



HAL
open science

Gradual Stabilization under τ -Dynamics

Karine Altisen, Stéphane Devismes, Anaïs Durand, Franck Petit

► **To cite this version:**

Karine Altisen, Stéphane Devismes, Anaïs Durand, Franck Petit. Gradual Stabilization under τ -Dynamics. Euro-Par 2016 - 22nd International Conference on Parallel and Distributed Computing, Aug 2016, Grenoble, France. pp.588-602, 10.1007/978-3-319-43659-3_43 . hal-01215190v4

HAL Id: hal-01215190

<https://hal.science/hal-01215190v4>

Submitted on 27 Jan 2016 (v4), last revised 2 Feb 2017 (v5)

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.

Public Domain

Gradual Stabilization under τ -Dynamics

Karine Altisen,[§] Stéphane Devismes,[§] Anaïs Durand,[§] and Franck Petit[†]

[§]VERIMAG UMR 5104, Université Grenoble Alpes, France

[†]LIP6 UMR 7606, INRIA, UPMC Sorbonne Universités, France

Abstract

We introduce the notion of *gradually stabilizing* algorithm as any self-stabilizing algorithm achieving the following additional feature. If at most τ *dynamic steps* occur starting from a legitimate configuration, a gradually stabilizing algorithm first quickly recovers to a configuration from which a specification offering a minimum quality of service is satisfied. It then gradually converges to specifications offering stronger and stronger safety guarantees until reaching a configuration (1) from which its initial (strong) specification is satisfied again, and (2) where it is ready to achieve gradual convergence again in case of up to τ new dynamic steps. By definition, a gradually stabilizing algorithm is also self-stabilizing. So, it still recovers within finite time (yet more slowly) after any other finite number of transient faults, including for example more than τ dynamic steps or other failure patterns such as memory corruptions, for example.

We illustrate this new property by considering three variants of a synchronization problem respectively called *strong*, *weak*, and *partial weak* unison. We propose a self-stabilizing algorithm which achieves gradual stabilization in the sense that after one dynamic step from a configuration which is legitimate for the strong unison, it immediately satisfies the specification of partial weak unison, then converges to the specification of weak unison in at most one round, and finally retrieves, after at most $(\mu + 1)\mathcal{D}_1 + 1$ additional rounds, a configuration (1) from which the specification of strong unison is satisfied, and (2) where it is ready to achieve gradual convergence again in case of another dynamic step. \mathcal{D}_1 is the diameter of the network after the dynamic step, and μ is a parameter satisfying $\mu \geq n + \#J$, where n is the initial number of processes and $\#J$ is an upper bound on the number of processes that join the system during a dynamic step.

Keywords: Self-stabilization, synchronization problems, unison, gradual stabilization, superstabilization, safe-convergence.

Contents

1	Introduction	3
2	Preliminaries	7
2.1	Processes	7
2.2	Executions	8
2.3	Static Steps	8
2.4	Dynamic Steps	9
2.5	Daemon	9
2.6	Functional Specification and Performances	9
3	Stabilization	10
3.1	Self-stabilization	10
3.2	Gradual Stabilization under τ -Dynamics	10
4	Unison	12
5	Necessary Condition	14
6	Self-Stabilizing Strong Unison	17
6.1	Algorithm \mathcal{WU}	17
6.2	Algorithm \mathcal{SU}	19
6.2.1	Correctness Proof	21
6.2.2	Complexity Analysis	24
7	Gradual Stabilization under 1-Dynamics for Strong Unison	25
7.1	Algorithm \mathcal{DSU}	25
7.2	Proof of Correctness	30
8	Conclusion	37
A	Results from [5]	41

1 Introduction

In 1974, Dijkstra [9] introduced *self-stabilization*, a general paradigm to enable the design of distributed systems tolerating *any* finite number of transient faults.¹ Consider the first configuration after all transient faults cease. This configuration is arbitrary, but no other transient faults will ever occur from this configuration. By abuse of language, this configuration is referred to as *arbitrary initial configuration* of the system in the literature. Then, a self-stabilizing algorithm (provided that faults have not corrupted its code) guarantees that starting from an arbitrary initial configuration, the system recovers *within finite time*, without any external intervention, to a so-called *legitimate configuration* from which its specification is satisfied. Thus, self-stabilization makes no hypotheses on the nature (*e.g.*, memory corruptions or topological changes) or extent of transient faults that could hit the system, and the system recovers from the effects of those faults in a unified manner. Such versatility comes at a price, *e.g.*, after transient faults cease, there is a finite period of time, called the *stabilization phase*, during which the safety properties of the system are violated. Hence, self-stabilizing algorithms are mainly compared according to their *stabilization time*, the maximum duration of the stabilization phase. For many problems, the stabilization time is significant, *e.g.*, for synchronization problems [2] and more generally for non-static problems [13] (such as token passing) the lower bound is $\Omega(\mathcal{D})$ rounds, where \mathcal{D} is the diameter of the network. By definition, the stabilization time is impacted by worst case scenarios. Now, in most cases, transient faults are sparse and their effect may be superficial. Recent research focuses on proposing self-stabilizing algorithms that also ensure drastically smaller convergence times in favorable cases.

Defining the number of faults hitting a network using some kind of Hamming distance (the minimal number of processes whose state must be changed in order to recover a legitimate configuration), variants of the self-stabilization paradigm have been defined, *e.g.*, a *time-adaptive* self-stabilizing algorithm [25] additionally guarantees a convergence time in $O(k)$ time-units when the initial configuration is at distance at most k from a legitimate configuration.

The property of *locality* consists in avoiding situations in which a small number of transient faults causes the entire system to be involved in a global convergence activity. Locality is, for example, captured by *fault containing* self-stabilizing algorithms [14], which ensure that when few faults hit the system, the faults are both spatially and temporally contained. “Spatially” means that if only few faults occur, those faults cannot be propagated further than a preset radius around the corrupted processes. “Temporally”

¹Transient faults have low frequency and results in perturbing the state of the system.

means quick stabilization when few faults occur.

Some other approaches consist in providing convergence times tailored by the type of transient faults. For example, a *superstabilizing algorithm* [11] is self-stabilizing and has two additional properties when transient faults are limited to a single topological change. Indeed, after adding or removing one link or process in the network, a superstabilizing algorithm recovers fast (typically $O(1)$ rounds), and a safety predicate, called a *passage predicate*, should be satisfied all along the stabilization phase.

Contribution. In this paper, we introduce a specialization of self-stabilization called *gradual stabilization*. A gradually stabilizing algorithm is a *self-stabilizing algorithm* with the following additional feature. If at most τ *dynamic steps*² occur starting from a legitimate configuration, a gradually stabilizing algorithm first quickly recovers to a configuration from which a specification offering a minimum quality of service is satisfied. It then gradually converges to specifications offering stronger and stronger safety guarantees until reaching a configuration (1) from which its initial (strong) specification is satisfied again, and (2) where it is ready to achieve gradual convergence again in case of up to τ new dynamic steps. Of course, the gradual stabilization makes sense only if the convergence to every intermediate weaker specification is fast.

We illustrate this new property by considering three variants of a synchronization problem respectively called *strong*, *weak*, and *partial weak* unison. In these problems, each process should maintain a local clock. We restrict here our study to periodic clocks, *i.e.*, all local clocks are integer variables whose domain is $\{0, \dots, \alpha - 1\}$, where $\alpha \geq 2$ is called the *period*. Each process should regularly increment its clock (modulo α) while fulfilling some safety requirements. The safety of strong unison imposes that at most two consecutive clock values exist in any configuration of the system. Weak unison only requires that the difference between clocks of every two neighbors is at most one increment. Finally, we defined partial weak unison as a property dedicated to dynamic systems. It only enforces the difference between clocks of neighboring processes present before the dynamic steps to remain at most one increment.

We propose a self-stabilizing strong unison algorithm which works with any period $\alpha > 4$ in any anonymous connected network. It assumes the knowledge of two values μ and β , where μ is any upper bound on n , and β should divide α and be greater than μ^2 . Our algorithm is designed in the locally shared memory model and assume the distributed unfair daemon, the most general daemon of the model. Its stabilization time is at most $n + (\mu + 1)\mathcal{D} + 1$ rounds, where n (resp. \mathcal{D}) is the size (resp. diameter) of the network.

We then slightly modify this algorithm to make it gradually stabilizing assuming at

²*N.b.*, a dynamic step is a step containing topological changes.

most one dynamic step. In particular, the parameter μ should now be greater than or equal to $n + \#J$, where $\#J$ is an upper bound on the number of processes that join the system during a dynamic step. Notice that these slight modifications lead to increase the stabilization time by one round. This new version is gradually stabilizing because after one dynamic step from a configuration which is legitimate for the strong unison, it immediately satisfies the specification of partial weak unison, then converges to the specification of weak unison in at most one round, and finally retrieves, after at most $(\mu + 1)\mathcal{D}_1 + 1$ additional rounds (where \mathcal{D}_1 is the diameter of the network after the dynamic step), a configuration (1) from which the specification of strong unison is satisfied, and (2) where it is ready to achieve gradual convergence again in case of another dynamic step. The dynamic step may contain several topological events, *i.e.*, link and/or process additions and/or removals. However, we require that, after those topological changes, the network should stay connected and if $\alpha > 4$, every process which joins the system should be linked to at least one process already in the system before the dynamic step. We show that this condition, called **UnderLocalControl**, is necessary to have the gradual convergence. However, the algorithm being self-stabilizing (by definition), it still converges to a legitimate configuration of the strong unison after the system suffers from arbitrary other kinds of transient fault including, for example, several dynamic steps that do not satisfy the **UnderLocalControl** condition. Now, in such cases, there is no safety guarantees during the stabilization phase.

Related Work. Gradual stabilization is related to two other stronger forms of self-stabilization, namely, *safe-converging self-stabilization* [19] and *superstabilization* [11]. The goal of a safely converging self-stabilizing algorithm is to first quickly (within $O(1)$ rounds is the usual rule) converge from an arbitrary configuration to a *feasible* legitimate configuration, where a minimum quality of service is guaranteed. Once such a feasible legitimate configuration is reached, the system continues to converge to an *optimal* legitimate configuration, where more stringent conditions are required. Hence, the aim of safe-converging self-stabilization is also to ensure a gradual convergence, but only for two specifications. However, this gradual convergence is stronger than ours as it should be ensured after any step of transient faults,³ while the gradual convergence of our property applies after dynamic steps only. Safe convergence is especially interesting for self-stabilizing algorithms that compute optimized data structures, *e.g.*, minimal dominating sets [19], approximately minimum weakly connected dominating sets [21], approximately minimum connected dominating sets [20, 22], and minimal (f, g) -alliances [7]. However, to the best of our knowledge, no safe-converging algorithm for non-static problems, such

³Such transient faults may include topological changes, but not only.

as unison for example, has been proposed until now.

In superstabilization, like in our approach, fast convergence and the passage predicate should be ensured only if the system was in a legitimate configuration before the topological change occurs. In contrast with our approach, superstabilization ensures fast convergence to the original specification. However, this strong property only considers one dynamic step consisting in only one topological event: the addition or removal of one link or process in the network. Again, superstabilization has been especially studied in the context of static problems, *e.g.*, spanning tree construction [11, 4, 3], and coloring [11]. However, notice that there exist few superstabilizing algorithms for non-static problems, such as mutual exclusion [16, 23].

We use the general term *unison* to name several close problems also known in the literature as *phase or barrier synchronization* problems. There exists many self-stabilizing algorithms for the strong as well as weak unison problem, *e.g.*, [15, 1, 17, 27, 18, 6, 28]. However, to the best of our knowledge, until now, no self-stabilizing solution for such problems addresses specific convergence properties in case of topological changes (in particular, no superstabilizing ones). Self-stabilizing strong unison was first considered in synchronous anonymous networks. Particular topologies were considered in [17] (rings) and [27] (trees). Gouda and Herman [15] proposed a self-stabilizing algorithm for strong unison working in anonymous synchronous systems of arbitrary connected topology. However, they considered unbounded clocks. A solution working with the same settings, yet implementing bounded clocks, is proposed in [1]. In [28], an asynchronous self-stabilizing strong unison algorithm is proposed for arbitrary connected rooted networks.

Johnen *et al.* investigated asynchronous self-stabilizing weak unison in oriented trees in [18]. The first self-stabilizing asynchronous weak unison for general graphs was proposed by Couvreur *et al.* [8]. However, no complexity analysis was given. Another solution which stabilizes in $O(n)$ rounds has been proposed by Boulinier *et al.* in [6]. Finally, Boulinier proposed in his PhD thesis a parametric solution which generalizes both the solutions of [8] and [6]. In particular, the complexity analysis of this latter algorithm reveals an upper bound in $O(\mathcal{D}.n)$ rounds on the stabilization time of the Couvreur *et al.*' algorithm.

Roadmap. The rest of the paper is organized as follows. In the next section, we define the computational model used in this paper. In Section 3, we recall the formal definition of self-stabilization, and introduce the notion of gradual stabilization. The three variants of the unison problem considered in this paper are defined in Section 4. In Section 5, we show that condition `UnderLocalControl` is necessary to obtain our gradually stabilizing solution. We present our self-stabilizing strong unison algorithm in Section 6.

The gradually stabilizing variant of this latter algorithm is proposed in Section 7. We make concluding remarks in Section 8. Some useful results from [5] are recalled in the appendix.

2 Preliminaries

We consider the *locally shared memory model* introduced by Dijkstra [9] enriched with the notion of topological changes. Thereupon, we follow an approach similar to the one used by Dolev in the context of superstabilization [10].

2.1 Processes

We consider distributed systems made of *anonymous* processes. The system *initially contains* $n > 0$ processes and its topology is connected, however it may suffer from topological changes along the time. Each process p can directly communicate with a subset $p.\mathcal{N}$ of other processes, called its *neighbors*. In our context, $p.\mathcal{N}$ can vary over time. Communications are assumed to be *bidirectional*, *i.e.*, for any two processes p and q , $q \in p.\mathcal{N} \Leftrightarrow p \in q.\mathcal{N}$ at any time. Communications are carried out by a finite set of locally shared variables at each process: each process can read its own variables and those of its (current) neighbors, but can only write into its own variables. The *state* of a process is the vector of values of its variables. We denote by \mathcal{S} the set of all possible states of a process.

Each process updates its variables according to a *local algorithm*. The collection of all local algorithms defines a *distributed algorithm*. In the distributed algorithm \mathcal{A} , the local algorithm of p consists of a finite set of *actions* of the following form:

$$\langle \text{label} \rangle :: \langle \text{guard} \rangle \rightarrow \langle \text{statement} \rangle$$

The *labels* are used to identify actions in the reasoning. The *guard* of an action is a Boolean predicate involving variables of p and its neighbors. The *statement* is a sequence of assignments on variables of p . If the guard of some action evaluates to true, then the action is said to be *enabled* at p . By extension, if at least one action is enabled at p , p is said to be enabled. An action can be executed only if it is enabled. In this case, the execution of the action consists in executing its statement, atomically.

