

Self-stabilizing Systems in Spite of High Dynamics

Karine Altisen, Stéphane Devismes
Univ. Grenoble Alpes, VERIMAG
Firstname.Lastname@univ-grenoble-alpes.fr

Anaïs Durand
Université de Clermont Auvergne, LIMOS
Anais.Durand@uca.fr

Colette Johnen
Univ. of Bordeaux, LaBRI
Colette.Johnen@labri.fr

Franck Petit
Sorbonne Université, LIP6
Franck.Petit@lip6.fr

ABSTRACT

We initiate research on self-stabilization in highly dynamic identified message passing systems where dynamics is modeled using time-varying graphs (TVGs). More precisely, we address the self-stabilizing leader election problem in three wide classes of TVGs: the class $\mathcal{TC}^{\mathcal{B}}(\Delta)$ of TVGs with temporal diameter bounded by Δ , the class $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ of TVGs with temporal diameter quasi-bounded by Δ , and the class $\mathcal{TC}^{\mathcal{R}}$ of TVGs with recurrent connectivity only, where $\mathcal{TC}^{\mathcal{B}}(\Delta) \subseteq \mathcal{TC}^{\mathcal{Q}}(\Delta) \subseteq \mathcal{TC}^{\mathcal{R}}$. We first study conditions under which our problem can be solved. We introduce the notion of size-ambiguity to show that the assumption on the knowledge of the number n of processes is central. Our results reveal that, despite the existence of unique process identifiers, any deterministic self-stabilizing leader election algorithm working in the class $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ or $\mathcal{TC}^{\mathcal{R}}$ cannot be size-ambiguous, justifying why our solutions for those classes assume the exact knowledge of n . We then present three self-stabilizing leader election algorithms for Classes $\mathcal{TC}^{\mathcal{B}}(\Delta)$, $\mathcal{TC}^{\mathcal{Q}}(\Delta)$, and $\mathcal{TC}^{\mathcal{R}}$, respectively. Our algorithm for $\mathcal{TC}^{\mathcal{B}}(\Delta)$ stabilizes in at most 3Δ rounds. In $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ and $\mathcal{TC}^{\mathcal{R}}$, stabilization time cannot be bounded, except for trivial specifications. However, we show that our solutions are speculative in the sense that their stabilization time in $\mathcal{TC}^{\mathcal{B}}(\Delta)$ is $O(\Delta)$ rounds.

KEYWORDS

self-stabilization, time-varying graphs, leader election, speculation

ACM Reference Format:

Karine Altisen, Stéphane Devismes, Anaïs Durand, Colette Johnen, and Franck Petit. 2021. Self-stabilizing Systems in Spite of High Dynamics. In *International Conference on Distributed Computing and Networking 2021 (ICDCN '21)*, January 5–8, 2021, Nara, Japan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3427796.3427838>

1 INTRODUCTION

Context. Starting from an arbitrary configuration, a *self-stabilizing algorithm* [18] makes a distributed system reach within finite time

This study has been partially supported by the French ANR projects ANR-16-CE40-0023 (DESCARTES) and ANR-16-CE25-0009-03 (ESTATE).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICDCN '21, January 5–8, 2021, Nara, Japan

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8933-4/21/01... \$15.00
<https://doi.org/10.1145/3427796.3427838>

a configuration from which its behavior is correct. Essentially, self-stabilizing algorithms tolerate *transient failures*, since by definition such failures last a finite time (as opposed to crash failures, for example) and their frequency is low (as opposed to intermittent failures). Indeed, the arbitrary initial configuration can be seen as the result of a finite number of transient faults, and after those faults cease, we can expect a sufficiently large time window without any fault so that the system recovers and then exhibits a correct behavior for a long time.

Even though self-stabilization is not inherently suited to handle other failure patterns, *a.k.a.*, *intermittent* and *permanent* failures, several works show that in many cases self-stabilization can still be achieved despite such faults occur. Indeed, strong forms of self-stabilization have been proposed to tolerate permanent failures, *e.g.*, *fault-tolerant self-stabilization* [6] to cope with process crashes, and *strict stabilization* [30] to withstand Byzantine failures. Furthermore, several self-stabilizing algorithms, *e.g.*, [17], withstand intermittent failures such as frequent lost, duplication, or reordering of messages, meaning their convergence is still effective despite such faults continue to often occur in the system. Hence, even if at the first glance guaranteeing a convergence property may seem to be contradictory with a high failure rate, the literature shows that self-stabilization may be a suitable answer even in such cases.

All these aforementioned works assume static communication networks. Nevertheless, self-stabilizing algorithms dedicated to arbitrary network topologies tolerate, up to a certain extent, some topological changes (*i.e.*, the addition or the removal of communication links or nodes). Precisely, if topological changes are eventually detected locally at involved processes and if the frequency of such events is low enough, then they can be considered as transient faults. Actually, a number of works, *e.g.*, [16, 19, 24] (*n.b.*, [16] deals with leader election), use this kind of argument to claim that they are suited for the dynamic context. Furthermore, several approaches, like *superstabilization* [23] and *gradual stabilization* [2], aims at additionally providing countermeasures to efficiently treat topological changes when they are both spatially and timely sparse. However, all these aforementioned approaches, *e.g.*, [2, 16, 19, 23, 24], become totally ineffective when the frequency of topological changes drastically increase, in other words when topological changes are intermittent rather than transient. Actually, in the intermittent case, the network dynamics should be no more considered as an anomaly but rather as an integral part of the system nature.

Several works aim at proposing a general graph-based model to capture the network dynamics. In [10], the network dynamics is represented as a sequence of digraphs called *evolving graphs*. In [13], the topological evolution of the network is modeled by a

(fixed) digraph where the nodes represent participating processes and the edges are communication links that may appear during the lifetime of the network. Each edge is labeled according to its presence during the lifetime of the network. Such digraphs are called *Time-Varying Graphs* (TVGs, for short).

In highly dynamic distributed systems, an expected property is *self-adaptiveness*, *i.e.*, the ability of a system to accommodate to sudden and frequent changes of its environment. By definition, achieving self-stabilization in highly dynamic networks is a suitable answer to self-adaptiveness. *Speculation* [29] is another possible approach for self-adaptiveness. Roughly speaking, it guarantees that the system satisfies its requirements for all executions, but also exhibits significantly better performances in a subset of more probable executions. The main idea behind speculation is that worst possible scenarios are often rare (even unlikely) in practice. So, a speculative algorithm is assumed to self-adapt its performances *w.r.t.* the “quality” of the environment, *i.e.*, the more favorable the environment is, the better the complexity of the algorithm should be. Interestingly, Dubois and Guerraoui [26] have investigated speculation in self-stabilizing, yet static, systems. They illustrate this property with a self-stabilizing mutual exclusion algorithm whose stabilization time is significantly better when the execution is synchronous.

Contribution. We initiate research on self-stabilization in *highly dynamic identified message passing systems where the dynamics is modeled using TVGs* to obtain solutions tolerating both transient faults and high dynamics. In our model, processes can only communicate through local broadcast primitives: at each round, every process can send a common message to its unknown set of current neighbors (if any).

We reformulate the definition of self-stabilization to accommodate TVGs, and investigate the self-stabilizing leader election problem. This problem is fundamental in distributed computing since it allows to synchronize and self-organize a network. Thus, leader election is a basic component in many protocols, *e.g.*, spanning tree constructions, broadcasts, and convergecasts.

