



HAL
open science

Being Efficient in Time, Space, and Workload: a Self-stabilizing Unison and its Consequences

Stéphane Devismes, David Ilcinkas, Colette Johnen, Frédéric Mazoit

► **To cite this version:**

Stéphane Devismes, David Ilcinkas, Colette Johnen, Frédéric Mazoit. Being Efficient in Time, Space, and Workload: a Self-stabilizing Unison and its Consequences. STACS 2025: 42nd International Symposium on Theoretical Aspects of Computer Science, Mar 2025, Jena, Germany. hal-04866194

HAL Id: hal-04866194

<https://hal.science/hal-04866194v1>

Submitted on 6 Jan 2025

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.



Distributed under a Creative Commons Attribution 4.0 International License

Being Efficient in Time, Space, and Workload: a Self-stabilizing Unison and its Consequences

Stéphane Devismes ✉ 

Laboratoire MIS, Université de Picardie, 33 rue Saint Leu - 80039 Amiens cedex 1, France

David Ilcinkas ✉ 

LaBRI, Université de Bordeaux, 351 cours de la Libération, F-33405 Talence cedex, France

Colette Johnen ✉ 

LaBRI, Université de Bordeaux, 351 cours de la Libération, F-33405 Talence cedex, France

Frédéric Mazoit ✉ 

LaBRI, Université de Bordeaux, 351 cours de la Libération, F-33405 Talence cedex, France

Abstract

We present a self-stabilizing algorithm for the unison problem which is efficient in time, workload, and space in a weak model. Precisely, our algorithm is defined in the atomic-state model and works in anonymous asynchronous connected networks in which even local ports are unlabeled. It makes no assumption on the daemon and thus stabilizes under the weakest one: the distributed unfair daemon.

In an n -node network of diameter D and assuming the knowledge $B \geq 2D + 2$, our algorithm only requires $\Theta(\log(B))$ bits per node and is fully polynomial as it stabilizes in at most $2D + 2$ rounds and $O(\min(n^2B, n^3))$ moves. In particular, it is the first self-stabilizing unison for arbitrary asynchronous anonymous networks achieving an asymptotically optimal stabilization time in rounds using a bounded memory at each node.

Furthermore, we show that our solution can be used to efficiently simulate synchronous self-stabilizing algorithms in asynchronous environments. For example, this simulation allows us to design a new state-of-the-art algorithm solving both the leader election and the BFS (Breadth-First Search) spanning tree construction in any identified connected network which, to the best of our knowledge, beats all existing solutions in the literature.

2012 ACM Subject Classification Theory of computation \rightarrow Distributed computing models; Theory of computation \rightarrow Distributed algorithms; Theory of computation \rightarrow Design and analysis of algorithms

Keywords and phrases Self-stabilization, unison, time complexity, synchronizer.

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23


Funding *David Ilcinkas, Colette Johnen, and Frédéric Mazoit:* This work was supported by the ANR project ENEDISC. *Colette Johnen and Stéphane Devismes:* This work was supported by the ANR project SkyData.



© Stéphane Devismes, David Ilcinkas, Colette Johnen and Frédéric Mazoit;
licensed under Creative Commons License CC-BY 4.0

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:17



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

1.1 Context

Self-stabilization is a general non-masking and lightweight fault tolerance paradigm [25, 3]. Precisely, a distributed system achieving this property inherently tolerates *any* finite number of transient faults.¹ Indeed, starting from an arbitrary configuration, which may be the result of such faults, a self-stabilizing system recovers within finite time, and without any external intervention, a so-called *legitimate configuration* from which it satisfies its specification.

In this paper, we consider the most commonly used model in the self-stabilizing area: the *atomic-state* model [25, 3]. In this model, the state of each node is stored into registers and these registers can be directly read by neighboring nodes. Furthermore, in one atomic step, a node can read its state and that of its neighbors, perform some local computation, and update its state accordingly. In the atomic-state model, asynchrony is materialized by an adversary called *daemon* that can restrict the set of possible executions. We consider here the weakest (i.e., the most general) daemon: the *distributed unfair daemon*.

Self-stabilizing algorithms are mainly compared according to their *stabilization time*, i.e., the worst-case time to reach a legitimate configuration starting from an arbitrary one. In the atomic-state model, stabilization time can be evaluated in terms of *rounds* and *moves*. Rounds [13] capture the execution time according to the speed of the slowest nodes. Moves count the number of local state updates. So, the move complexity is rather a measure of work than a measure of time. It turns out that obtaining efficient stabilization times both in rounds and moves is a difficult task. Usually, techniques to design an algorithm achieving a stabilization time polynomial in moves make its round complexity inherently linear in n , the number of nodes (see, e.g., [2, 23, 19]). Conversely, achieving the asymptotic optimality in rounds, usually $O(D)$ where D is the network diameter, commonly makes the stabilization time exponential in moves (see, e.g., [22, 31]). Surprisingly, Cournier, Rovedakis, and Villain [14] manage to prove the first *fully polynomial* (i.e., with $Poly(n)$ move and $Poly(D)$ round complexities) silent² self-stabilizing algorithm. Their algorithm builds a BFS (Breadth-First Search) spanning tree in any rooted connected network and they prove that it stabilizes in $O(n^6)$ moves and $O(D^2)$ rounds using $\Theta(\log B + \log \Delta)$ bits per node, where B is an upper bound on D and Δ is the maximum degree of the network.

Up to now, fully polynomial self-stabilizing algorithms have only been proposed (see [14, 21]) for so called *static* problems [34], such as spanning tree constructions and leader election, which compute a fixed object in finite time. In this paper, we propose an algorithm for a fundamental *dynamic* (i.e., non static) problem: the *asynchronous unison* (*unison* for short). It consists in maintaining a local clock at each node. The domain of clocks can be bounded (like everyday clocks) or infinite. The liveness property of the problem requests each node to increment its own clock infinitely often. Furthermore, the safety property of the unison requires the difference between the clocks of any two neighbors to always be at most one increment. The usefulness of the unison comes from the fact that asynchrony often makes fault tolerance very difficult in distributed systems. The impossibility of achieving consensus in an asynchronous system in spite of at most one process crash [30] is a famous example illustrating this fact. Thus, fault tolerance, and in particular self-stabilization, often requires some kind of barrier synchronization, which the unison provides, to control the

¹ A *transient fault* occurs at an unpredictable time, but does not result in a permanent hardware damage. Moreover, as opposed to intermittent faults, the frequency of transient faults is considered to be low.

² In the atomic-state model, a self-stabilizing algorithm is *silent* if all its executions terminate.

78 asynchronism of the system by making processes progress roughly at the same speed. Unison
 79 is thus a fundamental algorithmic tool that has numerous applications. Among others, it
 80 can be used to simulate synchronous systems in asynchronous environments [17], to free an
 81 asynchronous system from its fairness assumption (e.g., using the cross-over composition) [8],
 82 to facilitate the termination detection [9], to locally share resources [11], or to achieve
 83 infimum computations [10]. Thus, as expected, we also derive from our unison algorithm a
 84 *synchronizer* allowing us to obtain several new state-of-the-art self-stabilizing algorithms for
 85 various problems, including spanning tree problems and leader election.

86 **1.2 Related Work**

87 **Related Work on the Self-stabilizing Unison**

88 The first self-stabilizing asynchronous unison for general graphs was proposed by Couvreur,
 89 Francez, and Gouda [15] in the link-register model (a locally-shared memory model without
 90 composite atomicity [27, 26]). However, no complexity analysis was given. Another solution,
 91 which stabilizes in $O(n)$ rounds, is proposed by Boulinier, Petit, and Villain [11] in the
 92 atomic-state model assuming a distributed unfair daemon. Its move complexity is shown
 93 in [24] to be in $O(Dn^3 + \alpha n^2)$, where α is a parameter of the algorithm that should satisfy
 94 $\alpha \geq L - 2$, where L is the length of the longest hole in the network. In his PhD thesis,
 95 Boulinier proposes a parametric solution that generalizes the solutions of both [15] and [11].
 96 In particular, the time complexity analysis of this latter algorithm reveals an upper bound in
 97 $O(D \cdot n)$ rounds on the stabilization time of the atomic-state model version of the algorithm
 98 in [15]. Awerbuch, Kutten, Mansour, Patt-Shamir, and Varghese [4] propose a self-stabilizing
 99 unison that stabilizes in $O(D)$ rounds using an infinite state space. The move complexity of
 100 their solution is not analyzed. An asynchronous self-stabilizing unison algorithm is given
 101 in [23]. It stabilizes in $O(n)$ rounds and $O(\Delta \cdot n^2)$ moves using unbounded local memories.
 102 Emek and Keren [28] present in the stone age model a self-stabilizing unison that stabilizes
 103 in $O(B^3)$ rounds, where B is an upper bound on D known by all nodes. Their solution
 104 requires $\Theta(\log B)$ bits per node. Moreover, since node activations are required to be fair, the
 105 move complexity of their solution is unknown and may be unbounded.

