Silent Self-stabilizing BFS Tree Algorithms Revised

Stéphane Devismes, Colette Johnen

Silent Self-stabilizing BFS Tree Algorithms Revised

Stéphane Devismes

VERIMAG

Colette Johnen

Univ. Bordeaux, LaBRI, UMR 5800, F-33400 Talence, France

Abstract

In this paper, we revisit two fundamental results of the self-stabilizing literature about silent BFS spanning tree constructions: the Dolev *et al* algorithm and the Huang and Chen's algorithm. More precisely, we propose in the composite atomicity model three straightforward adaptations inspired from those algorithms. We then present a deep study of these three algorithms. Our results are related to both correctness (convergence and closure, assuming a distributed unfair daemon) and complexity (analysis of the stabilization time in terms of rounds and steps).

Keywords: Self-stabilization, BFS spanning tree, composite atomicity model, distributed unfair daemon, stabilization time, round and step complexity.

1 Introduction

Self-stabilization [1] is a versatile technique to withstand *any finite number* of transient faults in a distributed system: a self-stabilizing algorithm is able to recover a correct behavior in finite time, regardless of the *arbitrary* initial configuration of the system, and therefore, also after the occurrence of transient faults.

After the seminal work of Dijkstra, several self-stabilizing algorithms have been proposed to solve various tasks such as token circulations [2], clock synchronization [3], propagation of information with feedbacks [4], *etc.* Among the vast self-stabilizing literature, many works more precisely focus on the construction of distributed data structures, *e.g.*, minimal dominating sets [5], clustering [6], spanning trees [7]. Most of the self-stabilizing algorithms which construct distributed data structures actually achieve an additional property called *silence* [8]: a silent self-stabilizing algorithm converges within finite time to a configuration from which the value of all its communication variables are constant.

Related Works. We focus here on silent self-stabilizing spanning tree constructions, *e.g.*, [7, 9, 10, 11, 12, 13]. Spanning tree constructions are of major interest in networking, *e.g.*, they are often involved in the design of routing and broadcasting tasks. Moreover, (silent) self-stabilizing spanning tree constructions are widely used as a basic building blocks of more complex self-stabilizing

solutions. Indeed, *composition* is a natural way to design self-stabilizing algorithms [14] since it allows to simplify both the design and proofs of self-stabilizing algorithms. Various composition techniques have been introduced so far, *e.g.*, collateral composition [15], fair composition [16], and conditional composition [17]; and many self-stabilizing algorithms actually are made as a composition of a silent spanning tree algorithm and another algorithm designed for tree topologies. For example, collateral, fair, and conditional compositions are respectively used the design of the algorithms given in [18], [19], and [20]. Notably, the silence property is not mandatory in such designs, however it allows to write simpler proofs [21].

Many self-stabilizing spanning tree constructions have been proposed, *e.g.*, [7, 9, 10, 11, 12, 22, 13]. These constructions mainly differ by the type of tree they compute, *e.g.*, the tree can be arbitrary [7], depth-first [11], breadth-first [9, 10, 22, 13], *etc.* In this paper, we focus two particular *Breadth-First Search* (BFS) spanning tree constructions: the one of Huang and Chen [9], and the one of *Dolev et al* [10]. These two constructions are among the most commonly used in the self-stabilizing literature.[1] Indeed, these constructions cumulate several advantages:

1. Their design is simple.
2. The BFS spanning tree is really popular because of its minimum height.
3. They are silent. Notice, by contrast, that the solution given in [22] is not silent.
4. Despite their time complexity was not analyzed until now, they are commonly assumed to be asymptotically optimal in rounds, *i.e.*, $O(\mathcal{D})$ rounds, where $\mathcal{D}$ is the diameter of the network. Notice, by contrast, that the stabilization time of the solutions proposed in [22] and [13] are $O(n + \mathcal{D}^2)$ rounds and $O(\mathcal{D}^2)$ rounds, respectively.

More precisely, the Huang and Chen's algorithm [9] is written in the composite atomicity model. It assumes the processes have the knowledge of n, the size of the network. This assumption allows processes to have a bounded memory: $\Theta(\log n)$ bits are required per process. The algorithm is proven assuming a *distributed unfair daemon*, the most general scheduling assumption of the model. However, no complexity analysis is given about its stabilization time in steps or rounds, the two main complexity metrics of the model.

The Dolev *et al*'s algorithm [10] is written in the read/write atomicity model. This model is more general than the composite atomicity model. The algorithm does not assume any knowledge on any global parameter of the network, such as n for example. The counterpart being that there is no bound on process local memories. The algorithm is proven under the central fair assumption (*n.b.*, the notion of unfair daemon is meaningless in this model). Despite no complexity analysis is given in the paper, authors conjecture the stabilization time is asymptotically optimal, *i.e.*, $O(\mathcal{D})$ rounds. By definition, a straightforward translation of the algorithm of Dolev *et al* also works in the composite atomicity model assuming a distributed weakly fair daemon. However, an ad hoc proof is necessary if one want to establish its self-stabilization under a distributed unfair daemon. Notice that the algorithm of Dolev *et al* and its bounded memory variant are used in the design of several algorithms written in this composite atomicity model, *e.g.*, [23, 24]. Moreover, several algorithms are based on similar principles, *i.e.*, [25, 26]. Hence, a proof of its self-stabilization assuming a distributed unfair daemon is highly desirable.

[1] As a matter of facts, [9] and [10] are respectively cited 109 and 409 times in `Google Scholar`.

Contribution. In this paper, we study three silent self-stabilizing BFS spanning tree algorithms written in the composite atomicity model. The first algorithm, called $\texttt{Algo}_{\mathcal{U}}$, is the straightforward translation of the Dolev *et al*'s algorithm [10] into the composite atomicity model. The second one, $\texttt{Algo}_{\mathcal{B}}(D)$, is a variant of $\texttt{Algo}_{\mathcal{U}}$, where the process local memories are bounded. To that goal, the knowledge of some upper bound D on the network diameter is assumed. Finally, The third algorithm, noted $\texttt{Algo}_{\mathcal{HC}}(D)$, is a generalization of the Huang and Chen's algorithm [9], where the exact knowledge of n is replaced by the knowledge of some upper bound D on the network diameter (*n.b.*, by definition, n is a particular upper bound on the diameter).

The general purpose of this paper is twofold. First, we show the close relationship between these three algorithms. To see this, we propose a general and simple proof of self-stabilization for the three algorithms under the distributed unfair daemon. This proof implies that every executions of each of the three algorithms is finite in terms of steps. Moreover, notice that the proof shows in particular that the assumption on the exact knowledge of n in the initial algorithm of Huang and Chen was too strong. Second, the proof also validates the use of the Dolev *et al*'s algorithm and its bounded-memory variant in the composite atomicity model assuming any daemon (in particular, the unfair one).

Second, we propose a complexity analysis the stabilization time of these three algorithms in both steps and rounds. Our results are both positive and negative. First, we show that the stabilization time of $\texttt{Algo}_{\mathcal{U}}$ and $\texttt{Algo}_{\mathcal{B}}(D)$ is *optimal* in rounds by showing that in both cases the worst case is exactly $\mathcal{D}$ rounds. With few modifications our proof can be adapted for the read/write atomicity model, validating then the conjecture in [10] which claimed that the stabilization time the Dolev *et al*'s algorithm was asymptotically optimal in rounds.

We then establish a lower bound in $\Omega(D)$ rounds on the stabilization time of $\texttt{Algo}_{\mathcal{HC}}(D)$. Now, $\texttt{Algo}_{\mathcal{HC}}(n)$ is actually the algorithm of Huang and Chen. Thus, the algorithm of Huang and Chen stabilizes in $\Omega(n)$ rounds. This result may be surprising as until now this algorithm was conjectured to stabilize in $O(\mathcal{D})$ rounds. More precisely, this negative result is mainly due to the fact that two rules of the algorithm are not mutually exclusive, and when both rules are enabled at the same process p, the daemon may choose to activate any of them. Our lower bound is thus established when the daemon gives priority to one of the two rules (HC_1). Hence, to circumvent this problem we proposed a straightforward variant, noted $\texttt{Algo}_{\mathcal{FHC}}(D)$, where we give priority to the other rule (HC_2). We then establish that in this latter case the stabilization time becomes exactly $\mathcal{D}+1$ rounds in the worst case.

Finally, we consider the stabilization time in steps. Our results are all negative. Indeed, we first show that the stabilization time in steps of $\texttt{Algo}_{\mathcal{U}}$ cannot be bounded by any function depending on topological parameters. We then exhibit a lower bound exponential in $\mathcal{D}$ (the actual diameter of the network) on the stabilization time in steps which holds for both $\texttt{Algo}_{\mathcal{B}}(D)$ and $\texttt{Algo}_{\mathcal{HC}}(D)$. Notice, by contrast, that the stabilization time of the solutions proposed in [22] and [13] are $O(\Delta.n^3)$ steps and $O(n^6)$ steps, respectively.

Roadmap. The rest of the paper is organized as follows. The next section is dedicated to the description of the computation model and definitions. The formal codes of the three algorithms are given in Section 3. In Section 4, we propose our general proof of self-stabilization assuming a distributed unfair daemon. In Section 5, we analyze the stabilization time in rounds of the three solutions. We present an analysis of the stabilization time in steps of the three solutions in Section 6. Section 7 is dedicated to concluding remarks.

2 Preliminaries

2.1 Distributed Systems

We assume distributed systems of $n > 0$ interconnected processes. One process, called the *root*, is distinguished, the others are *anonymous*. The root of the system is simply denoted by R. Each process p can directly communicate with a subset $\mathcal{N}_p$ of other processes, called its *neighbors*. We assume bidirectional communications, *i.e.*, if $q \in \mathcal{N}_p$, then $p \in \mathcal{N}_q$. The topology of the system is a simple undirected connected graph $G = (V, E)$, where V is the set of processes and E is the set of edges, each edge being an unordered pair of neighboring processes. $\|p, q\|$ denotes the distance (the length of the shortest path) from p to q in G. We denote by $\mathcal{D}$ the diameter of G, *i.e.*, $\mathcal{D} = \max_{p,q \in V} \|p, q\|$.

2.2 Computational Model

We consider the *locally shared memory model with composite atomicity* introduced by Dijkstra [1], where each process communicates with its neighbors using a finite set of *locally shared variables*, henceforth called simply *variables*. Each process can read its own variables and those of its neighbors, but can write only to its own variables. Each process operates according to its (local) *program*. A *distributed algorithm* $\mathcal{A}$ consists of one local program $\mathcal{A}(p)$ per process p.

$\mathcal{A}(p)$ is given as a finite set of *rules*: $\{Label_i : Guard_i \rightarrow Action_i\}$. *Labels* are only used to identify rules in the reasoning. The *guard* of a rule in $\mathcal{A}(p)$ is a Boolean expression involving the variables of p and its neighbors. The *action* part of a rule in $\mathcal{A}(p)$ updates some variables of p. The *state* of a process in $\mathcal{A}$ is defined by the values of its variables in $\mathcal{A}$. A *configuration* of $\mathcal{A}$ is an instance of the states of every process in $\mathcal{A}$. $\mathcal{C}_\mathcal{A}$ is the set of all possible configurations of $\mathcal{A}$. (When there is no ambiguity, we omit the subscript $\mathcal{A}$.) A rule can be executed only if its guard evaluates to *true*; in this case, the rule is said to be *enabled*. A process is said to be enabled if at least one of its rules is enabled. We denote by $Enabled(\gamma)$ the subset of processes that are enabled in configuration γ. When the configuration is γ and $Enabled(\gamma) \neq \emptyset$, a *daemon* selects a non-empty set $\mathcal{X} \subseteq Enabled(\gamma)$; then every process of $\mathcal{X}$ *atomically* executes one of its enabled rule, leading to a new configuration γ', and so on. The transition from γ to γ' is called a *step* (of $\mathcal{A}$). The possible steps induce a binary relation over $\mathcal{C}_\mathcal{A}$, denoted by $\mapsto_\mathcal{A}$ (or, simply $\mapsto$, when it is clear from the context). An *execution* of $\mathcal{A}$ is a maximal sequence of its configurations $e = \gamma_0\gamma_1 \ldots \gamma_i \ldots$ 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 rule of $\mathcal{A}$ is enabled at any process. We denote by $\mathcal{E}_\mathcal{A}$ (or, simply $\mathcal{E}$, when it is clear from the context) the set of all possible executions of $\mathcal{A}$. The set of all executions starting from a particular configuration γ is denoted $\mathcal{E}(\gamma)$. Similarly, $\mathcal{E}(\mathcal{S})$ is the set of execution whose the initial configuration belongs to $\mathcal{S} \subseteq \mathcal{C}$.