We study self-stabilizing leader election in three wide classes of TVGs, respectively denoted by $\mathcal{TC}^B(\Delta)$, $\mathcal{TC}^Q(\Delta)$, and \mathcal{TC}^R , where $\mathcal{TC}^B(\Delta) \subseteq \mathcal{TC}^Q(\Delta) \subseteq \mathcal{TC}^R$: $\mathcal{TC}^B(\Delta)$ is the class of TVGs with temporal diameter bounded by Δ [27], $\mathcal{TC}^Q(\Delta)$ is the class of TVGs with temporal diameter quasi-bounded by Δ (introduced here), and \mathcal{TC}^R is the class of TVGs with recurrent connectivity [13]; this latter class is the most general infinite TVG class introduced so far [12, 13].¹

We first study conditions under which our problem can be solved. Actually, our results show that the assumption on the knowledge of the number n of processes is central. To see this, we introduce the notion of size-ambiguity, which formalizes the fact that some subsets of processes do not share enough initial knowledge on n to detect that the system is not limited to themselves. In other words, such an ambiguity comes from the fact that n is only partially known by the processes (*e.g.*, when processes only know an upper bound on n). Our results show that, despite the existence of unique process identifiers, any deterministic self-stabilizing leader election algorithm working

in the class $\mathcal{TC}^Q(\Delta)$ or \mathcal{TC}^R cannot be size-ambiguous. Hence, to make the problem solvable in those classes, we will assume each process knows the exact value of n .

We then propose self-stabilizing leader election algorithms for the three considered classes. In more detail, we present a self-stabilizing leader election algorithm for Class $\mathcal{TC}^B(\Delta)$ with a stabilization time of at most 3Δ rounds, assuming every process knows Δ , yet using no information on n . This in particular shows that our necessary condition is tight. Then, we propose a self-stabilizing leader election algorithm for Class $\mathcal{TC}^Q(\Delta)$ assuming every process knows Δ and n . In general, stabilization time cannot be bounded in $\mathcal{TC}^Q(\Delta)$; nevertheless we show that the algorithm is speculative since its stabilization time in $\mathcal{TC}^B(\Delta)$ is at most 2Δ rounds. Finally, we propose a self-stabilizing leader election algorithm for Class \mathcal{TC}^R , where only n is known, yet requiring unbounded local memories. Finding a self-stabilizing solution in this class was rather challenging, since there is no guarantee on message delivery timeliness at all (*n.b.*, by definition of the class, there is no bound on the temporal diameter). Again, in general, stabilization time cannot be bounded in \mathcal{TC}^R , yet we show that the algorithm is speculative since its stabilization time in $\mathcal{TC}^B(\Delta)$ is at most $\Delta + 1$ rounds.

Related Work. Ensuring convergence in highly dynamic networks regardless of the initial configuration may seem to be very challenging, even impossible in many cases [9]. However, there are a few works [8, 11, 22] that deal with this issue, yet in widely different models and assumptions than ours.

A recent work [8] deals with the self-stabilizing exploration of a *highly dynamic ring* by a cohort of synchronous robots equipped with visibility sensors, moving actuators, yet no communication capabilities. Note that, contrary to [8], the three classes studied in the present paper never enforce the network to have a particular topology at a given time.

In [11], Cai *et al.* tackles the self-stabilizing leader election problem in highly dynamic systems through the population protocol model. In this model, communications are achieved by atomic rendezvous between pair of *anonymous* processes, where ties are nondeterministically broken. The local broadcast primitive we use here is weaker. Moreover, authors assume global fairness, meaning that every configuration infinitely often reachable is infinitely often reached. We do not make such an assumption here. Actually, Cai *et al.* show that, in their model, self-stabilizing leader election is deterministically solvable if and only if the number of processes n is known, despite processes being *anonymous*. In our model, even with the knowledge of n , (deterministic) self-stabilizing leader election cannot be solved if processes are anonymous.² Moreover, our results show that (maybe surprisingly) even with process identifiers, the knowledge of n is necessary to solve self-stabilizing leader election in $\mathcal{TC}^Q(\Delta)$ and \mathcal{TC}^R .

Finally, Dolev *et al.* [22] assume the system is equipped with a routing algorithm, which allows any two processes to communicate, providing that the sender knows the identifier of the receiver. This blackbox protocol abstracts the dynamics of the system: the dynamics makes it fair lossy, non-FIFO, and duplication-prone. Moreover, the channel capacity is assumed to be bounded. Based on this weak

¹ Considering finite TVGs, *i.e.*, dynamic systems whose lifetime is limited, does not really make sense in the self-stabilizing context, since commonly self-stabilizing algorithms do not terminate [1, 20].

² Every static network is a very particular case of TVG, which belongs to all classes studied here, so the impossibility result of Angluin [3] still applies.

routing algorithm, they build a stronger routing protocol which is reliable, FIFO, and which prevents duplication. We should remark that techniques used here can be reengineered to implement their input black box routing protocol.

Roadmap. In Section 2, we present the computational model. In Section 3, we propose and justify our definition of self-stabilization for highly dynamic environments; we then study the impact of the knowledge of n on the solvability of the self-stabilizing leader election. In the three next sections, we present our algorithms. The last section is dedicated to conclusions and perspectives.

2 PRELIMINARIES

Time-varying Graphs. A *time-varying graph* (TVG for short) [13] is a tuple $\mathcal{G} = (V, E, \mathcal{T}, \rho)$ where V is a (static) set of nodes, E is a (static) set of arcs between pairwise nodes, \mathcal{T} is an interval over \mathbb{N}^* (the set of positive integers) called the *lifetime* of \mathcal{G} , and $\rho : E \times \mathcal{T} \rightarrow \{0, 1\}$ is the *presence* function that indicates whether or not a given arc exists at a given time. We denote by $o_{\mathcal{T}} = \min \mathcal{T}$ the first instant in \mathcal{T} .

From a global viewpoint, the evolution of \mathcal{G} is described as a sequence of digraphs, called *snapshots*: the *snapshot* of \mathcal{G} at time $t \in \mathcal{T}$ is the digraph $G_t = (V, \{e \in E : \rho(e, t) = 1\})$.

Let $[t, t'] \subseteq \mathcal{T}$. The *temporal subgraph* of \mathcal{G} for the interval $[t, t']$, denoted by $\mathcal{G}_{[t, t']}$, is the TVG $(V, E, [t, t'], \rho')$ where ρ' is ρ restricted to $[t, t']$. Roughly speaking, $\mathcal{G}_{[t, t']}$ is itself a TVG that reproduces all the interactions present in the original TVG \mathcal{G} , yet for the time window $[t, t']$.

A *journey* \mathcal{J} can be thought as a path over time from a source p_1 to a destination q_k , i.e., \mathcal{J} is a sequence $(e_1, t_1), (e_2, t_2), \dots, (e_k, t_k)$ where $\forall i \in \{1, \dots, k\}$, $e_i = (p_i, q_i) \in E$ satisfies $\rho(e_i, t_i) = 1$ and $i < k \Rightarrow q_i = p_{i+1} \wedge t_i < t_{i+1}$. We respectively denote by *departure*(\mathcal{J}) and *arrival*(\mathcal{J}) the *starting time* t_1 and the *arrival time* t_k of \mathcal{J} . A *journey from p to q* is a journey whose source is p and destination is q . Let $\mathcal{J}(p, q)$ be the set of journeys in \mathcal{G} from p to q . Let \rightsquigarrow be the binary relation over V such that $p \rightsquigarrow q$ if $p = q$ or there exists a journey from p to q in \mathcal{G} .

