in its subtree which is the home process of that package.*

Proof: By Lemmas 7 and 8, after $O(n)$ rounds, every process P holding an active package that does not have its home in the subtree of P satisfies $P.status = 0$ and no process copy this package.

Status 0 is propagated in $O(h)$ rounds (Action `Err`) up in the tree until reaching the root, or a process with status 1 and all its ancestors with status 1 as well as down to all descendants of P. After that, P and all its subtree reset its status to 1 in $O(h)$ rounds (Action `NewEpoch`).

Hence, within $O(n)$ rounds, all inconsistent active packages will be removed from the tree. $\square$

By Lemmas 5, 6, and 9, within $O(n)$ rounds from γ, $Error(P)$ is false forever for each process P. There may still exist processes with status 0, but in that case, by definition of $Error$, for any process P, we have: $P.status\{0,1,4\}$, $P.status = 0 \Rightarrow P = Root \lor P.par.status \in \{0,1\}$, $P.status = 1 \Rightarrow P = Root \lor P.par.status = 1$, $P.status = 4 \Rightarrow P \neq Root \land P.par.status \in \{1,4\}$. Hence, at the end of the broadcast of status 1 ($O(h)$ rounds), no process will have status 0. Thus, we have the following lemma:

Lemma 10 *Within $O(n)$ rounds from γ, $Error(P)$ is false and $P.status \in \{1,2,3,4\}$ forever, for each process P.*

From Lemma 10, we can deduce that the following invariant holds within $O(n)$ rounds from γ for all P.

1. $Error(P)$ is false and $P.status \in \{1,2,3,4\}$.

   That is, all initial errors will eventually be corrected.

2. If $P.status \in \{1,3\}$, then either $P = Root$ or $P.par.status = P.status$.

   The status values 1 and 3 indicate broadcast waves.

3. If $P.status \in \{2,4\}$, then $Q.status = P.status$ for all $Q \in Chldrn(P)$.

   The status values 2 and 4 indicate convergecast waves.

We now show that from any configuration where all previous invariants hold, infinitely many complete epochs are executed, each of them is performed in $O(n)$ rounds:

- If $Root.status = 4$, then all processes have status 4 and $Root$ initiates a broadcast of status 1 by Action `NewEpoch`.

- If $Root.status = 1$, then all processes P have either status 1 or 4. Moreover, $P.status = 1 \Rightarrow P = Root \lor P.par.status = 1$. So, the broadcast wave of status 1 is executed by all processes (Action `NewEpoch`), in at most h rounds.

- Then, a convergecast of status 2 is performed from the leaves (Action `ConvCast`) in at most h rounds.

- Once $Root.status = 2$, all processes have status 2. The flow of packages starts in parallel at processes of status 2.

- Then, by Claim 1, for every process P, while $P.status = 2$, if $P.up\_done$, every process Q in its subtree satisfies $Q.up\_done$. Moreover, $\neg P.started \Rightarrow \neg P.up\_done$. Now, by Action `CreateUpPkg`, the deepest node satisfying $P.status = 2$ and $\neg P.started$, eventually sets its Boolean $P.started$ to true and creates its own up-package. Then, the up-packages go up in the tree following the ordering on weight. So, every process P satisfies $P.up\_done$ after $O(n)$ rounds.

- When each process P satisfies $P.up\_done$, $Root$ eventually satisfies $Down\_Ready(Root)$ because every package it holds has its home in its subtree (Lemma 9), and so, is eventually copied by the children having the home in its subtree. Then, $Root$ initiates the broadcast of status 3 by Action `BroadCast`.

- When $Root.status = 3$, the all processes P have either status 2 or 3. Moreover, $P.status = 3 \Rightarrow P = Root \lor P.par.status 3$. So, status 3 is broadcast to the whole tree by Action `BroadCast`. As each node must wait for its down-package to become redundant before switching to status 3, this phase is performed in $O(n)$ rounds.

- Finally, once the status 3 reaches a leaf, the convergecast of status 4 is performed by Action `EndEpoch` in at most h rounds. $Root$ eventually has status 4 again.

Consider any epoch that starts from a configuration where all previous invariants hold. We define $\mathcal{S} = \{P : P.status \in \{1,2,3\}\}$. We call $\mathcal{S}$ the *active portion* of $\mathcal{T}$. The following invariants hold for all $P \in \mathcal{S}$.

4. If $P.status = 1$, then $P.started$ and $P.up\_done$ are false, and $P.up\_pkg = P.down\_pkg = \bot$.

   If the status of P is 1, then P has initialized its variables and has not yet begun the calculations of the epoch.

   Proof: By Claim 1, and definitions of $Error(P)$ and $Clean\_State(P)$. $\qquad\square$

5. If $P.up\_done$ then $P.started$, and $Q.up\_done$ for all $Q \in Chldrn(P)$.

   If there are no active packages in $\mathcal{T}_P$, then there are no active packages in $\mathcal{T}_Q$ for any child Q. Furthermore, the package whose home is P has already been created and copied up.

   Proof: $P.up\_done$ is initialized to false for all processes P during the broadcast wave of status 1. Then, all $P.up\_done$ are set to true in a bottom up fashion by Action `EndUpPkg`. $\qquad\square$

6. $P.up\_done$ if and only if there is no active package in $\mathcal{T}_P$.

   Proof: $P.up\_done$ is initialized to false for all processes P during the broadcast wave of status 1. Then, all $P.up\_done$ are set to true in a bottom up fashion by Action `EndUpPkg`. We can verify this claim by induction. $\qquad\square$

7. If P contains an active up-package, then there is some process $Q \in \mathcal{T} \cap \mathcal{S}$ that is the home process of that package.