A *configuration* γ_i of the system is a pair $(G_i, V_i \rightarrow \mathcal{S})$. $G_i = (V_i, E_i)$ is a simple undirected graph which represents the topology of the network in configuration γ_i , *i.e.*, V_i is the set of processes that are in the system in γ_i and $E_i \subseteq V_i \times V_i$ represents the communication links between processes of V_i in γ_i : $\forall p, q \in V_i, \{p, q\} \in E_i \Leftrightarrow p \in q.\mathcal{N}$

in γ_i . $V_i \rightarrow \mathcal{S}$ is a function which associates a state to any process of V_i . For sake of simplicity, we denote by $\gamma_i(p)$ the state of process $p \in V_i$ in configuration γ_i . Moreover, $\gamma_i(p).x$ denotes the value of the x -variable at process p in configuration γ_i . We denote by \mathcal{C} the set of all possible configurations.

2.2 Executions

The dynamicity and asynchronism of the system are materialized by an adversary, called *daemon*. To perform a *step* from a configuration γ_i , the daemon can

- activate some processes (of V_i) that are enabled in γ_i — each activated process executes one of its enabled action according to its own state and that of its neighbors in γ_i , and/or
- modify the topology.

Activation of enabled processes and/or topology modifications are done atomically, leading to a new configuration γ_{i+1} . The set of all possible steps induces a binary relation over configurations noted $\mapsto \subseteq \mathcal{C} \times \mathcal{C}$, in which we exclude empty steps of the form $\gamma_i \mapsto \gamma_i$.

We distinguish two kinds of steps, *i.e.*, \mapsto is partitioned into \mapsto_s and \mapsto_d . Relation \mapsto_s define all possible *static steps*, *i.e.*, all steps consisting in activations of enabled processes only. Relation \mapsto_d define all possible *dynamic steps*, *i.e.*, all steps containing topological changes (and possibly some process activations).

An *execution* is any sequence of configurations $\gamma_0, \gamma_1, \dots$ such that G_0 is connected and $\forall i \geq 0, \gamma_i \mapsto \gamma_{i+1}$. For sake of simplicity, we note $G_0 = G = (V, E)$; we also note \mathcal{D} the diameter of G and we recall that $|V| = n$. Moreover, we define \mathcal{E}^τ the set of maximal executions which contain at most τ dynamic steps. The set of all possible executions is therefore equal to $\mathcal{E} = \cup_{\tau \geq 0} \mathcal{E}^\tau$. Notice that $\forall i, j \in \mathbb{N}, i \leq j$ implies $\mathcal{E}^i \subseteq \mathcal{E}^j$. For any subset of configurations $X \subseteq \mathcal{C}$, we denote by \mathcal{E}_X^τ the set of all executions in \mathcal{E}^τ that start from a configuration of X .

2.3 Static Steps

Let γ_i be a configuration. The daemon can choose to make a static step from γ_i only if there exists at least one enabled process. In this case, it first selects a non-empty subset S of enabled processes. Next, every process $p \in S$ *atomically* executes one of its enabled actions, leading the system to a new configuration, say γ_{i+1} . In this case, $\gamma_i \mapsto_s \gamma_{i+1}$ with, in particular, $G_i = G_{i+1}$.

2.4 Dynamic Steps

Let $\gamma_i \mapsto_d \gamma_{i+1}$ be a dynamic step. We have in particular that $G_{i+1} \neq G_i$. Precisely, the step $\gamma_i \mapsto_d \gamma_{i+1}$ contains a finite number of topological events and maybe some process activations (like in static steps). Each topological event is of the following types.

A process p can *join* the system. This event is denoted by $join_p$ and triggers the atomic execution of a specific action, called *bootstrap*, which initializes the variables of p to a particular state, called *bootstate*. This bootstrap is executed without any communication. We denote by New_k the set of processes which are in bootstate in γ_k . When p joins the system in $\gamma_i \mapsto_d \gamma_{i+1}$, we have $p \in New_{i+1}$, but $p \notin New_i$. Moreover, until p executes its very first action, say in step $\gamma_x \mapsto \gamma_{x+1}$, it is still in bootstate. Hence $\forall j \in \{i+1, \dots, x\}, p \in New_j$, but $p \notin New_{x+1}$. We assume that there are at most $\#J$ joins in the system during a dynamic step.

A process can also *leave* the system. Finally, some communication links can *appear* or *disappear* between two different processes.

Several joins, leaves, as well as link appearances and disappearances can be made in the same step $\gamma_i \mapsto_d \gamma_{i+1}$.

2.5 Daemon

As previously explained, executions are driven by a daemon. We assume the daemon is *distributed* and *unfair*. In a static step, a distributed daemon must select at least one enabled process (maybe more). In a dynamic step, a distributed daemon can select 0, 1, or several enabled processes. An unfair daemon has no fairness constraint, *i.e.*, it might never select a process p during any step unless in the case of a static step from a configuration where p is the only enabled process. Moreover, at each configuration, an unfair daemon freely chooses between making a static or dynamic step, except if no more process is enabled; in this latter case, only a dynamic step containing no process activation can be chosen.

2.6 Functional Specification and Performances

A distributed algorithm \mathcal{A} is designed to ensure some functional properties called its *specification*. A specification SP is a predicate over \mathcal{E} .

We measure the time complexity of our algorithms in terms of *rounds* [12]. The first round of an execution $e = (\gamma_i)_{i \geq 0}$ is the minimal prefix e' of e such that every enabled process in γ_0 either executes an action or is *neutralized* (defined below). Let γ_j be the last configuration of e' , the second round of e is the first round of $e'' = (\gamma_i)_{i \geq j}$, and so forth.

Neutralized means that a process p is enabled in a configuration γ_i but either p is no more in the system in the next configuration γ_{i+1} , or p is not enabled in γ_{i+1} but does not execute any action during the step $\gamma_i \mapsto \gamma_{i+1}$.

3 Stabilization

3.1 Self-stabilization

Below we recall the definitions of some notions classically used in self-stabilization. Notice that all these notions are defined by only considering executions free of topological changes, yet starting from an arbitrary configuration. Indeed, self-stabilization considers the system immediately after the transient faults cease. So, the system is initially observed from an arbitrary configuration reached due to occurrence of transient faults (including some topological changes maybe), but from which no faults (in particular, no topological changes) will ever occur.

Let \mathcal{A} be a distributed algorithm. Let $X, Y \subseteq \mathcal{C}$ be two subsets of configurations. X is *closed* under \mathcal{A} if and only if $\forall \gamma, \gamma' \in \mathcal{C}, (\gamma \in X \wedge \gamma \mapsto_s \gamma') \Rightarrow \gamma' \in X$. Y *converges to X* under \mathcal{A} if and only if $\forall e \in \mathcal{E}_Y^0, \exists \gamma \in e$ such that $\gamma \in X$. \mathcal{A} *stabilizes from Y to a specification SP by X* if and only if

- X is closed under \mathcal{A} ,
- Y converges to X under \mathcal{A} ,
- and $\forall e \in \mathcal{E}_X^0, SP(e)$.

Moreover, the *convergence time in steps (resp. rounds) from Y to X* is the maximal number of steps (or rounds, respectively) to reach a configuration of X in over every execution of \mathcal{E}_Y^0 .

Self-stabilization has been defined by Dijkstra in 1974 [9] as follows: a distributed algorithm \mathcal{A} is *self-stabilizing* for a specification SP if and only if $\exists \mathcal{L} \subseteq \mathcal{C}$, \mathcal{A} stabilizes from \mathcal{C} to SP by \mathcal{L} .

\mathcal{L} (resp. $\mathcal{C} \setminus \mathcal{L}$) is then said to be a set of *legitimate configurations* (resp. *illegitimate configurations*) *w.r.t.* SP . The *stabilization time* of \mathcal{A} is then the convergence time from \mathcal{C} to \mathcal{L} .

3.2 Gradual Stabilization under τ -Dynamics

We now introduce a specialization of self-stabilization called *gradual stabilization under τ -dynamics* which requires that after at most τ dynamic steps from a legitimate config-

uration, the system re-stabilizes gradually to stronger and stronger specifications, until reaching a configuration (1) from which its initial (strong) specification is satisfied again, and (2) where it is ready to achieve gradual convergence again in case of up to τ new dynamic steps. Of course, the gradual stabilization makes sense only if the convergence to every intermediate weaker specification is fast.

Let $\tau \geq 0$. For a given execution $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}^\tau$, let $first(e)$ be the integer such that $\gamma_{first(e)}$ is the first configuration of e after the last topological change. Formally, $first(e)$ is the minimal index such that the suffix of e starting from $first(e)$ contains no dynamic step: $first(e) = \min\{i : (\gamma_j)_{j \geq i} \in \mathcal{E}^0\}$. For any subset of executions $E \subseteq \mathcal{E}^\tau$, let $FC(E) = \{\gamma_{first(e)} : e = (\gamma_i)_{i \geq 0} \in E\}$ be the set of all configurations that can be reached after the last topological changes in executions of E ($FC()$ stands for “First Configuration”).

Let SP_1, SP_2, \dots, SP_k , be an ordered sequence of specifications. Let B_1, B_2, \dots, B_k be (asymptotic) complexity bounds such that $B_1 < B_2 < \dots < B_k$.

A distributed algorithm \mathcal{A} is *gradually stabilizing under τ -dynamics* for $(SP_1 \bullet B_1, SP_2 \bullet B_2, \dots, SP_k \bullet B_k)$ if and only if $\exists \mathcal{L}_1, \dots, \mathcal{L}_k \subseteq \mathcal{C}$ such that

1. \mathcal{A} stabilizes from \mathcal{C} to SP_k by \mathcal{L}_k , i.e., \mathcal{A} is self-stabilizing for SP_k .
2. Starting from a legitimate configuration, after at most τ dynamic steps, \mathcal{A} gradually converges to every \mathcal{L}_i with $i \in \{1, \dots, k\}$, i.e., $\forall i \in \{1, \dots, k\}$, we have
 - \mathcal{A} stabilizes from $FC(\mathcal{E}_{\mathcal{L}_k}^\tau)$ to SP_i by \mathcal{L}_i , and
 - the convergence time in rounds from $FC(\mathcal{E}_{\mathcal{L}_k}^\tau)$ to \mathcal{L}_i is bounded by B_i .

Notice that, by definition, any gradually stabilizing algorithm is also a self-stabilizing algorithm for SP_k . Hence, the performances of any gradually stabilizing algorithm can be also evaluated at the light of its stabilization time.

Gradual stabilization is related to two other stronger forms of self-stabilization: *safe-converging self-stabilization* [19] and *superstabilization* [11].

The goal of a *safely converging self-stabilizing algorithm* is to first quickly (within $O(1)$ rounds is the usual rule) converge from any arbitrary configuration to a *feasible* legitimate configuration, where a minimum quality of service is guaranteed. Once such a feasible legitimate configuration is reached, the system continues to converge to an *optimal* legitimate configuration, where more stringent conditions are required. Hence, the aim of safe-converging self-stabilization is also to ensure a gradual convergence, but for two specifications. However, this gradual convergence should be ensured after any step of transient faults (such transient faults can include topological changes, but not only), while the gradual convergence of our property applies after dynamic steps only.

A *superstabilizing algorithm* is self-stabilizing and has two additional properties. In presence of a single topological change (adding or removing one link or process in the network), it recovers fast (typically $O(1)$), and a safety predicate, called a *passage* predicate, should be satisfied along the stabilization phase. Like in our approach, fast convergence, captured by the notion of *superstabilization time*, and the passage predicate should be ensured only if the system was in a legitimate configuration before the topological change occurs. In contrast with our approach, superstabilization only considers one dynamic step consisting in only one topological event: the addition or removal of one link or process in the network. A superstabilizing algorithm for a specification SP_1 can be seen as an algorithm which is gradually stabilizing under 1-dynamics for $(SP_0 \bullet 0, SP_1 \bullet f)$ where SP_0 is the passage predicate, f is the superstabilization time and the dynamic step consists of adding or removing one link or process in the network only.

Notice also that the bound B_k captures a complexity similar to the fault gap in *fault-containing algorithms* [14]: let $P1$ be a period of at most τ dynamic steps such that $P1$ starts in a legitimate configuration; B_k represents the necessary fault-free interval after $P1$ and before the next period of at most τ dynamic steps $P2$ so that the system becomes ready again to achieve gradual convergence after $P2$.

4 Unison

We consider several close synchronization problems included here under the general term of unison. In these problems, each process should maintain a local *clock*. We restrict here our study to periodic clocks, *i.e.*, all local clocks are integer variables whose domain is $\{0, \dots, \alpha - 1\}$, where $\alpha \geq 2$ is called the *period*. Each process should regularly increment its clock (modulo α) while fulfilling some safety requirements. Below we define three versions of the problem respectively named *strong*, *weak*, and *partial weak* unison.

Strong unison defined below is also known as the *phase* or *barrier* synchronization problem [26, 24]. In the definition, we use the following notation: for every configuration $\gamma_i = (G_i, V_i \rightarrow \mathcal{S})$, let $CV(\gamma_i) = \{\gamma_i(p).clock, p \in V_i\}$ be the set of clock values present in configuration γ_i .

Specification 1 (Strong Unison). *An execution $e = (\gamma_i)_{i \geq 0}$ satisfies the specification SP_{SU} of strong unison if and only if*

- *In any configuration $\gamma_i \in e$, there exists at most two different clock values, and if so, these two values are consecutive (modulo α), *i.e.*, $\forall \gamma_i \in e, |CV(\gamma_i)| \leq 2 \wedge (CV(\gamma_i) = \{x, y\} \Rightarrow x = (y + 1) \bmod \alpha \vee y = (x + 1) \bmod \alpha)$. (Safety)*

- *Every process increments its clock infinitely often in e , i.e., $\forall \gamma_i \in e, \forall p \in V_i, \exists j > i, (\forall k \in \{i+1 \dots j-1\}, p \in V_k \wedge \gamma_k(p).clock = \gamma_i(p).clock) \wedge (p \in V_j \wedge \gamma_j(p).clock = (\gamma_i(p).clock + 1) \bmod \alpha)$. (Liveness)*

The definition of *weak unison* below appeared first in [8] under the name of *asynchronous unison*.