As previously stated, each step from a configuration to another is driven by a daemon. In this paper we assume the daemon is *distributed* and *unfair*. "Distributed" means that while the configuration is not terminal, the daemon should select at least one enabled process, maybe more. "Unfair" means that there is no fairness constraint, *i.e.*, the daemon might never select an enabled process unless it is the only enabled process.

We say that a process p is *neutralized* during the step $\gamma_i \mapsto \gamma_{i+1}$ if p is enabled at γ_i and not enabled at γ_{i+1}, but does not execute any rule between these two configurations. An enabled process is neutralized if at least one neighbor of p changes its state between γ_i and γ_{i+1}, and this

change makes the guards of all rules of p *false*. To evaluate time complexity, we use the notion of *round*. This notion captures the execution rate of the slowest process in any execution. The first round of an execution e, noted e', is the minimal prefix of e in which every process that is enabled in the initial configuration either executes an action or becomes neutralized. Let e'' be the suffix of e starting from the last configuration of e'. The second round of e is the first round of e'', and so forth.

2.3 Self-Stabilization and Silence

We are interesting in algorithms which converge from an arbitrary configuration to a configuration where output variables define a specific data structure, namely a BFS spanning tree. Hence, we define a specification as a predicate $\mathbb{SP}$ on $\mathcal{C}$ which is *true* if and only if the outputs define the expected data structure.

Silent Self-stabilization is a particular form of self-stabilization defined by Dolev *et al* [8] as follows. A distributed algorithm $\mathcal{A}$ is *silent self-stabilizing w.r.t. specification* $\mathbb{SP}$ if following two conditions holds:

Termination: all executions of $\mathcal{A}$ are finite; and

Partial Correctness: all terminal configurations of $\mathcal{A}$ satisfy $\mathbb{SP}$.

In this context, the *stabilization time* is the maximum time (in steps or rounds) to reach a terminal configuration starting from any configuration.

3 Three Self-Stabilizing BFS Constructions

Below we give three self-stabilizing BFS constructions: $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, and $\texttt{Algo}_{\mathcal{HC}}(D)$.[2] In all these variants, the local program of the R just consists in the following constant: $d_R = 0$.

Each non-root process p maintains two variables: d_p and par_p. The domain of par_p is $\mathcal{N}_p$, the set of p's neighbors. The domain of d_p differs depending on the version:

- d_p is an unbounded positive integer in the first version, $\texttt{Algo}_{\mathcal{U}}$.
- $d_p \in [1..D]$ in the two other algorithms, $\texttt{Algo}_{\mathcal{B}}(D)$ and $\texttt{Algo}_{\mathcal{HC}}(D)$. The correctness of both $\texttt{Algo}_{\mathcal{B}}(D)$ and $\texttt{Algo}_{\mathcal{HC}}(D)$ will be established for any $D \geq \mathcal{D}$.

The three algorithms use the following three macros:

- $Min\_d(p) = \min\{d_q \mid q \in \mathcal{N}_p\}$
- $bestParent(p)$ is any neighbor q such that $d_q = Min\_d(p)$
- $update(p)$: $d_p \leftarrow Min\_d(p) + 1$; $par_p \leftarrow bestParent(p)$

Moreover, the following two predicates are used in the algorithms:

- $dOk(p) \equiv (d_p = Min\_d(p) + 1)$
- $parOk(p) \equiv (d_p = d_{par_p} + 1)$

The two variables are maintained using the actions defined below.

[2] $\mathcal{U}$, $\mathcal{B}$, $\mathcal{HC}$ respectively stand for *unbounded*, *bounded*, and *Huang-Chen*.

3.1 Actions of $\texttt{Algo}_{\mathcal{U}}$

The BFS algorithm in [10] is designed for the read/write atomicity model. The straightforward adaptation of this algorithm in the locally shared memory model is given below.

$$\begin{array}{lcll} U_1 & :: & \neg dOk(p) & \rightarrow \quad update(p) \\ U_2 & :: & dOk(p) \wedge \neg parOk(p) & \rightarrow \quad par_p \leftarrow bestParent(p); \end{array}$$

3.2 Actions of $\texttt{Algo}_{\mathcal{B}}(D)$

The following algorithm is a variant of $\texttt{Algo}_{\mathcal{U}}$, where the domain of the d-variable is now bounded by the input parameter D.

$$\begin{array}{lcll} B_1 & :: & Min\_d(p) < D \wedge \neg dOk(p) & \rightarrow \quad update(p) \\ B_2 & :: & Min\_d(p) < D \wedge dOk(p) \wedge \neg parOk(p) & \rightarrow \quad par_p \leftarrow bestParent(p); \\ B_3 & :: & Min\_d(p) = D \wedge (d_p \neq D) & \rightarrow \quad d_p \leftarrow D \end{array}$$

3.3 Actions of $\texttt{Algo}_{\mathcal{HC}}(D)$

Actions given below are essentially the same as those of the BFS algorithm given in [9]. Actually, they differ in two points from the version of [9].

- In [9], d-variables are defined in such way that $d_R = 1$ and $d_p \geq 2, \forall p \in V \setminus \{R\}$. We have changed the domain definition of the d-variables for sake of uniformity. However, this difference on the domain definitions has no impact on the behavior of the algorithm.
- Moreover, to be more general, we have replaced in the code the exact value n (the number of processes) by D.

$$\begin{array}{lcll} HC_1 & :: & \neg parOk(p) \wedge d_{par_p} < D & \rightarrow \quad d_p \leftarrow d_{par_p} + 1 \\ HC_2 & :: & d_{par_p} > Min\_d(p) & \rightarrow \quad update(p) \end{array}$$

Notice that HC_1 and HC_2 are not mutually exclusive, *i.e.*, in some configurations, both rules can be enabled at the same process. For instance, in the initial configuration of Figure 1 (page 12), rules HC_1 and HC_2 are enabled at process a. In this case, if a is selected by the daemon, the daemon also chooses which of the two rules is executed.

4 Correctness

In this section, we give a general proof which establishes the self-stabilization of the three algorithms under a distributed unfair daemon. The proof of correctness consists of the following two main steps:

Partial correctness (Theorems 1 and 2) which means that if an execution terminates, then the output of the terminal configuration is correct. (In our context, the output of the terminal configuration define a BFS spanning tree rooted at R.)

Termination (Theorem 4) which means that every possible execution is finite in terms of steps.

4.1 Partial Correctness

Below we establish the partial correctness of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, and $\texttt{Algo}_{\mathcal{HC}}(D)$ using three main steps: (1) we define a set of *legitimate* configurations (Definition 1), and show that (2) a BFS spanning tree is defined in each legitimate configuration (Theorem 1) and (3) every terminal configuration of each algorithm is legitimate (Theorem 2). additionally, we show that every legitimate configuration is terminal in any of the three algorithms (Theorem 3).

4.1.1 Legitimate Configurations

Definition 1 *A configuration is* legitimate *if and only if for every process p, we have $d_p = \|p, R\|$ and if $p \neq R$, then $d_p = d_{par_p} + 1$.*

Let $T_\gamma = (V, E_T)$, where $E_{T_\gamma} = \{\{p, q\} \in E \mid q \neq R \wedge par_q = p \text{ in } \gamma\}$.

Theorem 1 *T_γ is a BFS spanning tree in every legitimate configuration γ.*

Proof. Let γ be any legitimate configuration. We first demonstrate that T_γ is a spanning tree by showing the following two claims:

T_γ is acyclic: Assume, by contradiction, that T_γ contains a cycle $p_0, \ldots, p_k, p_0$. By definition, R is not involved into the cycle. Assume, without loss of generality, that for all $i > 0$, $par_{p_i} = p_{i-1}$ et $par_{p_0} = p_k$ in γ. From Definition 1, in γ we have $d_{p_i} > d_{p_{i-1}}$ and, by transitivity, $d_{p_k} > d_{p_0}$. Now, as $par_{p_0} = p_k$, we have $d_{p_0} > d_{p_k}$ by Definition 1, a contradiction.

$|E_{T_\gamma}| = n - 1$: First, by definition we have $|E_{T_\gamma}| \leq n - 1$. Now, if $|E_{T_\gamma}| < n - 1$, then there is at least one edge $\{p, q\}$ such that $par_p = q$ and $par_q = p$. This contradicts the fact that T_γ is acyclic. Hence, $|E_{T_\gamma}| = n - 1$.

We now show that T_γ is breadth-first. Let $p_0 = R, \ldots, p$ the unique path from R to p in the tree. By definition, for all $i > 0$, $par_{p_i} = p_{i-1}$. By Definition 1, we have $d_{p_i} = d_{p_{i-1}} + 1$ and, by transitivity, $d_{p_k} = d_{p_0} + k$. Moreover, $d_{p_0} = d_R = 0$. So, d_{p_k} is the length k of the path from R to p_k. Now, by Definition 1, d_{p_k} is also equal to $\|R, p_k\|$. Hence, the length k of the path from R to p_k is equal to the distance from R to p_k in G. □

4.1.2 Legitimacy of Terminal Configurations

Let $D \geq \mathcal{D}$. Let $\mathcal{TC}_{\texttt{Algo}_{\mathcal{U}}}$, $\mathcal{TC}_{\texttt{Algo}_{\mathcal{B}}(D)}$, and $\mathcal{TC}_{\texttt{Algo}_{\mathcal{HC}}(D)}$ be the set of terminal configurations of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, and $\texttt{Algo}_{\mathcal{HC}}(D)$, respectively.

Lemma 1 *Let γ be a configuration of $\mathcal{TC}_{\texttt{Algo}_{\mathcal{U}}}$. Let X be the largest distance value in γ. γ is a configuration of $\mathcal{TC}_{\texttt{Algo}_{\mathcal{B}}(X)}$.*

Proof. Every process p satisfies $dOk(p) \wedge parOk(p)$ in γ. So, the rule B_1 and B_2 of $\mathcal{TC}_{\texttt{Algo}_{\mathcal{B}}(X)}$ are disabled at p in γ. Moreover, by definition, $d_p \leq X$ and $parOk(p)$ in γ implies that $d_{par_p} < X$ in γ. So, $Min\_d(p) < X$ in γ and, consequently, B_3 of $\mathcal{TC}_{\texttt{Algo}_{\mathcal{B}}(X)}$ is disabled at p in γ. □

Lemma 2 *Let $D \geq \mathcal{D}$ and γ be a configuration. If γ be a terminal configuration of $\texttt{Algo}_{\mathcal{B}}(D)$, then γ is a terminal configuration of $\texttt{Algo}_{\mathcal{HC}}(D)$.*

Proof. We establish this lemma by showing its contrapositive. Let γ' be a non-terminal configuration of $\texttt{Algo}_{\mathcal{HC}}(D)$. There is a process p such that $(\neg parOk(p) \wedge d_{par_p} < D) \vee (d_{par_p} > Min\_d(p))$ holds in γ'.

Assume first that $\neg parOk(p) \wedge d_{par_p} < D$ holds in γ'. Then, $d_{par_p} < D$ implies that $Min\_d(p) < D$ holds. Thus either B_1 or B_2 is enabled in γ' because $\neg parOk(p)$ holds.

Assume then that $d_{par_p} > Min\_d(p)$. Then, as $d_{par_p} \leq D$ (by definition), we have $Min\_d(p) < D$. So, if $\neg dOk(p)$, then B_1 is enabled in γ'. Otherwise, we have $d_p = Min\_d(p) + 1$. So, $d_p \neq d_{par_p} + 1$ and, consequently, B_2 is enabled in γ'.

Hence, every non-terminal configuration of $\texttt{Algo}_{\mathcal{HC}}(D)$ is a non-terminal configuration of $\texttt{Algo}_{\mathcal{B}}(D)$. □

From the two previous lemmas, we have:

Corollary 1 $\mathcal{TC}_{\texttt{Algo}_{\mathcal{U}}} \subseteq \bigcup_{i=\mathcal{D}}^{\infty} \mathcal{TC}_{\texttt{Algo}_{\mathcal{B}}(i)} \subseteq \bigcup_{i=\mathcal{D}}^{\infty} \mathcal{TC}_{\texttt{Algo}_{\mathcal{HC}}(i)}$.

From the previous corollary, we know that it is sufficient to show that any configuration of $\bigcup_{i=\mathcal{D}}^{\infty} \mathcal{TC}_{\texttt{Algo}_{\mathcal{HC}}(i)}$ is legitimate to establish that any configuration of $\mathcal{TC}_{\texttt{Algo}_{\mathcal{U}}}$, $\bigcup_{i=\mathcal{D}}^{\infty} \mathcal{TC}_{\texttt{Algo}_{\mathcal{B}}(i)}$, and $\bigcup_{i=\mathcal{D}}^{\infty} \mathcal{TC}_{\texttt{Algo}_{\mathcal{HC}}(i)}$ are legitimate.

