



Self-stabilizing minimum-degree spanning tree within one from the optimal degree

Lélia Blin, Maria Gradinariu Potop-Butucaru, Stephane Rovedakis

► To cite this version:

Lélia Blin, Maria Gradinariu Potop-Butucaru, Stephane Rovedakis. Self-stabilizing minimum-degree spanning tree within one from the optimal degree. 23rd IEEE International Symposium on Parallel&Distributed Processing (IPDPS 2009), May 2009, Rome, Italy. pp.1-11, 10.1109/IPDPS.2009.5161042 . inria-00336713

HAL Id: inria-00336713

<https://inria.hal.science/inria-00336713>

Submitted on 4 Nov 2008

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

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

Self-stabilizing minimum-degree spanning tree within one from the optimal degree

Lélia Blin^{1,2} Maria Gradinariu Potop-Butucaru^{2,3} Stéphane Rovedakis¹

Abstract

We propose a self-stabilizing algorithm for constructing a Minimum-Degree Spanning Tree (**MDST**) in undirected networks. Starting from an arbitrary state, our algorithm is guaranteed to converge to a legitimate state describing a spanning tree whose maximum node degree is at most $\Delta^* + 1$, where Δ^* is the minimum possible maximum degree of a spanning tree of the network.

To the best of our knowledge our algorithm is the first self-stabilizing solution for the construction of a minimum-degree spanning tree in undirected graphs. The algorithm uses only local communications (nodes interact only with the neighbors at one hop distance). Moreover, the algorithm is designed to work in any asynchronous message passing network with reliable FIFO channels. Additionally, we use a fine grained atomicity model (i.e. the send/receive atomicity). The time complexity of our solution is $O(mn^2 \log n)$ where m is the number of edges and n is the number of nodes. The memory complexity is $O(\delta \log n)$ in the send-receive atomicity model (δ is the maximal degree of the network).

1 Introduction

The spanning tree construction is a fundamental problem in the field of distributed network algorithms being the building block for a broad class of fundamental services: communication primitives (e.g. broadcast or multicast), deadlock resolution or mutual exclusion. The new emergent distributed systems such as ad-hoc networks, sensor or peer-to-peer networks zoom on the cost efficiency of distributed communication structures and in particular spanning trees. The main issue in adhoc network systems for example is the communication cost. If the communication overlay contains nodes with large degree then undesirable effects can be observed: congestion, collisions or traffic burst. Moreover, these nodes are also the first targets in security attacks. In such case, the construction of reliable spanning trees in which the degree of a node is the lowest possible is needed.

Interestingly, in peer-to-peer networks the motivation for the construction of minimum degree trees is motivated by the nodes (users) welfare. Each communication on behalf of

¹Université d'Evry, IBISC, CNRS, France.

²Univ. Pierre & Marie Curie - Paris 6, LIP6-CNRS UMR 7606, France.

³INRIA REGAL, France.

other nodes in the network minimizes the possibility of a node to use the bandwidth for its own interests. As immediate consequence: nodes are incited to cheat on their real bandwidth or to have a free ridding behavior and illegally exploit the local traffic.

Self-stabilization as sub-field of fault-tolerance was first introduced in distributed computing area by Dijkstra in 1974 [6, 18]. A distributed algorithm is self-stabilizing if, starting from an arbitrary state, it guarantees to converge to a legitimate state in finite number of steps and to remain in a legitimate set of states thereafter. The property of self-stabilization enables a distributed algorithm to recover from a transient fault regardless of its initial state.

A broad class of self-stabilizing spanning trees have been proposed so far: by example [1, 7, 2] propose BFS trees, [5, 4] compute minimum diameter spanning trees, and [12] calculate minimum weight spanning trees. A detailed survey on the different self-stabilizing spanning tree schemes can be found in [18, 11]. Recently, [13] propose solutions for the construction of constant degree spanning trees in dynamic settings. To our knowledge there is no self-stabilizing algorithm for constructing minimum degree spanning trees.

This paper tackles the self-stabilizing construction of minimum-degree spanning tree in undirected graphs. More precisely, let $G = (V, E)$ be a graph. Our objective is to compute a spanning tree of G that has minimum degree among all spanning trees of G , where the degree of the tree is the maximum degree of its nodes. In fact, since this problem is NP-hard (by reduction from the Hamiltonian path problem), we are interested in constructing a spanning tree whose degree is within one from the optimal. This bound is achievable in polynomial-time as proved by Fürer and Raghavachari [8, 9] who describe an incremental sequential algorithm that constructs a spanning tree of degree at most $\Delta^* + 1$, where Δ^* is the degree of a MDST. Blin and Butelle proposed in [3] a distributed version of the algorithm by Fürer and Raghavachari. The algorithm in [3] uses techniques similar to those in [10] for controlling and updating fragment sub-trees of the currently constructed spanning trees. None of the previously mentioned solutions are self-stabilizing.

Our results We describe a self-stabilizing algorithm for constructing a minimum degree spanning tree in arbitrary networks. Our contribution is twofold. First, to the best of our knowledge our algorithm is the first self-stabilizing approximate solution for the construction of a minimum-degree spanning tree in undirected graphs. The algorithm converges to a legitimate state describing a spanning tree whose maximum node degree is at most $\Delta^* + 1$, where Δ^* is the minimum possible degree of a spanning tree of the network. Note that computing Δ^* is NP-hard. The algorithm uses only local communications — nodes interact only with their one hop neighbors — which makes it suitable for large scale and scalable systems. Our algorithm is designed to work in any asynchronous message passing network with reliable FIFO channels. Additionally, we use a fine grained atomicity model (i.e. *send/receive atomicity*) defined for the first time in [5]¹.