106 **Related Work on Simulations**

107 Simulation is a useful tool to simplify the design of algorithms. In self-stabilization, simulation
 108 has been mainly investigated to emulate schedulers or to port solutions from a strong
 109 computational model to a weaker one. Awerbuch [7] introduced the concept of *synchronizer*
 110 in a non-self-stabilizing context. A synchronizer simulates a synchronous execution of an input
 111 algorithm into an asynchronous environment. The first two self-stabilizing synchronizers have
 112 been proposed in [4] for message-passing systems. Both solutions achieve a stabilization time
 113 in $O(D)$ rounds. The first solution is based on the previously mentioned unison, also proposed
 114 in the paper, that uses an infinite state space. To solve this latter issue, they then propose
 115 to mix it with the reset algorithm of [5] applied on links of a BFS spanning tree computed
 116 in $O(D)$ rounds. This reset algorithm is devoted, and so limits the approach, to locally
 117 checkable and locally correctable problems, and the BFS spanning tree construction uses a
 118 finite yet unbounded number of states per node and requires the presence of a distinguished
 119 node (a root). Again, the move complexity of their solutions is not analyzed. Awerbuch
 120 and Varghese [6] propose, still in the message-passing model, two synchronizers: the *rollback*
 121 *compiler* and the *resynchronizer*. The resynchronizer additionally requires the input algorithm
 122 to be locally checkable and assumes the knowledge of a common upper bound \mathcal{D} on the

123 network diameter. Using the rollback, resp. the resynchronizer, method, a synchronous
 124 non-self-stabilizing algorithm can be turned into an asynchronous self-stabilizing algorithm
 125 that stabilizes in $O(T)$ rounds, resp. $O(T + D)$ rounds, using $\Omega(T \times S)$ space, resp. $\Theta(S)$
 126 space, per node where T , resp. S , is the execution time, resp. the space complexity, of the
 127 input algorithm. Again, the move complexity of these synchronizers is not analyzed. Now,
 128 the straightforward atomic-state model version of the rollback compiler is shown to achieve
 129 exponential move complexities in [21]. Finally, the synchronizer proposed in [21] works in the
 130 atomic-state model and achieves round and space complexities similar to those of the rollback
 131 compiler, but additionally offers polynomial move complexity. Hence, it allows to design
 132 fully polynomial self-stabilizing solutions for static problems, but still with an important
 133 memory requirement (using $\Omega(T \times S)$ space).

134 Simulation has been also investigated in self-stabilization to emulate other schedulers. For
 135 example, the *conflict manager* proposed in [32] allows to emulate an unfair locally central
 136 scheduler in fully asynchronous settings. Another example is fairness that can be enforced
 137 using a unison algorithm together with the *cross-over* composition [8].

138 Concerning now model simulations, Turau proposes in [35] a general procedure allowing
 139 to simulate any algorithm for the distance-two atomic-state model in the (classical) distance-
 140 one atomic-state model assuming that nodes have unique identifiers. Finally, simulation
 141 from the atomic-state model to the link-register one and from the link-register model to the
 142 message-passing one are discussed in [26].

143 1.3 Contributions

144 Fully Polynomial Self-stabilizing Unison

145 We propose a fully polynomial self-stabilizing bounded-memory unison in the atomic-state
 146 model assuming a distributed unfair daemon. It works in any anonymous network of arbitrary
 147 connected topology, and stabilizes in $O(D)$ rounds and $O(n^3)$ moves using $\Theta(\log B)$ bits per
 148 node, where $B \geq 2D + 2$ (see Table 1 below). To the best of our knowledge, our algorithm
 149 vastly improves on the literature as other self-stabilizing algorithms have at least one of the
 150 following drawbacks: an unbounded memory, an $\Omega(n)$ round complexity, a restriction on the
 151 daemon (synchronous, fair, ...). Note also that the computational model we use is at least
 152 as general as the *stone age* model of Emek and Wattenhofer [29]: it does not require any
 153 local port labeling at nodes, or knowing how many neighbors a node has.

154 Overall, our unison achieves outstanding performance in terms of time, workload, and
 155 space, which also makes it the first fully polynomial self-stabilizing algorithm for a *dynamic*
 156 problem.

157 Self-stabilizing Synchronizer

158 From our unison algorithm, we straightforwardly derive a self-stabilizing synchronizer that
 159 efficiently simulates synchronous executions of an input self-stabilizing algorithm in an
 160 asynchronous environment. More precisely, if the input algorithm Alg_I is silent, then the
 161 output algorithm $Sync(Alg_I)$ is silent as well and satisfies the same specification as Alg_I .
 162 The specification preservation property also holds for any algorithm, silent or not, solving
 163 a static problem. We analyze the complexity of this synchronizer and show that it mostly
 164 preserves the round and space complexities of the simulated algorithm (see Table 1 for
 165 details). This synchronizer is thus a powerful tool to ease the design of efficient *asynchronous*
 166 self-stabilizing algorithms. Indeed, for many tasks, the usual lower bound on the stabilization
 167 time in rounds is $\Omega(D)$. Now, thanks to our unison, one just has to focus on the design

168 of a *synchronous* $O(D)$ -round self-stabilizing algorithm to finally obtain an asynchronous
 169 self-stabilizing solution asymptotically optimal in rounds, with a low overhead in space
 170 ($\Theta(\log B)$ bits per node) and a polynomial move complexity (i.e., a fully polynomial solution).

171 The transformer of [21] has similar round and move complexities. But this algorithm and
 172 ours are incomparable as they make different trade-offs. This paper prioritizes memory over
 173 generality, while the transformer of [21] makes the opposite choice by prioritizing generality
 174 over memory. More precisely, the transformer of [21] can simulate any synchronous algorithm
 175 (not necessarily self-stabilizing), by storing its whole execution. It thus has a much larger
 176 space complexity than ours, which only stores two states of the simulated input algorithm. It
 177 turns out that the connections between our algorithm and the transformer of [21] are deeper
 178 than their move and round complexities. We further explain their similarities as well as their
 179 differences in Sections 3 and 4.5.

180 Implications of our Results

181 Using our synchronizer, one can easily obtain state-of-the-art (silent) self-stabilizing solutions
 182 for several fundamental distributed computing problems, e.g., BFS tree constructions, leader
 183 election, and clustering (see Table 1).

	Moves	Rounds	Space
Unison	$O(\min(n^2 B, n^3))$	$2D + 2$	$\lceil \log B \rceil + 2$
Synchronizer	$O(\min(n^2 B, n^3) + nT)$	$5D + 3T$	$2M + \lceil \log B \rceil + 2$

Problem	Moves	Rounds	Space
BFS tree in rooted networks	$O(n^3)$	$O(D)$	$\Theta(\log B + \log \Delta)$
BFS tree in identified networks	$O(n^3)$	$O(D)$	$\Theta(\log N)$
Leader election	$O(n^3)$	$O(D)$	$\Theta(\log N)$
$O(\frac{n}{k})$ -clustering	$O(n^3)$	$O(D)$	$\Theta(\log k + \log N)$

T and M are the synchronous time and space complexities of the input algorithm,
 and B and N are input parameters satisfying $B \geq 2D + 2$ and $N \geq n$.

■ **Table 1** Complexities of the Unison, the Synchronizer, and some consequences.

184 First, we obtain a new state-of-the-art asynchronous self-stabilizing algorithm for the BFS
 185 spanning tree construction in rooted and connected networks, by synchronizing the algorithm
 186 in [22] (which is a bounded-memory variant of the algorithm in [27]). This new algorithm
 187 converges in $O(n^3)$ moves and $O(D)$ rounds with $\Theta(\log B + \log \Delta)$ bits per node (the same
 188 round and space complexities as in [22]), where B is an upper bound on D and Δ is the
 189 maximum node degree. It improves both on the algorithm in [14], which only converges in
 190 $O(n^6)$ moves and $O(D^2)$ rounds, and on the algorithm in [21], which has similar complexities
 191 but uses $\Theta(B \cdot \log \Delta)$ bits per node.

192 In the following, we consider identified connected networks. In this setting, when nodes
 193 store identifiers, they usually know a bound k on the size of these identifiers. They thus
 194 know a bound $N = 2^k$ on n , and since N is a bound on D , we set $B = 2N + 2$.

195 In identified networks, a strategy to compute a BFS spanning tree is to compute a leader
 196 together with a BFS tree rooted at this leader. This is what the self-stabilizing algorithm
 197 in [33] actually does in a synchronous setting. Therefore, by synchronizing it, we obtain
 198 a new state-of-the-art asynchronous self-stabilizing algorithm for both the leader election

199 and the BFS spanning tree construction in identified and connected networks. This new
 200 algorithm converges in $O(n^3)$ moves and $O(D)$ rounds with $\Theta(\log N)$ bits per node (i.e.,
 201 the same round and space complexities as in [33]). To the best of our knowledge, no such
 202 efficient solutions exist until now in the literature. There are two incomparable asynchronous
 203 self-stabilizing algorithms that achieve an $O(D)$ round complexity [12, 1]. They operate in
 204 weaker models (resp. message-passing and link-register). However, their move complexity is
 205 not analyzed and the first one has a $\Theta(\log B \cdot \log N)$ space requirement (B being a known
 206 upper bound on D) while the second one uses an unbounded space.