Specification 2 (Weak Unison). *An execution $e = (\gamma_i)_{i \geq 0}$ satisfies the specification SP_{WU} of weak unison if and only if*

- *In any configuration $\gamma_i \in e$, the clocks of every two neighboring processes differ from at most one increment (modulo α), i.e., $\forall \gamma_i \in e, \forall p \in V_i, \forall q \in \gamma_i(p).\mathcal{N}, \gamma_i(p).clock = (\gamma_i(q).clock + 1) \bmod \alpha \vee \gamma_i(q).clock = (\gamma_i(p).clock + 1) \bmod \alpha$. (Safety)*
- *Every process increments its clock infinitely often in e , i.e., $\forall \gamma_i \in e, \forall p \in V_i, \exists j > i, (\forall k \in \{i+1 \dots j-1\}, p \in V_k \wedge \gamma_k(p).clock = \gamma_i(p).clock) \wedge (p \in V_j \wedge \gamma_j(p).clock = (\gamma_i(p).clock + 1) \bmod \alpha)$. (Liveness)*

Finally, in the context of dynamic systems, a straightforward variant of the weak unison is the following.

Specification 3 (Partial Weak Unison). *An execution $e = (\gamma_i)_{i \geq 0}$ satisfies the specification SP_{PU} of partial weak unison if and only if*

- *In any configuration $\gamma_i \in e$, the clocks of any two neighbors which are not in New_i differ from at most one increment (modulo α), i.e., $\forall \gamma_i \in e, \forall p \in V_i \setminus New_i, \forall q \in \gamma_i(p).\mathcal{N} \setminus New_i, \gamma_i(p).clock = (\gamma_i(q).clock + 1) \bmod \alpha \vee \gamma_i(q).clock = (\gamma_i(p).clock + 1) \bmod \alpha$. (Safety)*
- *Every process increments its clock infinitely often in e , i.e., $\forall \gamma_i \in e, \forall p \in V_i, \exists j > i, (\forall k \in \{i+1 \dots j-1\}, p \in V_k \wedge \gamma_k(p).clock = \gamma_i(p).clock) \wedge (p \in V_j \wedge \gamma_j(p).clock = (\gamma_i(p).clock + 1) \bmod \alpha)$. (Liveness)*

The property below sum up the straightforward relationship between the three variants of unison.

Property 1. $SP_{SU} \Rightarrow SP_{WU} \Rightarrow SP_{PU}$.

5 Necessary Condition

Through out this section, we assume the existence of a deterministic algorithm \mathcal{A} which is gradually stabilizing under 1-dynamics for $(SP_{\text{PU}} \bullet 0, SP_{\text{WU}} \bullet 1, SP_{\text{SU}} \bullet B)$ in any arbitrary anonymous network under the distributed unfair daemon, with $B > 1$ be any (asymptotic) complexity bound. Let $\mathcal{L}_{\text{SU}}^{\mathcal{A}}$ be the legitimate configurations of \mathcal{A} *w.r.t.* specification SP_{SU} .

The property given below states that, when $\alpha > 3$ and once a legitimate configuration of strong unison is reached, the system necessarily goes through a configuration where all clocks have the same value between any two increments at the same process.

Property 2. *Assume $\alpha > 3$. For every $(\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\text{SU}}^{\mathcal{A}}}^0$, for every process p , for every $k \in \{0, \dots, \alpha - 1\}$, for every $i \geq 0$, if p increments its clock from k to $(k + 1) \bmod \alpha$ in $\gamma_i \mapsto \gamma_{i+1}$ and $\exists j > i$ such that $\gamma_j(p).clock = (k + 2) \bmod \alpha$, then there exists $x \in \{i + 1, \dots, j - 1\}$, such that all clocks have value $(k + 1) \bmod \alpha$ in γ_x .*

Proof. Let $(\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\text{SU}}^{\mathcal{A}}}^0$ be an execution and p be a process. Let $k \in \{0, \dots, \alpha - 1\}$ and $i \geq 0$ such that p increments its clock from k to $(k + 1) \bmod \alpha$ in $\gamma_i \mapsto \gamma_{i+1}$ and $\exists j > i$ such that $\gamma_j(p).clock = (k + 2) \bmod \alpha$.

Assume that there is a process q such that $\gamma_i(q).clock = (k - 1) \bmod \alpha$. As the execution satisfies SP_{SU} , there exists a step after γ_i in which p increments, due to liveness; but due to safety, q necessarily increments at the same step. Now, because of the distributed unfair daemon, we can now build a possible step where p moves, but not q leading to a configuration where $q.clock = (k - 1) \bmod \alpha$ and $p.clock = (k + 1) \bmod \alpha$. Hence, there exists an execution starting from a configuration of $\mathcal{L}_{\text{SU}}^{\mathcal{A}}$ which does not satisfy SP_{SU} , a contradiction.

Hence, $\forall q \in V, \gamma_i(q).clock \in \{k, (k + 1) \bmod \alpha\}$, by the safety of SP_{SU} . Similarly to the previous case, while there are processes whose clock value is k , no process (in particular p) can increment its clock from $(k + 1) \bmod \alpha$ to $(k + 2) \bmod \alpha$. Hence, between γ_i and γ_j there exists a configuration where all processes have clock value $(k + 1) \bmod \alpha$. \square

In the following, we will establish that the property `UnderLocalControl` given below is a necessary condition for \mathcal{A} . The definition of `UnderLocalControl` uses the notion of *dominating set*: a *dominating set* of the graph $G = (V, E)$ is any subset D of V such that every node not in D is adjacent to at least one member of D . `UnderLocalControl` captures a condition on the network dynamics which is necessary to prevent a notable desynchronization of clocks. Namely, the network should stay connected and if $\alpha > 4$, every process that joins during the dynamic step $\gamma \mapsto_d \gamma'$ should be “under the control of” (that is, linked to) at least one process which exists in both γ and γ' .




Figure 1: Execution e'' in the proof of Theorem 1. The hachured nodes are in bootstate.

Definition 1 (UnderLocalControl). *UnderLocalControl holds if and only if for every execution $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\text{SU}}^A}^1$,*

1. $G_{\text{first}(e)}$ is connected, and
2. if $\alpha > 4$, then $V_{\text{first}(e)} \setminus \text{New}_{\text{first}(e)}$ is a dominating set.

Lemma 1. *For every execution $e \in \mathcal{E}_{\mathcal{L}_{\text{SU}}^A}^1$, $G_{\text{first}(e)}$ is connected.*

Proof. Assume, by the contradiction, that there is an execution $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\text{SU}}^A}^1$ such that $G_{\text{first}(e)}$ is disconnected. Let A and B be two connected components of $G_{\text{first}(e)}$. By definition, there exists $j \geq \text{first}(e)$ such that $\gamma_j \in \mathcal{L}_{\text{SU}}^A$ and A and B are defined in all configurations $(\gamma_i)_{i \geq j}$. From γ_j , all processes regularly increment their clocks in both A and B by the liveness property of strong unison. Now, as no process of B is linked to any process of A , the behavior of processes in B has no impact on processes in A and vice versa. So, liveness implies, in particular, that there always exists enabled processes in A . Consequently, there exists a possible execution of $\mathcal{E}_{\mathcal{L}_{\text{SU}}^A}^1$ prefixed by $\gamma_0 \dots \gamma_j$ where the distributed unfair daemon only selects processes in A from γ_j , hence violating the liveness property of strong unison, a contradiction. \square

Lemma 2. *If $\alpha > 4$, then for every execution $e \in \mathcal{E}_{\mathcal{L}_{\text{SU}}^A}^1$, $V_{\text{first}(e)} \setminus \text{New}_{\text{first}(e)}$ is a dominating set.*

Proof. We illustrate the following proof with Figure 1. Let $e \in (\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\text{SU}}^A}^1$. Let $x = \text{first}(e)$. Assume, by the contradiction, that $\alpha > 4$ and G_x is connected, but

$V_x \setminus New_x$ is not a dominating set. This implies that $\exists p \in New_x$ such that $\forall q \in \gamma_x(p).\mathcal{N}$, $q \in New_x$.

First, notice that every process among p and its neighbors are enabled in γ_x to take a clock value in $\{0, \dots, \alpha - 1\}$. Indeed, assume that the daemon makes a synchronous step from γ_x , then the step $\gamma_x \mapsto_s \gamma_{x+1}$ actually corresponds to a complete round (by definition) and so γ_{x+1} should be a legitimate configuration of weak unison. Let c be the clock value taken by p if p moves in $\gamma_x \mapsto_s \gamma_{x+1}$.

Consider now another execution e' in $\mathcal{E}_{\mathcal{L}_{SU}^A}^0$ (with no topological change) on a graph of at least two nodes which contains neither p nor its neighbors $\gamma_x(p).\mathcal{N}$. Strong unison is satisfied in e' and, as $\alpha > 4$, by Property 2, there is a configuration γ_S in e' where every clock equals $(c + 2) \bmod \alpha$. From γ_S , there is eventually a step in which at least one process increments its clock to $(c + 3) \bmod \alpha$. Assume not all processes are activated by the distributed unfair daemon during this step. Then, this step leads to a configuration γ_T where there is exactly two values of clock: $(c + 2) \bmod \alpha$ and $(c + 3) \bmod \alpha$, see Figure 1a.

Consider now another execution e'' having a prefix common to e' until γ_T . Assume that the daemon introduces a dynamic step at configuration γ_T . Assume that this step contains no process activation, but consists in adding p with the same neighborhood as well as two links from q , a neighbor of p , and two already existing nodes r_1 and r_2 , such that the clock of r_1 (resp. r_2) equals $(c + 2) \bmod \alpha$ (resp. $(c + 3) \bmod \alpha$) in γ_T , see Figure 1b. By definition, since strong unison is satisfied in γ_T (by assumption), the partial weak unison necessarily holds all along the suffix of e'' starting at γ_{T+1} .

Process p and its neighbors are in a situation similar to the one in γ_x , so they are enabled to take a clock value in $\{0, \dots, \alpha - 1\}$, in particular p is enabled to take value c . Assume that the daemon exactly selects p and its neighbors in the next step $\gamma_{T+1} \mapsto \gamma_{T+2}$. In γ_{T+2} (Figure 1c), the clock of r_1 and r_2 are respectively equal to $(c + 2) \bmod \alpha$ and $(c + 3) \bmod \alpha$, since they did not move; moreover, the clock of p is equal to c . Now, q also chooses a clock value in $\gamma_{T+1} \mapsto_s \gamma_{T+2}$ and that clock value should differ of at most one increment (mod α) from the clocks of p , r_1 , and r_2 since partial weak unison holds in γ_{T+1} and all subsequent configurations. As $\alpha > 3$, if the clock of q equals:

- c or $(c + 1) \bmod \alpha$, the difference between the clocks of q and r_2 is at least 2 increments (mod α),
- $(c + 2) \bmod \alpha$ or $(c + 3) \bmod \alpha$, the difference between the clocks of q and p is at least 2 increments (mod α),
- any value in $\{0, \dots, \alpha - 1\} \setminus \{c, (c + 1) \bmod \alpha, (c + 2) \bmod \alpha, (c + 3) \bmod \alpha\}$, the difference between the clocks of q and r_1 is at least 2 increments (mod α).

Hence, the safety of partial weak unison is necessarily violated in the configuration γ_{T+2} of e'' , a contradiction. \square

By Lemmas 1 and 2, follows:

Theorem 1. *An algorithm \mathcal{A} is gradually stabilizing under 1-dynamics for $(SP_{PU} \bullet 0, SP_{WU} \bullet 1, SP_{SU} \bullet B)$ in arbitrary anonymous networks under the distributed unfair daemon with a set of legitimate configurations w.r.t. specification SP_{SU} noted \mathcal{L}_{SU}^A only if *UnderLocalControl* holds.*

6 Self-Stabilizing Strong Unison

In this section, we propose an algorithm which is self-stabilizing for the strong unison problem in any arbitrary connected anonymous network. This algorithm works for any period $\alpha > 4$ and is based on an algorithm previously proposed by Boulinier in [5], this latter is self-stabilizing for the weak unison problem and works for any period $\beta > n^2$, where n is the number of processes. We first recall the algorithm of Boulinier, called here Algorithm \mathcal{WU} , in Subsection 6.1. Notice that the notation used in this algorithm will be also applicable to our algorithms. We give our self-stabilizing algorithm for the strong unison, Algorithm \mathcal{SU} , and its proof of correctness in Subsection 6.2.

6.1 Algorithm \mathcal{WU}

Algorithm \mathcal{WU} , see Algorithm 1 for its formal code, has been proposed by Boulinier in his PhD thesis [5]. Actually, it is a generalization of the self-stabilizing weak unison algorithm proposed by Couvreur *et al.* [8]. This algorithm being simply self-stabilizing, it only considers executions without any topological change, yet starting from arbitrary configuration. The topology of the network then consists in a connected graph $G = (V, E)$ of n nodes which is fixed all along the execution. Remind that \mathcal{D} is the diameter of G .

In Algorithm \mathcal{WU} , each process p is endowed with a clock variable $p.t \in \{0, \dots, \beta - 1\}$, where β is its period. β should be greater than n^2 . The algorithm also uses another constant, noted μ , which should satisfy $n \leq \mu \leq \frac{\beta}{2}$. The algorithm uses the following notations.

Notations. We define the *delay* between two integer values x and y by the function $d_\beta(x, y) = \min((x - y) \bmod \beta, (y - x) \bmod \beta)$. Then, let $\preceq_{\beta, \mu}$ be the relation such that for every two integer values x and y , $x \preceq_{\beta, \mu} y \equiv ((y - x) \bmod \beta) \leq \mu$.

The Algorithm. Two actions are used to maintain the clock $p.t$ at each process p . When the delay between $p.t$ and the clocks of some neighbors is greater than one, but the maximum delay is not too big (that is, does not exceed μ), then it is possible to “normally” converge, using Action $\mathcal{WU-N}$, to a configuration where the delay between those clocks is at most one by making increment the clocks of the most behind processes among p and its neighbors. Moreover, once legitimacy is achieved, p can “normally” increment its clock still using Action $\mathcal{WU-N}$ when it is on time or one increment late with all its neighbors. In contrast, if the delay is too big (that is, the delay between the clocks of p and one of its neighbors is more than μ) and the clock of p is not yet reset, then p should reset its clock to 0 using Action $\mathcal{WU-R}$.

Algorithm 1 \mathcal{WU} , for every process p

Parameters:

β : any positive integer such that $\beta > n^2$
 μ : any positive integer such that $n \leq \mu \leq \frac{\beta}{2}$

Variable:

$p.t \in \{0, \dots, \beta - 1\}$

Actions:

$\mathcal{WU-N} \quad :: \quad \forall q \in p.\mathcal{N}, p.t \preceq_{\beta, \mu} q.t \quad \rightarrow \quad p.t \leftarrow (p.t + 1) \bmod \beta$
 $\mathcal{WU-R} \quad :: \quad \exists q \in p.\mathcal{N}, d_{\beta}(p.t, q.t) > \mu \wedge p.t \neq 0 \quad \rightarrow \quad p.t \leftarrow 0$

From [5], we have the following theorem.