Secondly, our approach is based on the detection of fundamental cycles (i.e. cycles obtained by adding one edge to a tree) contrary to the technique used in [3] that perpetually updates membership information with respect to the different fragments (sub-trees)

¹The send/receive atomicity is the message passing counter-part of the read/write atomicity defined for the register model [18]

co-existing in the network. As a consequence, and in contrast with [3], our algorithm is able to decrease simultaneously the degree of each node of maximum degree. The time complexity of our solution is $O(mn^2 \log n)$ where m is the number of edges of the network and n the number of nodes. The memory complexity is $O(\log n)$ in a classical message passing model and $O(\delta \log n)$ in the send/received atomicity model (δ is the maximal degree of the network). The maximal length of messages used by our algorithm is $O(n \log n)$.

Paper road-map The paper is organized as follows. The next section introduces the model and the problem description. Section 3 presents the detailed description of our algorithm, then Section 4 and 5 contain the performance and complexity algorithm analysis respectively (only an abstract of correctness proof is given due to space constraints). The last section of the paper resumes the main results and outlines some open problems.

2 Model and notations

System model We borrow the model proposed in [5]. We consider an undirected connected network $G = (V, E)$ where V is the set of nodes and E is the set of edges. Nodes represent processors and edges represent communication links. Each node in the network has a unique identifier. For a node $v \in V$, we denote the set of its neighbors $\mathcal{N}(v) = \{u, \{u, v\} \in E\}$ ($\{u, v\}$ denotes the edge between the node u and its neighbor v). The size of the set $\mathcal{N}(v)$ is the degree of v . A spanning tree, T for G is a connected sub-graph of G such that $T = (V, E_T)$ and $E_T \subset E$, $|E_T| = |V| - 1$. The degree of a node u in T is denoted by $\deg_T(u)$, and we denote by $\deg(T)$ the degree of T , i.e., $\deg(T) = \max_v \deg_T(v)$.

We consider a static topology. The communication model is asynchronous message passing with FIFO channels (on each link messages are delivered in the same order as they have been sent). We use a refinement of this model: the send/receive atomicity ([5]). In this model each node maintains a local copy of the variables of its neighbors, refreshed via special messages (denoted in the sequel **InfoMsg**) exchanged periodically by neighboring nodes. A *local state* of a node is the value of the local variables of the node, the copy of the local variables of its neighbors and the state of its program counter. A *configuration* of the system is the cross product of the local states of all nodes in the system plus the content of the communication links. An *atomic step* at node p is an internal computation based on the current value of p 's local state and a single communication operation (send/receive) at p . An *execution* of the system is an infinite sequence of configurations, $e = (c_0, c_1, \dots, c_i, \dots)$, where each configuration c_{i+1} follows from c_i by the execution of a single atomic step.

Faults and self-stabilization In the sequel we consider the system can start in any configuration. That is, the local state of a node can be corrupted. Note that we don't make any assumption on the bound of corrupted nodes. In the worst case all the nodes in the system may start in a corrupted configuration. In order to tackle these faults we use self-stabilization techniques.

Definition 1 (self-stabilization) Let \mathcal{L}_A be a non-empty legitimacy predicate of an algorithm \mathcal{A} with respect to a specification predicate $Spec$ such that every configuration satisfying

\mathcal{L}_A satisfies Spec. Algorithm \mathcal{A} is self-stabilizing with respect to Spec iff the following two conditions hold:

- (i) Every computation of \mathcal{A} starting from a configuration satisfying \mathcal{L}_A preserves \mathcal{L}_A (closure).
- (ii) Every computation of \mathcal{A} starting from an arbitrary configuration contains a configuration that satisfies \mathcal{L}_A (convergence).

Minimum Degree Spanning Tree (MDST) A legitimate configuration for the MDST is a configuration that outputs an unique spanning tree of minimum degree. Since computing a Δ^* minimum degree spanning tree in a given network is NP-hard we propose in the following a self-stabilizing MDST with maximum node degree at most $\Delta^* + 1$.

3 A self-stabilizing MDST Algorithm

The main ingredients used by our MDST approximation algorithm are: (1) a module maintaining a spanning tree; (2) a module for computing the maximum node-degree of the spanning tree; (3) a module for computing the fundamental cycles and (4) a procedure for reducing the maximum node-degree of the spanning tree based on the fundamental cycles computed by the previous module.

One challenge is to design and run these modules concurrently, in a self-stabilizing manner. Note that the core of our algorithm is the reduction procedure that aims at repetitively reducing the degree of the spanning tree until getting a spanning tree of degree at most $\Delta^* + 1$.

The following section proposes the detailed description of our algorithm. The formal description of the algorithm can be found in Figure 2, with its sub-procedures in Figures 1 and 3².

3.1 Notations, Data Structures and Elementary procedures

Notations. Our algorithm makes use of the following notions (introduced first in [9]): *improving edge* and *blocking node* defined as follows. Let G be a network and T a spanning tree of G . Let $e = \{u, v\}$ be an edge of G that is not in T . The cycle C_e generated by adding e to T is called *fundamental*. Swapping e with any other edge of C_e (hence an edge in a spanning tree T) results in another spanning tree T' of G . Let w be a node of C_e , distinct from u and v , that has maximum degree in T among all nodes of C_e (i.e. $\deg_T(w) = \deg(T)$). If the exchange between e and an edge in C_e incident to w decreases by one the degree of w , then e is called an *improving edge*. An improving edge e satisfies

$$\deg_T(w) \geq \max(\deg_T(u), \deg_T(v)) + 2. \quad (1)$$

Let k be the degree of T (i.e., $k = \deg(T)$). If $\deg_T(u) = k - 1$ then u is called a *blocking node* for C_e . If u or v are blocking nodes for C_e then Eq. 1 states that e is not an improving edge for T .