207 Other memory-efficient fully polynomial self-stabilizing solutions can be easily obtained
 208 with our synchronizer, e.g., to compute the median or centers in anonymous trees by simulating
 209 algorithms proposed in [18]. Another application of our synchronizer is to remove fairness
 210 assumptions along with obtaining good complexities. For example, the silent self-stabilizing
 211 algorithm proposed in [16] computes a clustering of $O(\frac{n}{k})$ clusters in any rooted identified
 212 connected network. It assumes a distributed weakly fair daemon and its move complexity
 213 is unknown. With our synchronizer,³ we achieve a fully polynomial silent solution that
 214 stabilizes under the distributed unfair daemon and without the rooted network assumption,
 215 in $O(D)$ rounds, $O(n^3)$ moves, and using $\Theta(\log k + \log N)$ bits per node.

216 Note that, by using the compiler in [21], one can obtain similar time complexities for all
 217 the previous problems, but with a drastically higher space usage.

218 1.4 Roadmap

219 The rest of the paper is organized as follows. Section 2 is dedicated to the computational
 220 model and the basic definitions. We develop the links between the present paper and [21]
 221 in Section 3, and we present our algorithm in Section 4. We sketch its correctness and its
 222 time complexity in Section 5. In Section 6, the self-stabilizing synchronizer derived from our
 223 unison algorithm is presented and its complexity is also sketched. We conclude in Section 7.

224 2 Preliminaries

225 2.1 Networks

226 We model *distributed systems* as simple graphs, that is, pairs $G = (V, E)$ where V is a
 227 set of *nodes* and E is a set of *edges* representing communication links. We assume that
 228 communications are bidirectional. The set $N(p) = \{q \mid \{p, q\} \in E\}$ is the set of *neighbors*
 229 of p , with which p can communicate, and $N[p] = N(p) \cup \{p\}$ is the closed neighborhood of p .
 230 A *path* (from p_0 to p_l) of *length* l is a sequence $P = p_0 p_1 \cdots p_l$ of nodes such that consecutive
 231 nodes in P are neighbors. We assume that G is connected, meaning that any two nodes are
 232 connected by a path. We can thus define the *distance* $d(p, q)$ between two nodes p and q to
 233 be the minimum length of a path from p to q . The *diameter* D of G is then the maximum
 234 distance between nodes of G .

235 2.2 Computational Model: the Atomic-state Model

236 Our unison algorithm works in a variant of the *atomic-state model* in which each node holds
 237 locally shared registers, called *variables*, whose values constitute its *state*. The vector of all

³ Also replacing the spanning tree construction used in [16] by the new BFS tree construction of the previous paragraph.

238 node states defines a *configuration* of the system.

239 An algorithm consists of a finite set of *rules* of the form $label : guard \rightarrow action$. In
 240 the variant that we consider, a *guard* is a boolean predicate on the state of the node and
 241 on the set of states of its neighbors. The *action* changes the state of the node. To shorten
 242 guards and increase readability, priorities between rules may be set. A rule whose guard is
 243 *true* is *enabled*, and can be executed. By extension, a node with at least one enabled rule
 244 is also *enabled*, and $Enabled(\gamma)$ contains the enabled nodes in a configuration γ . Note that
 245 this model is quite weak. Indeed, in other variants, nodes may have, for example, distinct
 246 identifiers. In our case, the network is anonymous and since a node only accesses a set of
 247 states, it cannot even count how many neighbors it has.

248 An *execution* in this model is a maximal sequence of configurations $e = \gamma^0 \gamma^1 \dots \gamma^i \dots$
 249 such that for each transition (called *step*) $\gamma^i \mapsto \gamma^{i+1}$, there is a nonempty subset \mathcal{X}^i of
 250 $Enabled(\gamma^i)$ whose nodes simultaneously and atomically execute one of their enabled rules,
 251 leading from γ^i to γ^{i+1} . We say that each node of \mathcal{X}^i executes a *move* during $\gamma^i \mapsto \gamma^{i+1}$.
 252 Note that e is either infinite, or ends at a *terminal* configuration γ^f where $Enabled(\gamma^f) = \emptyset$.
 253 An algorithm with no infinite executions is *terminating* or *silent*.

254 A *daemon* \mathfrak{D} is a predicate over executions. An execution which satisfies \mathfrak{D} is said to be
 255 *an execution under \mathfrak{D}* . We consider the *synchronous daemon*, which is true if, at all steps,
 256 $\mathcal{X}^i := Enabled(\gamma^i)$, and the *fully asynchronous daemon*, also called *distributed unfair daemon*
 257 in the literature, which is always true. Note that under the distributed unfair daemon, a
 258 node may starve and may never be activated, unless it is the only enabled node.

259 In an execution, all the information in the states is not necessarily relevant for a problem.
 260 We thus use a *projection* to extract information (e.g., just an output boolean for the boolean
 261 consensus) from a node's state, and we canonically extend this projection to configurations
 262 and executions. A *specification* of a distributed problem is then a predicate over projected
 263 executions. A problem is *static* if its specification requires the projected executions to be
 264 constant, and it is *dynamic* otherwise.

265 An algorithm is *self-stabilizing* under a daemon \mathfrak{D} if, for every network and input
 266 parameters, there exists a set of *legitimate* configurations such that (1) the algorithm
 267 *converges*, i.e., every execution under \mathfrak{D} (starting from an arbitrary configuration) contains a
 268 legitimate configuration, and (2) the algorithm is *correct*, i.e., every execution under \mathfrak{D} that
 269 starts from a legitimate configuration satisfies the specification.

270 We consider three complexity measures: *space*, *moves* which model the total workload,
 271 and *rounds* which model an analogous of the synchronous time by taking the speed of the
 272 slowest nodes into account. As done in the literature on the atomic-state model, the space
 273 complexity is the maximum space used by one node to store its own variables. As explained
 274 before, a move is the execution of a rule by a node. To define the round complexity of an
 275 execution $e = \gamma^0 \gamma^1 \dots$, we first need to define the notion of *neutralization*: a node p is
 276 *neutralized* in $\gamma^i \mapsto \gamma^{i+1}$, if p is enabled in γ^i and not in γ^{i+1} , but it does not apply any rule
 277 in $\gamma^i \mapsto \gamma^{i+1}$. Then, the rounds are inductively defined as follows. The first round of an
 278 execution $e = \gamma^0 \gamma^1 \dots$ is the minimal prefix e' such that every node that is enabled in γ^0
 279 either executes a move or is neutralized during a step of e' . If e' is finite, then let e'' be
 280 the suffix of e that starts from the last configuration of e' ; the second round of e is the first
 281 round of e'' , and so on. For every $i > 0$, we denote by γ^{r_i} the last configuration of the i -th
 282 round of e , if it exists and is finite; we also conventionally let $\gamma^{r_0} = \gamma^0$. Consequently, $\gamma^{r_{i-1}}$
 283 is also the first configuration of the i -th round of e . The *stabilization time* of a self-stabilizing
 284 algorithm is the maximum time (in moves or rounds) over every execution possible under
 285 the considered daemon (starting from any initial configuration) to reach (for the first time) a

286 legitimate configuration.

287 **3 A Glimpse of our Research Process**

288 **3.1 An Unbounded Unison Algorithm**

289 We started this work on the bounded unison problem when we observed that an unbounded
290 solution can easily be derived from [21]. This can be seen as follows. The algorithm
291 given in [21] simulates a synchronous non self-stabilizing algorithm in an asynchronous
292 self-stabilizing setting. To do so, it uses a very natural idea. It stores, at each node, the
293 whole execution of the algorithm so far as a list of states. Given its list and the lists of its
294 neighbors, a given node can check for inconsistencies in the simulation and correct them.

295 Now if we implement this idea in an asynchronous algorithm which is not self-stabilizing,
296 then the length of the lists satisfy the unison property. Indeed, to compute its $(i + 1)$ -th
297 value, a node must wait for all its neighbors to have computed at least their i -th value.

298 Obviously, in a self-stabilizing setting, we cannot expect the length of the lists of the
299 nodes to initially satisfy the unison property. It turns out that the error recovery mechanism
300 in [21] not only solves the initial inconsistencies of the simulation, but also recovers the
301 unison property.

302 If we simulate an algorithm “that does nothing”, we can compress the lists by only storing
303 their lengths. We thus obtain a first (unbounded) unison algorithm, given below. Note that
304 although we describe the whole algorithm, the reader does not need to fully understand it.