The *temporal length* of a journey \mathcal{J} is equal to $\text{arrival}(\mathcal{J}) - \text{departure}(\mathcal{J}) + 1$. Let $t \geq o_{\mathcal{T}} - 1$. By extension, we define the *temporal distance* from p to q at t , denoted by $\hat{d}_{p,t}(q)$, as follows: $\hat{d}_{p,t}(q) = 0$, if $p = q$, $\hat{d}_{p,t}(q) = \min\{\text{arrival}(\mathcal{J}) - t : \mathcal{J} \in \mathcal{J}(p, q) \wedge \text{departure}(\mathcal{J}) > t\}$ otherwise (by convention, we let $\min \emptyset = +\infty$). Roughly speaking, the temporal distance from p to q at time t gives the minimum timespan for p to reach q after t . The *temporal diameter* at t is the maximum temporal distance between any two nodes at t .

We define $\text{ITVG}(\mathcal{G})$ to be the predicate that holds if \mathcal{T} is a right-open interval, in which case \mathcal{G} is said to be an *infinite* TVG; otherwise \mathcal{G} is called a *finite* TVG.

TVG Classes. Let $\mathcal{G} = (V, E, \mathcal{T}, \rho)$ be a TVG. We consider the following four TVG classes.

Class \mathcal{TC} (Connectivity over Time), also denoted by C_3 in [13]: every node can reach all the others at least once through a journey. Formally, $\mathcal{G} \in \mathcal{TC}$ if $\forall p, q \in V, p \rightsquigarrow q$.

Class \mathcal{TC}^R (Recurrent Connectivity), denoted by C_5 in [13]: at any point in time, every node can reach all the others through a journey. Formally, $\mathcal{G} \in \mathcal{TC}^R$ if $\text{ITVG}(\mathcal{G}) \wedge \forall t \in \mathcal{T}, \mathcal{G}_{[t, t+\infty)} \in \mathcal{TC}$.

Class $\mathcal{TC}^B(\Delta)$ with $\Delta \in \mathbb{N}^$ (Bounded Temporal Diameter), denoted by $\mathcal{TC}(\Delta)$ in [27]:* at any point in time, every node is at temporal distance at most Δ from all other nodes, i.e., the temporal diameter is bounded by Δ . Formally, $\mathcal{G} \in \mathcal{TC}^B(\Delta)$ if $\text{ITVG}(\mathcal{G}) \wedge \forall t \in \mathcal{T}, \mathcal{G}_{[t, t+\Delta)} \in \mathcal{TC}$.

Class $\mathcal{TC}^Q(\Delta)$ with $\Delta \in \mathbb{N}^$ (Quasi Bounded Temporal Diameter):* every node is infinitely often at temporal distance at most Δ from each other node. Formally, $\mathcal{G} \in \mathcal{TC}^Q(\Delta)$ if $\text{ITVG}(\mathcal{G}) \wedge \forall p, q \in V, \forall t \in \mathcal{T}, \exists t' \geq t - 1, \hat{d}_{p,t'}(q) \leq \Delta$.

Notice that, $\forall \Delta \in \mathbb{N}^*, \mathcal{TC}^B(\Delta) \subseteq \mathcal{TC}^Q(\Delta) \subseteq \mathcal{TC}^R \subseteq \mathcal{TC}$, by definition. Furthermore, we say that a TVG class C is *recurring* if C only contains infinite TVGs and, for every $\mathcal{G} \in C$ and every $t \geq o_{\mathcal{T}}, \mathcal{G}_{[t, t+\infty)} \in C$. (In other words, every recurring TVG class is suffix-closed.) The three classes we will consider hereafter (i.e., $\mathcal{TC}^R, \mathcal{TC}^B(\Delta), \mathcal{TC}^Q(\Delta)$) are recurring.

Computational Model. We consider the computational model defined in [5, 14]. We assume a *distributed system* made of a set of n *processes*, denoted by V . Each process has a local memory, a local sequential and deterministic algorithm, and message exchange capabilities. We assume that each process p has a unique identifier (ID for short). The identifier of p is denoted by $\text{id}(p)$ and taken in an arbitrary domain IDSET totally ordered by $<$. We assume that each identifier is stored using B bits. In the sequel, we denote by ℓ the process of minimum identifier. Processes are assumed to communicate by message passing through an interconnected network that evolves over the time. The topology of the network is then conveniently modeled by an infinite TVG $\mathcal{G} = (V, E, \mathcal{T}, \rho)$. Processes execute their local algorithms in *synchronous rounds*. For every $i > 0$, the communication network at Round i is defined by $G_{o_{\mathcal{T}}+i-1}$, i.e., the snapshot of \mathcal{G} after $i - 1$ instants elapse from the initial time $o_{\mathcal{T}}$. So, $\forall p \in V$, we denote by $\mathcal{N}(p)^i = \{q \in V : \rho((p, q), o_{\mathcal{T}} + i - 1) = 1\}$, the set of p 's neighbors at Round i . $\mathcal{N}(p)^i$ is assumed to be unknown by process p , whatever the value of i is.

A *distributed algorithm* \mathcal{A} is a collection of n local algorithms $\mathcal{A}(p)$, one per process $p \in V$ (n.b., different processes may have different codes). At any round, each process has a state. The *state* of each process $p \in V$ in \mathcal{A} is defined by the values of its variables in $\mathcal{A}(p)$. We denote by $\mathcal{S}_{\mathcal{A}}^V(p)$ the non-empty set of p 's possible local states in \mathcal{A} . Some variables may be constants in which case their values are predefined. A *configuration* of \mathcal{A} for V is a vector of n components (s_1, s_2, \dots, s_n) , where s_1 to s_n represent the states of the processes in V .

Let γ_0 be the initial configuration of \mathcal{A} for V . For any (synchronous) round $i \geq 1$, the system moves from the current configuration γ_{i-1} to some configuration γ_i , where γ_{i-1} (resp. γ_i) is referred to as the configuration *at the beginning of Round i* (resp. *at the end of Round i*). Such a move is atomically performed by every process $p \in V$ according to the following three steps, defined in its local algorithm $\mathcal{A}(p)$:

- (1) p sends a message consisting of all or a part of its local state in γ_{i-1} using the primitive $\text{SEND}()$,
- (2) using Primitive $\text{RECEIVE}()$, p receives all messages sent by processes in $\mathcal{N}(p)^i$, and
- (3) p computes its state in γ_i .

An *execution* of a distributed algorithm \mathcal{A} in the TVG $\mathcal{G} = (V, E, \mathcal{T}, \rho)$ is an infinite sequence of configurations $\gamma_0, \gamma_1, \dots$ of \mathcal{A} for V such

that $\forall i > 0$, γ_i is obtained by executing a synchronous round of \mathcal{A} on γ_{i-1} based on the communication network at Round i , *i.e.*, the snapshot $G_{o_{\mathcal{T}}+i-1}$.