Lemma 3 *Let $D \geq \mathcal{D}$. In any configuration of $\mathcal{TC}_{\texttt{Algo}_{\mathcal{HC}}(D)}$ we have $d_p \leq \|p, R\|$ for every process p.*

Proof. Let γ be any terminal configuration of $\texttt{Algo}_{\mathcal{HC}}(D)$. Assume, by the contradiction, that there is a process p such that $d_p > \|p, R\|$ in γ. Without loss of generality, assume p is one of the closest processes from R such that $d_p > \|p, R\|$ in γ. By definition, $\|p, R\| \leq \mathcal{D}$, and p has at least a neighbor, q, such that $\|q, R\| = \|p, R\| - 1 < \mathcal{D}$. By hypothesis, $d_q \leq \|q, R\| = \|p, R\| - 1$ in γ. So, we have $Min\_d(p) \leq \|p, R\| - 1 < \mathcal{D} \leq D$ in γ. Then, by definition, $d_{par_p} \geq Min\_d(p)$ in γ. Now, HC_2 is disabled at p in γ. So, $d_{par_p} = Min\_d(p)$ in γ. Consequently, $d_{par_p} \leq \|p, R\| - 1 < D$ in γ. Now, HC_1 is disabled at p in γ. So, $d_p = d_{par_p} + 1 \leq \|p, R\|$ in γ, a contradiction. □

Corollary 2 *Let γ be a configuration of $\mathcal{TC}_{\texttt{Algo}_{\mathcal{HC}}(D)}$ where $D \geq \mathcal{D}$. Every process $p \neq R$ satisfies $d_p = d_{par_p} + 1$ and $d_{par_p} = Min\_d(p)$ in γ.*

Proof. Let $p \neq R$ be process. By definition, $d_{par_p} \geq Min\_d(p)$ in γ. HC_2 is disabled at p in γ. So, $d_{par_p} = Min\_d(p)$ in γ. p has a neighbor such that $\|q, R\| \leq \mathcal{D} - 1$. So, we have $d_q \leq \mathcal{D} - 1$ in γ (by Lemma 3) and, consequently, $Min\_d(p) < D$ in γ. So, $d_{par_p} < D$ in γ. As HC_1 is disabled at p in γ, we have also $d_p = d_{par_p} + 1$. □

Lemma 4 *Let $D \geq \mathcal{D}$. Let γ be a configuration of $\mathcal{TC}_{\texttt{Algo}_{\mathcal{HC}}(D)}$. $d_p = \|p, R\|$ holds for every process p, in γ.*

Proof. We have $d_R = 0$ (by definition) and $d_q = d_{par_q} + 1$ for every process $q \neq R$ in γ (corollary 2). So, similarly to the proof of Theorem 1, we can establish that T_γ is spanning tree and for every process p, d_p is the length of the path in T_γ from R to p. So, we have $d_p \geq \|p, R\|$ in γ. Now, $d_p \leq \|p, R\|$ in γ, by Lemma 3. So, we conclude that $d_p = \|p, R\|$ in γ, for every process p.

□

From Corollaries 1,2, and Lemma 4, we can deduce the following theorem:

Theorem 2 *Let $D \geq \mathcal{D}$. Every terminal configuration of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$ is γ is legitimate.*

4.1.3 Legitimate Configurations are Terminal

Theorem 3 *Let $D \geq \mathcal{D}$. Every legitimate configuration is a terminal configuration of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, and $\texttt{Algo}_{\mathcal{HC}}(D)$, respectively.*

Proof. Let γ be a legitimate configuration. First, for every process p, $d_p = \|p, R\| \leq \mathcal{D}$. So, γ is a possible configuration of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, and $\texttt{Algo}_{\mathcal{HC}}(D)$, respectively.

Let p be a non-root process. We have $Min\_d(p) = \min\{d_q \mid q \in \mathcal{N}_p\} = \min\{\|q, R\| \mid q \in \mathcal{N}_p\} = \|p, R\| - 1 < \mathcal{D} \leq D$ in γ. So, B_3 is disabled at every non-root process in γ.

Moreover, $Min\_d(p) + 1 = \|p, R\| = d_p$ in γ. So, U_1 and B_1 are disabled at every non-root process in γ.

By definition, $parOk(p)$ holds in γ. So, U_2, B_2, and HC_1 are disabled at non-root process in γ.

Finally, $parOk(p)$ implies that $d_{par_p} = d_p - 1 = \|p, R\| - 1 = Min\_d(p)$ in γ. Hence, HC_2 is disabled at every non-root process in γ. □

4.2 Termination

In this subsection, we will establish that, under a distributed unfair daemon, all executions of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, and $\texttt{Algo}_{\mathcal{HC}}(D)$ are finite.

The lemma given below establish that we only need to prove that the number of d-variable updates is finite in any execution e of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$ to establish that e is finite.

Lemma 5 *Let $e = \gamma_0, \ldots \gamma_i, \gamma_{i+1}, \ldots$ be any execution of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$. If for every process p, d_p is modified only a finite number of time along e, then e is finite.*

Proof. Assume that every process p, d_p is modified only a finite number of time along e. Then, there exists $i \geq 0$ such that no d-variable is modified in the suffix $e' = \gamma_i\gamma_{i+1}\ldots$ of e. By definition of the three algorithms, only par-variables can be modified along e'. So the rules U_1 for $\texttt{Algo}_{\mathcal{U}}$, B_1 and B_3 for $\texttt{Algo}_{\mathcal{B}}(D)$, and HC_1 for $\texttt{Algo}_{\mathcal{HC}}(D)$ are not executed along e'. Now, by definition of the algorithms, in e', we have:

- Once the rule modifying par_p (Rule U_2, B_2, or HC_2) is disabled, it remains disabled forever by p, because the values of d-variables are constant (in particular, those of p and its neighbors).
- The rule modifying par_p (Rule U_2, B_2, or HC_2) becomes disabled immediately after p's execution.

Consequently, each process takes at most one step along e'; we conclude that the execution e is finite. □

Notation 1 *For every configuration γ, for any integer $k \geq 0$, we denote by $Set\_d_k(\gamma)$ the set of processes p such that $d_p = k$ in γ.*

Remark 1 *In every configuration γ, $Set\_d_0(\gamma) = \{R\}$, and $Set\_d_\ell(\gamma) \bigcap Set\_d_k(\gamma) = \emptyset$ for every $0 \leq \ell < k$. $V = \bigcup_{i=0}^{\infty} Set\_d_i(\gamma)$.*

The following lemma establishes that for every execution e of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$, there is a upper bound kb on the values taken in e by the d-variables of all processes.

Lemma 6 *Let $e = \gamma_0, \ldots \gamma_i, \gamma_{i+1}, \ldots$ be any execution of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$. $\exists kb \geq 0 \mid \forall j \geq 0, \forall \ell > kb$ we have $Set\_d_\ell(\gamma_j) = \emptyset$.*

Proof. By definition, the lemma is established by letting $kb = D$ if e is an execution of $\texttt{Algo}_{\mathcal{B}}(D)$ or $\texttt{Algo}_{\mathcal{HC}}(D)$.

Consider now the case where e is an execution of $\texttt{Algo}_{\mathcal{U}}$. VUB is the set of processes that have no upper bound on their distance value in e, formally, $VUB = \{p \in V \mid \forall k \geq 0, \exists j \geq 0, \exists \ell > k \mid p \in Set\_d_\ell(\gamma_j)\}$

Assume, by the contradiction, that VUB is not empty. The set $V \setminus VUB$ is not empty, by Remark 1. As the network is connected, there are two neighboring processes p and q such that $p \in VUB$ and $q \in V \setminus VUB$. By definition, $\exists x \geq 0$ such that $d_q \leq x$ in all configurations of e. Consequently, $Min\_d(p) \leq x$ in all configurations of e. So, we have $d_p \leq \max\{x+1, y\}$ in all configurations of e, where y be the initial value of d_p (according to the rule U_1). Consequently, $p \notin VUB$, a contradiction.

Hence, VUB is empty. Let ub_p the upper bound on the distance values taken by the process p in e. The lemma holds for $kb = \max_{p \in V}\{ub_p\}$. □

Below, we show that, for every $k \geq 0$, for every execution e of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$, if there is a suffix e' of e where every d-variable whose value is less than k is constant, then there is a suffix e'' of e' where no process switches its d-variable from any non-k value to k.

Lemma 7 *Let $e = \gamma_0, \ldots \gamma_i, \gamma_{i+1}, \ldots$ be an execution of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$. Let $k > 0$. If $\exists i_k \mid \forall j \geq i_k, \forall \ell \in [0..k-1]$ we have $Set\_d_\ell(\gamma_j) = Set\_d_\ell(\gamma_{i_k})$, then $\exists \ell \geq i_k \mid \forall j \geq \ell$ we have $Set\_d_k(\gamma_{j+1}) \subseteq Set\_d_k(\gamma_j)$.*

Proof. Let $\gamma_j \mapsto \gamma_{j+1}$ be any step in the suffix of e starting in γ_{i_k} where $Set\_d_k(\gamma_{j+1}) \not\subseteq Set\_d_k(\gamma_j)$. There is at least a process $p \neq R$ such that $p \notin Set\_d_k(\gamma_j) \wedge p \in Set\_d_k(\gamma_{j+1})$. In γ_j, we have $d_p > k$, as $Set\_d_\ell(\gamma_j) = Set\_d_\ell(\gamma_{j+1})$ $\forall \ell \in [0..k-1]$ and $p \notin Set\_d_k(\gamma_j)$. Moreover, p executes a rule in $\gamma_j \mapsto \gamma_{j+1}$. In the following, we prove that p will no more change its d-variable in e after this step.

- Consider first $\texttt{Algo}_{\mathcal{U}}$. We have $Min\_d(p) = k-1$ in γ_j. Moreover, by hypothesis, $Min\_d(p) = k-1$ forever from γ_{i_k} (so, in particular from γ_j). So, p will no more change its distance value after γ_{j+1}.
- Consider $\texttt{Algo}_{\mathcal{B}}(D)$. If $k = D$, then d_p should be greater than D in γ_j, a contradiction. So, $k < D$ and $d_p > k$ in γ_j. So p executes B_1 in $\gamma_j \mapsto \gamma_{j+1}$, and similarly to the previous case, $\gamma_j \mapsto \gamma_{j+1}$ is the only step in the suffix of e starting in γ_{i_k} where p sets d_p to k.
- Finally, consider $\texttt{Algo}_{\mathcal{HC}}(D)$. We have to study the two following cases:
  - Assume that p executes HC_2 to set par_p to q in $\gamma_j \mapsto \gamma_{j+1}$. By definition, $Min\_d(p) = k-1 = d_q$ holds in γ_{i_k} and all subsequent configurations. So, p is disabled forever from γ_{j+1}.
  - Assume that p executes HC_1 to set par_p to q in $\gamma_j \mapsto \gamma_{j+1}$ By definition, $Min\_d(p) \leq k-1 = d_q < D$ in γ_{i_k} and all subsequent configurations. So, until p next action, we have $parOk(p)$ and $d_p = k$. So, p next action is necessarily HC_2 to set d_p to a value smaller than k, a contradiction. So, p cannot execute any rule in the suffix starting from γ_{j+1}.

Hence, in the suffix of e starting in γ_{i_k}, there is at most n steps $\gamma_j \mapsto \gamma_{j+1}$ where $Set\_d_k(\gamma_{j+1}) \nsubseteq Set\_d_k(\gamma_j)$. □

Below, we show that, for every $k \geq 0$, for every execution e of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$, if eventually every d-variable whose value is less than k becomes constant, then eventually every d-variable whose value is k becomes constant.

Lemma 8 *Let $e = \gamma_0, \ldots \gamma_i, \gamma_{i+1}, \ldots$ be an execution of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, or $\texttt{Algo}_{\mathcal{HC}}(D)$. Let $k > 0$. If $\exists i_k \mid \forall j \geq i_k, \forall \ell \in [0..k-1]$ we have $Set\_d_\ell(\gamma_j) = Set\_d_\ell(\gamma_{i_k})$, then $\exists \ell \geq i_k \mid \forall j \geq \ell$ we have $Set\_d_k(\gamma_{j+1}) = Set\_d_k(\gamma_j)$.*

Proof. By lemma 7, there exists a suffix e' of e starting in γ_x, such that $\forall j \geq x$, we have $Set\_d_k(\gamma_{j+1}) \subseteq Set\_d_k(\gamma_j)$. During e', there is at most $|Set\_d_k(\gamma_{i_x})|$ steps $\gamma_j \mapsto \gamma_{j+1}$ where $Set\_d_k(\gamma_{j+1}) \neq Set\_d_k(\gamma_j)$. □

From Remark 1, Lemmas 6 and 8, we can deduce the following corollary.