305 Each node p has a status $p.s \in \{E, C\}$ (Error/Correct) and a time $p.t \in \mathbb{N}$. Given these
306 predicates,

$$\begin{aligned}
 307 \quad \text{root}(p) &:= (p.s = E \wedge \neg(\exists q \in N(p), q.s = E \wedge q.t < p.t)) \vee \\
 308 &\quad (p.s = C \wedge \exists q \in N(p), (q.t \geq p.t + 2)) \\
 309 \quad \text{activeRoot}(p) &:= \text{root}(p) \wedge (p.t > 0 \vee p.s = C) \\
 310 \quad \text{errProp}(p, i) &:= \exists q \in N(p), q.s = E \wedge q.t < i < p.t \\
 311 \quad \text{canClearE}(p) &:= p.s = E \wedge \forall q \in N(p), (|q.t - p.t| \leq 1 \wedge (q.t \leq p.t \vee q.s = C)) \\
 312 \quad \text{updatable}(p) &:= p.s = C \wedge (\forall q \in N(p), p.t \leq q.t \leq p.t + 1)
 \end{aligned}$$

313 the algorithm is defined by the following four rules

$$\begin{aligned}
 314 \quad R_R &: \text{activeRoot}(p) \longrightarrow p.t := 0 ; p.s := E \\
 R_P(i) &: \text{errProp}(p, i) \longrightarrow p.t := i ; p.s := E \\
 R_C &: \text{canClearE}(p) \longrightarrow p.s := C \\
 R_U &: \text{updatable}(p) \longrightarrow p.t := p.t + 1
 \end{aligned}$$

315 in which R_R has the highest priority, and $R_P(i)$ has a higher priority than $R_P(i')$ for $i < i'$.
316 The rules R_R , $R_P(i)$ and R_C are “error management” rules. Thus, once the algorithm has
317 stabilized, the status of all nodes is C and only R_U is applicable.

318 This unbounded self-stabilizing unison algorithm is not really interesting by itself. Indeed,
319 it converges in $2D + 2$ rounds in an asynchronous setting, but in this regard, the algorithm
320 in [4] converges twice as fast, is simpler and operates in the message-passing model, which is
321 more realistic. However, whereas nobody has been able to derive a bounded version of the
322 algorithm in [4], we hoped that this could be done with this new algorithm.

323 In the following subsections, we present a first very natural attempt, which ultimately
324 failed, and a more complex version, which we detail and prove in the next sections of the
325 paper.

326 **3.2 A Failed Bounded Unison Algorithm**

327 The most natural strategy to turn an unbounded unison into a bounded one is simply to
 328 count modulo a large enough fixed bound B . To outline this change of paradigm from an
 329 ever-increasing time to a circular clock, we rename the variable $p.t$ into $p.c$ for any node p .
 330 We thus modify the rule R_U as follows:

$$331 \quad \text{updatable}(p) := p.s = C \wedge (\forall q \in N(p), q.c \in \{p.c, p.c + 1 \bmod B\})$$

$$332 \quad R_U : \text{updatable}(p) \longrightarrow p.c := p.c + 1 \bmod B$$

334 At first glance, we do not need to modify the other rules as their purpose is only to
 335 correct errors, but this intuition is wrong. Indeed, when two neighboring nodes p and q are
 336 such that $p.s = q.s = C$, $p.c = 0$ and $q.c = B - 1$, they satisfy the unison property, but p
 337 can apply the rule R_R , although there are no errors to correct. The problem comes from
 338 the term $\exists q \in N(p), (q.c \geq p.c + 2)$ in the *root* predicate which should detect out-of-sync
 339 neighbors. Hence, we must at least modify this predicate as follows:

$$340 \quad \text{root}(p) := (p.s = E \wedge \neg(\exists q \in N(p), q.s = E \wedge q.c < p.c)) \vee$$

$$341 \quad (p.s = C \wedge \exists q \in N(p), (q.c \geq p.c + 2) \wedge \neg(p.c = 0 \wedge q.c = B - 1)).$$

343 Therefore, transforming the algorithm to implement this simple modulo- B idea is already
 344 not as straightforward as it may seem.

345 Moreover, even small modifications generally introduce new unforeseen behaviors, and
 346 the modified algorithm has no particular reasons to be efficient, or even correct. As a matter
 347 of fact, we failed to prove its correctness. To understand why, we must delve a bit into the
 348 proof scheme of [21].

349 An important observation is that rootless configurations (i.e., those without nodes
 350 satisfying the *root* predicate) satisfy the safety property of the unison. In [21], the correctness
 351 and the move complexity then follow from the key property that roots cannot be created,
 352 and that, in a “small” number of steps, at least one root disappears.

353 Sadly, this first attempt algorithm does not satisfy the “no root creations” property. To
 354 see this, consider a path $p - q - r$ and a configuration γ^a in which $p.c = q.c = B - 1$, $r.c = 3$,
 355 $p.s = q.s = C$ and $r.s = E$. In one step $\gamma^a \mapsto \gamma^b$,
 356 ■ p applies the rule R_U and thus, in γ^b , $p.c = 0$ and $p.s = C$
 357 ■ q applies the rule $R_P(4)$, and thus, in γ^b , $q.c = 4$.
 358 Therefore, in γ^b , $p.s = C$ and p has a neighbor q such that $q.c \geq p.c + 2$ and $q.c \neq B - 1$.
 359 Thus, p is a root in γ^b , although it is not one in γ^a .

360 Note that the fact that roots can be created is not necessarily a problem. Indeed, if only
 361 a finite number of them appears, we recover the correctness of the algorithm. We actually
 362 believe that, for B large enough, any node can become a root only once per execution, and
 363 this would most likely imply that the move complexity remains polynomial. But n roots may
 364 appear sequentially, which would lead to an $\Omega(n)$ round complexity.

365 At this point, we cannot rule out that this algorithm is correct and has good properties.
 366 However, because of these problems, we took another approach.

367 **3.3 Our Solution**

368 In the end, our solution is obtained by a rather limited modification of the previous algorithm:
 369 we extend the range of the counters $p.c$ to the interval $[-B, B)$, but we restrict their range
 370 to $[-B, 0)$ when $p.s = E$.

371 Actually, this modification prevents *all* root creations. But, as with the previous attempt,
 372 we must be extra careful even with the smallest change, as proofs can easily break. We thus
 373 present the whole algorithm and its proofs in more details in the next sections, and further
 374 highlight the differences with [21] in Section 4.5.

375 4 A Unison Algorithm

376 4.1 Data Structures

377 Let $B \geq 2D + 2$ be an integer. Each node p maintains a single variable $p.v \in \{(C, x) \mid x \in$
 378 $[-B, B)\} \cup \{(E, x) \mid x \in [-B, 0)\}$. In the algorithm, $p.s$ and $p.c$, the *status* and the *clock*
 379 of p , respectively denote the left and right part of $p.v$. An assignment to $p.s$ or $p.c$ modifies
 380 the corresponding field of $p.v$.

381 We define the unison increment $a \oplus_B 1$ as $(B - 1) \oplus_B 1 = 0$ and $a \oplus_B 1 = a + 1$ if
 382 $a \in [-B, B - 2]$. Two clocks are *synchronized* if they are at most one increment apart.
 383 We then define $a \oplus_B b$ as the result of b iterations of $\oplus_B 1$ over a . Note that, as hinted in
 384 Section 3.2, we also use the usual addition and subtraction.

385 4.2 Some Predicates

386 Apart from its state, a node p has only access to the set $\{q.v \mid q \in N(p)\}$ of its neighbors'
 387 variables. A guard should thus not contain a direct reference to a neighbor q of p . This may
 388 look like a problem for we have already used such references. Nevertheless, these uses are
 389 legitimate as, for any predicate Pred , the semantics of $\exists(s, c) \in \{q.v \mid q \in N(p)\}, \text{Pred}(s, c)$
 390 is precisely $\exists q \in N(p), \text{Pred}(q.s, q.c)$. We can similarly encode universal statements.