3 SELF-STABILIZATION IN HIGHLY DYNAMIC ENVIRONMENTS

Definition. In the following, we define a *specification* as a predicate over configuration sequences.

Self-stabilization has been originally defined for static networks. In the reference book of Dolev [20], self-stabilization is defined as follows. An algorithm is *self-stabilizing for a specification SP* if there exists a set of so-called *legitimate* configurations satisfying the following two properties: (1) every execution of the algorithm in the considered system eventually reaches a legitimate configuration (*Convergence*); and (2) every possible execution suffix starting from a legitimate configuration satisfies *SP* (*Correctness*). Below, we accommodate this definition to highly dynamic environments.

Definition 3.1 (Self-stabilization). An algorithm \mathcal{A} is *self-stabilizing* for the specification *SP* on a class C of infinite TVGs if for every set of processes V , there exists a subset of configurations \mathcal{L} of \mathcal{A} for V , called *legitimate configurations*, such that:

- (1) for every $\mathcal{G} \in C$ with set of processes V and every configuration γ of \mathcal{A} for V , every execution of \mathcal{A} in \mathcal{G} starting from γ contains a legitimate configuration $\gamma' \in \mathcal{L}$ (*Convergence*), and
- (2) for every $\mathcal{G} \in C$ with set of processes V , every $t \geq o_{\mathcal{T}}$, every legitimate configuration $\gamma \in \mathcal{L}$, and every execution e in $\mathcal{G}_{[t,+\infty)}$ starting from γ , *SP*(e) holds (*Correctness*).

The *length of the stabilization phase* of an execution e is the length of its maximum prefix containing no legitimate configuration. The *stabilization time* in rounds is the maximum length of a stabilization phase over all possible executions.

REMARK 1. *In the case of a recurring class of TVG, the definition of self-stabilization for an algorithm \mathcal{A} and a specification *SP* can be slightly simplified. Indeed, the correctness property can be equivalently rewritten as follows: given a set of processes V and a set of configurations \mathcal{L} on V , for every $\mathcal{G} \in C$ with set of processes V , every legitimate configuration $\gamma \in \mathcal{L}$, and every execution e in \mathcal{G} starting from γ , *SP*(e) holds (Recurring-Correctness).*

It is worth noting that Definition 3.1, as the one given in the reference book of Dolev [20], does not include the notion of *closure*: intuitively, a set of configurations \mathcal{S} is *closed* if every step of the algorithm starting in a configuration of \mathcal{S} leads to a configuration of \mathcal{S} ; see Definition 3.2 for a formal definition. Now, when dealing with high-level models (such as the atomic-state model), closure is most of the time present in definitions of self-stabilization. However, in the more practical message passing model, closure is usually simply given up; see, *e.g.*, [4, 25, 31]. Even if this absence is never motivated, this may be explained by the lack of functional significance of the closure property as compared to the convergence and correctness properties. Closure is rather a nice property that often helps to write elegant, and so simpler, proofs. Moreover, closure may be sometimes too restrictive, as we will show in Theorem 3.4 for example. Below, we reformulate closure in the context of TVGs.

Definition 3.2 (Closure). Let \mathcal{A} be a distributed algorithm, C be an infinite TVG class, V be a set of processes, and \mathcal{S} be a set of configurations of \mathcal{A} for V . \mathcal{S} is *closed* in C if for every $\mathcal{G} \in C$ with set of processes V , every $t \geq o_{\mathcal{T}}$, and every configuration $\gamma \in \mathcal{S}$, every execution of \mathcal{A} in $\mathcal{G}_{[t,+\infty)}$ starting from γ only contains configurations of \mathcal{S} .

REMARK 2. *Again, when the considered class of TVGs is recurring, the definition of closure can be slightly simplified. If \mathcal{A} is a distributed algorithm, C is a recurring TVG class, V is a set of processes, and \mathcal{S} is a set of configurations of \mathcal{A} for V , then \mathcal{S} is closed in C if for every $\mathcal{G} \in C$ with set of processes V and every configuration $\gamma \in \mathcal{S}$, every execution of \mathcal{A} in \mathcal{G} starting from γ only contains configurations of \mathcal{S} .*

Self-stabilizing Leader Election. The *leader election* problem consists in distinguishing a single process in the system. In identified networks, the election usually consists in making the processes agree on one of the identifiers held by processes. The identifier of the elected process is then stored at each process p in an output variable, denoted here by $lid(p)$. We call *fake ID* any value $v \in IDSET$ (recall that *IDSET* is the definition domain of the identifiers) such that v is not assigned as a process identifier in the system, *i.e.*, there is no process $p \in V$ such that $id(p) = v$. In the self-stabilizing context, the output variables lid may be initially corrupted; in particular some of them may be initially assigned to *fake IDs*. Despite such fake IDs, the goal of a self-stabilizing algorithm is to make the system converge to a configuration from which a unique process is forever adopted as leader by all processes, *i.e.*, $\exists p \in V$ such that $\forall q \in V, lid(q) = id(p)$ forever. Hence, the leader election specification SP_{LE} can be formulated as follows: a sequence of configurations $\gamma_0, \gamma_1, \dots$ satisfies SP_{LE} if and only if $\exists p \in V$ such that $\forall i \geq 0, \forall q \in V$, the value of $lid(q)$ in configuration γ_i is $id(p)$. In the sequel, we say that an algorithm is a *self-stabilizing leader election algorithm* for the class of infinite TVG C if it is self-stabilizing for SP_{LE} on C .

Knowledge of n and Closure in $\mathcal{TC}^{\mathcal{B}}(\Delta)$. We now advocate that closure of legitimate configurations may be cumbersome in $\mathcal{TC}^{\mathcal{B}}(\Delta)$ since to achieve it, any (deterministic) self-stabilizing leader election algorithm somehow requires the exact knowledge of n (the number of processes), *i.e.*, even partial knowledge such as an upper bound on n is not sufficient; see Theorem 3.4. To that goal, we need to first define what we mean by *not exactly knowing n* .

When an algorithm \mathcal{A} uses the number of its processes, this means that this information is given as an input in the local state of each process. So, the definition of the set of possible local states of each process is adjusted according to the size of the system it belongs to. Conversely, if an algorithm \mathcal{A} does not know its exact size, this means that there are sizes of systems that cannot be distinguished by part of its processes using their local inputs (and so their possible local states). More precisely, for a given set of processes V executing \mathcal{A} , there should exist a size $k < |V|$ for which the processes of any k -subset U of V do not share enough initial information to distinguish whether the system is made of the process-set V or U . Below, we formalize this intuitive idea by the notion of *size-ambiguity*.

Definition 3.3 (Size-Ambiguity). Let V be a set of processes and $k \in \mathbb{N}$. A distributed algorithm \mathcal{A} is (k, V) -*ambiguous* if $0 < k < |V|$ and for every $U \subset V$ such that $|U| = k$ and every $p \in U$, $S_{\mathcal{A}}^U(p)$

is well-defined and $S_{\mathcal{A}}^U(p) = S_{\mathcal{A}}^V(p)$. We simply say that \mathcal{A} is *size-ambiguous* if there exists V and k such that \mathcal{A} is (k, V) -ambiguous.