Corollary 3 *For every process p, d_p can be modified only a finite number of time in e.*

By Lemma 5, Corollary 3, follows:

Theorem 4 *Under a distributed unfair daemon, all executions of $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, and $\texttt{Algo}_{\mathcal{HC}}(D)$ are finite.*

5 Stabilization Time in Rounds

In this section, we study the stabilization time in rounds of the three algorithms presented in Section 3. Throughout this section we will use the notion of *attractor* defined below.

Let $\mathcal{A}$ be a distributed algorithm. Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be two subsets of $\mathcal{C}$, the set of all possible configurations of $\mathcal{A}$. $\mathcal{C}_2$ is an *attractor* for $\mathcal{C}_1$ (under $\mathcal{A}$) if the following conditions hold:

Convergence: $\forall e = \gamma_0, \gamma_1, \ldots \in \mathcal{E}(\mathcal{C}_1), \exists i \geq 0 \mid \gamma_i \in \mathcal{C}_2$.

Closure: $\forall e = \gamma_0, \gamma_1, \ldots \in \mathcal{E}(\mathcal{C}_2), \forall i \geq 0 \mid \gamma_i \in \mathcal{C}_2$.

The following predicate is useful to establish a sequence of attractors.

$$Pred\_correct\_node(p, i) \equiv (\|R, p\| \leq i \Rightarrow (d_p = \|R, p\| = d_{par_p} + 1))$$

For every $i \geq 0$, $Att(i)$ is the set of configurations, where every process $p \neq R$ satisfies $Pred\_correct\_node(p, i)$.

In any configuration of $Att(\mathcal{D})$, every process $p \neq R$ satisfies $(d_p = \|R, p\| = d_{par_p} + 1)$, moreover $d_r = 0$, by definition. So, all configurations of $Att(\mathcal{D})$ are legitimate. Furthermore, every legitimate configuration is terminal in all of the three algorithms (Theorem 3). Hence, the stabilization time of any of the three algorithms is bounded by the maximum number of rounds it requires to reach any configuration of $Att(\mathcal{D})$ starting from any arbitrary configuration.

5.1 Lower Bound in $\Omega(D)$ Rounds for $\texttt{Algo}_{\mathcal{HC}}(D)$

We first show that the stabilization time in rounds of $\texttt{Algo}_{\mathcal{HC}}(D)$ actually depends on the size of the domain of the d-variables. Hence, we can conclude that $\texttt{Algo}_{\mathcal{HC}}(n)$, *i.e.*, the algorithm proposed in [9], stabilizes in $\Omega(n)$ rounds, where n is the number of processes.

Let $k \geq 1$. We now exhibit a possible execution of $\texttt{Algo}_{\mathcal{HC}}(2k)$ which stabilizes in $k+1$ rounds in the 3-nodes graph given in Figure 1 (its diameter is 2). Notice that this execution requires $2k$ steps.

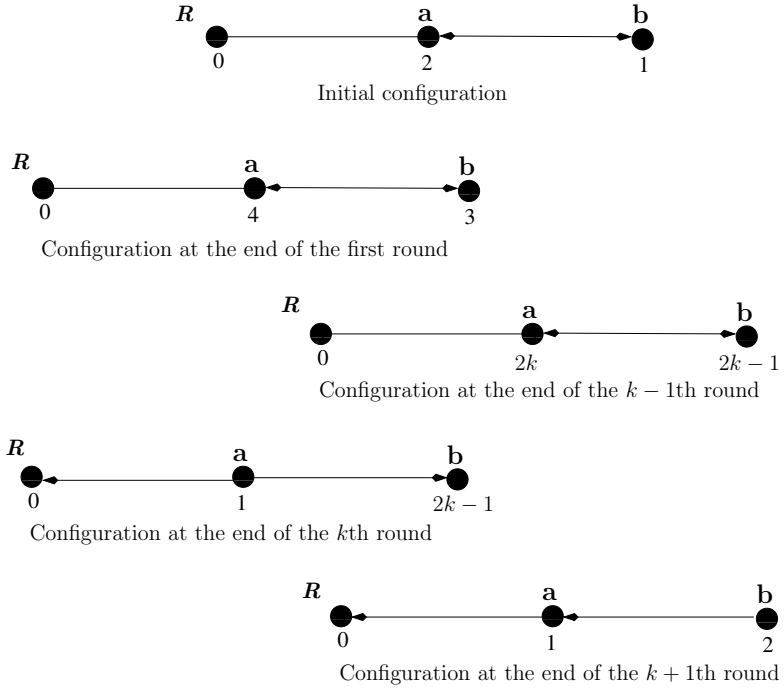

Figure 1: Execution of $\texttt{Algo}_{\mathcal{HC}}(2k)$ which converges in $k+1$ rounds

- At the beginning of the ith round with $i \in [1, k-1]$, processes a and b are enabled. In the first step of these rounds, b executes HC_1. During the second step (the last step of these rounds) the node a executes HC_1.
- At the beginning of the kth round, only process a is enabled. During the only step of this round, a executes HC_2 and gets its terminal state.
- At the beginning of the $k+1$th round, only process b is enabled. During the only step of this round, process b executes and gets its terminal state.

This example can be generalized to any number of processes $n \geq 3$. Just construct a network G of n processes by adding $n-3$ processes to the network given in Figure 1; those $n-3$ processes being only neighbors of R. Since the state of R is constant, these $n-3$ processes have no impact on the behavior of a and b. Hence, the previous execution is a possible execution prefix in G which contains $\Omega(D)$ rounds.

Hence, the stabilization time of $\texttt{Algo}_{\mathcal{HC}}(D)$ is $\Omega(D)$ rounds.

The lower bound on the stabilization time is mainly due to the fact that rules HC_1 and HC_2 are not mutually exclusive. Hence, when both are enabled at the same process p, the daemon may choose to activate any of them. Our lower bound is then established when the daemon makes priority on HC_1.

In the following subsection, we show that this lower bound can be easily circumvented to obtain the stabilization time in $\Theta(\mathcal{D})$ rounds.

5.2 Fast Implementation of $\texttt{Algo}_{\mathcal{HC}}(D)$

5.2.1 Algorithm $\texttt{Algo}_{\mathcal{FHC}}(D)$

Below, we propose a variant of $\texttt{Algo}_{\mathcal{HC}}(D)$ where we have modified HC_1 into FHC_1, so that FHC_1 and HC_2 are now mutually exclusive. The modification of HC_1 into FHC_1 gives it priority on HC_2. In the following, this variant will be denoted by $\texttt{Algo}_{\mathcal{FHC}}(D)$ and called *fast implementation of* $\texttt{Algo}_{\mathcal{HC}}(D)$.

$$\begin{array}{rcll} FHC_1 & :: & \neg parOk(p) \wedge d_{par_p} < D \wedge d_{par_p} = Min\_d(p) & \rightarrow \quad d_p \leftarrow d_{par_p} + 1 \\ HC_2 & :: & d_{par_p} > Min\_d(p) & \rightarrow \quad update(p) \end{array}$$

The lemma given below show the close relationship between $\texttt{Algo}_{\mathcal{FHC}}(D)$ and $\texttt{Algo}_{\mathcal{B}}(D)$.

Lemma 9 *If $\gamma \mapsto \gamma'$ is a step of $\texttt{Algo}_{\mathcal{FHC}}(D)$ containing execution of rules $HC2$ only, then $\gamma \mapsto \gamma'$ is a possible step of $\texttt{Algo}_{\mathcal{B}}(D)$.*

Proof. Let $\gamma \mapsto \gamma'$ be any step of $\texttt{Algo}_{\mathcal{FHC}}(D)$ containing execution of rules $HC2$ only. Consider any process p that moves during $\gamma \mapsto \gamma'$. So, p performs HC_2 during $\gamma \mapsto \gamma'$ and we have $Min\_d(p) < d_{par_p} \leq D$ in γ.

If $\neg dOk(p)$ is true in γ, then B_1 is enabled at p in γ. Now, the action part of B_1 and HC_2 are identical.

Conversely, assume that $dOk(p)$ is true in γ. During $\gamma \mapsto \gamma'$, p does not modify d_p, however par_p is set to $bestParent(p)$. Now, $\neg parOk(p)$ is true in γ, so B_2 is enabled while B_1 is not. In this case, the action part of B_2 has the same effect as the action part of HC_2.

Hence, in any case, $\gamma \mapsto \gamma'$ is a possible step of $\texttt{Algo}_{\mathcal{B}}(D)$. □

5.2.2 Upper Bound on Stabilization Time in Rounds of $\texttt{Algo}_{\mathcal{FHC}}(D)$

The $Pred\_UB\_d(p,i)$ and $Pred\_correct\_d(p,i)$ predicate defined below are used to establish a sequence of $\mathcal{D}+1$ attractors under $\texttt{Algo}_{\mathcal{FHC}}(D)$ (with $D \geq \mathcal{D}$) ending in the set of the terminal configurations.

$$Pred\_UB\_d(p,i) \equiv (\|R,p\| > i \Rightarrow (d_p > i \vee (d_p = i \wedge (\exists q \in \mathcal{N}_p \mid d_q \leq i+1))))$$

$Pred\_UB\_d(p,i)$ means that if a process p is at distance larger than i from R, then either d_p should be also larger than i, or d_p should be equal to i and a neighbor of p should have its distance to R smaller than or equal to $i+1$.

$$Pred\_correct\_d(p,i) \equiv (\|R,p\| \leq i \Rightarrow d_p = \|R,p\|)$$

$Pred\_correct\_d(p,i)$ means that if a process p is at least at distance i from R, then its distance value should be correct, *i.e.*, d_p is equal to its distance to the R.

Below, we define some useful subsets of configurations.

- Let $Att\_UB(i)$ be the set of configurations, where every process $p \neq R$ satisfies $Pred\_UB\_d(p,i)$.
- Let $Att\_dist(i)$ be the set of configurations, where every process p satisfies $Pred\_correct\_d(p,i)$.
- Let $Att\_HC(i) = Att\_dist(i) \cap Att\_UB(i)$.

Notice that $Att\_HC(0)$ is the set of all possible configurations.

Observation 1 *Let p be an process such that $\|p,R\| > i+1$. So, we have $i+2 \leq \mathcal{D}$. Let γ be a configuration of $Att\_HC(i)$. By definition of $Att\_HC(i)$, we have $d_p \geq i$ and $d_q \geq i$ for every $q \in \mathcal{N}_p$ in every execution from γ. Consequently, $Min\_d(p) \geq i$ along any execution from γ.*

Lemma 10 *Assume that $Att\_HC(i)$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ with $D \geq \mathcal{D}$ and $0 \leq i < \mathcal{D}$. Let γ be a configuration of $Att\_HC(i)$. Let $\gamma \mapsto \gamma'$ be a possible step where process p moves. $Pred\_UB\_d(p,i+1)$ holds in γ'.*

Proof. Let p be a process such that $\|p,R\| > i+1$ (the other case is trivial). We have $d_p \geq i+1$ in γ' according to observation 1.
If $d_p > i+1$ in γ', then $Pred\_UB\_d(p,i+1)$ holds. Otherwise, we have $d_p = i+1$ in γ'. Let $q \in \mathcal{N}_p$ such that $d_q = Min\_d(p)$ in γ. We have $d_q = i$ in γ. By definition of $Att\_HC(i)$, we have $Min\_d(p) \leq i+2$ in γ'. Hence, $Pred\_UB\_d(p,i+1)$ holds in γ'. □

Lemma 11 *If $Att\_HC(i)$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ with $D \geq \mathcal{D}$ and $0 \leq i < \mathcal{D}$, then $Att\_UB(i+1)$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ from $Att\_HC(i)$ which is reached within at most one round from $Att\_HC(i)$.*

Proof.

Closure. Let $\gamma \mapsto \gamma'$ be a possible step from any configuration γ of $Att\_HC(i)$. We show that for every process $p \neq R$, if $Pred\_UB\_d(p,i+1)$ holds in γ, then $Pred\_UB\_d(p,i+1)$ holds in γ'. Assume $\|p,R\| > i+1$ (the other case is trivial). Assume that p does not move during the step; otherwise $Pred\_UB\_d(p,i+1)$ holds in γ' according to Lemma 10. If $d_p > i+1$ in γ, $Pred\_UB\_d(p,i+1)$ holds in γ'. Assume now that in γ, $d_p = i+1$ and p has a neighbor q such that $d_q \leq i+2$. While $d_p = i+1$, $Min\_d(q) \leq i+1$, so $d_q \leq i+2$. Hence, we can conclude that $Pred\_UB\_d(p,i+1)$ still holds in γ'.