Theorem 2. *Algorithm \mathcal{WU} is self-stabilizing for $SP_{\mathcal{WU}}$ (specification of weak unison) by the set of legitimate configurations*

$$\mathcal{L}_{\mathcal{WU}} = \{\gamma \in \mathcal{C} : \forall p \in V, \forall q \in \gamma(p).\mathcal{N}, d_{\beta}(\gamma(p).t, \gamma(q).t) \leq 1\}$$

in an arbitrary connected network assuming a distributed unfair daemon.

Its stabilization time is at most $n + \mu\mathcal{D}$ rounds, where n (resp. \mathcal{D}) is the size (resp. diameter) of the network and μ is a parameter satisfying $n \leq \mu \leq \frac{\beta}{2}$.

By definition, $\mathcal{D} < n$, consequently we have:

Remark 1. *Once Algorithm \mathcal{WU} has stabilized, the delay between t -clocks of any two arbitrary far processes is at most $n - 1$, the size of the network.*

Complexity Analysis. Let C_{μ} be the set of configurations where the delay between two neighboring clocks is at most μ . Below, we prove in Lemma 3 (resp. Lemma 4) a bound on the time required to ensure that all t -variables have incremented k times which holds since the system has reached a configuration of C_{μ} (resp. $\mathcal{L}_{\mathcal{WU}}$).

Lemma 3. $\forall k \geq 1, \forall e \in \mathcal{E}_{C_\mu}^0$, every process p increments $p.t$ executing \mathcal{WU} - N at least k times every $\mu\mathcal{D} + k$ rounds, where \mathcal{D} is the diameter of the network.

Proof. Let $k \geq 1$. Let $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}_{C_\mu}^0$. Using Lemma 30, $\forall i \geq 0$, there is a function f on processes such that $\forall p \in V, f(\gamma_i, p) \bmod \beta = \gamma_i(p).t$ and $\forall p \in V, \forall q \in p.\mathcal{N}, d_\beta(\gamma_i(p).t, \gamma_i(q).t) \leq \mu$. Hence, $\forall p, q \in V, |f(\gamma_i, p) - f(\gamma_i, q)| \leq \mu\mathcal{D}$.

For every $i \geq 0$, we note $f_{\gamma_i}^{\min} = \min\{f(\gamma_i, x) : x \in V\}$. Action \mathcal{WU} - N is enabled in γ_i at every process $x \in V$ for which $\gamma_i(x).t = f(\gamma_i, x) = f_{\gamma_i}^{\min}$. So, after one round, every such a process x has incremented its t -variable (executing action \mathcal{WU} - N) at least once. Let γ_j be the first configuration after one round. Then, $f_{\gamma_j}^{\min} \geq f_{\gamma_i}^{\min} + 1$. We now consider γ_d to be the first configuration after $\mu\mathcal{D} + k$ rounds, starting from γ_i . Using the same arguments as for γ_j inductively, we have $f_{\gamma_d}^{\min} \geq f_{\gamma_i}^{\min} + \mu\mathcal{D} + k$ (*).

Let p be a process in V . By definitions of f and $f_{\gamma_i}^{\min}$, we have that $f_{\gamma_i}^{\min} \leq f(\gamma_i, p) \leq f_{\gamma_i}^{\min} + \mu\mathcal{D}$ (**). Assume now that p increments $\#incr < k$ times $p.t$ between γ_i and γ_d . Then

$$\begin{aligned} f(\gamma_d, p) = f(\gamma_i, p) + \#incr &< f(\gamma_i, p) + k \text{ (assumption on } \#incr) \\ &\leq f_{\gamma_i}^{\min} + \mu\mathcal{D} + k, \text{ by (**)} \\ &\leq f_{\gamma_d}^{\min}, \text{ by (*)} \end{aligned}$$

So, p satisfies $f(\gamma_d, p) < f_{\gamma_d}^{\min}$, a contradiction. \square

Lemma 4. $\forall k \geq 1, \forall e \in \mathcal{E}_{\mathcal{L}_{\mathcal{WU}}}^0$, every process p increments its clock $p.t$ executing action \mathcal{WU} - N at least k times every $\mathcal{D} + k$ rounds, where \mathcal{D} is the diameter of the network.

Proof. The proof of this lemma is exactly the same as the one of Lemma 3, yet replacing C_μ by $\mathcal{L}_{\mathcal{WU}}$ and $\mu\mathcal{D}$ by \mathcal{D} . \square

Some other useful results from [5] about Algorithm \mathcal{WU} are recalled in Appendix A.

6.2 Algorithm \mathcal{SU}

In this subsection, we still assume a non-dynamic context (no topological change) and we use the notations defined in Subsection 6.1. Algorithm \mathcal{SU} is a straightforward adaptation of Algorithm \mathcal{WU} . More precisely, Algorithm \mathcal{SU} maintains two clocks at each process p . The first one, $p.t \in \{0, \dots, \beta - 1\}$, is called the *internal clock* and is maintained exactly as in Algorithm \mathcal{WU} . Then, $p.t$ is used as an internal pulse machine to increment a second, yet actual, clock of Algorithm \mathcal{SU} $p.c \in \{0, \dots, \alpha - 1\}$, also referred to as *external clock*.

6.2.1 Correctness Proof

We first define a set of legitimate configurations *w.r.t.* specification $SP_{\mathcal{SU}}$ (Definition 2). Then, we prove the closure and convergence *w.r.t.* those legitimate configurations (see Lemmas 5 and 6). Afterwards, we prove the correctness *w.r.t.* specification $SP_{\mathcal{SU}}$ in any execution starting in a legitimate configuration, namely, safety is shown in Lemma 10 and liveness is proven in Lemma 11.

Definition 2 (Legitimate Configurations of \mathcal{SU} *w.r.t.* $SP_{\mathcal{SU}}$). *A configuration γ of \mathcal{SU} is legitimate w.r.t. $SP_{\mathcal{SU}}$ if and only if*

1. $\forall p \in V, \forall q \in \gamma(p).\mathcal{N}, d_\beta(\gamma(p).t, \gamma(q).t) \leq 1.$
2. $\forall p \in V, \gamma(p).c = \left\lfloor \frac{\alpha}{\beta} \gamma(p).t \right\rfloor.$

We denote by $\mathcal{L}_{\mathcal{SU}}$ the set of legitimate configurations of \mathcal{SU} w.r.t. $SP_{\mathcal{SU}}$.

By definition, $\mu \geq n > 0$, hence from Definition 2, follows.

Remark 4. *In any legitimate configuration $\gamma \in \mathcal{L}_{\mathcal{SU}}, \forall p, q \in V, d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu.$*

Lemma 5 (Closure). *$\mathcal{L}_{\mathcal{SU}}$ is closed under \mathcal{SU} .*

Proof. First, from Theorem 2, note that the set of legitimate configurations defined for \mathcal{WU} is also closed for \mathcal{SU} . Hence we only have to check closure for the second constraint of Definition 2, the one on c -variables .

Let $\gamma \in \mathcal{L}_{\mathcal{SU}}$ be a legitimate configuration of \mathcal{SU} and let $\gamma \mapsto_s \gamma'$ be a static step of \mathcal{SU} . Let $p \in V$. As $\gamma \in \mathcal{L}_{\mathcal{SU}}, \gamma(p).c = \left\lfloor \frac{\alpha}{\beta} \gamma(p).t \right\rfloor$. Either p does not execute any action during step $\gamma \mapsto_s \gamma'$, or p executes \mathcal{SU} -N or \mathcal{SU} -R. These two actions update $p.c$ according to the new value of $p.t$. Hence $\gamma'(p).c = \left\lfloor \frac{\alpha}{\beta} \gamma'(p).t \right\rfloor$. \square

Lemma 6 (Convergence). *\mathcal{C} (the set of all possible configurations) converges to $\mathcal{L}_{\mathcal{SU}}$ under \mathcal{SU} .*

Proof. From Theorem 2, the set of legitimate configurations for \mathcal{WU} is also reached in a finite number of steps for \mathcal{SU} . Hence, we only have to check that the second constraint (the one on c -variables) is also achievable within a finite number of steps.

Again using Theorem 2, liveness of Specification $SP_{\mathcal{WU}}$ is ensured by \mathcal{WU} and therefore by \mathcal{SU} . Hence, after stabilization, each process p updates its internal clock $p.t$ within a finite time; meanwhile $p.c$ is also updated to $\left\lfloor \frac{\alpha}{\beta} p.t \right\rfloor$. \square

Lemmas 7, 8 and 9 are technical results on the values of t - and c - variables that will be used to prove that the safety of Specification $SP_{\mathcal{SU}}$ is achieved in any execution that




Figure 2: Relationship between variables t and c .

starts from a legitimate configuration. For all these lemmas, we assume that α, β, K are positive numbers that satisfies the constraint declared on the **Parameters** section of Algorithm \mathcal{SU} , namely $\beta = K\alpha$.

Lemma 7. *Let $x \in \{0, \dots, \alpha - 1\}$ and $\xi \in \{0, \dots, \frac{\beta}{\alpha} - 1\}$. The following equality holds:*
 $\left\lfloor \frac{\alpha}{\beta} (x \frac{\beta}{\alpha} + \xi) \right\rfloor = x.$

Proof. Let $x \in \{0, \dots, \alpha - 1\}$ and $\xi \in \{0, \dots, \frac{\beta}{\alpha} - 1\}$. As $\left\lfloor \frac{\alpha}{\beta} (x \frac{\beta}{\alpha} + \xi) \right\rfloor = \left\lfloor x + \frac{\alpha}{\beta} \xi \right\rfloor = x + \left\lfloor \frac{\alpha}{\beta} \xi \right\rfloor$ and $\xi \in \{0, \dots, \frac{\beta}{\alpha} - 1\}$, we have that $0 \leq \frac{\alpha}{\beta} \xi < 1$ and therefore $\left\lfloor \frac{\alpha}{\beta} \xi \right\rfloor = 0$. \square

We apply Lemma 7 by instantiating the value of the internal clock t with $x \frac{\beta}{\alpha} + \xi$. Since the value of the external clock c is computed as $\left\lfloor \frac{\alpha}{\beta} t \right\rfloor$ in Algorithm 2, we have $c = x$. Now, if we chose β (period of internal clocks) such that it can be written as $\beta = K\alpha$ with K a positive integer, the value of $c = \left\lfloor \frac{\alpha}{\beta} t \right\rfloor$ is always a non negative integer which evolves according to $t = c \frac{\beta}{\alpha} + \xi$ as shown in Figure 2.

Lemma 8. *Let $x_1, x_2 \in \{0, \dots, \alpha - 1\}$ and $\xi_1, \xi_2 \in \{0, \dots, \frac{\beta}{\alpha} - 1\}$. The following assertion holds: $x_1 \frac{\beta}{\alpha} + \xi_1 \leq x_2 \frac{\beta}{\alpha} + \xi_2 \Rightarrow x_1 \leq x_2$*

Proof. Let $x_1, x_2 \in \{0, \dots, \alpha - 1\}$ and $\xi_1, \xi_2 \in \{0, \dots, \frac{\beta}{\alpha} - 1\}$. If $x_1 \frac{\beta}{\alpha} + \xi_1 \leq x_2 \frac{\beta}{\alpha} + \xi_2$, then we have that $x_1 - x_2 \leq (\xi_1 - \xi_2)/K$. As $\xi_1, \xi_2 \in \{0, \dots, K - 1\}$, we have $(\xi_2 - \xi_1)/K \leq 1 - 1/K$. By transitivity, we obtain that $x_1 - x_2 \leq 1 - 1/K < 1$. As x_1 and x_2 are natural integers, so is their difference; this proves that $x_1 - x_2 \leq 0$ \square

Again, Lemma 8 will be used with the internal clock $t = c \frac{\beta}{\alpha} + \xi$: this establishes the monotonic relation between internal and external clocks.

Lemma 9. *Let $t_1, t_2 \in \{0, \dots, \beta - 1\}$. The following assertion holds:*

$$\forall d < K, \quad d_\beta(t_1, t_2) \leq d \Rightarrow d_\alpha\left(\left\lfloor \frac{\alpha}{\beta} t_1 \right\rfloor, \left\lfloor \frac{\alpha}{\beta} t_2 \right\rfloor\right) \leq 1$$

Proof. Let $t_1, t_2 \in \{0, \dots, \beta - 1\}$ such that $d_\beta(t_1, t_2) \leq d$. We write t_1 and t_2 as $t_1 = x_1K + \xi_1$ and $t_2 = x_2K + \xi_2$ where $x_1, x_2 \in \{0, \dots, \alpha - 1\}$ (resp. $\xi_1, \xi_2 \in \{0, \dots, K - 1\}$) are the quotients (resp. remainders) of the Euclidean division of t_1, t_2 by K . From Lemma 7, we have that $\lfloor t_1/K \rfloor = x_1$ and $\lfloor t_2/K \rfloor = x_2$.

Assume, by contradiction, that $d_\alpha(x_1, x_2) > 1$. By definition, this means that $\min((x_1 - x_2) \bmod \alpha, (x_2 - x_1) \bmod \alpha) > 1$. This implies that both $(x_1 - x_2) \bmod \alpha > 1$ and $(x_2 - x_1) \bmod \alpha > 1$. As $d_\beta(t_1, t_2) \leq d$, $\min((t_1 - t_2) \bmod \beta, (t_2 - t_1) \bmod \beta) \leq d$. Without loss of generality, assume that $(t_1 - t_2) \bmod \beta \leq d$. There are two cases:

1. If $t_1 \geq t_2$, then $(t_1 - t_2) \bmod \beta = t_1 - t_2$. So, $t_1 - t_2 \leq d$.

Now, as $t_1 \geq t_2$, $x_1 \geq x_2$ by Lemma 8. Hence $x_1 - x_2 = (x_1 - x_2) \bmod \alpha > 1$. As x_1 and x_2 are natural numbers, this implies that $x_1 - x_2 \geq 2$. We rewrite the inequality as $x_1K + \xi_1 - x_2K - \xi_2 \geq 2K + \xi_1 - \xi_2$. Since $\xi_1, \xi_2 \in \{0, \dots, K - 1\}$, we have $-K < \xi_1 - \xi_2 < K$ and therefore $x_1K + \xi_1 - x_2K - \xi_2 > K > d$. Hence, $t_1 - t_2 > d$, a contradiction.