Consider now a few examples. First, if each process has a constant input whose value is the number n of processes in the system (*i.e.*, each process "exactly knows n "), then from our definition, the algorithm is not size-ambiguous since, in this case, the set of possible local states of any process differs from one size of system to another, at least because of the input storing n . Conversely, if the processes do not know the exact number of processes but its parity, then we can choose any set V of at least three processes and any positive value $k < |V|$ with same parity as $|V|$: for every subset U of V such that $|U| = k$, the constant input giving the parity will be the same at each process of U whether running its algorithm in a TVG with process-set V or U . Consequently, every process $p \in U$ will have the same set of possible local states in both TVGs; hence the size-ambiguity. Similarly, an algorithm is size-ambiguous if each process p only knows an upper bound $N_p \geq 2$ on the number of processes in the TVG (*n.b.*, processes may not know the same bound) since the property can be achieved with any set V of at least two processes and any value k such that $0 < k < |V|$.

THEOREM 3.4. *Let \mathcal{A} be a deterministic self-stabilizing leader election algorithm for $\mathcal{TC}^{\mathcal{B}}(\Delta)$ (with $\Delta \geq 2$), V be a set of processes, \mathcal{L} be a set of legitimate configurations of \mathcal{A} for V , and $k \in \mathbb{N}$. \mathcal{L} is not closed in $\mathcal{TC}^{\mathcal{B}}(\Delta)$ if \mathcal{A} is (k, V) -ambiguous.*

PROOF. Let $n = |V|$ and $V = \{p_0, \dots, p_{n-1}\}$. Assume, by the contradiction, that \mathcal{L} is closed in $\mathcal{TC}^{\mathcal{B}}(\Delta)$. Let $\mathcal{G} = (V, E, \mathcal{T}, \rho)$ be an infinite TVG such that

- (1) $E = \{(p_i, p_j) : p_i, p_j \in V \wedge i \neq j\}$
- (2) $\forall t \geq o_{\mathcal{T}}, \forall (p_i, p_j) \in E, \rho((p_i, p_j), t) = 1$ if and only if either t is odd, or $i \notin \{\frac{t}{2} \bmod n, \dots, (\frac{t}{2} + n - k - 1) \bmod n\}$ and $j \notin \{\frac{t}{2} \bmod n, \dots, (\frac{t}{2} + n - k - 1) \bmod n\}$.

Notice first that, $\forall t \geq o_{\mathcal{T}}$, the snapshot G_t of \mathcal{G} is fully connected when t is odd. Consequently, \mathcal{G} belongs to $\mathcal{TC}^{\mathcal{B}}(\Delta)$, with $\Delta \geq 2$. Then, by definition, we have:

Claim 1: *For every $x \in \{0, \dots, n-1\}$ and every $i \geq 0$, in the snapshot $G_{t_{x,i}}$ of \mathcal{G} at time $t_{x,i} = 2((i + o_{\mathcal{T}}).n + x)$, the set $V \setminus \{p_x, \dots, p_{(x+n-k-1) \bmod n}\}$ is fully connected and all processes in the set $\{p_x, \dots, p_{(x+n-k-1) \bmod n}\}$ are isolated.*

Let $\gamma \in \mathcal{L}$ and $p_\ell \in V$ be the elected process in γ . $\forall i \geq 0$, we inductively define Configuration γ^i as follows. $\gamma^0 = \gamma$. $\forall i > 0$, γ^i is the configuration at the end of the first round of the execution of \mathcal{A} in $\mathcal{G}_{[t_{\ell,i}, +\infty)}$ starting from γ^{i-1} . Since $\mathcal{TC}^{\mathcal{B}}(\Delta)$ is recurring, we can use the closure of \mathcal{L} to show by induction Claim 2 below (*n.b.*, Claim 2 is the only result of the proof where closure of \mathcal{L} is used).

Claim 2: $\forall i \geq 0$, γ^i is legitimate and $\forall p_j \in V, \text{lid}(p_j) = p_\ell$ in γ^i .

Let $V^- = V \setminus \{p_\ell, \dots, p_{(\ell+n-k-1) \bmod n}\}$ and $E^- = \{(p_i, p_j) : p_i, p_j \in V^- \wedge i \neq j\}$. Let $\mathcal{G}^- = (V^-, E^-, \mathcal{T}^-, \rho^-)$ be the infinite TVG having k processes such that $\forall t \geq o_{\mathcal{T}}, \forall (p_i, p_j) \in E^-$, we have $\rho((p_i, p_j), t) = 1$. In other word, \mathcal{G}^- is a static fully connected network. Consequently, \mathcal{G}^- in particular belongs to $\mathcal{TC}^{\mathcal{B}}(\Delta)$ with $\Delta \geq 2$ (actually, it also belongs to $\mathcal{TC}^{\mathcal{B}}(\Delta)$ with $\Delta = 1$). Let γ_0^- be the configuration of \mathcal{A} for V^- where each process has the same state as in the configuration γ^0 (such a configuration exists by definition of (k, V) -ambiguity). We now consider the execution $e = \gamma_0^-, \dots, \gamma_i^-$, ...

of \mathcal{A} in \mathcal{G}^- starting from the configuration γ_0^- . Claim 3 below can be proven by induction on i .

Claim 3: $\forall i \geq 0$, the state of each process in V^- in γ_i^- is the same as in γ^i .

By Claims 2 and 3, for every process p_j in V^- , in every configuration γ_i^- , we have $\text{lid}(p_j) = p_\ell \notin V^-$, *i.e.*, $\text{lid}(p_j)$ is a fake ID (for V^-). Hence, no suffix of e satisfies SP_{LE} . As a consequence, \mathcal{A} is not a self-stabilizing leader election algorithm for $\mathcal{TC}^{\mathcal{B}}(\Delta)$ (with $\Delta \geq 2$), a contradiction. \square

REMARK 3. *The condition $\Delta \geq 2$ is necessary in Theorem 3.4, indeed if $\Delta = 1$, there is a trivial deterministic self-stabilizing leader election algorithm for $\mathcal{TC}^{\mathcal{B}}(\Delta)$ that does not need information on n and has a closed set of legitimate configurations: it simply consists of all processes sending their own IDs at each round; since $\Delta = 1$, all processes learn the exact set of all IDs present in the network at each round and just have to choose, e.g., the smallest one, $\text{id}(\ell)$. The legitimate configurations are then all configurations where every process p satisfies $\text{lid}(p) = \text{id}(\ell)$.*

According to Theorem 3.4, the set of legitimate configurations of our solution for $\mathcal{TC}^{\mathcal{B}}(\Delta)$ (Algorithm 1) is not closed, since by making no assumption on n , this algorithm is size-ambiguous. The contrapositive of Theorem 3.4 is given in Corollary 3.5. This latter justify the need of the exact knowledge of the number of processes to obtain a closed set of legitimate configurations in a deterministic self-stabilizing leader election algorithm for $\mathcal{TC}^{\mathcal{B}}(\Delta)$, with $\Delta \geq 2$.