Convergence. We now show that for every process $p \neq R$, $Pred\_UB\_d(p,i+1)$ becomes true within at most one round from any configuration γ of $Att\_HC(i)$. Assume $\|p,R\| > i+1$ (the other case is trivial).

- If $d_p > i+1$ in γ, then $Pred\_UB\_d(p,i+1)$ holds.
- If $d_p = i$ in γ, then p is enabled while $d_p = i$ because $Min\_d(p) \geq i$ forever from γ (Observation 1). So, p moves during the first round from γ, and we are done, by Lemma 10.

- Assume that $d_p = i + 1$ in γ.
  - If p moves during the first round from γ, then $Pred\_UB\_d(p, i + 1)$ holds after the step, by Lemma 10.
  - Assume that p does not move during the first round from γ.
    * If a neighbor of p, q, moves during a step of the round, then after this step $Pred\_UB\_d(p, i + 1)$ holds because $d_q \leq i + 2$.
    * Assume that neither p nor its neighbors move during the round. So the value of d_q is less than or equal to $i + 2$ in γ, for all $q \in \mathcal{N}_p$. Indeed, if some neighbor of p, q, satisfies $d_q > i + 2$ in γ, then q stay enabled along the round from γ, because of the state of p. This contradicts the definition of round. Hence, the value of d_q is less than or equal to $i + 2$ in γ, for all $q \in \mathcal{N}_p$, and consequently $Pred\_UB\_d(p, i + 1)$ holds in γ.

□

Lemma 12 *If $Att\_HC(i)$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ with $D \geq \mathcal{D}$ and $0 \leq i < \mathcal{D}$, then $Att\_dist(i+1)$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ from $Att\_HC(i)$ which is reached within at most one round from any configuration of $Att\_HC(i)$.*

Proof. Let p_{i+1} be a process at distance $i + 1$ of R. p_{i+1} has at least a neighbor p_i such that $\|p_i, R\| = i$. Let $\gamma \in Att\_HC(i)$. By definition of $Att\_HC(i)$ and as $Att\_HC(i)$ is an attractor, we can deduce that $d_{p_i} = i$ and $\forall q \in \mathcal{N}_{p_{i+1}}, d_q \geq i$ forever from γ. So, from γ $Min\_d(p_{i+1}) = i$ forever.

Consequently, if $d_{p_{i+1}} \neq i + 1$ in γ, then p_{i+1} is enabled to execute FHC_1 or HC_2 to set $d_{p_{i+1}}$ to $i + 1$.

Moreover, if $d_{p_{i+1}} = i + 1$ in γ, then p_{i+1} cannot modify $d_{p_{i+1}}$ in any step from γ.

Hence, $Att\_dist(i + 1)$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ from $Att\_HC(i)$ which is reached within at most one round from any configuration of $Att\_HC(i)$. □

From the two previous lemmas, we can deduce the following corollary.

Corollary 4 *If $Att\_HC(i)$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ with $D \geq \mathcal{D}$ and $0 \leq i < \mathcal{D}$, then $Att\_HC(i+1)$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ from $Att\_HC(i)$ which is reached within at most one round from any configuration of $Att\_HC(i)$.*

The previous corollary establishes that after at most $\mathcal{D}$ rounds, the distance value in every process is accurate forever. We now show one additional is necessary to fix the *par*-variables.

Lemma 13 *$Att(\mathcal{D})$ is an attractor under $\texttt{Algo}_{\mathcal{FHC}}(D)$ (with $D \geq \mathcal{D}$) from $Att\_HC(\mathcal{D})$ which is reached within at most one round from any configuration of $Att\_HC(\mathcal{D})$.*

Proof. In any configuration of $Att\_HC(\mathcal{D})$, $d_p = Min\_d(p) + 1$ holds forever for every process p. So, the distance value of any process stays unchanged along any execution of $\texttt{Algo}_{\mathcal{FHC}}(D)$ from a configuration of $Att\_HC(\mathcal{D})$.

Let γ be a configuration of $Att\_HC(\mathcal{D})$ where $d_{par_p} \neq Min\_d(p)$. The rule HC_2 is enabled at p until p executes it. After the execution of this rule, we have $d_p = d_{par_p} + 1$. As no process changes its distance value in $Att\_HC(\mathcal{D})$, p is become disabled forever.

Hence, we conclude that $Att(\mathcal{D})$ is an attractor which is reached within at most one round from $Att\_HC(\mathcal{D})$. □

From Corollary 4 and Lemma 13, we have the following theorem:

Theorem 5 *For every $D \geq \mathcal{D}$, the stabilization time of $\texttt{Algo}_{\mathcal{FHC}}(D)$ is at most $\mathcal{D}+1$ rounds.*

5.2.3 Lower Bound on Stabilization Time in Rounds of $\texttt{Algo}_{\mathcal{FHC}}(D)$

Below, we show that the upper bound given in the previous theorem is exact when $D = \mathcal{D}$: $\forall \mathcal{D} \geq 2$, there exists an execution of $\texttt{Algo}_{\mathcal{HC}}(\mathcal{D})$ in a graph of diameter $\mathcal{D}$ that stabilizing in $\mathcal{D}+1$ rounds.

We consider any graph $G = (V, E)$ of $\mathcal{D}+2$ nodes of diameter $\mathcal{D} \geq 1$, where

- $V = \{p_0 = R, \ldots, p_{\mathcal{D}+1}\}$, and
- $E = \{\{p_i, p_{i+1}\} | i \in [0..\mathcal{D}]\} \cup \{\{p_{\mathcal{D}+1}, p_{\mathcal{D}-1}\}\}$.

We consider a synchronous execution (*i.e.* an execution where the distributed unfair daemon activates all enabled processes at each step) which starts from the following initial configuration:

- $d_R = 0$,
- $\forall i \in [1..\mathcal{D}-2], par_{p_i} = p_{i-1} \wedge d_{p_i} = \mathcal{D}$,
- $par_{p_{\mathcal{D}-1}} = p_{\mathcal{D}+1} \wedge d_{p_{\mathcal{D}-1}} = \mathcal{D}$,
- $par_{p_{\mathcal{D}}} = p_{\mathcal{D}+1} \wedge d_{p_{\mathcal{D}}} = \mathcal{D}$, and
- $par_{p_{\mathcal{D}+1}} = p_{\mathcal{D}-1} \wedge d_{p_{\mathcal{D}}} = \mathcal{D}-1$.

An example of initial configuration is given in Figure 2. Notice that in a synchronous execution, every round lasts one step.

- At each round $i \in [1..\mathcal{D}-2]$, p_i executes FHC_1 to change d_{p_i} to i.
- At the $\mathcal{D}-1$th round, $p_{\mathcal{D}-1}$ executes HC_2 to set $d_{p_{\mathcal{D}-1}}$ to $\mathcal{D}-1$ and $par_{p_{\mathcal{D}-1}}$ to $p_{\mathcal{D}-2}$.
- At the $\mathcal{D}$th round, $p_{\mathcal{D}+1}$ executes FHC_1 to set $d_{p_{\mathcal{D}+1}}$ to $\mathcal{D}$.
- At the $\mathcal{D}+1$th round, $p_{\mathcal{D}}$ executes HC_2 to set $par_{p_{\mathcal{D}}}$ to $p_{\mathcal{D}-1}$.

Hence, we can conclude with the theorem below.

Theorem 6 *The worst case stabilization time of $\texttt{Algo}_{\mathcal{FHC}}(\mathcal{D})$ is $\mathcal{D}+1$ rounds.*

5.3 Algorithms $\texttt{Algo}_{\mathcal{U}}$ and $\texttt{Algo}_{\mathcal{B}}(D)$

We now establish that the stabilization time of both $\texttt{Algo}_{\mathcal{U}}$ and $\texttt{Algo}_{\mathcal{B}}(D)$ is exactly $\mathcal{D}$ rounds in the worst case.

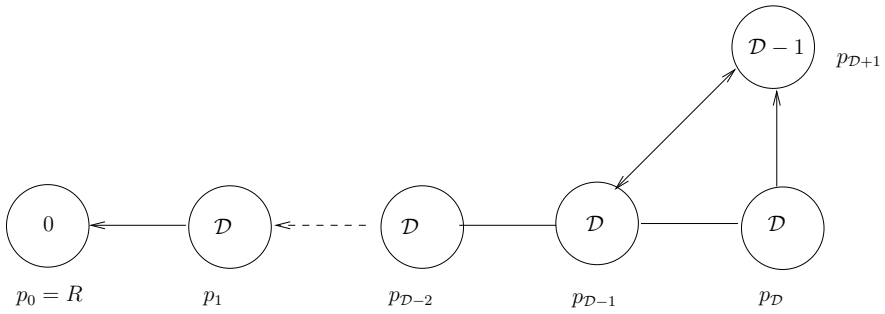

Figure 2: Initial configuration of a synchronous execution of $\texttt{Algo}_{\mathcal{FHC}}(\mathcal{D})$ which stabilizes in $\mathcal{D}+1$ rounds

5.3.1 Upper bound on the Stabilization Time in Rounds for both $\texttt{Algo}_{\mathcal{U}}$ and $\texttt{Algo}_{\mathcal{B}}(D)$

We first establish that the stabilization time of both $\texttt{Algo}_{\mathcal{U}}$ and $\texttt{Algo}_{\mathcal{B}}(D)$ is at most $\mathcal{D}$ rounds in the worst case. To that goal, we use the predicate $Pred\_SUB\_d(p,i)$ defined below:

$$Pred\_SUB\_d(p,i) \equiv (\|R,p\| > i \Rightarrow d_p > i)$$

$Pred\_SUB\_d(p,i)$ means that d_p must be larger than i if the process p is at a distance larger than i from R.

We will also use the following sets:

- Let $Att\_SUB(i)$ be the set of configurations where every process p satisfies $Pred\_SUB\_d(p,i)$).
- Let $Att\_B(i) = Att(i) \cap Att\_SUB(i)$.

Notice that all configurations belong to $Att\_SUB(0)$.

Lemma 14 *If $Att\_B(i)$ is an attractor under $\texttt{Algo}_{\mathcal{U}}$ (resp. $\texttt{Algo}_{\mathcal{B}}(D)$ where $D \geq \mathcal{D}$) with $0 \leq i < \mathcal{D}$, then $Att(i+1)$ is an attractor under $\texttt{Algo}_{\mathcal{U}}$ (resp. $\texttt{Algo}_{\mathcal{B}}(D)$) from $Att\_B(i)$ which is reached within at most one round from any configuration of $Att\_B(i)$.*

Proof. Let p_{i+1} be a process at distance $i+1$ of R. By definition, p_{i+1} has at least one neighbor p_i at distance i of R. As $Att\_B(i)$ is an attractor, from any configuration of $Att\_B(i)$, the three following conditions hold forever: (i) $d_{p_i} = i$, (ii) $Min\_d(p_{i+1}) = i < D$, and (iii) for every process q, $d_q = i \Rightarrow \|q,R\| = i$.

Let $\gamma \mapsto \gamma'$ be a possible step such that γ is a configuration of $Att\_B(i)$. We first show that for every process p, if $Pred\_correct\_node(p_{i+1}, i+1)$ holds in γ, then $Pred\_correct\_node(p_{i+1}, i+1)$ holds in γ'. From γ, $d_{par_{p_{i+1}}} = i$, so $d_{par_{p_{i+1}}}$ is no more modified by (iii). In γ, $d_{p_{i+1}} = i+1$ and $d_{p_{i+1}}$ is no more modified by (ii). Hence, $d_{p_{i+1}} = d_{par_{p_{i+1}}} + 1$ and p is disabled forever from γ. Hence, $Pred\_correct\_node(p_{i+1}, i+1)$ still holds in γ'.

We now show that for every process p, $Pred\_correct\_node(p_{i+1}, i+1)$ becomes true within at most one round from any configuration γ of $Att\_B(i)$. Assume that $d_{p_{i+1}} \neq i+1$ or $d_{par_{p_{i+1}}} \neq i$ in γ. Then, p_{i+1} is enabled in $\texttt{Algo}_{\mathcal{U}}$ (resp. $\texttt{Algo}_{\mathcal{B}}(D)$) until it executes an action, by (i) and (ii). Moreover, after p_{i+1} move, we have $d_{p_{i+1}} = i+1 = d_{par_{p_{i+1}}} + 1$, by (i) and (ii). Hence, $Pred\_correct\_node(p_{i+1}, i+1)$ becomes true within at most one round from γ. □

Lemma 15 *If $Att\_B(i)$ is an attractor under $\texttt{Algo}_{\mathcal{U}}$ (resp. $\texttt{Algo}_{\mathcal{B}}(D)$ where $D \geq \mathcal{D}$) with $0 \leq i < \mathcal{D}$, then $Att\_SUB(i+1)$ is an attractor under $\texttt{Algo}_{\mathcal{U}}$ (resp. $\texttt{Algo}_{\mathcal{B}}(D)$) from $Att\_B(i)$ which is reached within at most one round from any configuration of $Att\_B(i)$.*

Proof. Let p be a process such that $\|p, R\| > i+1$. In this case, we have $i+2 \leq \mathcal{D} \leq D$.