²In the proposed algorithms \oplus concatenates lists

```

1 Update_State( $u, \deg_u$ )
2 use rules 1 and 2 to correct the local state of  $v$ 
3  $\text{edge\_status}_v[u] \leftarrow \text{is\_tree\_edge}(v, u); \deg_v[u] \leftarrow \deg_u;$ 
4  $\text{color\_tree}_v \leftarrow \text{degree\_stabilized}(v);$ 
5 Action_on_Cycle( $\text{idblock}, y, \text{path}, s$ )
6 if  $\text{idblock} = \text{NIL}$  then
7   let  $\text{d\_path} = \max\{\deg_u : u \in \text{path}\}$ 
8   if  $\text{dmax}_v = \text{d\_path}$  then
9     if  $\max\{\deg_v, \deg_{v,y}\} = \text{dmax}_v - 1$  then
10       Deblock( $y, s$ );
11     else
12       if  $\max\{\deg_v, \deg_{v,y}\} < \text{dmax}_v - 1$  then
13         let  $w = \min\{\text{ID}_u : \deg_u =$ 
14            $\text{d\_path}, u \in \text{path}\} \text{ and } (w, z) \in \text{path}$ 
15         Improve( $y, \deg_w, (w, z), \text{path}$ );
16   else if  $\text{idblock} \in \text{path}$  then
17     if  $\max\{\deg_v, \deg_{v,y}\} = \text{dmax}_v - 1$  then
18       Deblock( $y, s$ );
19     else
20       if  $\max\{\deg_v, \deg_{v,y}\} < \text{dmax}_v - 1$  then
21         let  $e = (\text{idblock}, z) \in \text{path}$ 
22         Improve( $y, \deg_{\text{idblock}}, e, \text{path}$ );
23 Broadcast( $\text{idblock}, s$ )
24  $\forall w \in N(v), w \neq s, \text{edge\_status}_v[w],$ 
25   send  $\langle \text{Deblock}, \text{idblock} \rangle$  to  $w$ .
26 Cycle_Search( $\text{idblock}$ );

26 Improve( $y, \deg, e, \text{path}$ )
27 send  $\langle \text{Remove}, (v, y), \deg, e, \text{path} \rangle$  to  $\text{head}(\text{path})$ 
28 Deblock( $y, s$ )
29 if  $\deg_v \geq \deg_y$  then Broadcast( $v, s$ );
30 if  $\deg_y \geq \deg_v$  then send  $\langle \text{Deblock}, y \rangle$  to  $y$ 
31 Reverse_Orientation(( $x, y$ ),  $\deg, (w, z), u, \text{path}$ )
32 let  $\text{path} = \text{list}_1 \oplus v \oplus \text{list}_2$  and  $p = \text{head}(\text{list}_2)$ .
33 if  $u = \text{parent}_v$  then
34   if  $\neg \text{edge\_status}_v[p]$  then Reverse_Aux( $p$ );
35   else
36      $\text{parent}_v \leftarrow p; \text{distance}_v \leftarrow \text{distance}_p + 1;$ 
37   send  $\langle \text{Remove}, (x, y), \deg, (w, z), \text{path} \rangle$  to  $p$ 
38 else
39   let  $q = \text{head}(\text{reverse}(\text{list}_1))$ .
40   if  $\neg \text{edge\_status}_v[q]$  then Reverse_Aux( $q$ );
41   else
42      $\text{parent}_v \leftarrow q; \text{distance}_v \leftarrow \text{distance}_q + 1;$ 
43   send  $\langle \text{Back}, (x, y), \text{reverse}(\text{list}_1) \rangle$  to  $q$ 
44 Reverse_Aux( $u$ )
45 send from  $u$   $\langle \text{Reverse}, v \rangle$  to  $\text{parent}_u$ , then wait
and treat only  $\text{InfoMsg}_v$  until message
 $\langle \text{Reverse}, v \rangle$  was received.
46 target_remove( $v, \deg\_max, (w, z)$ )  $\equiv (v = w \vee$ 
 $v = z) \wedge (\deg_w = \deg\_max \vee \deg_z = \deg\_max)$ 
 $\wedge (\text{edge\_status}_w[z] = \text{true});$ 
47 source_remove( $v, (x, y)$ )  $\equiv (v = x \vee v = y);$ 

```

Figure 1: Algorithm's procedures at node v

Variables. This short section lists all the variables used by the algorithm. For any node $v \in V(G)$, $N(v)$ (we assume an underlying self-stabilizing protocol that regularly updates the neighbors set) denotes the set of all neighbors of v in the network G , and $\text{ID}_v \in \mathbb{N}$ is the unique identifier of v . Nodes repeatedly send their variables to each of their neighbors u via a **InfoMsg** and update their local variables upon reception of this type of message from a neighbors. Each node v maintains the following variables:

- **Integer type variables:** root_v : the ID of the root of the spanning tree (computed by node v); parent_v : the parent ID of v in the spanning tree; distance_v : the distance of v to the root of the tree; dmax_v : Local estimation of k , it is updated upon the reception of a **InfoMsg** message. A change on k is detected via the color_tree_v variable. This information is further disseminated in the network via **InfoMsg** messages; \deg_v : the degree of v in the tree;
- **Boolean type variables:** $\text{edge_status}_v[u]$: is true when $\{u, v\}$ is an edge of the tree; color_tree_v : used to track any change on dmax_v ;

Predicates. The predicates at node v are the following:

- $\text{better_parent}(v) \equiv \bigvee_{u \in N(v)} \text{root}_v > \text{root}_u$

```

1 Do forever: send  $\text{InfoMsg}_v$  to all  $u \in N(v)$ 
2 Upon receipt of  $\text{InfoMsg}_u$  from  $u$ :  $\text{Update\_State}(u, \deg_u)$ ;  $\text{Cycle\_Search}(\text{NIL})$ ;
3 Upon receipt of  $\langle \text{Remove}, (x, y), \text{deg\_max}, (w, z), \text{path} \rangle \wedge \text{locally\_stabilized}(v)$  from  $u$ :
4 if  $\text{target\_remove}(v, \text{deg\_max}, (w, z))$  then
5    $\text{Reverse\_Orientation}((x, y), \text{deg\_max}, (w, z), u, \text{path})$ ;  $\text{color\_tree}_v \leftarrow \neg \text{color\_tree}_v$ ;
6 else
7   if  $\text{source\_remove}(v, (x, y))$  then send  $\langle \text{UpdateDist}, (w, z), \text{distance}_y + 1 \rangle$  to  $\text{parent}_v$ ;
8      $\text{parent}_v \leftarrow y$ ;  $\text{distance}_v \leftarrow \text{distance}_{\text{parent}_v} + 1$ ;
9   else let  $\text{path} = \text{list}_1 \oplus v \oplus \text{list}_2$  and  $p = \text{head}(\text{list}_2)$ .
10     if  $w, z \notin \text{list}_2$  then
11       if  $\neg \text{edge\_status}_v[u]$  then  $\text{Reverse\_Aux}(\text{head}(\text{list}_2))$ ;
12       else  $\text{parent}_v \leftarrow \text{head}(\text{list}_2)$ ;  $\text{distance}_v \leftarrow \text{distance}_{\text{parent}_v} + 1$ ;
13       send  $\langle \text{Remove}, (x, y), \text{deg\_max}, (w, z), \text{path} \rangle$  to  $\text{head}(\text{list}_2)$ 
14 send  $\text{InfoMsg}_v$  to all  $w \in N(v)$ 
15 Upon receipt of  $\langle \text{Back}, (x, y), \text{path} \rangle \wedge \text{locally\_stabilized}(v)$  from  $u$ :
16 if  $\text{source\_remove}(v, (x, y))$  then send  $\langle \text{UpdateDist}, (w, z), \text{distance}_y + 1 \rangle$  to  $\text{parent}_v$ ;
17    $\text{parent}_v \leftarrow y$ ;  $\text{distance}_v \leftarrow \text{distance}_{\text{parent}_v} + 1$ ;
18 else let  $\text{path} = \text{list}_1 \oplus v \oplus \text{list}_2$  and  $p = \text{head}(\text{list}_2)$ .
19   if  $\neg \text{edge\_status}_v[u]$  then  $\text{Reverse\_Aux}(p)$ ; else  $\text{parent}_v \leftarrow p$ ;  $\text{distance}_v \leftarrow \text{distance}_p + 1$ ;
20   send  $\langle \text{Back}, (x, y), \text{path} \rangle$  to  $p$ 
21 send  $\text{InfoMsg}_v$  to all  $w \in N(v)$ 
22 Upon receipt of  $\langle \text{Deblock}, \text{idblock} \rangle \wedge \text{locally\_stabilized}(v)$  from  $u$ :  $\text{Broadcast}(\text{idblock}, u)$ ;
23 Upon receipt of  $\langle \text{Reverse}, \text{target} \rangle \wedge \text{locally\_stabilized}(v)$  from  $u$ :
24 if  $\text{target} \neq v$  then send  $\langle \text{Reverse}, \text{target} \rangle$  to  $\text{parent}_v$ ;  $\text{parent}_v \leftarrow u$ ;
25 Upon receipt of  $\langle \text{UpdateDist}, (w, z), \text{dist} \rangle \wedge \text{locally\_stabilized}(v)$ :
26 if  $v = w \vee v = z$  then  $\text{distance}_v \leftarrow \text{dist} + 1$ ;
27  $\forall u \in N(v), u \neq \text{parent}_v \wedge \text{edge\_status}_v[u]$ , send  $\langle \text{UpdateDist}, (w, z), \text{distance}_v \rangle$ .

```

Figure 2: Algorithm at node v

- $\text{coherent_parent}(v) \equiv (\text{parent}_v \in N(v) \cup \{v\}) \wedge (\text{root}_v = \text{root}_{\text{parent}_v})$
- $\text{coherent_distance}(v) \equiv (\text{parent}_v \neq \text{ID}_v \wedge \text{distance}_v = \text{distance}_{\text{parent}_v} + 1) \vee (\text{parent}_v = \text{ID}_v \wedge \text{distance}_v = 0)$
- $\text{new_root_candidate}(v) \equiv \neg \text{coherent_parent}(v) \vee \neg \text{coherent_distance}(v)$
- $\text{is_tree_edge}(v, u) \equiv \text{parent}_v = \text{ID}_u \vee \text{parent}_u = \text{ID}_v$
- $\text{tree_stabilized}(v) \equiv \neg \text{better_parent}(v) \wedge \neg \text{new_root_candidate}(v)$
- $\text{degree_stabilized}(v) \equiv \bigwedge_{u \in N(v)} \text{dmax}_v = \text{dmax}_u$
- $\text{color_stabilized}(v) \equiv \bigwedge_{u \in N(v)} \text{color_tree}_v = \text{color_tree}_u$
- $\text{locally_stabilized}(v) \equiv \text{tree_stabilized}(v) \wedge \text{color_stabilized}(v)$

Procedures. The procedures at node v are the following:

- $\text{change_parent_to}(v, u) \equiv \text{root}_v \leftarrow \text{root}_u$; $\text{parent}_v \leftarrow \text{ID}_u$; $\text{distance}_v \leftarrow \text{distance}_u + 1$
- $\text{create_new_root}(v) \equiv \text{root}_v \leftarrow \text{ID}_v$; $\text{parent}_v \leftarrow \text{ID}_v$; $\text{distance}_v \leftarrow 0$