2. If $t_1 < t_2$, then $(t_1 - t_2) \bmod \beta = \beta + t_1 - t_2$. So, $\beta + t_1 - t_2 \leq d$.

Now, as $t_1 < t_2$, $x_1 \leq x_2$ by Lemma 8. Hence $(x_1 - x_2) \bmod \alpha = \alpha + x_1 - x_2 > 1$. As x_1 and x_2 are natural numbers, this implies that $\alpha + x_1 - x_2 \geq 2$. We rewrite the inequality as $\beta + x_1K + \xi_1 - x_2K - \xi_2 \geq 2K + \xi_1 - \xi_2$. Since $\xi_1, \xi_2 \in \{0, \dots, K - 1\}$, we have $-K < \xi_1 - \xi_2 < K$ and therefore $\beta + x_1K + \xi_1 - x_2K - \xi_2 > K > d$. Hence, $\beta + t_1 - t_2 > d$, a contradiction. □

As previous lemmas, Lemma 9 will be used with the internal clock $t = c \frac{\beta}{\alpha} + \xi$: it expresses that once internal clocks have stabilized at a delay smaller than d , external clocks are at delay smaller than 1. We now prove that Algorithm 2 achieves the safety and liveness properties of SP_{SU} in any execution starting from a legitimate configuration.

Lemma 10 (Safety). *Every execution $e \in \mathcal{E}_{\text{LSU}}^0$ satisfies the safety of SP_{SU} .*

Proof. Let $\gamma \in \mathcal{L}_{\text{SU}}$: the delay (β) between any two internal clocks t in γ is upper bounded by $n - 1$ and for any process, $p \in V$, $\gamma(p).c = \left\lfloor \frac{\alpha}{\beta} \gamma(p).t \right\rfloor$. Hence, using Lemma 9 with $d = n - 1 < K$, we have $\forall p, q \in V$, $d_\alpha(\gamma(p).c, \gamma(q).c) \leq 1$. As $\alpha > 4$, this proves that the variables c in γ have at most two different consecutive values.

Finally, as the set $\mathcal{L}_{\mathcal{SU}}$ is closed (Lemma 5), we are done. \square

Lemma 11 (Liveness). *Every execution $e \in \mathcal{E}_{\mathcal{LSU}}^0$ satisfies the liveness of $SP_{\mathcal{SU}}$.*

Proof. Let $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{LSU}}^0$. Let p be a process. γ_0 is a legitimate configuration of \mathcal{WU} so p increments infinitely often $p.t$ using Action $\mathcal{SU}-N$ (by Theorem 2 and Remark 3). So $p.t$ goes through each integer value between 0 and $\beta - 1$ infinitely often (in increasing order). Hence, by Lemma 7, $p.c$ is incremented infinitely often and goes through each integer value between 0 and $\alpha - 1$ (in increasing order). \square

Proof of Theorem 3. Lemma 5 (closure), Lemma 6 (convergence), Lemmas 10 and 11 (correctness) prove that Algorithm \mathcal{SU} is self-stabilizing for $SP_{\mathcal{SU}}$ in any arbitrary connected anonymous network assuming a distributed unfair daemon. \square

6.2.2 Complexity Analysis

We now give some complexity results about Algorithm \mathcal{SU} . Precisely, a bound on the stabilization time of \mathcal{SU} is given in Theorem 4. Then, a delay between any two consecutive clocks increments, which holds once \mathcal{SU} has stabilized, is given in Theorem 5.

Theorem 4. *The stabilization time of \mathcal{SU} to $\mathcal{L}_{\mathcal{SU}}$ is at most $n + (\mu + 1)\mathcal{D} + 1$ rounds, where n (resp. \mathcal{D}) is the size (resp. diameter) of the network.*

Proof. Let $(\gamma_i)_{i \geq 0} \in \mathcal{E}^0$. The behavior of the t -variables in \mathcal{SU} is similar to that of \mathcal{WU} (Remark 3), which stabilizes in at most $n + \mu\mathcal{D}$ rounds (see Theorems 11 and 12) to weak unison. So, in $n + \mu\mathcal{D}$ rounds, the delay between the t -clocks of any two arbitrary far processes is at most $n - 1$ (Remark 1). If c -variables are well-calculated according to t -variables, *i.e.*, if $c = \left\lfloor \frac{\alpha}{\beta}t \right\rfloor$, then the delay between the c -clocks of any two arbitrary far processes is at most 1 (Lemma 9). In at most $\mathcal{D} + 1$ additional rounds, each process executes $\mathcal{SU}-N$ (Lemma 4) and updates its c -variable according to its t -variable. Hence, in at most $n + (\mu + 1)\mathcal{D} + 1$ rounds, the system reaches a legitimate configuration. \square

Theorem 5. *After convergence of \mathcal{SU} to $\mathcal{L}_{\mathcal{SU}}$, each process p increments its clock $p.c$ at least once every $\mathcal{D} + \frac{\beta}{\alpha}$ rounds, where \mathcal{D} is the diameter of the network.*

Proof. If \mathcal{DSU} converged to $\mathcal{L}_{\mathcal{SU}}^d$, by Remark 3 and Lemma 4, after $\mathcal{D} + \frac{\beta}{\alpha}$ rounds, p increments $p.t$ at least $\frac{\beta}{\alpha}$ times. Now, by Lemma 7, if a t -variable is incremented $\frac{\beta}{\alpha}$ times, then its corresponding c -variable is incremented once. \square

7 Gradual Stabilization under 1-Dynamics for Strong Unison

We now propose Algorithm \mathcal{DSU} (Algorithm 3), a variant of Algorithm \mathcal{SU} . \mathcal{DSU} is still self-stabilizing for strong unison and achieves a gradual convergence after one dynamic step. Precisely, after any dynamic step which fulfills condition `UnderLocalControl` (*n.b.*, this step may include several topological events), \mathcal{DSU} maintains clocks almost synchronized during the convergence to strong unison since it immediately satisfies partial weak unison, then converges in at most one round to weak unison, and finally re-stabilizes to strong unison.

Notice that, after one dynamic step, the graph contains at most $n + \#\mathcal{J}$ processes, by definition. Moreover, we denote by \mathcal{D}_1 the diameter of the new graph.

7.1 Algorithm \mathcal{DSU}

Our solution withstands one dynamic step, which may include several topological events (*i.e.* link or process additions or removals). However, according to Theorem 1, such a dynamic step should satisfy Condition `UnderLocalControl`. Namely, the graph should stay connected and, as $\alpha > 4$, every process that joins during the dynamic step $\gamma \mapsto_d \gamma'$ should be linked to at least one process which exists in both γ and γ' .

The following theorem allows to simplify proofs and explanations.

Theorem 6. *Let X be a closed set of configurations. If `UnderLocalControl` holds, then*

$$\forall \gamma_i \in \mathcal{C}, (\exists \gamma_j \in X \mid \gamma_j \mapsto_d \gamma_i) \Leftrightarrow (\exists \gamma_k \in X \mid \gamma_k \mapsto_{d_{\text{only}}} \gamma_i)$$

where $\mapsto_{d_{\text{only}}}$ is the set of all dynamic steps containing no process activation.

Proof. Let $\gamma_i \in \mathcal{C}$ such that $\gamma_j \mapsto_d \gamma_i$ with $\gamma_j \in X$. If $\gamma_j \mapsto_{d_{\text{only}}} \gamma_i$, we are done. Otherwise, let A be the non-empty subset of processes that are activated in $\gamma_j \mapsto_d \gamma_i$. There exists $\gamma_j \mapsto_s \gamma_u$, where A is activated. As X is closed, $\gamma_u \in X$. Moreover, $\forall x \in G_j \cap G_i$, $x \in G_u$ (since $G_u = G_j$) and $\gamma_u(x) = \gamma_i(x)$. Let $\gamma_u \mapsto_{d_{\text{only}}} \gamma_k$ such that $G_k = G_i$. $\forall x \in G_j \cap G_i$, $x \in G_k$ (since $G_k = G_i$) and $\gamma_k(x) = \gamma_u(x) = \gamma_i(x)$. Moreover, $\forall x \in G_i \setminus G_j$, $x \in G_k$ (since $G_k = G_i$) and $\gamma_k(x) = \gamma_i(x)$ because in both cases, x is in bootstate. Hence, $\gamma_k = \gamma_i$, and we are done.

The second part of the assertion is trivial since, by definition, $\mapsto_{d_{\text{only}}} \subseteq \mapsto_d$. \square

Since, by definition, $\mathcal{L}_{\text{SU}}^d$ is closed, the previous theorem applies: the set of configurations reachable from $\mathcal{L}_{\text{SU}}^d$ after one dynamic step (which may also include process




Figure 3: Delay between neighboring t -clocks may become greater than one after adding a link (e.g., the dashed one). The value of c - (resp. t -) variable is in the upper (resp. lower) part of the node.




Figure 4: Delay between neighboring t -clocks remains bounded by one after removing processes and/or links.

activations) is the same as the one reachable from $\mathcal{L}_{\text{SU}}^d$ after one dynamic step made of topological events only. At the light of this result, we only consider this latter kind of dynamic steps in the following explanations.

Consider first link additions only. Adding a link can break the safety of weak unison on internal clocks, see for example Figure 3. Indeed, adding a link may create a delay between two (new) neighboring t -clocks greater than one. Nevertheless, the delay between any two t -clocks remains bounded by $n - 1$ and, consequently, no process will reset its t -clock (Figure 3 shows a worst case). Moreover, c -clocks still satisfies strong unison immediately after the link addition. Besides, since increments are constrained by neighboring clocks, adding links only reinforces those constraints. Thus, the delay between internal clocks of arbitrary far processes remains bounded by $n - 1$, and so strong unison remains satisfied in all subsequent steps in this case. Consider again example in Figure 3: before the dynamic step, p_{n-1} had only to wait until p_{n-2} increments $p_{n-2}.t$ in order to be able to increment its own t -clock; yet after the step, it has also to wait for p_0 .

Assume now a dynamic step containing only process and link removals. Due to Condition `UnderLocalControl`, the network remains connected. Hence, constraints between (still existing) neighbors are maintained: the delay between t -clocks of two neighbors remains bounded by one, see example on Figure 4. So, weak unison on t -clocks remains satisfied and so is strong unison on c -clocks.

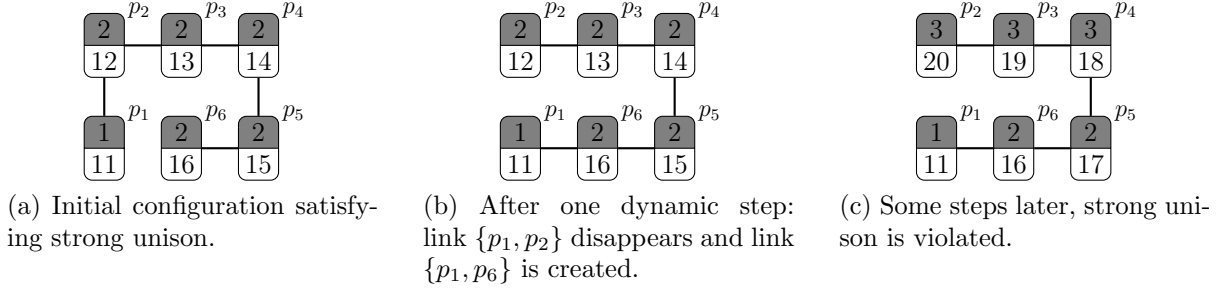


Figure 5: Example of execution where one link is added and another is removed: $\mu = 6$, $\alpha = 7$, and $\beta = 42$.

Consider now a more complex scenario, where the dynamic step contains link additions as well as process and/or link removals. Figure 5 shows an example of such a scenario, where safety of strong unison is violated. As above, the addition of link $\{p_1, p_6\}$ in Figure 5b leads to a delay between t -clocks of these two (new) neighbors which is greater than one (here 5). However, the removal of link $\{p_1, p_2\}$ relaxes the neighborhood constraint on p_2 : p_2 can now increment without waiting for p_1 . Consequently, executing Algorithm \mathcal{SU} does not ensure that the delay between t -clocks of any two arbitrary far processes remains bounded by $n - 1$, *e.g.*, in Figure 5c, the delay between p_1 and p_2 is 9 while $n - 1 = 5$. Since c -clock values are computed from t -clock values, we also cannot guarantee that there is at most two consecutive c -clock values in the system, see in Figure 5c $p_1.c = 1$, $p_6.c = 2$, and $p_2.c = 3$.

Again, in the worst case scenario, after a dynamic step, the delay between two neighboring t -clocks is bounded by $n - 1$. Moreover, t -clocks being computed like in Algorithm \mathcal{WU} , we can use two of its useful properties (see [5]): (1) when the delay between every pair of neighboring t -clocks is at most μ with $\mu \geq n$, the delay between these clocks remains bounded by μ because processes never reset; (2) furthermore, from such configurations, the system converges to a configuration from which the delay between the t -clocks of every two neighbors is at most one. So, keeping $\mu \geq n$, processes will not reset after one dynamic step and the delay between any two neighboring t -clocks will monotonically decrease from at most $n - 1$ to at most one. Consequently, the delay between any two neighboring c -clocks (which are computed from t -clocks) will stay less than or equal to one, *i.e.*, weak unison will be satisfied, all along the convergence to strong unison.

Consider now a process p that joins the system. The event $join_p$ occurs and triggers the specific action *bootstrap* that sets both the clocks $p.t$ and $p.c$ to a specific *bootstate* value, noted \perp . Note that by definition and from the previous discussion, the system immediately satisfies partial weak unison since it only depends on processes that were in the system before the dynamic step. Now, to ensure that weak unison holds within a

round, we add the action $\mathcal{DSU}\text{-}J$ which is enabled as soon as the process is in bootstate. This action initializes the two clocks of p according to the clock values in its neighborhood. Precisely, the value of $p.t$ can be chosen among the non- \perp values in its neighborhood, and such values exist by Condition UnderLocalControl . We choose to set $p.t$ to the minimum non- \perp t -clock value in its neighborhood, using the function MinTime_p given below.

$$\begin{aligned} \text{MinTime}_p &= 0 \text{ if } \forall q \in p.\mathcal{N}, q.t = \perp; \\ &\quad \min\{q.t : q \in p.\mathcal{N} \wedge q.t \neq \perp\} \text{ otherwise.} \end{aligned}$$

The value of $p.c$ is then computed according to the value of $p.t$. Notice that MinTime_p returns 0 when p and all its neighbors have their respective t -clock equal to \perp . This ensures that Algorithm \mathcal{DSU} remains self-stabilizing (in particular, if the system starts in a configuration where all t -clocks are equal to \perp).

To prevent the unfair daemon from blocking the convergence to a configuration containing no \perp values, we should also forbid processes with non- \perp t -clock values to increment while there are t -clock with \perp -values in their neighborhood. So, we define the predicate Locked which holds for a given process p when either $p.t = \perp$, or at least one of its neighbor q satisfies $q.t = \perp$. We then enforce the guard of both normal and reset actions, so that no Locked -process can execute them. See actions $\mathcal{DSU}\text{-}N$ and $\mathcal{DSU}\text{-}R$. This will ensure that t -clocks are initialized first by Action $\mathcal{DSU}\text{-}J$, before any value in their neighborhood increments.