391 As a matter of fact, we use the following shortcuts to increase readability:

Shortcut1 $\exists q \in N(p), \text{Pred}(q.s, q.c) := \exists(s, c) \in \{q.v \mid q \in N(p)\}, \text{Pred}(s, c)$

Shortcut2 $\forall q \in N(p), \text{Pred}(q.s, q.c) := \forall(s, c) \in \{q.v \mid q \in N(p)\}, \text{Pred}(s, c)$

392 Below, we define the predicates used by our algorithm.

$$\begin{aligned} \text{root}(p) := & (p.s = E \wedge \neg(\exists q \in N(p), q.s = E \wedge q.c < p.c)) \vee \\ & (p.s = C \wedge \exists q \in N(p), (q.c \geq p.c + 2) \wedge \neg(p.c = 0 \wedge q.c = B - 1)) \end{aligned}$$

$$\text{activeRoot}(p) := \text{root}(p) \wedge (p.c \neq -B \vee p.s = C)$$

$$\text{errorPropag}(p, i) := i < 0 \wedge \exists q \in N(p), q.s = E \wedge q.c < i < p.c$$

$$\text{canClearE}(p) := p.s = E \wedge \forall q \in N(p), (|q.c - p.c| \leq 1 \wedge (q.c \leq p.c \vee q.s = C))$$

$$\text{updatable}(p) := p.s = C \wedge \forall q \in N(p), q.c \in \{p.c, p.c \oplus_B 1\}$$

393 A node p is a *root* if $\text{root}(p)$. An *error rule* is either the rule R_R or a rule $R_P(i)$.

394 4.3 The Algorithm

395 A unison algorithm is rarely used alone. It is merely a tool to drive another algorithm. It
 396 thus makes sense that our algorithm depends on some properties which are external to the
 397 unison algorithm and its variables. Our algorithm uses a predicate P_{aux} which is not yet
 398 defined. As a matter of fact, its influence on the complexity analysis of the algorithm is very
 399 limited. To prove the correctness of the unison, we set $P_{\text{aux}} = \text{true}$, and we specialize P_{aux}
 400 differently in Section 6 when using our algorithm as a synchronizer.

$$\begin{array}{lll}
R_R : & \text{activeRoot}(p) & \longrightarrow p.c := -B ; p.s = E \\
R_P(i) : & \text{errorPropag}(p, i) & \longrightarrow p.c := i ; p.s = E \\
R_C : & \text{canClearE}(p) & \longrightarrow p.s := C \\
R_U : & \text{updatable}(p) \wedge P_{\text{aux}}(p) & \longrightarrow p.c := p.c \oplus_B 1
\end{array}$$

The rule R_R has the highest priority, and $R_P(i)$ has a higher priority than $R_P(i')$ for $i < i'$.

4.4 An Overview of the Algorithm

Contrary to [4] which proceeds by only locally synchronizing out-of-sync clocks, i.e., the clocks of two neighboring nodes that differ by at least two increments, we organize the synchronizations in *error broadcasts*. Every node p involved in such a broadcast is *in error* and its status is $p.s = E$. Otherwise, it is *correct* and $p.s = C$.

If p is correct, in sync with its neighbors, and if its clock $p.c$ is a local minimum, then p can apply the rule R_U to increment its clock.

There is a *cliff* between r and one of its neighbors p if their clocks are out-of-sync and $p.c > r.c$. If r is correct and has a cliff with a neighbor, then r is said to be a *root* and should initiate an error broadcast by applying the rule R_R , which respectively sets $r.c$ to $-B$ and $r.s$ to E .

If there is a cliff between r and p , r is in error, and $r.c < -1$, then p should propagate the broadcast by applying the rule R_P which sets $p.c$ to $r.c + 1$ and $p.s$ to E . If p has several such neighbors r , it applies R_P according to the one with the minimum clock.

As a consequence, any node p in error with $p.c > -B$ should have at least one neighbor in error with a smaller clock. This way, the structure of an error broadcast is a *dag* (directed acyclic graph). We therefore extend the definition of root to include nodes in error with no “parents” in the broadcast dag.

Note that a node may decrease its clock multiple times using R_P , and in doing so may consecutively join several error dags or several parts of them. This way, nodes reduce the height of the error dags, which is a key element to achieve the $O(D)$ -round complexity. Furthermore, any node in error eventually has a clock smaller than $-B + D$ and all cliffs are eventually destroyed.

Finally, if p is in error, is not involved in any cliff (in which case an error must be propagated), and if all its neighbors with larger clocks are correct, then the broadcast from p is finished, and p can apply the rule R_C to switch back $p.s$ to C .

A key element to bound the move complexity is that a dag built during an error broadcast is cleaned from the larger clocks to the smaller, but nodes previously in the dag resume the “normal” increments (using the rule R_U) in the reverse order (i.e., from the smaller clocks to the larger). Indeed, a non-root node in an error broadcast is one increment ahead of its parents in the dag and so has to wait for their increment before being able to perform one itself. Hence, the first node in the dag that makes a R_U move after an error broadcast is its root.

4.5 Some Subtleties

Some statements and the corresponding proof arguments are very similar to the ones of [21] (rather its arXiv version [20]). However, the fact that the algorithm and its data structures are different imply that proofs are indeed different. As a matter of fact, we have tried but failed to unify both algorithms into a *natural* more general one.

Below, we outline subtleties which are specific to our algorithm.

- 442 ■ Since nodes in error are restricted to negative clocks, it is natural to expect that legitimate
 443 configurations require all clocks to be non-negative. This would suggest a $\Theta(B)$ round
 444 complexity, which is weaker than what we claim. But this intuition is false. For example,
 445 the configuration where all nodes are correct and all clocks are set to $-B$ is legitimate.
 446 This is one of the reasons for our $O(D)$ round complexity.
- 447 ■ In the unbounded unison algorithm above which we derive from [21], whenever two
 448 neighboring nodes p and q are such that $q.s = E$ and $p.c \geq q.c + 2$, the node p can always
 449 apply a rule $R_P(i)$. In our algorithm, this is not the case when $q.c = -1$. This could
 450 introduce unexpected behaviors which could impact the complexities of our algorithm, or
 451 in the worst case, lead to deadlocks. We thus have to deal with this slight difference in
 452 the proofs.
- 453 ■ In [21], the proofs heavily rely on the fact that the counters increase when applying the
 454 rule R_U while they decrease when applying the rules R_R and $R_P(i)$. This monotony
 455 property is however not true in our setting. More generally, having two addition operators
 456 $+$ and \oplus_B requires special care throughout the proofs.
- 457 ■ Finally, to bound the memory, the maximum clock is $B - 1$, after which clocks go back
 458 to 0. Notice that to ensure the liveness property of the unison, we must have $B \geq 2D + 2$
 459 (an example of deadlock is presented for $B = 2D + 1$ in Subsection 5.2).

460 5 Self-Stabilization and Complexity of the Unison Algorithm

461 As already mentioned, the unison algorithm corresponds to $P_{\text{aux}} = \text{true}$. However, since
 462 most proofs are valid regardless of the definition of P_{aux} , we only specify it when needed.
 463 We define the *legitimate* configurations as the configurations without roots. Let $e = \gamma^0 \gamma^1 \dots$
 464 be an execution. We respectively denote by $p.s^i$ and $p.c^i$ the value of $p.s$ and $p.c$ in γ^i .

465 5.1 Convergence and Move Complexity of the Unison Algorithm

466 Although it is tedious, it is straightforward to prove, by case analysis, that roots cannot be
 467 created. Since roots are obstructions to legitimate configurations, it is natural to partition
 468 the steps of e into *segments* such that each step in which at least one root disappears is the
 469 last step of a segment. There are thus at most n segments with roots, which constitute the
 470 *stabilization phase*, and at most one root-less segment. We now show that the stabilization
 471 phase is finite by providing a (finite) bound on its move complexity.

472 In the following, s is any segment of the stabilization phase. The key fact is that in s , a
 473 node p in error cannot apply the rule R_U until the end of s . We prove this by induction on
 474 $p.c$. If $p.c = -B$, then p is a root, and the only rule that p can apply is R_C , which removes
 475 its root status. The base case thus follows. Now let p be in error with $p.c > -B$. If p does
 476 not move in s , then our claim holds. Otherwise let $\gamma^a \mapsto \gamma^b$ be the first step in which p
 477 moves. If p applies the rule R_R , then $p.c^b = -B$, and for the remainder of s , the claim holds
 478 by induction. Otherwise, p has a neighbor q such that $q.s^a = E$ and $q.c^a < p.c^a < 0$. By
 479 induction, $q.c$ cannot increase until the end of s . As long as $p.c > q.c$, p cannot apply the rule
 480 R_U and if, at some point, $p.c \leq q.c$, then p must have applied an error rule, thus decreasing
 481 its clock, at which point the claim holds by induction.

482 Since roots cannot be created, the number of R_R -moves is at most n . Moreover, since
 483 between two R_C -moves, there has to be at least one error move (R_R - or R_P -move), we have
 484 $\#R_C\text{-moves} \leq n + \#R_R\text{-moves} + \#R_P\text{-moves} \leq 2n + \#R_P\text{-moves}$. We thus only need to
 485 bound the number of R_U -moves and R_P -moves.

486 We now bound the number of R_U moves by a node in s . If a node q does not move
 487 between γ^a and γ^b in s with $a < b$, then a neighbor p can apply the rule R_U at most twice,
 488 to go from $q.c^a - 1$ to $q.c^a + 1$. More generally, if $p.c^b \geq p.c^a + 2 + i$ (we really mean the
 489 $+$ operator and not the \oplus_B operator), then every neighbor q of p must increase its clock by
 490 at least i between γ^a and γ^b . By induction on d , if $p.c^b = p.c^a + 2d + i$, then every node q
 491 at distance d from p increases its clock by at least i between γ^a and γ^b . Since roots cannot
 492 increase their clocks, this implies that $p.c^b \leq p.c^a + 2D$.