COROLLARY 3.5. *Let \mathcal{A} be a deterministic self-stabilizing leader election algorithm for $\mathcal{TC}^{\mathcal{B}}(\Delta)$ (with $\Delta \geq 2$), V be a set of processes, and \mathcal{L} be a set of legitimate configurations of \mathcal{A} for V . If \mathcal{L} is closed in $\mathcal{TC}^{\mathcal{B}}(\Delta)$, then \mathcal{A} should not be size-ambiguous.*

REMARK 4. *The scheme used in the proof of Theorem 3.4 can be adapted to handle other problems consisting in computing a constant output whose value depends on the set of processes. For example, one can show that no deterministic self-stabilizing size-ambiguous algorithm for $\mathcal{TC}^{\mathcal{B}}(\Delta)$ can both compute the exact number of processes and achieve the closure of its legitimate configurations.*

Knowledge of n and Closure in $\mathcal{TC}^{\mathcal{Q}}(\Delta)$. We now show that every execution of a self-stabilizing algorithm for a recurring specification in $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ necessarily converges to a closed set of legitimate configurations; see Theorem 3.8. Consequently, no deterministic self-stabilization leader election algorithm for $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ can be size-ambiguous (Theorem 3.10 and Corollary 3.11); justifying why algorithms in Sections 5 and 6 assume the exact knowledge of n .

Definition 3.6 (Recurring Specification). We say that a specification SP is *recurring* if for every sequence of configurations $\gamma_0, \gamma_1, \dots$, $SP(\gamma_0, \gamma_1, \dots) \Rightarrow (\forall i \geq 0, SP(\gamma_i, \gamma_{i+1}, \dots))$.

SP_{LE} (as most of specifications used in self-stabilization) is a recurring specification.

Definition 3.7 (Sequential Composition). Let $\mathcal{G} = (V, E, \mathcal{T}, \rho)$ be an infinite TVG and $\mathcal{G}' = (V', E', [a, b], \rho')$ be a finite TVG. The *sequential composition* of \mathcal{G}' and \mathcal{G} , denoted by $\mathcal{G}' \triangleright \mathcal{G}$, is the infinite TVG $\mathcal{G}'' = (V'', E'', \mathcal{T}'', \rho'')$ such that $V'' = V \cup V'$, $E'' = E \cup E'$, $\mathcal{T}'' = [a, +\infty)$, and $\forall e \in E''$,

- $\forall t \in [a, b], \rho''(e, t) = 1$ if and only if $e \in E' \wedge \rho'(e, t) = 1$, and
- $\forall t > b, \rho''(e, t) = 1$ if and only if $e \in E \wedge \rho(e, o_{\mathcal{T}} + t - b - 1) = 1$.

PROPERTY 1. Let $\mathcal{G} = (V, E, \mathcal{T}, \rho) \in \mathcal{TC}^Q(\Delta)$ and $\mathcal{G}' = (V', E', \mathcal{T}', \rho')$ be a finite TVG. If $V' \subseteq V$, $\mathcal{G}' \triangleright \mathcal{G} \in \mathcal{TC}^Q(\Delta)$.

THEOREM 3.8. Let SP be a recurring specification, \mathcal{A} be a self-stabilizing algorithm for SP on $\mathcal{TC}^Q(\Delta)$, and V be a set of processes. There exists a set of legitimate configurations of \mathcal{A} for V which is closed in $\mathcal{TC}^Q(\Delta)$.

PROOF. Assume, by the contradiction, that every set of legitimate configurations of \mathcal{A} for V is not closed in $\mathcal{TC}^Q(\Delta)$. Let \mathcal{L} be the set of legitimate configurations of \mathcal{A} for V defined as follows: for every configuration γ of \mathcal{A} for V , $\gamma \in \mathcal{L}$ if and only if for every $\mathcal{G} \in \mathcal{TC}^Q(\Delta)$ with set of processes V and every execution e of \mathcal{A} in \mathcal{G} starting from γ , $SP(e)$ holds. Since \mathcal{L} is not closed in $\mathcal{TC}^Q(\Delta)$, there exists $\gamma_0 \in \mathcal{L}$, $\mathcal{G} = (V, E, \mathcal{T}, \rho) \in \mathcal{TC}^Q(\Delta)$ with V as set of processes, and an execution $\gamma_0, \dots, \gamma_i, \dots$ in \mathcal{G} starting from γ_0 which contains a configuration $\gamma_i \notin \mathcal{L}$. By definition of \mathcal{L} , there exists $\mathcal{G}' \in \mathcal{TC}^Q(\Delta)$ with set of processes V and an execution e' in \mathcal{G}' starting from γ_i such that $\neg SP(e')$ (otherwise γ_i should be in \mathcal{L}). Now, $\mathcal{G}_{[o_{\mathcal{T}}, o_{\mathcal{T}} + i - 1]} \triangleright \mathcal{G}' \in \mathcal{TC}^Q(\Delta)$, by Property 1. Consequently, $\gamma_0, \dots, \gamma_{i-1}, e'$ is an execution of \mathcal{A} in $\mathcal{TC}^Q(\Delta)$ that starts from γ_0 and violates SP since $\neg SP(e')$ and SP is recurring. By the correctness property of the self-stabilizing definition (see Remark 1), γ_0 cannot be a legitimate configuration, a contradiction. \square

COROLLARY 3.9. Let \mathcal{A} be any self-stabilizing leader election algorithm for $\mathcal{TC}^Q(\Delta)$ and V be a set of processes. There exists a set of legitimate configurations of \mathcal{A} for V which is closed in $\mathcal{TC}^Q(\Delta)$.

Since $\mathcal{TC}^B(\Delta) \subseteq \mathcal{TC}^Q(\Delta)$, from Corollaries 3.5 and 3.9, we deduce Theorem 3.10 and Corollary 3.11:

THEOREM 3.10. No deterministic self-stabilizing leader election algorithm for $\mathcal{TC}^Q(\Delta)$, with $\Delta \geq 2$, can be size-ambiguous.

COROLLARY 3.11. No deterministic self-stabilizing leader election algorithm for \mathcal{TC}^R can be size-ambiguous.

REMARK 5. Like in Remark 4, one can show, for example, that no deterministic self-stabilizing size-ambiguous algorithm for $\mathcal{TC}^Q(\Delta)$ (resp. \mathcal{TC}^R) can compute the exact number of processes.

4 CLASS $\mathcal{TC}^B(\Delta)$ WITH Δ KNOWN

Overview of Algorithm 1. Each process p maintains two variables: the output $lid(p)$ will eventually contain the ID of the leader; $tll(p)$ represents the *degree of mistrust* of p in $lid(p)$ and allows to eliminate messages containing fake IDs. The value $tll(p)$ increases at each round if p does not receive a message; otherwise it is updated thanks to the received messages. The value of $tll(p)$ can increase up to $2\Delta - 1$. Process p never increases $tll(p)$ from $2\Delta - 1$ to 2Δ ; instead it locally *resets* and declares itself as the leader: $lid(p) := id(p)$ and $tll(p) := 0$ (see Line u1).

At each round i , p first sends its leader ID together with its degree of mistrust; see Line 2. Then, p selects the received message $\langle id, tll \rangle$

which is minimum using the lexicographic order (*i.e.*, the message the lowest ID and with the lowest tll to break ties, if any; see Line 7). If id is smaller than $lid(p)$, p updates its leader $lid(p)$; see Line 8. If $id = lid(p)$, it updates the $tll(p)$ by taking the smallest value between $tll(p)$ and tll (in this way, p may decrease its mistrust in $lid(p)$; see Line 11). In either case, $tll(p)$ is then incremented if $lid(p) \neq id(p)$. Finally, if $lid(p) \geq id(p)$, p systematically resets; see Line 15. So, if p believes to be the leader at the end of Round i (*i.e.*, $lid(p) = id(p)$), then it sends its own ID together with a degree of mistrust 0 at the beginning of the next round, $i + 1$.