Finally, notice that all the previous explanation relies on the fact that, once the system recovers from process additions (*i.e.*, once no \perp value remains), the algorithm behaves exactly the same as Algorithm \mathcal{SU} . Hence, it has to match the assumptions made for \mathcal{SU} . In particular, the assumptions on α and β remain the same. But the constraint on μ has to be adapted, since μ should be greater than or equal to the actual number of processes in the network and this number may increase. Now, the number of processes added in a dynamic step is bounded by $\#J$. So, we require μ to be greater than or equal to $n + \#J$.

We now consider the example of execution of Algorithm \mathcal{DSU} given in Figure 6. This execution starts in a configuration satisfying strong unison, see Figure 6a. Then, one dynamic step happens (step (a) \mapsto (b)), where a process p_6 joins the system. We now try to delay as long as possible the execution of $\mathcal{DSU}\text{-}J$ by p_6 . In configuration (b), p_3 and p_5 , the new neighbors of p_6 , are locked. They will remain disabled until p_6 executes $\mathcal{DSU}\text{-}J$. p_1 and p_4 execute $\mathcal{DSU}\text{-}N$ in (b) \mapsto (c). Then, p_4 is disabled because of p_5 and p_1 executes $\mathcal{DSU}\text{-}N$ in (c) \mapsto (d). In configuration (d), p_1 is from now on disabled: p_1 must wait until p_2 and p_4 get t -clock value 7. p_6 is the only enabled process, so the distributed unfair

Algorithm 3 *DSU*, for every process p

Parameters: α : any positive integer such that $\alpha > 4$ μ : any positive integer such that $\mu \geq n + \#J$ β : any positive integer such that $\beta > \mu^2$, and $\exists K$ such that $K > \mu$ and $\beta = K\alpha$ **Variables:** $p.c \in \{0, \dots, \alpha - 1\} \cup \{\perp\}$ $p.t \in \{0, \dots, \beta - 1\} \cup \{\perp\}$ **Predicates:** $Locked_p \equiv p.t = \perp \vee \exists q \in p.\mathcal{N}, q.t = \perp$ $NormalStep_p \equiv \neg Locked_p \wedge \forall q \in p.\mathcal{N}, p.t \preceq_{\beta, \mu} q.t$ $ResetStep_p \equiv \neg Locked_p \wedge (\exists q \in p.\mathcal{N}, d_\beta(p.t, q.t) > \mu \wedge p.t \neq 0)$ $JoinStep_p \equiv p.t = \perp$ **Actions:** $DSU-N \quad :: \quad NormalStep_p \quad \rightarrow \quad p.t \leftarrow (p.t + 1) \bmod \beta$ $p.c \leftarrow \left\lfloor \frac{\alpha}{\beta} p.t \right\rfloor$ $DSU-R \quad :: \quad ResetStep_p \quad \rightarrow \quad p.t \leftarrow 0$ $p.c \leftarrow 0$ $DSU-J \quad :: \quad JoinStep_p \quad \rightarrow \quad p.t \leftarrow MinTime_p$ $p.c \leftarrow \left\lfloor \frac{\alpha}{\beta} p.t \right\rfloor$ $bootstrap \quad :: \quad join_p \quad \rightarrow \quad p.t \leftarrow \perp$ $p.c \leftarrow \perp$




Figure 6: Example of execution where the daemon delays the first step of a new process: $\mu = 6$, $\alpha = 6$, and $\beta = 42$.

daemon has no other choice: it selects p_6 to initialize its variables executing $DSU-J$ in (d) \rightarrow (e).

7.2 Proof of Correctness

Self-stabilization w.r.t. SP_{SU} .

Remark 5. In DSU , if all t -variables have values different from \perp , predicates $JoinStep$ and $Locked$ are false. Furthermore, no action can assign \perp to t . Consequently,

- when all t -variables have values different from \perp ,
- as far as no topological change occurs,

Algorithms DSU and SU are syntactically identical. This implies in particular, that the set of executions \mathcal{E}^0 of SU and the set of executions \mathcal{E}_{nobot}^0 of DSU , where $nobot = \{\gamma \in \mathcal{C} : \forall p \in V, \gamma(p.t) \neq \perp\}$, are identical too.

Definition 3 (Legitimate Configurations of DSU w.r.t. SP_{SU}). A configuration γ of DSU is legitimate w.r.t. SP_{SU} if and only if

- $\forall p \in V, \gamma(p).t \neq \perp$
- $\forall p \in V, \forall q \in \gamma(p).\mathcal{N}, d_\beta(\gamma(p).t, \gamma(q).t) \leq 1$
- $\forall p \in V, \gamma(p).c = \lfloor \frac{\alpha}{\beta} \gamma(p).t \rfloor$.

We denote by \mathcal{L}_{SU}^d the set of legitimate configurations of DSU w.r.t. SP_{SU} .

Remark 6. As $\mathcal{L}_{\text{SU}}^d$ restricts the values of t to non- \perp values, we trivially have the equivalence between $\mathcal{L}_{\text{SU}}^d$ and \mathcal{L}_{SU} : for every configuration γ , we have $\gamma \in \mathcal{L}_{\text{SU}}^d \Leftrightarrow \gamma \in \mathcal{L}_{\text{SU}}$.

Lemma 12 (Closure of $\mathcal{L}_{\text{SU}}^d$ under \mathcal{DSU}). *The set of configurations $\mathcal{L}_{\text{SU}}^d$ is closed under \mathcal{DSU} .*

Proof. Let $\gamma \in \mathcal{L}_{\text{SU}}^d$ be a legitimate configuration of \mathcal{DSU} and let $\gamma \mapsto_s \gamma'$ be a static step of \mathcal{DSU} , from configuration γ . By Remark 6, γ is also in \mathcal{L}_{SU} . From Remark 5, the step $\gamma \mapsto_s \gamma'$ is also a step under \mathcal{SU} since γ contains no \perp -value and Lemma 5 applies (\mathcal{L}_{SU} is closed under \mathcal{SU}). Therefore, γ' is also in \mathcal{L}_{SU} and in $\mathcal{L}_{\text{SU}}^d$ as well, using again Remark 6. \square

Lemma 13. *For any execution $(\gamma_i)_{i \geq 0} \in \mathcal{E}^0$ under \mathcal{DSU} , $\exists j \geq 0$ such that $\forall k \geq j$, $\forall p \in V$, $\gamma_k(p).t \neq \perp$.*

Proof. Let $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}^0$. For any $i \geq 0$, we note $\text{Bottom}(\gamma_i) = \{p \in V : \gamma_i(p).t = \perp\}$. As actions $\mathcal{DSU-N}$, $\mathcal{DSU-R}$ and $\mathcal{DSU-J}$ do not create \perp values, $\forall i > 0$, $\text{Bottom}(\gamma_i) \subseteq \text{Bottom}(\gamma_{i-1})$. Now, assume by contradiction that $\exists p \in V$ such that $\forall i \geq 0$, $p \in \text{Bottom}(\gamma_i)$. There is a configuration γ_s , $s \geq 0$, from which no \perp -value disappears anymore, *i.e.*, $\forall p \in V$, $p \in \text{Bottom}(\gamma_s) \Rightarrow \forall i \geq s, p \in \text{Bottom}(\gamma_i)$.

If $\text{Bottom}(\gamma_s) = V$, every process is enabled for action $\mathcal{DSU-J}$. So, the unfair daemon selects at least one process to execute action $\mathcal{DSU-J}$ and sets its t -variable to a value different from \perp , a contradiction with the definition of γ_s .

Hence there is at least one process that is not in $\text{Bottom}(\gamma_s)$. Again, if the only enabled processes are in $\text{Bottom}(\gamma_s)$, then the unfair daemon has no other choice but selecting one of them, a contradiction. So, $\forall i \geq s$, there exists a process that is enabled in γ_i but which is not in $\text{Bottom}(\gamma_i)$. Remark that this implies in particular that e is an infinite execution (no terminal configuration reached).

Now, let consider the subgraph G' of G induced by $V \setminus \text{Bottom}(\gamma_s)$. G' is composed of a finite number of connected components and, as e is infinite, there is an infinite number of actions of e executed in (at least) one of these components. Let $G'' = (V'', E'')$ be such a connected component.

Let $e' = (\gamma'_i)_{i \geq 0}$ be the projection of e on G'' and t -variable: $\forall i \geq 0$, $\forall x \in V''$, $\gamma'_i(x).t = \gamma_i(x).t$. We construct $e'' = (\gamma''_j)_{j \geq 0}$ from e' by removing duplicate configurations with the following inductive schema:

- $\gamma''_0 = \gamma'_0$,
- and, $\forall j > 0$, if $\gamma''_0 \dots \gamma''_j$ represents $\gamma'_0 \dots \gamma'_k$ without duplicate configurations, $\gamma''_{j+1} = \gamma'_{\text{next}}$, where $\text{next} = \min\{l > k : \gamma'_l \neq \gamma'_k\}$. (Notice that next is always defined as there is an infinite number of actions executed in G'' .)

Let $L = \{p \in V'' : \exists q \in \text{Bottom}(\gamma_s), \{p, q\} \in E\}$ be the set of processes that are neighbors of a $\text{Bottom}(\gamma_s)$ process in G . As G is connected, L is not empty. Furthermore, during the execution e , Locked holds forever for processes in L , hence are disabled. As a consequence, in execution e'' , no process in L can execute a static step.

Now, from Remark 3 and 5, and since γ_0'' contains no \perp value, e'' is also an execution of \mathcal{WU} in graph G'' . The fact that existing processes (from the non-empty set L) never increment their clocks during an infinite execution e'' of \mathcal{WU} is a contradiction with the liveness of unison (Specification 2) and Theorem 2 which states that \mathcal{WU} is self-stabilizing for unison under an unfair daemon. \square

Lemma 14 (Convergence to $\mathcal{L}_{\text{SU}}^d$). \mathcal{C} (the set of all possible configurations) converges under \mathcal{DSU} to the set of legitimate configurations $\mathcal{L}_{\text{SU}}^d$.

Proof. Let $(\gamma_i)_{i \geq 0} \in \mathcal{E}^0$ under \mathcal{DSU} . Using Lemma 13, $\exists j \geq 0$ such that $\forall k \geq j, \forall p \in V, \gamma_k(p).t \neq \perp$. After γ_j , the execution of the system, $(\gamma_k)_{k \geq j}$, is also a possible execution of \mathcal{SU} (see Remark 5). Hence, it converges to a configuration γ_k ($k \geq j$) in \mathcal{L}_{SU} (Lemma 6). So, using Remark 6, $\gamma_k \in \mathcal{L}_{\text{SU}}^d$. \square

Lemma 15 (Correctness of SP_{SU} under \mathcal{DSU}). For every execution $e \in \mathcal{E}_{\mathcal{L}_{\text{SU}}^d}^0$ under \mathcal{DSU} , $SP_{\text{SU}}(e)$.

Proof. From Remark 5, every execution in $\mathcal{E}_{\mathcal{L}_{\text{SU}}^d}^0$ under \mathcal{DSU} is also an execution in $\mathcal{E}_{\mathcal{L}_{\text{SU}}}^0$ under \mathcal{SU} . Therefore, the correctness is proven in Lemmas 10 and 11. \square

Using Lemmas 12, 14 and 15, we can deduce the following theorem:

Theorem 7 (Self-stabilization of \mathcal{DSU} w.r.t. strong unison). Algorithm \mathcal{DSU} is self-stabilizing for SP_{SU} in any arbitrary connected anonymous network assuming a distributed unfair daemon.

Theorem 8 states the stabilization time of \mathcal{DSU} .

Theorem 8. The stabilization time of \mathcal{DSU} to $\mathcal{L}_{\text{SU}}^d$ is at most $n + (\mu + 1)\mathcal{D} + 2$, where n (resp. \mathcal{D}) is the size (resp. diameter) of the network, and μ is a parameter satisfying $\mu \geq n$.

Proof. Let $(\gamma_i)_{i \geq 0} \in \mathcal{E}^0$. If there are some processes p such that $\gamma_0(p).t = \perp$, \mathcal{DSU} -J is continuously enabled at p . So, after at most one round $p.t \neq \perp$. Afterwards, the behavior of the algorithm is similar to the one of \mathcal{SU} , which stabilizes in at most $n + (\mu + 1)\mathcal{D} + 1$ rounds (see Theorem 4). Hence, in at most $n + (\mu + 1)\mathcal{D} + 2$ rounds, the system reaches a legitimate configuration. \square

Immediate stabilization to SP_{PU} after one dynamic step.

Definition 4 (Legitimate Configurations of \mathcal{DSU} w.r.t. SP_{PU}). A configuration γ of \mathcal{DSU} is legitimate w.r.t. SP_{PU} if and only if

- a. $\forall p \in V, \gamma(p).t = \perp \Rightarrow (\exists q \in \gamma(p).\mathcal{N}, \gamma(q).t \neq \perp)$.
- b. $\forall p \in V, \forall q \in \gamma(p).\mathcal{N}, (\gamma(p).t \neq \perp \wedge \gamma(q).t \neq \perp) \Rightarrow (d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu)$.
- c. $\forall p, q \in V, (\gamma(p).t \neq \perp \wedge (\exists x \in \gamma(p).\mathcal{N}, \gamma(x).t = \perp) \wedge \gamma(q).t \neq \perp \wedge (\exists y \in \gamma(q).\mathcal{N}, \gamma(y).t = \perp)) \Rightarrow (d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu)$.
- d. $\forall p \in V, \gamma(p).t \neq \perp \Rightarrow \gamma(p).c = \lfloor \frac{\alpha}{\beta} \gamma(p).t \rfloor$.

We denote by \mathcal{L}_{PU}^d the set of legitimate configurations of \mathcal{DSU} w.r.t. SP_{PU} .

Lemma 16 (Closure of \mathcal{L}_{PU}^d under \mathcal{DSU}). The set of configurations \mathcal{L}_{PU}^d is closed under \mathcal{DSU} .

Proof. Let $\gamma \in \mathcal{L}_{PU}^d$ be a legitimate configuration of \mathcal{DSU} and let $\gamma \mapsto_s \gamma'$ be a computation step of \mathcal{DSU} , from configuration γ . In γ , action $\mathcal{DSU-R}$ is disabled for all processes: a process can only execute action $\mathcal{DSU-N}$ or $\mathcal{DSU-J}$ depending whether its clock is \perp or not.