493 From this “linear” bound, we now derive a “circular” bound which takes into account the
 494 fact that the clock of a node may decrease while applying the rule R_U (from $B - 1$ to 0).
 495 In the worst case, p could apply the rule R_U $2D$ times to reach $p.c = B - 1$, then apply
 496 R_U once so that $p.c = 0$, then reapply R_U $2D$ more times (recall that $B \geq 2D + 2$). To
 497 summarize, p may apply R_U at most $4D + 1$ times in s . This gives an $O(n^2D)$ bound on the
 498 number of R_U -moves done during the stabilization phase.

499 We now focus on the rule R_P in s . If a node p_0 applies a rule R_P in a step $\gamma^{j_1} \mapsto \gamma^{j_1+1}$ of
 500 s , it does so to “connect” to a neighbor p_1 which is already in error. Now p_1 may be in error
 501 in γ^{j_1} because it has applied a rule R_P in another step $\gamma^{j_2} \mapsto \gamma^{j_2+1}$ of s with $j_2 < j_1$, to
 502 connect to a neighbor p_2 , and so on. This defines a *causality chain* $p_0 \cdots p_l$ for some l . Since,
 503 according to the key fact, rules R_P and R_U do not alternate in s , a node cannot appear
 504 twice in the causality chain, thus $l < n$. Moreover, when considering a maximal causality
 505 chain, $p_l.c^{j_l}$ is either the value of $p_l.c$ at the beginning of s , or $-B$ if p_l has applied the
 506 rule R_R . The clock $p_0.c$ can thus take at most $n(n + 1)$ distinct values in s , which implies
 507 that the rule R_P is applied at most $O(n^2)$ times in s by a given node. This gives an overall
 508 $O(n^4)$ -bound on the number of rules R_P . Note that a more careful analysis gives an overall
 509 bound of $O(n^3)$ on the total number of R_P -moves.

510 We can also easily obtain a bound that involves B . Indeed, a node p has at most B
 511 R_P -moves and $4D + 1 = O(B)$ R_U -moves in s . This gives an $O(n^2B)$ -bound on the number
 512 of moves. To summarize, the stabilization phase terminates after at most $O(\min(n^3, n^2B))$
 513 moves.

514 Note that any configuration γ with at least one root contains at least one enabled node.
 515 Indeed, if any two neighboring clocks are at most one increment apart, then any root is in
 516 error, and the rule R_C is enabled at any node p in error with $p.c$ maximum. Otherwise, there
 517 exist two neighbors p and q such that $p.c$ and $q.c$ are more than one increment apart. We
 518 choose them with $q.c < p.c$ and $q.c$ minimum. $q.c$ being minimum, we can show that either q
 519 is a root that is enabled for R_R , or p can apply the rule R_P because q is in error and satisfies
 520 $q.c \leq -B + D < -1$. Thus, the last configuration of the stabilization phase is legitimate.

521 Also, note that since roots cannot be created, being legitimate is a closed property,
 522 meaning that in a step $\gamma^a \mapsto \gamma^b$, if γ^a is legitimate, then so is γ^b .

523 5.2 Correctness of the Unison Algorithm

524 We now show that any legitimate configuration γ satisfies the safety property of the unison.
 525 First, γ cannot contain nodes in error, because any such node p with $p.c$ minimum would be
 526 a root. Moreover, if the clocks of two correct neighbors differ by more than one increment,
 527 then the node with the smaller clock is a root.

528 To prove the liveness property of the unison, we set $P_{\text{aux}} = \text{true}$ in this paragraph. In
 529 legitimate configurations, since neighboring clocks differ by at most one increment, any two
 530 clocks differ by at most D increments. And since $B \geq 2D + 2$, there exists $c \in [0, B)$ which
 531 is not the clock of any node. This implies that there exists at least one node p whose clock is
 532 not the increment of any other clock. Thus, p satisfies *updatable* and can apply R_U . This

533 proves that at least one node applies R_U infinitely often, and thus so do all nodes. Observe
 534 that $B \geq 2D + 2$ is tight. Indeed, when $B = 2D + 1$, the configuration of the cycle p_0 ,
 535 p_1, \dots, p_{2D} in which all nodes are correct and $p_i.c = i$, is legitimate but is terminal.

536 5.3 Round Complexity of the Unison Algorithm

537 We claim that $\gamma^{r_{2D+2}}$ contains no roots and so is legitimate. Recall that for all $i \geq 1$, Round i
 538 is $\gamma^{r_{i-1}} \dots \gamma^{r_i}$. We suppose that all γ^{r_i} with $i \leq 2D + 1$ contain roots otherwise our claim
 539 directly holds.

540 We now study the first $D + 1$ rounds. Let r be any root in $\gamma^{r_{D+1}}$. Since there are no
 541 root creations, r is already a root in γ^0 . By the end of the first round (using the rule R_R if
 542 needed), $r.c = -B$ and $r.s = E$. Now, since r is still a root in $\gamma^{r_{D+1}}$, it cannot make a move
 543 in the meantime, and its state does not change until $\gamma^{r_{D+1}}$. Furthermore, every neighbor
 544 p of r such that $p.c > -B + 1$ can apply the rule R_P . So, by the end of Round 2, and as
 545 long as r does not increment its clock, $p.c \leq -B + 1$. By induction on the distance $d(p, r)$
 546 between p and r , we can prove that $p.c^{r_{D+1}} \leq -B + d(p, r)$ for every node p and every root
 547 r in $\gamma^{r_{D+1}}$.

548 We claim that $\gamma^{r_{D+1}}$ does not contain any *cliff*, i.e., a pair (q, p) of neighboring nodes
 549 whose clocks are out-of-sync and such that $q.c < p.c$. Suppose that (q, p) is a cliff in $\gamma^{r_{D+1}}$.
 550 The node q is in error as otherwise it would be a root not in error, which, as already mentioned,
 551 is impossible from γ^{r_1} . Moreover, we can prove by induction on $q.c$ that there is a root r in
 552 $\gamma^{r_{D+1}}$ such that $q.c \geq -B + d(q, r)$. Since $p.c \geq q.c + 2$, we have $p.c > -B + d(p, r)$, which
 553 contradicts the result of the previous paragraph.

554 We now consider the next $D + 1$ rounds. Since $\gamma^{r_{D+1}}$ contains no nodes which can apply
 555 R_R , and no cliffs, nodes can only apply the rules R_U or R_C . Furthermore, among nodes in
 556 error, those with the largest clock can apply the rule R_C , which implies that roots no longer
 557 exist by the end of Round $2D + 2$, and thus $\gamma^{r_{2D+2}}$ is legitimate.

558 6 A Synchronizer

559 Let us consider a variant of the atomic-state model which is at least as expressive as the
 560 model of our unison algorithm. This means that, in this model, we should be able to encode
 561 the shortcuts `Shortcut1` and `Shortcut2` (defined page 9).

562 In this model, let Alg_I be any silent algorithm which is self-stabilizing with a projection
 563 $proj$ for a static specification SP under the synchronous daemon. Using folklore ideas (see,
 564 e.g., [4] and [28]), we define in this section a synchronizer which uses our unison to transform
 565 Alg_I into an algorithm $Sync(Alg_I)$ which “simulates” synchronous executions of Alg_I in an
 566 asynchronous environment under a distributed unfair daemon.

567 6.1 The Synchronized Algorithm

568 On top of its unison variables, each node p stores two states of Alg_I , in the variables $p.old$
 569 and $p.curr$. These variables ought to contain the last two states of p in a synchronous
 570 execution of Alg_I . When p applies the rule R_U , it also computes a next state of Alg_I . It
 571 does so by applying the function $\widehat{Alg_I}$ which selects $p.curr$ and, for each neighbor q , the
 572 variable $q.curr$ if $p.c = q.c$, and $q.old$ if $q.c = p.c \oplus_B 1$, and applies Alg_I on these values. We
 573 thus modify the rule R_U in the following way:

$$R_U : updatable(p) \wedge P_{aux}(p) \longrightarrow p.old := p.curr; p.curr := \widehat{Alg_I}(p); p.c := p.c \oplus_B 1.$$

574 The folklore algorithm corresponds to the case when $P_{\text{aux}}(p)$ is always *true*. In this case,
 575 the clocks of the unison constantly change. Thus, even if Alg_I is silent, its simulation is not.
 576 To obtain a silent simulation, we devise another strategy by defining $P_{\text{aux}}(p)$ as follows.

$$P_{\text{aux}}(p) = (\widehat{Alg_I}(p) \neq p.\text{curr}) \vee (\exists q \in N(p), q.c = p.c \oplus_B 1).$$

577 We define the legitimate configurations of $Sync(Alg_I)$ to be its terminal configurations.
 578 In the next sections, we sketch the proof that $Sync(Alg_I)$ is self-stabilizing for the same
 579 specification as Alg_I . As a matter of fact, our result is more general as the silent assumption
 580 is not necessary (we need a different definition for the legitimate configurations though).