   Proof: By Lemma 9, there is some process $Q \in \mathcal{T}$ that is the home process of that package. Assume, by the contradiction, that $Q \notin \mathcal{S}$. Before $Q.status = 4$, Q switches its status from 1 to 2, 2 to 3 and finally 3 to 4. When, Q switches its status to 3, all its ancestors, in particular $Root$, also have the status 3. Before switching to 3, $Root$ switches its status to 2. Now, $Root$ switches its status to 2, only if $Root.up\_done$. So, when Q switches its status to 3, there is no up-package in the network by Claim 6. Following a similar reasoning, there is no down-package in the path from the root to P. Hence, there is no package in the network having Q as its home by Lemma 9, a contradiction. $\qquad\square$

8. If $P.started$ is false, then there is no active up-package in $\mathcal{S}$ whose home process is P.

   Proof: P resets $P.started$ to false during the broadcast wave of status 1. Then, $P.started$ remains false until P creates its up-package. $\qquad\square$

9. If $P \in \mathcal{S}$ and $P.started$ is true, then there is at most one active package whose home process is P.

   Proof: P will create its own package only once, and once a package has been copied, previous copies become redundant. $\qquad\square$

10. If $Q \in Chldrn(P)$ and if $P.up\_pkg \neq \bot$, then either $Q.up\_pkg.weight \geq P.up\_pkg.weight$ or $Q.up\_done$.

    Proof: Min-heap order is maintained, so that up-packages reach $Root$ in weight order. $\qquad\square$

11. Let p be the number of processes P such that $P.started$ is false or $P.status = 4$, and q be the number of active up-packages. Then $p + q + counter = n$, the size of $\mathcal{T}$.

    Proof: At the beginning of each epoch, $p = n$ and $q = counter = 0$. Each time a process executes Action `CreateUpPkg`, p is decremented and q is incremented. Each time $Root$ executes Action `CreateDownPkg`, q is decremented and $counter$ is incremented. At the end of the epoch, $p = q = 0$ and $counter = n$. $\qquad\square$

12. If $P.weight$ is the i^{th} smallest weight in $\mathcal{T}$, then $i > counter$ if and only if either $P.started$ is false, or there is an active up-package whose home process is P.

    Proof: Follows from the two previous claims. $\qquad\square$

13. If $P.started$ is true and $P.rank$ is not the correct rank of P, then there is an active package in $\mathcal{S}$ whose home process is P.

    Proof: From the previous claim, the active up-package whose home is P will cause the creation of a down-package whose home is P with the correct rank value. $\qquad\square$

14. If $P.status = 3$, then $P.started$ and $P.up\_done$ are true, $P.up\_pkg$ and $P.down\_pkg$ are both redundant, and $P.rank$ has the correct value.

If the status of P is 3, then P has completed its role in the epoch.

Proof: Before switching to 3, P switches to status 2, hence, $P.up\_done$ is true. Moreover, if $P.up\_done$, then $P.started$ and $P.up\_pkg$ are redundant. P received status 3, only if there is no active down-package in the path from the root to P, so $P.down\_pkg$ is redundant too. Finally, by Claim 13, $P.rank$ has the correct value. □

As infinitely many complete epochs are executed, and during these epochs, all processes switch to status 3, by Claim 14 we have the following theorem:

Theorem 11 RANK *is self-stabilizing, computes the ranking of all processes in $O(n)$ rounds from an arbitrary initial configuration, and works under the weakly fair scheduler.*

References

[BDN95]   Brian Bourgon, Ajoy Kumar Datta, and Viruthagiri Natarajan. A self-stabilizing ranking algorithm for tree structured networks. In *Proceedings of the First Workshop on Self-Stabilizing Systems (WSS'95)*, pages 23–28, 1995.

[BDV05]   D. Bein, A.K. Datta, and V. Villain. Snap-stabilizing optimal binary-search-tree. *Proceedings of the 7-th International Symposium on Self-Stabilizing Systems*, 2005.

[DDH$^+$11]   Ajoy K. Datta, Stéphane Devismes, Karel Heurtefeux, Lawrence L. Larmore, and Yvan Rivierre. Self-stabilizing small k-dominating sets. Technical report, VERIMAG, 2011. http://www-verimag.imag.fr/TR/TR-2011-6.pdf.

[DGS96]   S. Dolev, MG Gouda, and M Schneider. Memory requirements for silent stabilization. In *PODC '96: Proceedings of the fifteenth annual ACM symposium on Principles of distributed computing*, pages 27–34, 1996.

[Dij74]   EW Dijkstra. Self stabilizing systems in spite of distributed control. *Communications of the Association of Computing Machinery*, 17:643–644, 1974.

[Dol00]   S Dolev. *Self-Stabilization.* The MIT Press, 2000.

[FEP$^+$06]   Paola Flocchini, Antonio Mesa Enriques, Linda Pagli, Giuseppe Prencipe, and Nicola Santoro. Point-of-failure shortest-path rerouting: Computing the optimal swap edges distributively. *IEICE Transactions*, 89-D(2):700–708, 2006.

[HM01]   T Herman and T Masuzawa. A stabilizing search tree with availability properties. In *Fifth International Symposium on Autonomous Decentralized Systems (ISADS 2001)*, pages 398–405, 2001.

[HP01]   T. Herman and I. Pirwani. A composite stabilizing data structure. *5th International Workshop on Self-Stabilizing Systems (WSS 2001), Lecture Notes in Computer Science LNCS 2194, Springer Verlag*, pages 167–182, 2001.

[Tel00]   G Tel. *Introduction to distributed algorithms (2nd Ed.).* Cambridge University Press, 2000.