- a. Let $p \in V$ such that $\gamma'(p).t = \perp$. As no action can set $p.t$ to \perp , $\gamma(p).t = \perp$ and by Definition 4, $\exists q \in \gamma(p).\mathcal{N}$ such that $\gamma(q).t \neq \perp$. $Locked_q$ holds in γ (because of p). Hence, q is disabled in γ and $\gamma(q).t = \gamma'(q).t \neq \perp$.
- b. Let $p \in V$ and $q \in \gamma(p).\mathcal{N}$ such that $\gamma'(p).t \neq \perp$ and $\gamma'(q).t \neq \perp$.
 1. If $\gamma(p).t \neq \perp$ and $\gamma(q).t \neq \perp$, as $\gamma \in \mathcal{L}_{PU}^d$, $d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu$. Now, p and q can only execute action $\mathcal{DSU-N}$ during $\gamma \mapsto_s \gamma'$. If both p and q , or none of them, execute action $\mathcal{DSU-N}$, the delay between $p.t$ and $q.t$ remains the same. If only one of them, say p , executes action $\mathcal{DSU-N}$, $p.t \preceq_{\beta, \mu} q.t$ holds in γ . So, either $\gamma(p).t = \gamma(q).t$ and $d_\beta(\gamma'(p).t, \gamma'(q).t) = 1 \leq \mu$, or the increment of $p.t$ decreases the delay between $p.t$ and $q.t$ and again we have $d_\beta(\gamma'(p).t, \gamma'(q).t) \leq \mu$.
 2. If $\gamma(p).t = \perp$ and $\gamma(q).t \neq \perp$, as $\gamma \in \mathcal{L}_{PU}^d$, $\exists x \in \gamma(p).\mathcal{N}$ such that $\gamma(x).t \neq \perp$. We choose x as such a neighbor of p , with minimum value for t , i.e., $\gamma(x).t = MinTime_p$ in γ . Hence, $d_\beta(\gamma(x).t, \gamma(q).t) \leq \mu$ because of Definition 4.c: q and x have a (common) neighbor p whose t -variable equals \perp . q is disabled in γ because of p ($Locked_q$ holds in γ), hence $\gamma(q).t = \gamma'(q).t$. As $\gamma'(p).t = \gamma(x).t$ (since p executes action $\mathcal{DSU-J}$), $d_\beta(\gamma'(p).t, \gamma'(q).t) \leq \mu$.

3. If $\gamma(p).t \neq \perp$ and $\gamma(q).t = \perp$, similar to case 2.
 4. If $\gamma(p).t = \perp$ and $\gamma(q).t = \perp$, as $\gamma \in \mathcal{L}_{\text{PU}}^d$, $\exists x \in p.\mathcal{N}$ such that $\gamma(x).t = \text{MinTime}_p \neq \perp$ in γ and $\exists y \in p.\mathcal{N}$ such that $\gamma(y).t = \text{MinTime}_q \neq \perp$ in γ . Because of Definition 4.c, $d_\beta(\gamma(x).t, \gamma(y).t) \leq \mu$ because they have neighbors whose t -variables equal \perp (p and q , respectively). Since p and q execute action $\mathcal{DSU}\text{-}J$, $\gamma'(p).t = \gamma(x).t$ and $\gamma'(q).t = \gamma(y).t$, so $d_\beta(\gamma'(p).t, \gamma'(q).t) \leq \mu$.
- c. Let $p, q \in V$ such that $\gamma'(p).t \neq \perp$, $\exists x \in \gamma'(p).\mathcal{N}$ with $\gamma'(x).t = \perp$, $\gamma'(p).t \neq \perp$, and $\exists y \in \gamma'(q).\mathcal{N}$ with $\gamma'(y).t = \perp$.

As no action can set variable t to \perp , $\gamma(x).t = \perp$ and $\gamma(y).t = \perp$.

1. If $\gamma(p).t \neq \perp$ and $\gamma(q).t \neq \perp$, as $\gamma \in \mathcal{L}_{\text{PU}}^d$, $d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu$. Now, p and q are disabled in γ (Locked_p , Locked_q) because of x and y , respectively. Hence, $d_\beta(\gamma'(p).t, \gamma'(q).t) \leq \mu$.
 2. If $\gamma(p).t = \perp$ and $\gamma(q).t \neq \perp$, as $\gamma \in \mathcal{L}_{\text{PU}}^d$, $\exists x' \in \gamma(p).\mathcal{N}$ such that $\gamma(x').t = \text{MinTime}_p \neq \perp$ in γ . Hence, $d_\beta(\gamma(x').t, \gamma(q).t) \leq \mu$ because they have neighbors whose t -variables equal \perp (p and y , respectively). q is disabled in γ (Locked_q) because of y : $\gamma(q).t = \gamma'(q).t$. And $\gamma'(p).t = \gamma(x').t$ since p executes action $\mathcal{DSU}\text{-}J$. So $d_\beta(\gamma'(p).t, \gamma'(q).t) \leq \mu$.
 3. If $\gamma(p).t \neq \perp$ and $\gamma(q).t = \perp$, similar to case 2.
 4. If $\gamma(p).t = \perp$ and $\gamma(q).t = \perp$, as $\gamma \in \mathcal{L}_{\text{PU}}^d$, $\exists x' \in \gamma(p).\mathcal{N}$ such that $\gamma(x').t = \text{MinTime}_p \neq \perp$ in γ and $\exists y' \in \gamma(p).\mathcal{N}$ such that $\gamma(y').t = \text{MinTime}_q \neq \perp$ in γ . Hence, $d_\beta(\gamma(x').t, \gamma(y').t) \leq \mu$ because they have neighbors whose t -variables equal \perp (p and q , respectively). $\gamma'(p).t = \gamma(x').t$ and $\gamma'(q).t = \gamma(y').t$ since p and q execute action $\mathcal{DSU}\text{-}J$. So $d_\beta(\gamma'(p).t, \gamma'(q).t) \leq \mu$.
- d. Let $p \in V$ such that $\gamma'(p).t \neq \perp$. Two cases are possible: either p executes no action and the constraint between $p.t$ and $p.c$ is preserved, or p executes an action $\mathcal{DSU}\text{-}N$. In the latter case, the assignment of the action ensures the constraint.

□

Lemma 17 (Safety of SP_{PU} in $\mathcal{E}_{\mathcal{L}_{\text{PU}}^d}^0$). *Every execution $e \in \mathcal{E}_{\mathcal{L}_{\text{PU}}^d}^0$ satisfies the safety of SP_{PU} .*

Proof. Let $\gamma \in \mathcal{L}_{\text{PU}}^d$. By Definition 4, $\forall p \in V, \forall q \in \gamma(p).\mathcal{N}, (\gamma(p).t \neq \perp \wedge \gamma(q).t \neq \perp) \Rightarrow d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu$. Furthermore, $\forall p \in V, \gamma(p).t \neq \perp \Rightarrow \gamma(p).c = \left\lfloor \frac{\alpha}{\beta} \gamma(p).t \right\rfloor$. Hence, using Lemma 9 with $d = \mu < K$, $\forall p \in V, \forall q \in \gamma(p).\mathcal{N}, (\gamma(p).t \neq \perp \wedge \gamma(q).t \neq \perp) \Rightarrow d_\alpha(\gamma(p).c, \gamma(q).c) \leq 1$. Finally, as the set $\mathcal{L}_{\text{PU}}^d$ is closed (Lemma 16), we are done. □

Lemma 18 (Liveness of SP_{PU} in $\mathcal{E}_{\mathcal{L}_{\text{PU}}^d}^0$). *Every execution $e \in \mathcal{E}_{\mathcal{L}_{\text{PU}}^d}^0$ satisfies the liveness of SP_{PU} .*

Proof. Let $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\text{PU}}^d}^0$. Using Lemma 13, $\exists i \geq 0$ such that $\forall j \geq i, \forall p \in V, \gamma_j(p).t \neq \perp$. The legitimate configurations $\mathcal{L}_{\text{PU}}^d$ are closed (Lemma 16), so γ_i is a legitimate configuration *w.r.t.* SP_{PU} and is also a legitimate configuration of \mathcal{WU} so p increments infinitely often $p.t$ executing action $\mathcal{DSU}\text{-}N$ (Lemma 27). Furthermore, actions $\mathcal{DSU}\text{-}R$ and $\mathcal{DSU}\text{-}J$ are disabled. So $p.t$ goes through each integer value between 0 and $\beta - 1$ infinitely often (in increasing order). Hence, by Lemma 7, $p.c$ is incremented infinitely often and goes through each integer value between 0 and $\alpha - 1$ (in increasing order). \square

Lemma 19. *\mathcal{DSU} converges from $\mathcal{L}_{\text{PU}}^d$ to $\mathcal{L}_{\text{WU}}^d$ in a finite time. The convergence time is at most one round.*

Proof. Let $(\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\text{PU}}^d}^0$. $\forall p \in V$, such that $\gamma_0(p).t = \perp$, action $\mathcal{DSU}\text{-}J$ is continuously enabled at p . By Lemma 13, $\exists i > 0$, such that $\forall j \geq i, \forall p \in V, \gamma_j(p).t \neq \perp$. Hence, $\exists j \in \{1, \dots, i\}$ such that p executes action $\mathcal{DSU}\text{-}J$ during $\gamma_{j-1} \mapsto \gamma_j$.

So, in at most one round, the system reaches a configuration $\gamma_k, k \geq 0$, such that $\forall p \in V_k, \gamma_k(p).t \neq \perp$. Now, $\gamma_0 \in \mathcal{L}_{\text{PU}}^d$ and $\mathcal{L}_{\text{PU}}^d$ is closed under \mathcal{DSU} (Lemma 16), so $\gamma_k \in \mathcal{L}_{\text{PU}}^d$. As there is no t -variable with \perp value in $\gamma_k, \gamma_k \in \mathcal{L}_{\text{WU}}^d$. \square

Stabilization to SP_{WU} in at most one round after one dynamic step.

Definition 5 (Legitimate Configurations of \mathcal{DSU} *w.r.t.* SP_{WU}). *A configuration γ of \mathcal{DSU} is legitimate w.r.t. SP_{WU} if and only if*

- $\forall p \in V, \gamma(p).t \neq \perp$.
- $\forall p \in V, \forall q \in \gamma(p).\mathcal{N}, d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu$.
- $\forall p \in V, \gamma(p).c = \lfloor \frac{\alpha}{\beta} \gamma(p).t \rfloor$.

We denote by $\mathcal{L}_{\text{WU}}^d$ the set of legitimate configurations of \mathcal{DSU} w.r.t. SP_{WU} .

Lemma 20 (Closure of $\mathcal{L}_{\text{WU}}^d$ under \mathcal{DSU}). *The set of configurations $\mathcal{L}_{\text{WU}}^d$ is closed under \mathcal{DSU} .*

Proof. Notice that the first and third constraints of Definition 5 are closed, as for $\mathcal{L}_{\text{SU}}^d$. As, (1) for every configuration γ , we have that $\gamma \in \mathcal{L}_{\text{WU}}^d \Rightarrow \gamma \in C_\mu$, (2) C_μ is closed under \mathcal{WU} (see Lemma 25) and (3) every static step of \mathcal{DSU} from $\mathcal{L}_{\text{WU}}^d$ is also a static step for \mathcal{WU} (see Remarks 3 and 5), the second constraint is also closed under \mathcal{DSU} . \square

Lemma 21 (Safety of $SP_{\mathcal{WU}}$ in $\mathcal{E}_{\mathcal{L}_{\mathcal{WU}}^d}^0$). *Every execution $e \in \mathcal{E}_{\mathcal{L}_{\mathcal{WU}}^d}^0$ satisfies the safety of $SP_{\mathcal{WU}}$.*

Proof. Let $\gamma \in \mathcal{L}_{\mathcal{WU}}^d$, $p \in V$, and $q \in \gamma(p).\mathcal{N}$. We have that $d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu$. and $\gamma(p).c = \left\lfloor \frac{\alpha}{\beta} \gamma(p).t \right\rfloor$. Using Lemma 9 with $d = \mu < K$, we obtain that $d_\alpha(\gamma(p).c, \gamma(q).c) \leq 1$. Finally, as the set $\mathcal{L}_{\mathcal{WU}}^d$ is closed (Lemma 20), we are done. \square

Lemma 22 (Liveness of $SP_{\mathcal{WU}}$ in $\mathcal{E}_{\mathcal{L}_{\mathcal{WU}}^d}^0$). *Every execution $e \in \mathcal{E}_{\mathcal{L}_{\mathcal{WU}}^d}^0$ satisfies the liveness of $SP_{\mathcal{WU}}$.*

Proof. Let $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\mathcal{WU}}^d}^0$. Let p be a process. γ_0 is a legitimate configuration of \mathcal{WU} so p increments infinitely often $p.t$ executing $\mathcal{DSU-N}$ (see Lemma 27). Furthermore, actions $\mathcal{DSU-R}$ and $\mathcal{DSU-J}$ are disabled. So $p.t$ goes through each integer value between 0 and $\beta - 1$ infinitely often (in increasing order). Hence, by Lemma 7, $p.c$ is incremented infinitely often and goes through each integer value between 0 and $\alpha - 1$ (in increasing order). \square

Gradual stabilization after one dynamic step. We now consider executions in $\mathcal{E}_{\mathcal{L}_{\mathcal{SU}}^d}^1$ where `UnderLocalControl` holds.

Lemma 23. *Let $\gamma_i \in \mathcal{L}_{\mathcal{SU}}^d$ be a legitimate configuration under \mathcal{DSU} , $\gamma_i \mapsto_d \gamma_{i+1}$ be a dynamic step, such that $\forall p \in V_{i+1}, \gamma_{i+1}(p).t = \perp \Rightarrow (\exists q \in \gamma_{i+1}(p).\mathcal{N}, \gamma_{i+1}(q).t \neq \perp)$, then $\gamma_{i+1} \in \mathcal{L}_{\mathcal{PU}}^d$.*

Proof. Let $\gamma_i \mapsto_d \gamma_{i+1}$ be a dynamic step such that $\gamma_i \in \mathcal{L}_{\mathcal{SU}}^d$ and $\forall p \in V_{i+1}, \gamma_{i+1}(p).t = \perp \Rightarrow (\exists q \in \gamma_{i+1}(p).\mathcal{N}, \gamma_{i+1}(q).t \neq \perp)$. By Theorem 6 and as $\mathcal{L}_{\mathcal{SU}}^d$ is closed (Lemma 12), the set of reachable configurations from $\mathcal{L}_{\mathcal{SU}}^d$ after one dynamic step satisfying `UnderLocalControl` is the same than after one dynamic step satisfying `UnderLocalControl` without process activations. Hence, we consider here that $\gamma_i \mapsto_d \gamma_{i+1}$ contains no process activations.

By Definition 3, $\forall p \in V_i, \gamma_i(p).t \neq \perp, \gamma_i(p).c = \left\lfloor \frac{\alpha}{\beta} \gamma_i(p).t \right\rfloor$, and $\forall q \in \gamma_i(p).\mathcal{N}, d_\beta(\gamma_i(p).t, \gamma_i(q).t) \leq 1$. Hence $\forall p, q \in V_i, d_\beta(\gamma_i(p).t, \gamma_i(q).t) \leq \mu$.