The reset mechanism allows to remove all fake IDs within at most 2Δ rounds. From that time, the lower ID process, ℓ , satisfies $(lid(\ell), tll(\ell)) = (id(\ell), 0)$ forever (Corollary 4.2). So, after 2Δ rounds, ℓ sends $\langle id(\ell), 0 \rangle$ at each round and all processes will receive messages $\langle id(\ell), d \rangle$, with $d \leq \Delta < 2\Delta$ (since $\Delta \in \mathbb{N}^*$), at least every Δ rounds since the temporal diameter is upper bounded by Δ . Thus, within at most Δ additional rounds, they will all adopt ℓ as leader and never more reset, ensuring that ℓ will remain the leader forever (Lemma 4.4). Hence, Algorithm 1 is a self-stabilizing leader election for $\mathcal{TC}^B(\Delta)$ and its stabilization time is at most 3Δ rounds (Corollary 4.5).

Algorithm 1: Self-stabilizing leader election for $\mathcal{TC}^B(\Delta)$, for each process p .

Inputs:
 $\Delta \in \mathbb{N}^*$: upper bound on the temporal diameter
 $id(p) \in IDSET$: ID of p

Local Variables:
 $lid(p) \in IDSET$: ID of the leader
 $tll(p) \in \{0, \dots, 2\Delta - 1\}$: degree of mistrust in $lid(p)$

Macros:
 $updateTTL(v)$:

```

u1:   if  $v \geq 2\Delta$  then // Reset
u2:   |    $lid(p) := id(p)$ 
u3:   |    $tll(p) := 0$ 
u4:   |   else if  $lid(p) \neq id(p)$  then  $tll(p) := v$ 

```

1: **Repeat Forever**
2: SEND($\langle lid(p), tll(p) \rangle$)
3: mailbox := RECEIVE()
4: if mailbox = \emptyset then
5: | $updateTTL(tll(p) + 1)$
6: else
7: | $\langle lid, tll \rangle := \min\{\text{messages in mailbox}\}$
8: | if $lid < lid(p)$ then
9: | | $lid(p) := lid$
10: | | $updateTTL(tll + 1)$
11: | else if $lid = lid(p)$ then
12: | | $updateTTL(\min(tll(p), tll) + 1)$
13: | else
14: | | $updateTTL(tll(p) + 1)$
15: if $lid(p) \geq id(p)$ then // Reset
16: | $lid(p) := id(p)$
17: | $tll(p) := 0$

Self-stabilization and Complexity. First, by definition of the algorithm, the next remark follows.

REMARK 6. Since the end of the first round, $\forall p \in V$, we have $lid(p) \leq id(p) \wedge (lid(p) = id(p) \Rightarrow tll(p) = 0)$.

Lemma 4.1 below can be proven by induction on i .

LEMMA 4.1. *Let f be a fake ID. For every $i \geq 1$, at the beginning of Round i , $\forall p \in V, lid(p) = f \Rightarrow ttl(p) \geq i - 1$.*

Lemma 4.1 implies that for every $i > 0$ and every fake ID f , $\forall p \in V, lid(p) = f \Rightarrow ttl(p) \geq i$ at the end of Round i . We define a *quasi-legitimate* configuration of Algorithm 1 as any configuration where $lid(\ell) = id(\ell)$ and $ttl(\ell) = 0$ and there is no fake ID in the system (i.e., $\forall p \in V, lid(p)$ is not a fake ID). So, from Lemma 4.1 and thanks to the reset mechanism of Algorithm 1, we deduce the following corollary.

COROLLARY 4.2. *At the end of Round 2Δ , the configuration is quasi-legitimate.*

The proof of the next lemma consists in showing that for every set of processes V , the set of quasi-legitimate configurations of Algorithm 1 for V is closed in $\mathcal{TC}^{\mathcal{B}}(\Delta)$.

LEMMA 4.3. *Let e be an execution of Algorithm 1 in an arbitrary TVG that starts from a quasi-legitimate configuration. The configuration reached at the end of every round of e is quasi-legitimate.*

PROOF. Consider any step from γ to γ' such that γ is quasi-legitimate. First, since γ contains no fake ID, no message containing a fake ID can be sent in the step from γ to γ' , and γ' contains no fake ID too. Moreover, $id(\ell)$ is the smallest non-fake ID. So, $\forall p \in V, lid(p) \geq id(\ell)$ in γ' . By Remark 6, we conclude that $lid(\ell) = id(\ell)$ and $ttl(\ell) = 0$ in γ' . Hence, γ' is quasi-legitimate. \square

A process p has a *legitimate state* iff $lid(p) = id(\ell)$, $ttl(p) \leq \Delta$, and $p = \ell \Rightarrow ttl(p) = 0$. We define a *legitimate configuration* of Algorithm 1 as any configuration where every process has a legitimate state. By definition, every legitimate configuration is also quasi-legitimate.

LEMMA 4.4. *Let \mathcal{G} be a TVG of Class $\mathcal{TC}^{\mathcal{B}}(\Delta)$, $t \geq o_{\mathcal{T}}$, and e be an execution of Algorithm 1 in $\mathcal{G}_{[t,+\infty)}$ starting in a quasi-legitimate configuration. For every $r \geq \Delta$, the configuration at the end of Round r in e is legitimate.*

PROOF. First, remark that for every $j > 0$, the communication network at Round j in e is G_{t+j-1} . Then, the proof of the lemma is based on the claim below.

Claim (*): *for every $i \geq 0$, $d \geq 0$, every process p such that $\hat{d}_{\ell, t+i-1}(p) \leq d$ satisfies: $\forall j \in \{1, \dots, \Delta - \hat{d}_{\ell, t+i-1}(p) + 1\}$, $lid(p) = id(\ell)$ and $ttl(p) \leq \hat{d}_{\ell, t+i-1}(p) + j - 1$ at the beginning of Round $(i + j + \hat{d}_{\ell, t+i-1}(p))$ of e .*

Let $r \geq \Delta \in \mathbb{N}^*$. We now apply Claim (*) to $d = \Delta$ so that every process p is taken into account by the claim: with $i = r - \Delta$, $j = \Delta - \hat{d}_{\ell, t+i-1}(p) + 1$, we obtain that $lid(p) = id(\ell)$ and $ttl(p) \leq \Delta$ at the beginning of Round $r + 1$; in addition, $ttl(\ell) = 0$ at the beginning of Round $r + 1$, by Remark 6. Hence, the configuration at the end of Round r is legitimate. \square

As a direct consequence of Corollary 4.2 and Lemma 4.4, we obtain the convergence.