Messages. The messages used by our algorithm are the following:

- $\langle \text{InfoMsg}, \text{root}_v, \text{parent}_v, \text{distance}_v, \text{dmax}_v, \text{deg}_v, \text{edge_status}_v[u], \text{color_tree}_v \rangle$. This message is used to gossip the local variables of a node.
- $\langle \text{Search}, \text{init_edge}, \text{idblock}, \text{path} \rangle$: it is used to find the fundamental cycle induced by a non tree edge, where *init_edge* is the initiate non tree edge, *idblock* is the ID of a blocking node and *path* is the fundamental cycle.
- $\langle \text{Remove}, \text{init_edge}, \text{deg_max}, \text{target}, \text{path} \rangle$: it is used to reduce the degree of a maximum degree node by deleting an adjacent edge then it changes edge orientation in the fundamental cycle, where *init_edge* is the initiate non tree edge, *target* is the edge to be deleted, *deg_max* is the maximum degree of *target* extremities and *path* is the fundamental cycle.
- $\langle \text{Back}, \text{init_edge}, \text{path} \rangle$: change edge orientation in the fundamental cycle after an improvement, where *init_edge* is the initiate non tree edge and *path* is the fundamental cycle.
- $\langle \text{Deblock}, \text{idblock} \rangle$: it is used to change the state of a blocking node, where *idblock* is the ID of a blocking node.
- $\langle \text{Reverse}, \text{target} \rangle$: it is used to change the edge orientation in the fundamental cycle where *target* is the ID of the node when the change orientation stops.
- $\langle \text{UpdateDist}, \text{target}, \text{dist}, \text{path} \rangle$: update the distance of fundamental cycle nodes, where *target* is the edge when the update stops, *dist* is the distance from the tree root of the sender and *path* is the fundamental cycle.

Note that the *path* information is never stored at a node, therefore it is not listed as local variable of a particular node. However this information is carried by different messages. Therefore the complexity of our solution in the length of the network buffers is at least $O(n \log n)$ (the maximal length of the *path* chain) as explained in the complexity section.

3.2 Elementary building blocks

In this section, we provide first a detailed description of the underlying modules for our degree reduction algorithm, namely: a spanning tree module, a module that detects fundamental cycles and finally a module that computes and disseminates the maximal degree of the spanning tree. We conclude the section by the detailed presentation of the degree reduction module.

3.2.1 Spanning tree module

The algorithm described below is a simplification of the BFS algorithm proposed in [1]. Each node v maintains three variables: the local known ID of the tree root, a pointer to v 's parent and the distance of v to the spanning tree root. The output tree is rooted at the node having the minimum root value.

The algorithm uses two rules formally specified below. The first rule, "correction parent", enables the update of the locally known root: for a node v if a neighbor u has a lower root ID than v then v changes its root to u 's root, and u becomes the parent of v . If several neighbors verify this property then v will choose as parent the node with the minimal ID. This choice is realized by the argmin operation. The second rule, "correction root", creates a root if the neighborhood of a node has an incoherent state. In a coherent state, the distance

node-degree of the current spanning tree, can be done using a Propagation of Information with Feedback (PIF) protocol [16, 17]. That is, the protocol selects in the feedback phase at each level from the leaves to the root the maximum node-degree (by using deg_v). In the propagation phase that can be piggybacked on the **InfoMsg** messages the root disseminates the new maximum degree to all the nodes in the spanning tree. Upon the reception of **InfoMsg** a node updates its $dmax_v$ variable.

Note that the maximal degree of the tree changes during the execution of the algorithm via the degree reduction procedure. When a node of maximum degree has its degree decreased by one, the algorithm uses $color_tree_v$ variable to detect this change (see line 5, Figure 2). Additionally, a node that detects an incoherence in its neighborhood related to the maximum degree blocks the construction of the minimum-degree spanning tree until the neighborhood becomes locally stabilized (all neighbors have the same value for their $color_tree$ variable).

3.2.4 Degree reduction module

For all edges that do not belong to the current tree T our algorithm proceeds as follows. For each of these edges e , one of its adjacent nodes computes the fundamental cycle C_e , and checks whether e is an improving edge, and whether the degree of C_e is equal to $\deg(T)$. If the two latter conditions are satisfied, then e is swapped with one edge of C_e incident to a node of maximum degree in T . This operation is repeated until there is no more improving edge e with degree of C_e equal to the degree of the current tree. If a blocking node for C_e is encountered, then the algorithm tries to reduce the degree of the blocking node, recursively, in a way similar to the one used for decreasing the degree of the nodes of maximum degree. When no more improvement is possible then the algorithm stops. Note that [9] proved that the degree of a tree for which no improvement is possible is at most $\Delta^* + 1$, where Δ^* is the degree of a MDST of G . In the following we describe in details the above reduction process.




Figure 4: Degree reduction process

Cycle_Search: Repeatedly each non tree edge looks for its fundamental cycle by using the procedure **Cycle_Search** is described in Section 3.2. When a non tree edge $e = \{u, v\}$ discovers its fundamental cycle (Figure 3, line 8) it checks whether it is an improving edge in procedure **Action_on_Cycle**.

Action_on_Cycle: If the fundamental cycle contains no maximum degree node, no improvement is possible for this fundamental cycle. Otherwise (line 8) if the fundamental cycle also contains no blocking node (line 6) an improvement is possible (procedure **Improve** is called, lines 11-14). If there are blocking nodes the procedure **Deblock** is started (lines 9-

10). The procedure tries to improve the blocking node (lines 18-21), either via the procedure **Improve** or via the procedure **Deblock** (lines 15-17).