In any configuration of $Att\_B(i)$, we have $d_p > i$ and $d_q > i$ for any neighbor q of p by definition of $Att\_B(i)$. So, starting from any configuration γ of $Att\_B(i)$, $Min\_d(p) > i$ holds forever.

Hence, if $d_p > i+1$ in γ, then $d_p > i+1$ forever from γ, which implies that if $Att\_SUB(i+1)$ holds in γ, then $Att\_SUB(i+1)$ holds forever from γ.

Assume now that $d_p = i+1$ in γ. Then, $d_p = i+1 < D$ and, as $Min\_d(p) > i$ holds forever from γ, p is continuously enabled from γ until it executes either U_1 in $\texttt{Algo}_{\mathcal{U}}$, or $B_j, j \in \{1,2\}$ in $\texttt{Algo}_{\mathcal{B}}(D)$. After p move, $d_p \geq i+1$. Hence, $Att\_SUB(i+1)$ holds within at most one round from γ. □

Corollary 5 *If $Att\_B(i)$ is an attractor under $\texttt{Algo}_{\mathcal{U}}$ (resp. of $\texttt{Algo}_{\mathcal{B}}(D)$ where $D \geq \mathcal{D}$) with $0 \leq i < \mathcal{D}$, then $Att\_B(i+1)$ is an attractor under $\texttt{Algo}_{\mathcal{U}}$ (resp. $\texttt{Algo}_{\mathcal{B}}(D)$) from $Att\_B(i)$ which is reached within at most one round from any configuration of $Att\_B(i)$.*

From the previous corollary and owing the fact that $Att\_B(\mathcal{D}) = Att(\mathcal{D})$, we can deduce the following theorem:

Theorem 7 *The stabilization time of $\texttt{Algo}_{\mathcal{U}}$ and $\texttt{Algo}_{\mathcal{B}}(D)$ (for every $D \geq \mathcal{D}$) is at most $\mathcal{D}$ rounds.*

5.3.2 Lower Bound on the Stabilization Time in Rounds for both $\texttt{Algo}_{\mathcal{U}}$ and $\texttt{Algo}_{\mathcal{B}}(D)$

Below, we show that the upper bound is exact for both $\texttt{Algo}_{\mathcal{U}}$ and $\texttt{Algo}_{\mathcal{B}}(D)$ when $D \geq \mathcal{D}$.

Consider first $\texttt{Algo}_{\mathcal{U}}$. Let $G = (V, E)$ be any line graph of diameter $\mathcal{D}$, *i.e.*, $V = \{p_0 = R, p_1, \ldots, p_{\mathcal{D}}\}$ and $E = \{\{p_i, p_{i+1}\} \mid i \in [0..\mathcal{D}-1]\}$. Consider the initial configuration where $d_{p_0} = 0$ and $\forall i \in [1..\mathcal{D}]$, $d_{p_i} = X$, where $X > \mathcal{D}$. (Par-variables have arbitrary values). Consider a synchronous execution starting from that initial configuration. Then, at each round i, with $i \in [1..\mathcal{D}]$, p_i executes U_1 to definitely set d_{p_i} to i and par_{p_i} to p_{i-1}. Moreover, $\forall j \in [i+1..\mathcal{D}]$, p_j increments d_{p_j} by U_1. Hence, after $\mathcal{D}$ rounds the system is in a terminal configuration.

Consider now $\texttt{Algo}_{\mathcal{B}}(D)$ with any value $D \geq \mathcal{D}$. Consider the same graph as for $\texttt{Algo}_{\mathcal{FHC}}(\mathcal{D})$. However, we consider now a synchronous execution starting from any configuration where:

- $d_R = 0$,
- $\forall i \in [1..\mathcal{D}+1], d_{p_i} = D$,
- $par_{p_{\mathcal{D}}} = p_{\mathcal{D}+1}$, and
- $par_{p_{\mathcal{D}+1}} = p_{\mathcal{D}}$.

An example of initial configuration is given in Figure 3. The synchronous execution starting from that configuration then works as follows:

- In round i, with $i \in [1..\mathcal{D}-1]$, only process p_i moves. It executes rule B_1 to set d_{p_i} to i and par_{p_i} to p_{i-1}.

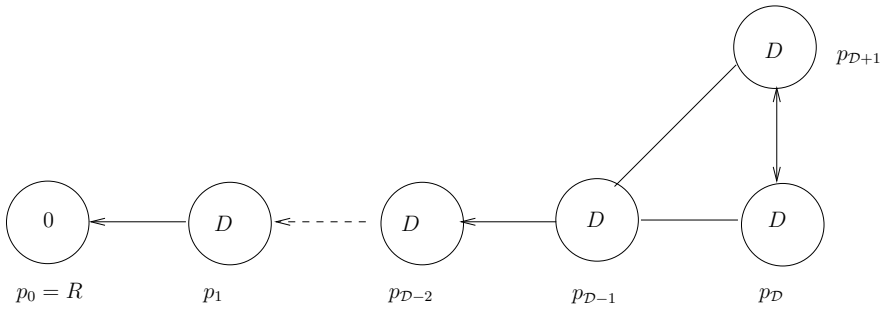

Figure 3: Initial configuration of a synchronous execution of $\texttt{Algo}_{\mathcal{B}}(D)$ which stabilizes in $\mathcal{D}$ rounds

- In round $\mathcal{D}$, only $p_{\mathcal{D}}$ and $p_{\mathcal{D}+1}$ move. Two cases are possibles. Either $D > \mathcal{D}$ and they both execute B_1 to set $d_{p_{\mathcal{D}}}$ (resp. $d_{p_{\mathcal{D}+1}}$) to $\mathcal{D}$ and $par_{p_{\mathcal{D}}}$ (resp. $par_{p_{\mathcal{D}+1}}$) to $p_{\mathcal{D}-1}$. Or, $D = \mathcal{D}$ and they both execute B_2 to set $par_{p_{\mathcal{D}}}$ and $par_{p_{\mathcal{D}+1}}$ to $p_{\mathcal{D}-1}$.

Hence, we can conclude with the theorem below.

Theorem 8 *The worst case stabilization time of* $\texttt{Algo}_{\mathcal{U}}$ *and* $\texttt{Algo}_{\mathcal{B}}(D)$ *(with* $D \geq \mathcal{D}$*) is* $\mathcal{D}$ *rounds.*

6 Stabilization Time in Steps

In this section, we propose a step complexity analysis of the three algorithms presented in Section 3.

6.1 A General Bound

The theorem below exhibits a trivial upper bound on the stabilization time in steps of every self-stabilizing algorithm working under an unfair daemon.

Theorem 9 *Let* $\mathcal{A}$ *be any self-stabilizing algorithm under an unfair daemon,*[3] *the stabilization time of* $\mathcal{A}$ *is less than or equal to* $\prod_{p \in V} |S_p| - 2$ *steps, where* S_p *is the set of possible states of* p*, for every process* p*.*

Proof. First, the number of possible configurations of $\mathcal{A}$ is $\prod_{p \in V} |S_p|$. Let e be any execution of $\mathcal{A}$. $\mathcal{A}$ being self-stabilizing, e contain a maximal prefix of finite size $e' = \gamma_i, \gamma_{i+1} \dots$ where its specification is not achieved. Let e'' such that $e = e'e''$.

Assume, by the contradiction, that $\exists k, \ell$ such that $i \leq k < \ell$ and $\gamma_k = \gamma_\ell$. Then, $(\gamma_{k+1}, \dots, \gamma_\ell)^\infty$ is an infinite execution of $\mathcal{A}$ under the unfair daemon that never stabilized. So, $\mathcal{A}$ is not self-stabilizing under an unfair daemon, a contradiction.

Hence, all configurations of e' are distinct. Moreover, $|e''| \geq 1$ and e' and e'' have no common configuration. Hence, e' contains at most $\prod_{p \in V} |S_p| - 1$ configurations, and so at most $\prod_{p \in V} |S_p| - 2$ steps. □

The previous theorem is useless when considering algorithms where at least one variable as an infinite domain, *e.g.*, $\texttt{Algo}_{\mathcal{U}}$. Now, for $\texttt{Algo}_{\mathcal{B}}(D)$ and $\texttt{Algo}_{\mathcal{HC}}(D)$, the theorem claims that their respective stabilization times are less than or equal to $(n-1)^{\Delta.D}$ steps. This upper bound may appear to be overestimated at the first glance. However, we will see in the next subsections that those algorithms are exponential in steps in the worst case.

[3] The daemon can be central or distributed.

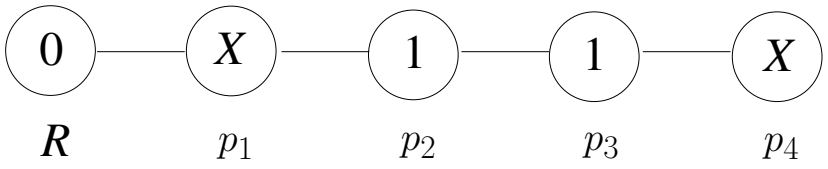

Figure 4: Possible initial Configuration of the line of 5 nodes

6.2 Algorithm $\texttt{Algo}_{\mathcal{U}}$

Here we consider the unbounded version given in Subsection 3.1. The following theorem shows that the step complexity of $\texttt{Algo}_{\mathcal{U}}$ cannot be bounded by any function depending on topological parameters, *e.g.*, n, N, $\mathcal{D}$, or $D\ldots$

Theorem 10 *Let f be any function mapping graphs to integers. There exists a graph G and an execution e of $\texttt{Algo}_{\mathcal{U}}$ in G such that e stabilizes in more than $f(G)$ steps.*

(The following proof is illustrated with Figure 4.)
Proof. Consider a line graph G of 5 nodes, where R is an extremity, *i.e.*, $G = \{R, p_1, p_2, p_3, p_4\}$ and $E = \{\{R, p_1\}\} \cup \{\{p_i, p_{i+1}\}, i \in [1..3]\}$. Let $X \geq f(G) + 1$. Assume an initial configuration, where $d_{p_1} = d_{p_4} = X$ and $d_{p_2} = d_{p_3} = 1$. (The initial values of *par* variables are arbitrary and, by definition, $d_R = 0$.) Initially, all processes, except R, are enabled. Assume that p_2 moves, then in the next configuration, d_{p_2} takes value 2 and all processes, except R and p_2, are enabled. Assume that p_3 moves, then in the next configuration, d_{p_3} takes value 3 and all processes, except R and p_3, are enabled. By alternating activations of p_2 and p_3 the system reaches in $X \geq f(G) + 1$ steps a configuration where $d_R = 0$, $d_{p_1} = d_{p_4} = X$, and $d_{p_2} = d_{p_3} = X + 1$. □

6.3 Algorithm $\texttt{Algo}_{\mathcal{HC}}(D)$

In this subsection, we establish that the stabilization times in steps of both $\texttt{Algo}_{\mathcal{HC}}(D)$ and $\texttt{Algo}_{\mathcal{B}}(D)$ are exponential in the worst case. The lowed bound is based on a family of graphs called $\mathcal{G}_k$. For every $k \geq 0$, the graph $\mathcal{G}_k$ contains $4k + 3$ processes and has a diameter of $2k + 3$.

Definition 2 (Graph $\mathcal{G}_1$) *Let $\mathcal{G}_1 = (V_1, E_1)$ be the undirected graph, where*

- $V_1 = \{f.0, e.1, f.1, h.0, g.1, h.1, R\}$ *and*
- $E_1 = \{\{R, h.0\}, \{h.0, f.0\}, \{f.0, e.1\}, \{e.1, f.1\}, \{f.1, h.1\}, \{g.1, e.1\}\}$.

We now consider three classes of configurations for the graph $\mathcal{G}_1$. In all consider configurations:

- the distance value of $g.1$, $h.1$, and $h.0$ are $z - 1$,
- $d_{e.1} = z$ if and only if $par_{e.1} = g.1$, and
- $d_{f.i} = z$ if and only if $par_{f.i} = h.i$, for $i \in [0, 1]$.