Now, after the dynamic step, in γ_{i+1} , the state of processes that are in $V_i \cap V_{i+1}$ remains the same. So, $\forall p \in V_{i+1}, \forall q \in \gamma_{i+1}(p).\mathcal{N}$, if $\gamma_{i+1}(p).t \neq \perp$ and $\gamma_{i+1}(q).t \neq \perp$, then $p, q \in V_i \cap V_{i+1}$. As $d_\beta(\gamma_i(p).t, \gamma_i(q).t) \leq \mu$ (Remark 4), $d_\beta(\gamma_{i+1}(p).t, \gamma_{i+1}(q).t) \leq \mu$.

Finally, $\forall p \in V_{i+1}$, if $\gamma_{i+1}(p).t \neq \perp, p \in V_i \cap V_{i+1}$ so $\gamma_{i+1}(p).c = \gamma_i(p).c = \left\lfloor \frac{\alpha}{\beta} \gamma_i(p).t \right\rfloor = \left\lfloor \frac{\alpha}{\beta} \gamma_{i+1}(p).t \right\rfloor$.

Hence, $\gamma_{i+1} \in \mathcal{L}_{\mathcal{PU}}^d$. \square

Lemma 24. *\mathcal{DSU} converges from $\mathcal{L}_{\mathcal{WU}}^d$ to $\mathcal{L}_{\mathcal{SU}}^d$ in a finite time. The convergence time is at most $(\mu + 1)\mathcal{D}_1 + 1$ rounds.*

Proof. Let $e = (\gamma_i)_{i \geq 0} \in \mathcal{E}_{\mathcal{L}_{\text{WU}}^d}^0$. The behavior of the algorithm is similar to the one of \mathcal{WU} (Remarks 3 and 5). Furthermore, $\forall p \in V, \forall q \in p.\mathcal{N}, d_\beta(\gamma_0(p).t, \gamma_0(q).t) \leq \mu$, so $\gamma_0 \in C_\mu$. By Lemma 29, in a finite time, $\forall p \in V, \forall q \in p.\mathcal{N}, d_\beta(\gamma_0(p).t, \gamma_0(q).t) \leq 1$. This convergence lasts at most $\mu\mathcal{D}_1$ rounds (Theorem 12).

The liveness of weak unison is ensured in e (Lemma 22), so each process eventually increments its clock executing \mathcal{DSU} - N and updates its c -variable. By Lemma 4, the c -variables are well computed according to t -variables in at most $\mathcal{D}_1 + 1$ additional rounds. Hence, in at most $(\mu + 1)\mathcal{D}_1 + 1$ rounds, the system reaches a $\mathcal{L}_{\text{SU}}^d$. \square

Since `UnderLocalControl` imposes that the system remains connected after one dynamic step, we can use Theorem 7 and Lemmas 20-19 to establish the following theorem:

Theorem 9. *If `UnderLocalControl` is satisfied then \mathcal{DSU} is gradually stabilizing under 1-dynamics for $(SP_{\text{PU}} \bullet 0, SP_{\text{WU}} \bullet 1, SP_{\text{SU}} \bullet (\mu + 1)\mathcal{D}_1 + 2)$, where \mathcal{D}_1 (resp. $n + \#J$) is the diameter (resp. an upper bound on the size) of the network after the dynamic step and μ is a parameter satisfying $\mu \geq n + \#J$.*

Theorem 10 establish a bound on how many rounds are necessary to ensure that a given process increments its c -clock after the convergence to legitimate configurations *w.r.t.* SP_{SU} (resp. SP_{WU}).

Theorem 10. *After convergence of \mathcal{DSU} to $\mathcal{L}_{\text{WU}}^d$ (resp. $\mathcal{L}_{\text{SU}}^d$), each process p increments its clock $p.c$ at least once every $\mu\mathcal{D}_1 + \frac{\beta}{\alpha}$ rounds (resp. $\mathcal{D}_1 + \frac{\beta}{\alpha}$ rounds).*

Proof. By Remarks 3 and 5, we can use results on \mathcal{WU} for \mathcal{DSU} . If \mathcal{DSU} converged to a configuration $\gamma \in \mathcal{L}_{\text{WU}}^d$, then $\gamma \in C_\mu$. So, by Lemma 3, after $\mu\mathcal{D}_1 + \frac{\beta}{\alpha}$ rounds, p increments $p.t$ at least $\frac{\beta}{\alpha}$ times. Now, by Lemma 7, if t -variable is incremented $\frac{\beta}{\alpha}$ times, c -variable is incremented once.

If \mathcal{DSU} converged to $\mathcal{L}_{\text{SU}}^d$, the result of Theorem 5 can be applied (Remark 5). So, after $\mathcal{D}_1 + \frac{\beta}{\alpha}$ rounds, p increments $p.c$ at least once. \square

8 Conclusion

The apparent seldomness of superstabilizing solutions for non-static problems, such as unison, may suggest the difficulty of obtaining such a strong property and if so, make our notion of gradual stabilization very attractive compared to merely self-stabilizing solutions. For example, in our unison solution, gradual stabilization ensures that processes remain “almost” synchronized during the convergence phase started after one dynamic step. Hence, it is worth investigating whether this new paradigm can be applied to other, in particular non-static, problems.

Concerning our unison algorithm, the graceful recovery after one dynamic step comes at the price of slowing down the clock increments. The question of limiting this drawback remains open.

Finally, it would be interesting to address in future work gradual stabilization for non-static problems in context of more complex dynamic patterns.

References

- [1] Anish Arora, Shlomi Dolev, and Mohamed G. Gouda. Maintaining digital clocks in step. *Parallel Processing Letters*, 1:11–18, 1991.
- [2] Baruch Awerbuch, Shay Kutten, Yishay Mansour, Boaz Patt-Shamir, and George Varghese. Time optimal self-stabilizing synchronization. In *STOC*, pages 652–661, 1993.
- [3] Lélia Blin, Maria Potop-Butucaru, and Stephane Rovedakis. A super-stabilizing $\log(n)\log(n)$ -approximation algorithm for dynamic steiner trees. *Theor. Comput. Sci.*, 500:90–112, 2013.
- [4] Lélia Blin, Maria Gradinariu Potop-Butucaru, Stephane Rovedakis, and Sébastien Tixeuil. Loop-free super-stabilizing spanning tree construction. In *SSS*, pages 50–64, 2010.
- [5] Christian Boulinier. *L’Unisson*. PhD thesis, Université de Picardie Jules Vernes, France, 2007.
- [6] Christian Boulinier, Franck Petit, and Vincent Villain. When graph theory helps self-stabilization. In *PODC*, pages 150–159, 2004.
- [7] Fabienne Carrier, Ajoy Kumar Datta, Stéphane Devismes, Lawrence L. Larmore, and Yvan Rivierre. Self-stabilizing (f,g)-alliances with safe convergence. *J. Parallel Distrib. Comput.*, 81-82:11–23, 2015.
- [8] Jean-Michel Couvreur, Nissim Francez, and Mohamed G. Gouda. Asynchronous unison (extended abstract). In *ICDCS*, pages 486–493, 1992.
- [9] Edsger W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11):643–644, 1974.
- [10] Shlomi Dolev. *Self-stabilization*. MIT Press, 2000.

- [11] Shlomi Dolev and Ted Herman. Superstabilizing protocols for dynamic distributed systems. *Chicago J. Theor. Comput. Sci.*, 1997, 1997.
- [12] Shlomi Dolev, Amos Israeli, and Shlomo Moran. Uniform Dynamic Self-Stabilizing Leader Election. *IEEE Trans. Parallel Distrib. Syst.*, 8(4):424–440, 1997.
- [13] Christophe Genolini and Sébastien Tixeuil. A lower bound on dynamic k-stabilization in asynchronous systems. In *SRDS*, page 212, 2002.
- [14] Sukumar Ghosh, Arobinda Gupta, Ted Herman, and Sriram V. Pemmaraju. Fault-containing self-stabilizing distributed protocols. *Distributed Computing*, 20(1):53–73, 2007.
- [15] Mohamed G. Gouda and Ted Herman. Stabilizing unison. *Inf. Process. Lett.*, 35(4):171–175, 1990.
- [16] Ted Herman. Superstabilizing mutual exclusion. *Distributed Computing*, 13(1):1–17, 2000.
- [17] Shing-Tsaan Huang and Tzong-Jye Liu. Four-state stabilizing phase clock for unidirectional rings of odd size. *Inf. Process. Lett.*, 65(6):325–329, 1998.
- [18] Colette Johnen, Luc Onana Alima, Ajoy Kumar Datta, and Sébastien Tixeuil. Optimal snap-stabilizing neighborhood synchronizer in tree networks. *Parallel Processing Letters*, 12(3-4):327–340, 2002.
- [19] Hirotsugu Kakugawa and Toshimitsu Masuzawa. A self-stabilizing minimal dominating set algorithm with safe convergence. In *IPDPS*, pages 8.–, 2006.
- [20] Sayaka Kamei, Tomoko Izumi, and Yukiko Yamauchi. An asynchronous self-stabilizing approximation for the minimum connected dominating set with safe convergence in unit disk graphs. In *SSS*, pages 251–265, 2013.
- [21] Sayaka Kamei and Hirotsugu Kakugawa. A self-stabilizing approximation algorithm for the minimum weakly connected dominating set with safe convergence. In *WRAS*, pages 57–67, 2007.
- [22] Sayaka Kamei and Hirotsugu Kakugawa. A self-stabilizing 6-approximation for the minimum connected dominating set with safe convergence in unit disk graphs. *Theor. Comput. Sci.*, 428:80–90, 2012.
- [23] Yoshiaki Katayama, Eiichiro Ueda, Hideo Fujiwara, and Toshimitsu Masuzawa. A latency optimal superstabilizing mutual exclusion protocol in unidirectional rings. *J. Parallel Distrib. Comput.*, 62(5):865–884, 2002.

- [24] Sandeep S. Kulkarni and Anish Arora. Multitolerant barrier synchronization. *Inf. Process. Lett.*, 64(1):29–36, 1997.
- [25] Shay Kutten and Boaz Patt-Shamir. Stabilizing time-adaptive protocols. *Theor. Comput. Sci.*, 220(1):93–111, 1999.
- [26] Jayadev Misra. Phase synchronization. *Inf. Process. Lett.*, 38(2):101–105, 1991.
- [27] Florent Nolot and Vincent Villain. Universal self-stabilizing phase clock protocol with bounded memory. In *IPCCC*, pages 228–235, 2001.
- [28] Chi-Hung Tzeng, Jehn-Ruey Jiang, and Shing-Tsaan Huang. Size-independent self-stabilizing asynchronous phase synchronization in general graphs. *J. Inf. Sci. Eng.*, 26(4):1307–1322, 2010.

A Results from [5]

In this section, we recall some useful technical results from [5] about Algorithm \mathcal{WU} . Algorithm \mathcal{WU} is an instance of the parametric algorithm GAU in [5]: $\mathcal{WU} = GAU(\beta, 0, \mu)$.

The following five lemmas (25-29) are used to establish the self-stabilization of \mathcal{WU} for $SP_{\mathcal{WU}}$ by the set of legitimate configurations $\mathcal{L}_{\mathcal{WU}}$ such that $\gamma \in \mathcal{L}_{\mathcal{WU}}$ if and only if $\forall p \in V, \forall q \in \gamma(p).\mathcal{N}, d_\beta(\gamma(p).t, \gamma(q).t) \leq 1$.

The proof of self-stabilization is divided into several steps. The first step (Lemma 26) consists in showing the convergence of \mathcal{WU} from \mathcal{C} to C_μ , where C_μ is the set of configurations where the delay between the clocks of two neighbors is at most μ , *i.e.*,

$$C_\mu = \{\gamma \in \mathcal{C} : \forall p \in V, \forall q \in \gamma(p).\mathcal{N}, d_\beta(\gamma(p).t, \gamma(q).t) \leq \mu\}$$

C_μ is shown to be closed under \mathcal{WU} in Lemma 25. (Notice that $\mathcal{L}_{\mathcal{WU}} \subseteq C_\mu$.) The liveness part of $SP_{\mathcal{WU}}$ (the clock $p.t$ of every process p goes through each value in $\{0, \dots, \beta - 1\}$ in increasing order infinitely often) is shown for every execution starting from C_μ in Lemma 27.

Lemma 25 (Property 8 in [5]). *C_μ is closed under \mathcal{WU} .*

Lemma 26 (Theorem 56 in [5]). *If $n \leq \mu < \frac{\beta}{2}$, then $\forall e \in \mathcal{E}^0, \exists \gamma \in e$ such that $\gamma \in C_\mu$.*

Lemma 27 (Theorem 21 in [5]). *If $\beta > n^2$, then $\forall e \in \mathcal{E}_{C_\mu}^0, e$ satisfies the liveness part of $SP_{\mathcal{WU}}$.*

Then, the second step consists of showing closure of $\mathcal{L}_{\mathcal{WU}}$ under \mathcal{WU} (Lemma 28) and the convergence from C_μ to $\mathcal{L}_{\mathcal{WU}}$ (Lemma 29). Regarding the correctness, the safety part of $SP_{\mathcal{WU}}$ (two neighboring clocks differ from at most 1) is ensured by definition of $\mathcal{L}_{\mathcal{WU}}$, whereas the liveness part is already ensured by Lemma 27. Precisely:

Lemma 28 (Property 2 in [5]). *$\mathcal{L}_{\mathcal{WU}}$ is closed under \mathcal{WU} .*

Lemma 29 (Theorems 29 in [5]). *If $\beta > n^2$ and $\mu < \frac{\beta}{2}$, then $\forall e \in \mathcal{E}_{C_\mu}^0, \exists \gamma \in e$ such that $\gamma \in \mathcal{L}_{\mathcal{WU}}$.*

Some performances of Algorithm \mathcal{WU} are already recalled in Theorems 11 and 12 (page 41).

Theorem 11 (Theorem 61 in [5]). *If $n \leq \mu < \frac{\beta}{2}$, the convergence time of \mathcal{WU} from \mathcal{C} to C_μ is at most n rounds.*

Theorem 12 (Theorems 20 and 28 in [5]). *If $\beta > n^2$ and $\mu < \frac{\beta}{2}$, the convergence time of \mathcal{WU} from C_μ to $\mathcal{L}_{\mathcal{WU}}$ is at most $\mu\mathcal{D}$ rounds.*

Finally, Lemma 30 below is a technical result about the values of t -variables.

Lemma 30 (Theorem 20, Property 27, and Lemma 22 in [5]). *If $\beta > n^2$ and $\beta > 2\mu$, then $\forall e = (\gamma_i)_{i \geq 0} \in \mathcal{E}_{C_\mu}^0$, there exists a function f on processes such that*

- $\forall i \geq 0, \forall p \in V, f(\gamma_i, p) \bmod \beta = \gamma_i(p).t,$
- *and* $\forall i \geq 0, \forall p, q \in V, |f(\gamma_i, p) - f(\gamma_i, q)| = d_\beta(\gamma_i(p).t, \gamma_i(q).t).$