Improve: The procedure **Improve** sends a **Remove** message (lines 26-27), which is propagated along the fundamental cycle (Figure 2, lines 9-13). When the **Remove** message reaches its destination (edge $e' = \{w, z\}$), the degree and the status of e' are checked. If the maximum degree or the edge status have changed, then the **Remove** message is discarded. Otherwise, the status of edge e' is modified (it becomes a non-tree edge), and a message is sent along the cycle for correcting the parent orientation (procedure **Reverse_Orientation** explained below). When **Remove** reaches the edge e that initiated the process then e is added to the tree (Figure 2, lines 7-8). In the case when **Remove** message meets a deleted edge on its path, it carries on as if the deleted edge would be still alive.

Reverse_Orientation: After the remove of an edge e' , the orientation of the fundamental cycle must be corrected (see Figure 5). This is achieved via the circulation of two messages: **Back** and **Remove**. The procedure **Reverse_Orientation** deletes e' and checks following the orientation of e' whether a **Back** or a **Remove** message must be used to correct the orientation. If e' is oriented opposite to the direction followed by **Remove** message, then a **Remove** message is used (Figure 5(a)) otherwise a **Back** message is used (Figure 5(b)). The result of **Reverse_Orientation** on the path from e' to e on the fundamental cycle of e can be seen in Figure 5(c).

Note that a **Remove** (or **Back**) message propagated along the fundamental cycle may meet an edge deleted by another improvement (see message **Reverse** Figure 5). This implies that the fundamental cycle of e has changed. In this case it is necessary to finish the improvement operation otherwise the tree partitions.

Deblock: To perform an edge exchange with a blocking node $w \in \{u, v\}$, our algorithm starts by reducing by one the degree of w . For this purpose, w broadcasts a **Deblock** message in all its sub-tree (see procedure **Deblock**). When a descendant w' of w , adjacent to a non tree edge f , receives a **Deblock** message, it proceeds with the discovery of the fundamental cycle C_f of f as described in Section 3.2. If this fundamental cycle C_f enables to decrease the degree of the blocking node w , then w is not anymore blocking, and the swap e with e' can be performed. However, if the fundamental cycle C_f does not enable to decrease the degree of the blocking node w , then the procedure carries on recursively by broadcasting a **Deblock** messages for w' .




Figure 5: Illustration of **Reverse_Orientation** and **Reverse_Aux** procedures: (a) when $ID_x > ID_y$, (b) when $ID_x < ID_y$, (c) the resulted network given by (a) and (b).

Note that in order to maintain the tree stabilized we must update the distance of nodes on the path reversed after the remove of e' , as their distances to the tree root has changed. Therefore an `UpdateDist` message is diffused for all children on the reversed path.

4 Algorithm correctness

In the sequel we address different aspects related to the corruption of the local state of a node.

`InfoMsg` messages ensure the coherence between the copy of a node variables and their original values. Therefore, even if a node maintains corrupted copies their values will be corrected as soon as the node receives `InfoMsg` messages from each of its neighbors.

The correctness of the computation and the dissemination of the maximum degree of the current tree follow directly from the correctness of the self-stabilizing PIF scheme [17].

The reduction procedure and the search of fundamental cycles are frozen at a node v until the neighborhood of v is locally stabilized. That is, modules executed at an arbitrary node are locally ordered on priority bases: the module with maximal priority being the spanning tree construction, followed by the degree reduction module.

In the following we prove that the execution of the reduction procedure is blocked for a finite time (i.e. the time needed to the algorithm to compute a spanning tree). That is, we prove that the algorithm computes in a self-stabilizing manner a spanning tree in a finite number of steps.

Lemma 1 *Starting from an arbitrary configuration, eventually all nodes share the same root value.*

Proof. Nodes periodically exchange `InfoMsg` messages with their neighborhood. Therefore, even if the local copy of the neighbors variables are not identical with the original ones, upon the reception of a `InfoMsg` message a node corrects its local copies (via the procedure `Update_State`). Assume w.r.g. two different root values coexist in the network. Let $v_1 < v_2$ be these values. Let k be the number of nodes having their root variable setted to v_2 . Since the network is connected there are two neighboring nodes p_1 and p_2 such that the local root of p_1 is v_1 and the local root of p_2 is v_2 . After the reception of an `InfoMsg` from p_1 , p_2 detects its inconsistency and changes the local root value to v_1 and the parent variable to p_1 by executing the rule “correction parent”. Hence the number of root variables setted to v_2 decreases by 1. Recursively, the number of root variables setted to v_2 falls to 0.

Note that the root values are not modified by the other procedures of the algorithm. Overall, all nodes in the network will have the same root value. \square

Lemma 2 *Starting from an arbitrary configuration, eventually an unique spanning tree is alive in the network.*

Proof. Assume the contrary: two or more spanning trees are created. This leads to either the existence of two or more different root values or the presence of at least one cycle. Since

the values of the root variable are totally ordered then by Lemma 1 one of those values will eventually won the contest which invalidates the multiple roots assumption.

Assume there is cycle in the network. Note that there is an unique root and each node has a distance to the root which is the distance of the parent plus one. In a cycle there is at least one node such that the `coherent_distance` predicate returns false. This node will eventually execute the “correction root” rule. According to the proof of Lemma 1 the new root will run a competition with the existing root and the root with the lowest ID remains root while the other one modifies its root and distance variables. \square

Note that the minimal degree reduction procedure interfere with the tree maintenance via the modification of the distance variable in the edge reversal process. The messages of type `UpdateDist` are in charge of correcting the distance value and the parent value. During the correction of a path in the tree the messages related to the tree reduction are frozen until the correction is completed. Therefore, after the edge reversal completes the tree is coherent.

Now we prove that our algorithm maintains the same performance as algorithm in [9] and converges to a legitimate configuration verifying the hypothesis of Theorem 1 ([9]).