The three classes of configurations are defined as follows where $x \geq 1$ and $z > 1$:

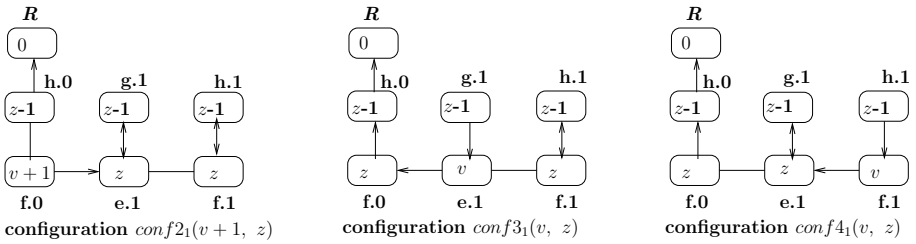

Figure 5: Examples of configurations of $\mathcal{G}_1$

- In the configurations of $conf2_1(x, z)$, the distance value of $e.1$ and $f.1$ is z, and the distance value of $f.0$ is x.
- In the configurations of $conf3_1(x, z)$, the distance value of $e.1$ is x, and the distance value of $f.1$ and $f.0$ is z.
- In the configurations of $conf4_1(x, z)$, the distance value of $e.1$ and $f.0$ is z, and the distance value of $f.1$ is x.

Except otherwise mentioned, all other variables have arbitrary values. Notice that we have $conf2(z, z) = conf3_1(z, z) = conf4_1(z, z)$. An illustrative example of these three types of configurations of $\mathcal{G}_1$ is given in Figure 5.

Observation 2 *Let v, z, D be three integers such that $1 \geq v < z \leq D$.*

- *From any configuration of $conf3_1(v, z)$, a configuration of $conf2_1(v + 1, z)$ is reachable in a single step of $\texttt{Algo}_{\mathcal{HC}}(D)$, where $e.1$ and $f.0$ execute HC_2.*
- *From a configuration of $conf2_1(v, z)$, a configuration of $conf3_1(v+1, z)$ is reached in a single step of $\texttt{Algo}_{\mathcal{HC}}(D)$, where $e.1$ and $f.0$ execute HC_2.*

Notation 2 *Let v and z be two integers such that $1 \leq v \leq z$ and $z > 1$. Let $nbSteps(v, z, 1)$ be the maximal number of steps of $\texttt{Algo}_{\mathcal{HC}}(D)$ (with $D \geq z$) to reach a configuration of $conf3_1(z, z)$ from a configuration of $conf3_1(v, z)$.*

Observation 3 *Let v and z be two integers such that $1 \leq v \leq z$ and $z > 1$. We have $nbSteps(v + 2, z + 2, 1) = nbSteps(v, z, 1)$.*

Lemma 16 *In $\mathcal{G}_1$, for every $1 \leq v \leq z - 2$, there is a execution $e_1(k)$ of $\texttt{Algo}_{\mathcal{HC}}(D)$ (with $D \geq z$), starting in a configuration of $conf3_1(v, z)$ and where only rules HC_2 are executed, which reaches a configuration of $conf3_1(v + 2, z)$ in at least 2 steps.*

Proof. Immediate from Observation 2. □

Corollary 6

- *If $1 \leq v \leq z$ and $z > 1$ then $nbSteps(v, z + 2, 1) \geq z - v$.*

- *Let $k \geq 1$. In $\mathcal{G}_1$, there is a execution of $\texttt{Algo}_{\mathcal{HC}}(D)$, with $D \geq 2k+3$, which starts in a configuration of $conf3_1(1, 2k+3)$, contains only executions of rules HC_2, and reaches a configuration of $conf3_1(2k+3, 2k+3)$ in at least $2k+2$ steps.*
- $nbSteps(1, 5, 1) = 4$.

The following definition generalizes Definition 2.

Definition 3 (Graph $\mathcal{G}_{i+1}$) *Let $\mathcal{G}_{i+1} = (V_{i+1}, E_{i+1})$ be the undirected graph, where*

- $V_{i+1} = V_i \cup \{e.i+1, f.i+1, g.i+1, h.i+1\}$ *and*
- $E_{i+1} = E_i \cup E'_{i+1}$, *where* $E'_{i+1} = \{\{f.i, e.i{+}1\}, \{g.i{+}1, e.i{+}1\}, \{e.i{+}1, f.i{+}1\}, \{f.i{+}1, h.i{+}1\}\}$.

We mainly consider four classes of configurations for any graph $\mathcal{G}_{i+1}$. In all consider configurations:

- the distance value of $g.i+1$ and $h.i+1$ is $z-1$,
- for every $j \in [0, i+1]$, $d_{e.j} = z$ if and only if $par_{e.j} = g.j$, and
- for every $j \in [0, i+1]$, $d_{f.j} = z$ if and only if $par_{f.j} = h.j$.

The four classes of configurations are then defined as follows where $x \geq 1$ and $z > 1$:

- In the configurations of $conf1_{i+1}(x, z)$, the configuration of the subgraph $\mathcal{G}_i$ belongs to $conf3_i(x, z)$, the distance value of $e.i+1$ is x and the distance value of $f.i+1$ is z.
- In the configurations of $conf2_{i+1}(x, z)$, the configuration of the subgraph $\mathcal{G}_i$ belongs to $conf4_i(x, z)$, the distance value of $e.i+1$ and $f.i+1$ is z.
- In the configurations of $conf3_{i+1}(x, z)$, the configuration of the subgraph $\mathcal{G}_i$ belongs to $conf3_i(z, z)$, the distance value of $e.i+1$ is x, and the distance value of $f.i+1$ is z.
- In the configurations of $conf4_{i+1}(x, z)$, the configuration of the subgraph $\mathcal{G}_i$ belongs to $conf3_i(z, z)$, the distance value of $e.i+1$ is z, and the distance value of $f.i+1$ is x.

Except otherwise mentioned, all other variables have arbitrary values. Notice that we have $conf1_{i+1}(z, z) = conf2_{i+1}(z, z)$ and $conf3_{i+1}(z, z) = conf4_{i+1}(z, z)$. Some illustrative examples are given in Figures 6.

Observation 4 *Let v, z, D be three integers such that $1 \leq v < z \leq D$.*

- *From any configuration of $conf3_{i+1}(v, z)$, a configuration of $conf2_{i+1}(v+1, z)$ is reachable in a single step of $\texttt{Algo}_{\mathcal{HC}}(D)$, where $e.i+1$ and $f.i$ execute HC_2.*
- *From a configuration of $conf2_{i+1}(v, z)$, a configuration of $conf1_{i+1}(v+1, z)$ is reached in a single step of $\texttt{Algo}_{\mathcal{HC}}(D)$, where $e.i+1$, $e.i$, and $f.i$ execute HC_2.*

The notation below generalizes Notation 2.

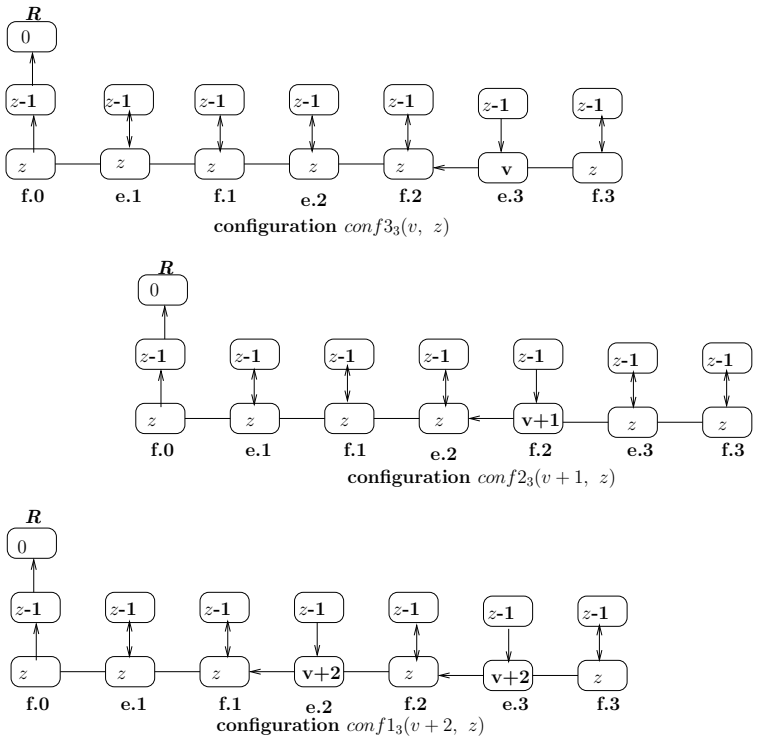

Figure 6: Some configurations of $\mathcal{G}_3$

Notation 3 *Let v and z be two integers such that $v \leq z$. Let $i \geq 1$. Let $nbSteps(v, z, i)$ be the maximal number of steps of $\texttt{Algo}_{\mathcal{HC}}(D)$ (with $D \geq z$) to reach a configuration of $conf3_i(z, z)$ from a configuration of $conf3_i(v, z)$ in the graph $\mathcal{G}_j$ with $j \geq i$.*

Observation 5 *Let v and z be two integers such that $1 \leq v \leq z$ and $1 < z$. We have $nbSteps(v + 2, z + 2, i) = nbSteps(v, z, i)$ and $nbSteps(v, z, i + 1) \geq nbSteps(v, z, i)$.*

Lemma 17 *Let z and D be two integers such that $z \leq D$. Let $i \geq 1$. In the graph $\mathcal{G}_{j+1}$ with $j \geq i$, for every $1 \leq v \leq z - 2$, there is an execution of $\texttt{Algo}_{\mathcal{HC}}(D)$ starting from a configuration of $conf3_{i+1}(v, z)$, where only rules HC_2 are executed, which reaches a configuration of $conf3_{i+1}(v + 2, z)$ in at least $nbSteps(v + 2, z, i) + 2$ steps.*

Proof. From a configuration of $conf3_{i+1}(v, z)$, a configuration of $conf1_{i+1}(v + 2, z)$ is reached in two steps of $\texttt{Algo}_{\mathcal{HC}}(D)$ where only rules HC_2 are executed, by Observation 4. From a configuration of $conf1_{i+1}(v + 2, z)$, a configuration of $conf3_{i+1}(v + 2, z)$ is reached in at least $nbSteps(v + 2, z, i)$ steps of $\texttt{Algo}_{\mathcal{HC}}(D)$ where processes of the subgraph $\mathcal{G}_i$ only execute rules HC_2 (according to the definition of $nbSteps(v, z, i)$). □

From Observation 5 and Lemma 17, we can deduce the following corollary.

Corollary 7

- *Let $1 \leq v \leq z$ and $1 < z$. $nbSteps(v, z + 2, i + 1) \geq 2 + nbSteps(v, z, i) + nbSteps(v, z, i + 1)$.*

- *Let $j \geq k$. In the graph $\mathcal{G}_j$, there is a execution $e_k(j)$ of $\texttt{Algo}_{\mathcal{HC}}(D)$, with $D \geq 2j+3$, which starts in a configuration of $conf3_k(1, 2j+3)$, contains only executions of rules HC_2, and reaches a configuration of $conf3_k(2j+3, 2j+3)$.*

Notation 4 *$nbTotal(2k+3) = \sum_{\ell=1}^{k} nbSteps(1, 2k+3, \ell)$.*

Definition 4 *let $k \geq 1$. Let e^k be the execution of $\texttt{Algo}_{\mathcal{HC}}(D)$, with $D \geq 2k+3$, in the graph $\mathcal{G}_k$ defined as follows: e^k is the concatenation of $e_1(k) \dots e_k(k)$.*

By definition, e^k contains at least $nbTotal(2k+3)$ steps, moreover those steps are only made of rules HC_2's executions.

Theorem 11 *For all $k > 1$, $nbTotal(2k+3) = 2.nbTotal(2k+1) + 2k + nbSteps(3, 2k+3, k)$.*

Proof. $nbTotal(2k+3) = \sum_{\ell=1}^{k} nbSteps(1, 2k+3, \ell)$.
We have:

- $nbSteps(1, 2k+3, 1) = 2 + nbSteps(1, 2k+1, 1)$, by Corollary 6.
- $nbSteps(1, 2k+3, k) = 2 + nbSteps(1, 2k+1, k-1) + nbSteps(3, 2k+3, k)$, by Corollary 7 and Observation 5.
- $nbSteps(1, 2k+3, \ell) = 2 + nbSteps(1, 2k+1, \ell-1) + nbSteps(1, 2k+1, \ell)$ for $\ell \in [2, k-1]$, by Corollary 7.

So, we can conclude that $nbTotal(2k+3) = \sum_{\ell=1}^{k-1}(2.nbSteps(1, 2k+1, \ell)) + 2k + nbSteps(3, 2k+3, k)$.
□

The following corollary establishes a lower bound on the number of steps of e^k which is exponential on the graph diameter: $2^{\frac{\mathcal{D}-1}{2}}$.