581 6.2 Convergence and Move Complexity of the Synchronized Algorithm

582 In everyday life, we have a distinction between the value of a clock (modulo 24 hours) and
 583 the time. Both are obviously linked. We would like to make a similar distinction here. Let
 584 $e = \gamma^0 \gamma^1 \dots$ be an execution of legitimate configurations. Since γ^0 is legitimate, every two
 585 neighboring clocks differ by at most one increment.

586 Since $B \geq 2D + 2$, as already mentioned, at least one element of $[0, B)$ is the clock of no
 587 nodes in γ^0 . This implies that there is a node x such that $x.c^0 \oplus_B 1$ is not the clock of any
 588 node. We extend this local synchronization property by uniquely defining $time(p)^0 \in [-D, 0]$
 589 by (1) $time(x)^0 = 0$, (2) if $p.c^0 = q.c^0$, then $time(p)^0 = time(q)^0$, and (3) if $p.c \oplus_B 1 = q.c$,
 590 then $time(p)^0 = time(q)^0 - 1$. Moreover, we also define $time(p)^{i+1} = time(p)^i + 1$ if p applies
 591 R_U in $\gamma^i \mapsto \gamma^{i+1}$, and $time(p)^{i+1} = time(p)^i$ otherwise.

592 For any $i, j \geq 0$ such that $time(p)^j = i$, we set $\mathbf{st}_p^i := p.\text{curr}^j$. When \mathbf{st}_p^i is defined for
 593 all p , let λ^i be the configuration in which the state of each node p is \mathbf{st}_p^i . A careful analysis
 594 shows that, by definition of P_{aux} , λ^i exists as soon as some \mathbf{st}_p^i does, and $\Lambda = \lambda^0 \lambda^1 \dots$ is
 595 precisely the synchronous execution of Alg_I from λ^0 .

596 Suppose that T is a bound on the number of rounds that Alg_I needs to reach silence.
 597 Thus, $\Lambda = \lambda^0 \dots \lambda^H$ for some $H \leq T$. In the simulation phase, a node makes at most D
 598 moves to have a non-negative time, and then at most T moves to finish the simulation.
 599 Together with the stabilization time of the unison, our simulated algorithm is also silent with
 600 an $O(\min(n^3, n^2 B) + nD + nT) = O(\min(n^3, n^2 B) + nT)$ move complexity.

601 6.3 Correctness of the Synchronized Algorithm

602 In $Sync(Alg_I)$, we define the *restriction* $rest(s)$ of the state s of any node p to be $p.\text{curr}$,
 603 and we canonically extend $rest$ to configurations and executions. Let us consider a legitimate
 604 configuration γ of $Sync(Alg_I)$. This configuration is terminal, and therefore there exists a
 605 unique execution e of $Sync(Alg_I)$ starting at γ (the one restricted to γ alone). Besides, since
 606 γ is terminal, its restriction is terminal too (for Alg_I). Therefore $rest(\gamma)$ is legitimate, and
 607 $proj(rest(e))$ satisfies the specification SP . Hence, the algorithm $Sync(Alg_I)$ also satisfies
 608 SP (for the projection $proj \circ rest$).

609 6.4 Round Complexity of the Synchronized Algorithm

610 The round complexity is analyzed by considering two stages: a first stage to have all times
 611 non-negative, and a second stage to have all times equal to H .

612 To give an intuition of our proof, as it is the more complex, we first consider the second
 613 stage. Figure 1 is an illustration of the following explanation. Suppose that all times are 0
 614 in γ^0 , and only s_1 is such that $\mathbf{st}_{s_1}^0 \neq \mathbf{st}_{s_1}^1$. In the first round of the synchronous execution,

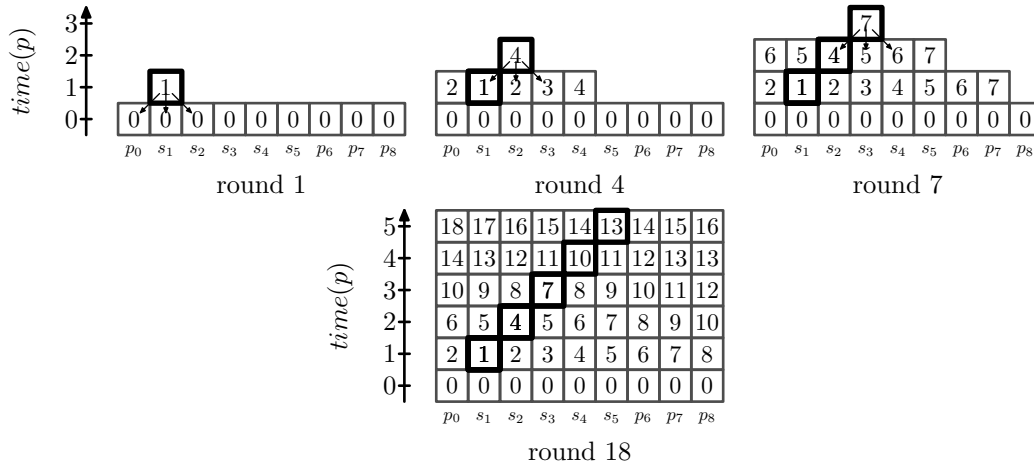


Figure 1 The intuition of the round complexity for the second stage.

615 s_1 applies R_U , and then, after each new round, nodes at distance 1 from s_1 , then 2, and so
 616 on will increase their time to 1. Now suppose that only $s_2 \in N[s_1]$ is such that $st_{s_2}^1 \neq st_{s_2}^2$.
 617 As soon as all nodes in $N[s_2]$ have a time of 1, s_2 applies R_U . This happens at Round 3 if
 618 $s_2 = s_1$ and at Round 4 otherwise. After this, after each new round, nodes at distance 1
 619 from s_2 , then 2, and so on will increase their time to 2. If we consider some $s_3 \in N[s_2]$, and
 620 so on, then s_i increases its time to i at Round at most $3i - 2$, and all nodes do so at Round
 621 at most $3i + D - 2$. If nodes increase their time earlier, this only speed up the process.

622 Now, by definition of H , there is a node s_H whose state changes between λ_{H-1} and λ_H .
 623 If the states of all nodes in $N[p]$ were the same in λ_{H-2} and λ_{H-1} , then s_H would not have
 624 changed its state between λ_{H-1} and λ_H . There thus exists $s_{H-1} \in N[s_H]$ that changes its
 625 state between λ_{H-2} and λ_{H-1} . By repeating this process, we can prove that, unless $H = 0$,
 626 Λ has a *starting sequence* that is a sequence $s_1 \dots s_H$ verifying $st_{s_i}^{i-1} \neq st_{s_i}^i$ for $1 \leq i \leq H$,
 627 and $s_{i-1} \in N[s_i]$ for $1 < i \leq H$. We can then prove that, if all nodes have a positive time at
 628 Round X , then the algorithm becomes silent after at most $X + 3H + D - 2$ rounds.

629 Using similar ideas, we can prove that all times are non-negative after at most $X = 2D$
 630 rounds. Taking into account the $3D + 2$ rounds of the stabilization phase, we obtain an
 631 overall $5D + 3T$ round complexity to reach the silence from any configuration.

7 Conclusion

633 We propose the first fully polynomial self-stabilizing unison algorithm for anonymous
 634 asynchronous bidirectional networks of arbitrary connected topology, and use it to obtain new
 635 state-of-the-art algorithms for various problems such as BFS constructions, leader election,
 636 and clustering.

637 A challenging perspective would be to generalize our approach to weaker models such
 638 as the message passing or the link-register models. We would also be curious to know the
 639 properties of the algorithm proposed in Section 3.2. Thirdly, although we could not do it, it
 640 would be nice to unify our result with that of [21] in a satisfactory manner. Finally, it would
 641 be interesting to know whether or not constant memory can be achieved by an asynchronous
 642 self-stabilizing unison for arbitrary topologies.