Theorem 1 (Fürer and Raghavachari [9]) *Let T be a spanning tree of degree Δ of a graph G . Let Δ^* be the degree of a minimum-degree spanning tree. Let S be the set of vertexes of degree Δ in T . Let B be an arbitrary subset of vertexes of degree $\Delta - 1$ in T . Let $S \cup B$ be removed from the graph, breaking the tree T into a forest F . Suppose G satisfies the condition that there are no edges between different trees in F . Then $\Delta \leq \Delta^* + 1$.*

We call *improvement* a step of our algorithm in which the degree of at least one node of maximum degree is decreased by 1, without increasing the degree of a node which already has the maximum degree, nor the degree of a blocking node.

Lemma 3 *When applying procedure **Improve**, our algorithm makes an improvement.*

Proof. Let $e = \{u, v\}$ be an improving edge for the current tree T of maximum degree k . Before adding e to T in the procedure **Improve**, some extremity of the edge e sends a **Remove** message along C_e to delete some edge $e' = \{x, y\}$ of the cycle, adjacent to a node of degree k or $k - 1$. This message **Remove** is routed along C_e . When it reaches one extremity of edge e' , here are two cases:

Case 1: e' is still a tree edge of degree k or $k - 1$. Then, according to procedure **Improve**, edge e' is removed from T , and u or v is informed by a **Back** message that e can be added to the tree T . In this case, we have an improvement.

Case 2: e' is no more a tree edge of degree k or $k - 1$. This implies that the degree of C_e has been decreased (by a concurrent improving edge), yielding an improvement. In this case, according to procedure **Improve**, the message **Remove** is discarded, for preserving the spanning tree property. In both cases, an improvement occurred. \square

A blocking node w in the current tree T is said *eventually non blocking* if it is non blocking, or there is an improving edge $e = \{u, v\}$ connecting two nodes u and v in the sub-tree T_w of T rooted at w , such that $w \in C_e$, and u and v are eventually non blocking.

Lemma 4 *Let w be a blocking node which can be made non blocking. Then procedure **Deblock**(w) reduces the degree of w by 1, without increasing the degree of neither a node with maximum degree, nor a blocking node.*

Proof. The proof proceeds by induction on the number of calls to procedure **Deblock**. If at the first call to **Deblock** there is an improving edge in the sub-tree of w (according to [9] it is sufficient to look for improving edge in the sub-tree of w) then an improvement is performed by the procedure **Improve** reducing the degree of w according to Lemma 3. If there is no improving edge but at least one blocking node adjacent to a non tree edge in the sub-tree of w which contains w in its fundamental cycle, then the procedure **Deblock** is recursively called by these ones according to procedure **Deblock_Node**. So if an improvement is possible, this improvement is propagated to w by a sequence of improvements. Note that each improvement can only increase the degree of nodes with a degree $\leq k - 2$ according to Lemma 3, so procedure **Deblock**(w) can not increase the degree of neither a node with maximum degree, nor a blocking node. \square

Theorem 2 *The algorithm returns a spanning tree of G with degree at most $\Delta^* + 1$.*

Proof. In the following we prove that a legitimate configuration verifies the hypothesis of Theorem 1. That is, when our algorithm reaches a legitimate configuration, there is no more possible improvement. Assume the contrary. We assume there is a tree structure T with maximum degree k . Suppose there is a possible improvement for T , by performing the swap between the edges $e_1 \notin T$ and $e_2 \in T$ to obtain the tree $T' = T \cup \{e_1\} \setminus \{e_2\}$ and assume algorithm does not perform this improvement. This implies that edge e_1 has a degree equal to k , otherwise as showed by Lemmas 3 and 4 the node of minimum id of e_1 will perform an improvement with procedure **Improve** or run **Deblock** to reduce the degree of e_1 , according to procedure **Action_on_Cycle**. This contradicts the fact that e_1 can be used to perform an improvement, and so if there is an improvement for T then the algorithm will perform this improvement. \square

5 Complexity issues and Discussions

The following lemma states that the time complexity is polynomial while the memory complexity is logarithmic in the size of the network.

Lemma 5 *Our algorithm converges to a legitimate state for the MDST problem in $O(mn^2 \log n)$ rounds using $O(\delta \log n)$ bits memory in the send/receive model³, where δ is the maximal degree of the network.*

Proof. The algorithm is the composition of three layers: the spanning tree construction, the maximum degree computation and the tree degree reduction layer. The first and the second layer have respectively a time complexity of $O(n^2)$ rounds [1, 11] and $O(n)$ rounds [16]. The time complexity is dominated by the last layer. According to [9], there are $O(n \log n)$

³In the classical message passing model the memory complexity is $O(\log n)$

phases. A phase consists of decreasing by one the maximum degree in the tree (except for the last phase). In fact, let k the degree of the initial tree, there are $O(n/k)$ maximum degree nodes. As $2 \leq k \leq n - 1$, summing up the harmonic series corresponding to the k values leads to $O(n \log n)$ phases. During each phase a non tree edge $\{u, v\}$ performs the following operations: finding its fundamental cycle and achieving an improvement (if it is an improving edge). To find its fundamental cycle, $\{u, v\}$ uses the procedure **Cycle_Search** which propagates a **Search** message via a DFS traversal of the current tree. So the first operation requires at most $O(n)$ rounds. For the second operation there are two possible cases: a direct and an indirect improvement. A direct improvement is performed if $\{u, v\}$ is an improving edge, in this case a **Remove** message is sent and follows the fundamental cycle to make the different changes, requiring $O(n)$ rounds. An indirect improvement is performed when u or v are blocking node ($\{u, v\}$ is a blocking edge). In the worst case, the blocking edges chain is of size at most $m - (n - 1)$. Therefore $m - (n - 1)$ exchanges are necessary to reduce the tree degree. This requires $O(mn)$ rounds (according to the direct improvement). So as the second operation needs at most $O(mn)$ rounds, the third layer requires $O(mn^2 \log n)$ rounds.