Corollary 8 *For all $k \geq 1$, $nbTotal(2k+3) \geq 2^{k-1}.nbSteps(1, 5, 1) \geq 2^{k+1}$.*

We now propose a tighter bound on $nbTotal(2k+3)$.

Lemma 18 *For all $k \geq 1$, $nbSteps(3, 2k+3, k) \geq 2.(2^k - 1)$.*

Proof. By induction.

Base Case: Let $k = 1$. We have $nbSteps(3, 5, 1) \geq 2 = 2.(2^1 - 1)$, by Corollary 6.

Induction Hypothesis: Assume that $nbSteps(3, 2k+3, k) \geq 2.(2^k - 1)$.

Induction Step: By Corollary 7, we have:

$$nbSteps(3, 2k+5, k+1) \geq 2 + nbSteps(3, 2k+3, k) + nbSteps(3, 2k+3, k+1)$$

By Observation 5, we have:

$$nbSteps(3, 2k+5, k+1) \geq 2 + 2.nbSteps(3, 2k+3, k) \geq 2 + 4.(2^k - 1) = 2.(2^{k+1} - 1)$$

□

Lemma 19 *For all $k \geq 1$, $nbTotal(2k+3) \geq (2k+2)(2^k-1)$.*

Proof. By induction.

Base Case: Let $k = 1$. $nbTotal(5) = nbSteps(1, 5, 1) = 4$, by Corollary 6.

Induction Hypothesis: Assume that $nbTotal(2k+3) \geq (2k+2)(2^k-1)$ for $k \geq 1$.

Induction Step:

- $nbTotal(2k+5) = 2.nbTotal(2k+3) + 2k + 2 + nbSteps(3, 2k+5, k)$, by Theorem 11.
- $nbTotal(2k+5) \geq 2.(2k+2)(2^k-1) + 2k + 2 + nbSteps(3, 2k+5, k)$, by induction hypothesis.
- $nbTotal(2k+5) \geq 2.(2k+2)(2^k-1) + 2k + 2 + 2.(2^k-1)$, by Lemma 18.
- $nbTotal(2k+5) \geq (2k+2)(2^{k+1}-2) + 2.2^{k+1} + 2k$.
- $nbTotal(2k+5) \geq (2k+4)2^{k+1} - 2.(2k+2) + 2k$.

So, we conclude that $nbTotal(2k+5) \geq (2k+4)2^{k+1} - (2k+4)$.

□

Theorem 12 *Let $n \geq 7$. Let k the maximum integer such that $n = 4k+3+y$ with $y \geq 0$. For every $D \geq 2k+3$, there is an execution of $\texttt{Algo}_{\mathcal{HC}}(D)$ which stabilizes in at least $(2k+2)(2^k-1)$ steps containing only executions of rules HC_2 in an n-node graph of diameter at most $2k+4$.*

Proof. Let $\mathcal{G}_k = (V_k, E_k)$. Let $\mathcal{G}'_k = (V_k \cup v_1, \ldots, v_y, E_k \cup \{\{v_i, R\}, i \in [1..y]\})$. Since $\mathcal{G}_k$ has diameter $2k+3$, $\mathcal{G}'_k$ has at most diameter $2k+4$. Since $\mathcal{G}_k$ contains $4k+3$ nodes, $\mathcal{G}'_k$ contains n nodes. Finally, nodes $v_1, \ldots, v_y$ are only neighbors of R whose state is constant. So, $v_1, \ldots, v_y$ have no impact on the behavior of nodes of $\mathcal{G}_k$. Hence, we can apply Lemma 19 and we are done. □

Corollary 9 *For every n-node graph G, $\texttt{Algo}_{\mathcal{HC}}(n)$,* i.e.*, the algorithm proposed in [9], stabilizes in $\Omega(2^{\mathcal{D}})$ steps, where $\mathcal{D}$ is the diameter of G.*

6.4 Algorithm $\texttt{Algo}_{\mathcal{B}}(D)$

Theorem 12 exhibits an execution exponential in steps where only rules HC_2 are executed. So, this execution is also an execution of $\texttt{Algo}_{\mathcal{FHC}}(D)$ (*i.e.*, the fast implementation of $\texttt{Algo}_{\mathcal{HC}}(D)$). Moreover, this is also a definition of $\texttt{Algo}_{\mathcal{B}}(D)$, by Lemma 9). Hence, we can conclude with the following theorem:

Theorem 13 *Let $n \geq 7$. Let k the maximum integer such that $n = 4k+3+y$ with $y \geq 0$. For every $D \geq 2k+3$, there is an execution of $\texttt{Algo}_{\mathcal{B}}(D)$ (resp. $\texttt{Algo}_{\mathcal{FHC}}(D)$) which stabilizes in at least $(2k+2)(2^k-1)$ steps in an n-node graph of diameter at most $2k+4$.*

7 Conclusion and Perspective

In this paper, we revisited two fundamental results of the self-stabilizing literature [9, 10]. More precisely, we proposed three silent self-stabilizing BFS spanning tree algorithms working in the composite atomicity model inspired from the solutions proposed in [9, 10]: Algorithms $\texttt{Algo}_{\mathcal{U}}$, $\texttt{Algo}_{\mathcal{B}}(D)$, and $\texttt{Algo}_{\mathcal{HC}}(D)$. We then presented a deep study of these algorithms. Our results are related to both correctness and complexity.

Concerning the correctness part, we proposed in particular a new, simple, and general proof scheme to show the convergence of silent algorithms under the distributed unfair daemon. We believe that our approach, based on process partitioning, is versatile enough to be applied in the convergence proof of many other silent algorithms.

Concerning the complexity part, our analysis notably shows that the Huang and Chen's algorithm [9] stabilizes in $\Omega(n)$ rounds (where n is the size of the network), while it confirms that the stabilization time in rounds of the Dolev *et al*'s algorithm [10] is optimal (exactly $\mathcal{D}$ rounds in the worst case). Finally, our analysis reveals that the stabilization time in steps of $\texttt{Algo}_{\mathcal{U}}$ cannot be bounded, while the stabilization time of both $\texttt{Algo}_{\mathcal{B}}(D)$ and $\texttt{Algo}_{\mathcal{HC}}(D)$ can be exponential in $\mathcal{D}$, the diameter of the network. Our results must be put in perspective with the complexities of the silent BFS construction proposed in [13], which stabilizes in $O(\mathcal{D}^2)$ rounds and $O(n^6)$ steps, respectively. This suggests the existence of a trade-off between the complexity in rounds and steps for the silent construction of a BFS tree. This conjecture would have to be investigated in future works.

References

[1] Edsger W. Dijkstra. Self-Stabilizing Systems in Spite of Distributed Control. *Commun. ACM*, 17:643–644, 1974.

[2] Shing-Tsaan Huang and Nian-Shing Chen. Self-stabilizing depth-first token circulation on networks. *Distributed Computing*, 7(1):61–66, 1993.

[3] Jean-Michel Couvreur, Nissim Francez, and Mohamed G. Gouda. Asynchronous unison (extended abstract). In *Proceedings of the 12th International Conference on Distributed Computing Systems, Yokohama, Japan, June 9-12, 1992*, pages 486–493. IEEE Computer Society, 1992.

[4] Alain Bui, Ajoy Kumar Datta, Franck Petit, and Vincent Villain. Optimal PIF in tree networks. In Yuri Breitbart, Sajal K. Das, Nicola Santoro, and Peter Widmayer, editors, *Distributed Data & Structures 2, Records of the 2nd International Meeting (WDAS 1999), Princeton, USA, May 10-11, 1999*, volume 6 of *Proceedings in Informatics*, pages 1–16. Carleton Scientific, 1999.

[5] Hirotsugu Kakugawa and Toshimitsu Masuzawa. A self-stabilizing minimal dominating set algorithm with safe convergence. In *20th International Parallel and Distributed Processing Symposium (IPDPS 2006), Proceedings, 25-29 April 2006, Rhodes Island, Greece*. IEEE, 2006.

[6] Eddy Caron, Ajoy Kumar Datta, Benjamin Depardon, and Lawrence L. Larmore. A self-stabilizing k-clustering algorithm for weighted graphs. *J. Parallel Distrib. Comput.*, 70(11):1159–1173, 2010.

[7] NS Chen, HP Yu, and ST Huang. A self-stabilizing algorithm for constructing spanning trees. *Information Processing Letters*, 39:147–151, 1991.

[8] Shlomi Dolev, Mohamed G. Gouda, and Marco Schneider. Memory requirements for silent stabilization. *Acta Informatica*, 36(6):447–462, 1999.

[9] Shing-Tsaan Huang and Nian-Shing Chen. A self-stabilizing algorithm for constructing breadth-first trees. *Information Processing Letters*, 41(2):109–117, 1992.

[10] Shlomi Dolev, Amos Israeli, and Shlomo Moran. Self-stabilization of dynamic systems assuming only read/write atomicity. *Distributed Computing*, 7(1):3–16, 1993.

[11] Zeev Collin and Shlomi Dolev. Self-stabilizing depth-first search. *Inf. Process. Lett.*, 49(6):297–301, 1994.

[12] Adrian Kosowski and Lukasz Kuszner. A self-stabilizing algorithm for finding a spanning tree in a polynomial number of moves. In *6th International Conference Parallel Processing and Applied Mathematics, (PPAM'05), Springer LNCS 3911*, pages 75–82, 2005.

[13] Alain Cournier, Stephane Rovedakis, and Vincent Villain. The first fully polynomial stabilizing algorithm for BFS tree construction. In *the 15th International Conference on Principles of Distributed Systems (OPODIS'11), Springer LNCS 7109*, pages 159–174, 2011.

[14] G Tel. *Introduction to distributed algorithms*. Cambridge University Press, Cambridge, UK, Second edition 2001.

[15] Mohamed G. Gouda and Ted Herman. Adaptive programming. *IEEE Trans. Software Eng.*, 17(9):911–921, 1991.

[16] Shlomi Dolev. *Self-Stabilization*. MIT Press, 2000.

[17] Ajoy Kumar Datta, Shivashankar Gurumurthy, Franck Petit, and Vincent Villain. Self-stabilizing network orientation algorithms in arbitrary rooted networks. *Stud. Inform. Univ.*, 1(1):1–22, 2001.

[18] Lin Fei, Sun Yong, Ding Hong, and Ren Yizhi1. Self stabilizing distributed transactional memory model and algorithms. *Journal of Computer Research and Development*, 51(9):2046, 2014.

[19] Franck Butelle, Christian Lavault, and Marc Bui. A uniform self-stabilizing minimum diameter tree algorithm (extended abstract). In Jean-Michel Hélary and Michel Raynal, editors, *Distributed Algorithms, 9th International Workshop, WDAG '95, Le Mont-Saint-Michel, France, September 13-15, 1995, Proceedings*, volume 972 of *Lecture Notes in Computer Science*, pages 257–272. Springer, 1995.

[20] Alain Cournier, Stéphane Devismes, and Vincent Villain. Snap-stabilizing detection of cutsets. In David A. Bader, Manish Parashar, Sridhar Varadarajan, and Viktor K. Prasanna, editors, *High Performance Computing - HiPC 2005, 12th International Conference, Goa, India, December 18-21, 2005, Proceedings*, volume 3769 of *Lecture Notes in Computer Science*, pages 488–497. Springer, 2005.

[21] Ajoy Kumar Datta, Lawrence L. Larmore, Stéphane Devismes, Karel Heurtefeux, and Yvan Rivierre. Self-stabilizing small k-dominating sets. *IJNC*, 3(1):116–136, 2013.

[22] Alain Cournier, Stéphane Devismes, and Vincent Villain. Light enabling snap-stabilization of fundamental protocols. *ACM Transactions on Autonomous and Adaptive Systems*, 4(1), 2009.

[23] Anish Arora and Mohamed G. Gouda. Distributed reset. *IEEE Trans. Computers*, 43(9):1026–1038, 1994.

[24] Mehmet Hakan Karaata. A stabilizing algorithm for finding biconnected components. *J. Parallel Distrib. Comput.*, 62(5):982–999, 2002.

[25] Ajoy Kumar Datta, Lawrence L. Larmore, and Priyanka Vemula. An O(n)-time Self-stabilizing Leader Election Algorithm. *J. Parallel Distrib. Comput.*, 71(11):1532–1544, 2011.

[26] Christian Glacet, Nicolas Hanusse, David Ilcinkas, and Colette Johnen. Disconnected components detection and rooted shortest-path tree maintenance in networks. In *the 16th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS'14), Springer LNCS 8736*, pages 120–134, 2014.