643 ——— **References** ———

- 644 1 S. Aggarwal and S. Kutten. Time optimal self-stabilizing spanning tree algorithms. In *13th*
645 *Foundations of Software Technology and Theoretical Computer Science, (TSTTCS'93)*, volume
646 761, pages 400–410, 1993. doi:10.1007/3-540-57529-4_72.
- 647 2 K. Altisen, A. Cournier, S. Devismes, A. Durand, and F. Petit. Self-stabilizing leader
648 election in polynomial steps. *Information and Computation*, 254(3):330–366, 2017. doi:
649 10.1016/j.ic.2016.09.002.
- 650 3 K. Altisen, S. Devismes, S. Dubois, and F. Petit. *Introduction to Distributed Self-Stabilizing*
651 *Algorithms*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool, 2019.
652 doi:10.2200/S00908ED1V01Y201903DCT015.
- 653 4 B. Awerbuch, S. Kutten, Y. Mansour, B. Patt-Shamir, and G. Varghese. Time optimal self-
654 stabilizing synchronization. In *25th Annual Symposium on Theory of Computing, (STOC'93)*,
655 pages 652–661, 1993. doi:10.1145/167088.167256.
- 656 5 B. Awerbuch, B. Patt-Shamir, and G. Varghese. Self-stabilization by local checking and
657 correction. In *32nd Annual Symposium of Foundations of Computer, Science (FOCS'91)*,
658 pages 268–277, 1991. doi:10.1109/SFCS.1991.185378.
- 659 6 B. Awerbuch and G. Varghese. Distributed program checking: a paradigm for building self-
660 stabilizing distributed protocols. In *32nd Annual Symposium on Foundations of Computer*
661 *Science, (FOCS'91)*, pages 258–267, 1991. doi:10.1109/SFCS.1991.185377.
- 662 7 Baruch Awerbuch. Complexity of network synchronization. *J. ACM*, 32(4):804–823, 1985.
663 doi:10.1145/4221.4227.
- 664 8 J. Beauquier, M. Gradinariu, and C. Johnen. Cross-over composition - enforcement of fairness
665 under unfair adversary. In *5th International Workshop on Self-Stabilizing Systems, (WSS'01)*,
666 pages 19–34, 2001. doi:10.1007/3-540-45438-1_2.
- 667 9 L. Blin, C. Johnen, G. Le Boudier, and F. Petit. Silent anonymous snap-stabilizing termination
668 detection. In *41st International Symposium on Reliable Distributed Systems, (SRDS'22)*, pages
669 156–165, 2022. doi:10.1109/SRDS55811.2022.00023.
- 670 10 C. Boulinier and F. Petit. Self-stabilizing wavelets and rho-hops coordination. In *22nd IEEE*
671 *International Symposium on Parallel and Distributed Processing, (IPDPS'08)*, pages 1–8, 2008.
672 doi:10.1109/IPDPS.2008.4536130.
- 673 11 C. Boulinier, F. Petit, and V. Villain. When graph theory helps self-stabilization. In *23rd*
674 *Annual Symposium on Principles of Distributed Computing, (PODC'04)*, pages 150–159, 2004.
675 doi:10.1145/1011767.1011790.
- 676 12 J. Burman and S. Kutten. Time optimal asynchronous self-stabilizing spanning tree. In *21st*
677 *International Symposium on Distributed Computing, (DISC'07)*, volume 4731, pages 92–107,
678 2007. doi:10.1007/978-3-540-75142-7_10.
- 679 13 A. Cournier, A. K. Datta, F. Petit, and V. Villain. Snap-stabilizing PIF algorithm in arbitrary
680 networks. In *22nd International Conference on Distributed Computing Systems (ICDCS'02)*,
681 pages 199–206, 2002. doi:10.1109/ICDCS.2002.1022257.
- 682 14 A. Cournier, S. Rovedakis, and V. Villain. The first fully polynomial stabilizing algorithm for
683 BFS tree construction. *Information and Computation*, 265:26–56, 2019. doi:10.1016/j.ic.
684 2019.01.005.
- 685 15 J.-M. Couvreur, N. Francez, and M. G. Gouda. Asynchronous unison (extended abstract). In
686 *12th International Conference on Distributed Computing Systems, (ICDCS'92)*, pages 486–493,
687 1992. doi:10.1109/ICDCS.1992.235005.
- 688 16 A. K. Datta, S. Devismes, K. Heurtefeux, L. L. Larmore, and Y. Rivierre. Competitive
689 self-stabilizing k -clustering. *Theoretical Computer Science*, 626:110–133, 2016. doi:10.1016/
690 j.tcs.2016.02.010.
- 691 17 A. K. Datta, S. Devismes, and L. L. Larmore. A silent self-stabilizing algorithm for the
692 generalized minimal k -dominating set problem. *Theoretical Computer Science*, 753:35–63,
693 2019. doi:10.1016/j.tcs.2018.06.040.

- 694 18 A. K. Datta and L. L. Larmore. Leader election and centers and medians in tree networks. In
695 *15th International Symposium on Stabilization, Safety, and Security of Distributed Systems,*
696 *(SSS'13)*, pages 113–132, 2013. doi:10.1007/978-3-319-03089-0_9.
- 697 19 S. Devismes, D. Ilcinkas, and C. Johnen. Optimized silent self-stabilizing scheme for tree-based
698 constructions. *Algorithmica*, 84(1):85–123, 2022. doi:10.1007/s00453-021-00878-9.
- 699 20 S. Devismes, D. Ilcinkas, C. Johnen, and F. Mazoit. Making local algorithms efficiently
700 self-stabilizing in arbitrary asynchronous environments. *CoRR*, abs/2307.06635, 2023. doi:
701 10.48550/arXiv.2307.06635.
- 702 21 S. Devismes, D Ilcinkas, C. Johnen, and F. Mazoit. Asynchronous self-stabilization made
703 fast, simple, and energy-efficient. In *43rd Symposium on Principles of Distributed Computing,*
704 *(PODC'24)*, pages 538–548, 2024. doi:10.1145/3662158.3662803.
- 705 22 S. Devismes and C. Johnen. Silent self-stabilizing BFS tree algorithms revisited. *Journal on*
706 *Parallel Distributed Computing*, 97:11–23, 2016. doi:10.1016/j.jpdc.2016.06.003.
- 707 23 S. Devismes and C. Johnen. Self-stabilizing distributed cooperative reset. In *39th International*
708 *Conference on Distributed Computing Systems, (ICDCS'19)*, pages 379–389, 2019. doi:
709 10.1109/ICDCS.2019.00045.
- 710 24 S. Devismes and F. Petit. On efficiency of unison. In *4th Workshop on Theoretical Aspects*
711 *of Dynamic Distributed Systems, (TADDS'12)*, pages 20–25, 2012. doi:10.1145/2414815.
712 2414820.
- 713 25 E. W. Dijkstra. Self-stabilization in spite of distributed control. *Communications of the ACM*,
714 17(11):643–644, 1974. doi:10.1145/361179.361202.
- 715 26 S. Dolev. *Self-Stabilization*. MIT Press, 2000.
- 716 27 S. Dolev, A. Israeli, and S. Moran. Self-stabilization of dynamic systems assuming only
717 read/write atomicity. *Distributed Computing*, 7(1):3–16, 1993. doi:10.1007/BF02278851.
- 718 28 Y. Emek and E. Keren. A thin self-stabilizing asynchronous unison algorithm with applications
719 to fault tolerant biological networks. In *40nd Symposium on Principles of Distributed*
720 *Computing, (PODC'21)*, pages 93–102, 2021. doi:10.1145/3465084.3467922.
- 721 29 Y. Emek and R. Wattenhofer. Stone age distributed computing. In *32nd Symposium on*
722 *Principles of Distributed Computing, (PODC'13)*, pages 137–146, 2013. doi:10.1145/2484239.
723 2484244.
- 724 30 M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with
725 one faulty process. *Journal of the ACM*, 32(2):374–382, 1985. doi:10.1145/3149.214121.
- 726 31 C. Glacet, N. Hanusse, D. Ilcinkas, and C. Johnen. Disconnected components detection
727 and rooted shortest-path tree maintenance in networks. *Journal of Parallel and Distributed*
728 *Computing*, 132:299–309, 2019. doi:10.1016/j.jpdc.2019.05.006.
- 729 32 Maria Gradinariu and Sébastien Tixeuil. Conflict managers for self-stabilization without
730 fairness assumption. In *27th IEEE International Conference on Distributed Computing Systems*
731 *(ICDCS 2007), June 25-29, 2007, Toronto, Ontario, Canada*, page 46. IEEE Computer Society,
732 2007. doi:10.1109/ICDCS.2007.95.
- 733 33 A. Kravchik and S. Kutten. Time optimal synchronous self stabilizing spanning tree. In
734 *27th International Symposium on Distributed Computing, (DISC'13)*, pages 91–105, 2013.
735 doi:10.1007/978-3-642-41527-2_7.
- 736 34 S. Tixeuil. *Vers l'auto-stabilisation des systèmes à grande échelle*. Habilitation à diriger des
737 recherches, Université Paris Sud - Paris XI, 2006. URL: [https://tel.archives-ouvertes.
738 fr/tel-00124848/file/hdr_final.pdf](https://tel.archives-ouvertes.fr/tel-00124848/file/hdr_final.pdf).
- 739 35 V. Turau. Efficient transformation of distance-2 self-stabilizing algorithms. *Journal of Parallel*
740 *and Distributed Computing*, 72(4):603–612, 2012. doi:10.1016/j.jpdc.2011.12.008.