In the following we analyze the memory complexity of our solution. Each node maintains a constant number of local variables of size $O(\log n)$ bits. However, due to specificity of our model (the send/receive model) the memory complexity including the copies of the local neighborhood is $O(\delta \log n)$ where δ is the maximal degree of the network. \square

Note that the messages exchanged during the execution of our algorithm carry information related to the size of fundamental cycles. Therefore, the buffers length complexity of our solution is $O(n \log n)$.

6 Conclusion and open problems

We proposed a self-stabilizing algorithm for constructing a minimum-degree spanning tree in undirected networks. Starting from an arbitrary state, our algorithm is guaranteed to converge to a legitimate state describing a spanning tree whose maximum node degree is at most $\Delta^* + 1$, where Δ^* is the minimum possible maximum degree of a spanning tree of the network.

To the best of our knowledge our algorithm is the first self-stabilizing solution for the construction of a minimum-degree spanning tree in undirected graphs. The algorithm can be easily adapted to large scale systems since it uses only local communications and no node centralizes the computation of the MDST. That is, each node exchanges messages with its neighbors at one hop distance and the computations are totally distributed. Additionally the algorithm is designed to work in any asynchronous message passing network with reliable FIFO channels. The time complexity of our solution is $O(mn^2 \log n)$ where n and m are respectively the number of nodes and edges in the network. The memory complexity is $O(\delta \log n)$ in the send/receive atomicity model.

Several problems remain opened. First, the computation of a minimum-degree spanning tree in directed topologies seems to be a promising research direction with a broad area

of application (i.e. sensor networks, adhoc network, robot networks). Second, the new emergent networks need scalable solutions able to cope with nodes churn. An appealing research direction would be to design a super-stabilizing algorithm [18] for approximating a MDST.

References

- [1] Y. Afek and S. Kutten and M. Yung. Memory-efficient self stabilizing protocols for general networks. *4th International Workshop on Distributed Algorithm, WDAG, Springer LNCS volume 486*, 15-28, 1991.
- [2] A. Arora and M.G. Gouda. Distributed Reset. *In Proceedings of the Tenth Conference on Foundations of Software Technology and theoretical Computer Science, FSTTCS*, 17-19, 1990.
- [3] L. Blin. and F. Butelle. The first approximated distributed algorithm for the minimum degree spanning tree problem on general graphs. *Int. J. Found. Comput. Sci.*, 15(3):507-516, 2004.
- [4] F. Butelle and C. Lavault and M. Bui. A Uniform Self-Stabilizing Minimum Diameter Tree Algorithm (Extended Abstract). *Distributed Algorithms, 9th International Workshop (WDAG)*, 257-272, 1995.
- [5] J. Burman and S. Kutten. Time Optimal Asynchronous Self-stabilizing Spanning Tree. *Distributed Computing, 21st International Symposium (DISC)* 92-107, 2007.
- [6] E.W. Dijkstra. Self-stabilizing systems in spite of Distributed Control. *Communications of the ACM*, 17(11):643-644, 1974.
- [7] S. Dolev and A. Israeli and S. Moran. Self-Stabilization of Dynamic Systems Assuming only Read/Write Atomicity. *in ninth annual ACM symposium on Principles of distributed computing (PODC)*, 103-117, 1990.
- [8] M. Fürer and B. Raghavachari. Approximating the minimum degree spanning tree to within one from the optimal degree. *In the Proc. of the 3rd ACM-SIAM Symp. on Discr. Algo. (SODA)*, 317-324, 1992.
- [9] M. Fürer and B. Raghavachari. Approximating the minimum degree steiner tree to within one of optimal. *Journal of Algorithms*, 17:409-423, 1994.
- [10] R.G. Gallager and P.A. Humblet and P.M. Spira. A Distributed Algorithm for Minimum-Weight Spanning Trees. *ACM Trans. Program. Lang. Syst.*, 5(1):66-77, 1983.
- [11] F.C. Gärtner. A Survey of Self-Stabilizing Spanning-Tree Construction Algorithms. *TR,EPFL*, October 2003.
- [12] L. Higham and Z. Liang. Self-Stabilizing Minimum Spanning Tree Construction on Message-Passing Networks. *Distributed Computing, 15th International Conference (DISC)*, 194-208, 2001.
- [13] T. Héroult and Pierre Lemarinier and Olivier Peres and Laurence Pilard and Joffroy Beauquier. A Model for Large Scale Self-Stabilization. *IPDPS* 2007.
- [14] Silvia Bianchi and Ajoy Kumar Datta and Pascal Felber and Maria Gradinariu. Stabilizing Peer-to-Peer Spatial Filters. *ICDCS* 2007.
- [15] Baruch Awerbuch and Rafail Ostrovsky. Memory-efficient and self-stabilizing network RESET (extended abstract). *PODC '94: Proceedings of the thirteenth annual ACM symposium on Principles of distributed computing*.
- [16] Lélia Blin and Alain Cournier and Vincent Villain. An Improved Snap-Stabilizing PIF Algorithm. *In 6th Symposium on Self-Stabilizing Systems (SSS), LNCS 2704*, pages 199-214, 2003.
- [17] Alain Cournier and Ajoy Kumar Datta and Franck Petit and Vincent Villain. Optimal snap-stabilizing PIF algorithms in un-oriented trees. *J. High Speed Networks*, volume 14, number 2, pages 185-200, 2005.
- [18] Shlomi Dolev. *Self-Stabilization*. MIT Press, 2000.