COROLLARY 4.5. *Let \mathcal{G} be a TVG of $\mathcal{TC}^{\mathcal{B}}(\Delta)$. For every $i \geq 3\Delta$, at the end of Round i of any execution of Algorithm 1 in \mathcal{G} , the configuration is legitimate.*

LEMMA 4.6. *Let \mathcal{G} be a TVG of $\mathcal{TC}^{\mathcal{B}}(\Delta)$, $t \geq o_{\mathcal{T}}$, and e be an execution of Algorithm 1 for $\mathcal{G}_{[t,+\infty)}$ starting in a legitimate configuration. For every $r \in \{1, \dots, \Delta - 1\}$, the configuration e has reached at the end of Round r satisfies $lid(p) = id(\ell)$ and $ttl(p) \leq \Delta + r$, for every process p .*

PROOF. First, the lemma trivially holds for $\Delta \leq 1$. So, we now show by induction on r that the lemma holds in the case where $\Delta > 1$. For the base case, at the beginning of Round 1, $\forall p \in V, lid(p) = id(\ell)$ and $ttl(p) \leq \Delta$ as the first configuration of e is legitimate. According to the algorithm, at the end of Round 1, $\forall p \in V, lid(p) = id(\ell)$ and $ttl(p) \leq \Delta + 1 < 2\Delta$ (since $\Delta > 1$).

Induction step: let $i \in \{2, \dots, \Delta - 1\}$. At the end of Round $i - 1$, hence at the beginning of Round i , $\forall p \in V, lid(p) = id(\ell)$ and $ttl(p) \leq \Delta + i - 1 < 2\Delta - 1$, by induction hypothesis. According to the algorithm, and since $\Delta + i < 2\Delta$, no process can reset. Hence, at the end of Round i , $\forall p \in V, lid(p) = id(\ell)$ and $ttl(p) \leq \Delta + i$, and we are done. \square

Theorem 4.7 below is a direct consequence of Lemmas 4.4 and 4.6.

THEOREM 4.7. *For every $\mathcal{G} = (V, E, \mathcal{T}, \rho) \in \mathcal{TC}^{\mathcal{B}}(\Delta)$, for every legitimate configuration γ of Algorithm 1 for V , the execution of Algorithm 1 in \mathcal{G} starting from γ satisfies SP_{LE} .*

By Corollary 4.5 and Theorem 4.7, we have the next corollary.

COROLLARY 4.8. *Algorithm 1 is a self-stabilizing leader election algorithm for $\mathcal{TC}^{\mathcal{B}}(\Delta)$. Its stabilization time is at most 3Δ rounds. It requires $O(B + \log \Delta)$ bits per process and messages of size $O(B + \log \Delta)$ bits.*

5 CLASS $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ WITH Δ AND n KNOWN

Overview of Algorithm 2. Each process p collects IDs in its variable $members(p)$. Actually, $members(p)$ is a (FIFO) queue containing at most n pairs $\langle id, t \rangle$, where id is an identifier and t is a timestamp, i.e., an integer value less than or equal to Δ . (We denote by $members(p)[id]$ the timestamp associated to the identifier id belonging to $members(p)$.)

At each round i , p sends all pairs $\langle id, t \rangle$ of $members(p)$ such that $t < \Delta$ at the end of Round $i - 1$ (Line 2). (The timestamps allow to eventually remove all fake IDs.) Then, p updates $members(p)$ by calling function *insert* on each received pair $\langle id, t \rangle$ such that $id \neq id(p)$ (Lines 4-5).

The insertion function *insert* works as follows: if id already appears in $members(p)$, then the old pair tagged with id is removed first from the queue (Line i1), and in either case, $\langle id, t \rangle$ is appended at the tail of the queue (Lines i1 and i4). In particular, since the size of $members(p)$ is limited, if the queue is full, its head is removed to make room for the new value (Line i3). Using this FIFO mechanism, initial spurious values eventually vanish from $members(p)$.

After all received pairs have been managed, the timestamps of all pairs in the queue are incremented (Line 6) and then, $\langle id(p), 0 \rangle$ is systematically inserted at the tail of the queue (Line 7). This mechanism ensures two main properties. First, every timestamp associated to a fake ID in a variable $members$ is eventually forever greater than or equal to Δ (Lemma 5.1); and consequently, eventually no message containing fake IDs is sent (Corollary 5.2). Second, by

definition of $\mathcal{TC}^Q(\Delta)$, for every two distinct processes p and q , there are journeys of length at most Δ infinitely often, so each process p regularly receives messages containing $id(q)$ with timestamps smaller than Δ . Thus, eventually $members(p)$ exactly contains all IDs of the networks (Lemma 5.4). Now, at the end of each round, p updates its leader variable with the smallest ID in $members(p)$ (Line 7). Hence, the process of lowest ID, ℓ , is eventually elected.

Algorithm 2: Self-stabilizing leader election for $\mathcal{TC}^Q(\Delta)$ for each process p .

Inputs:

- $n \in \mathbb{N}$: number of processes
- $\Delta \in \mathbb{N}^*$: recurrent bound on the temporal distance between processes
- $id(p) \in IDSET$: ID of p

Local Variables:

- $members(p)$: queue of at most n elements contains pairs $\langle id, t \rangle \in IDSET \times \{0, \dots, \Delta\}$
- $lid(p) \in IDSET$: ID of the leader

Macros:

$insert(p, \langle id, t \rangle)$:

```

11: if  $\exists t', \langle id, t' \rangle \in members(p)$  then remove  $\langle id, t' \rangle$  from
    members(p) ; push  $\langle id, \min(t, t') \rangle$  at the tail of members(p)
12: else
13:   if  $|members(p)| = n$  then remove the head of members(p)
14:   push  $\langle id, t \rangle$  at the tail of members(p)

```

Repeat Forever

```

1: SEND( $\{\langle id, t \rangle \in members(p) : t < \Delta\}$ )
2: mailbox := RECEIVE()
3: for all pair  $\langle id, t \rangle$  in a message of mailbox do
4:   if  $id \neq id(p)$  then  $insert(p, \langle id, t \rangle)$ 
5: for all  $\langle id, t \rangle \in members(p) : t < \Delta$  do  $members(p)[id] ++$ 
6:    $insert(p, \langle id(p), 0 \rangle)$  ;  $lid(p) := \min\{id : \langle id, \_ \rangle \in members(p)\}$ 
7:

```

Self-stabilization. Lemma 5.1 below can be proven by induction on i .

LEMMA 5.1. *Let f be a fake ID. For every $i \geq 1$, at the beginning of Round i , the following property holds: $\forall p \in V$ if f is in $members(p)$, then $members(p)[f] \geq i - 1$.*

Since a process p does not send a pair $\langle id, t \rangle$ of $members(p)$ with $t \geq \Delta$, we have:

COROLLARY 5.2. *In any round $\Delta + i$ with $i \geq 1$, no process receives a message containing fake IDs.*

LEMMA 5.3. $\forall p, q \in V$, if $id(q)$ is inserted into $members(p)$ during Round $\Delta + i$ with $i \geq 1$, $id(q)$ remains into $members(p)$ forever.

PROOF. If an ID id is in $members(p)$, it can only be removed from $members(p)$ if function $insert(p, \langle id', t \rangle)$ is called and one of the following two situations occurs:

- Line 11: if $id = id'$ but in this case id is immediately added at the tail of $members(p)$,
- Line 13: if $id \neq id'$, id is the head of the queue, and the size of $members(p)$ is already n .

After id is inserted (at tail) into $members(p)$, it requires the insertion of n different IDs that are not into $members(p)$ (and that are different from id) in order to get id at the head of the queue and remove it. If id is inserted during Round $\Delta + i$, it is not a fake ID and the only

other IDs that can be inserted into $members(p)$ are IDs of processes in V since p will not receive any fake ID (Corollary 5.2). Thus, at most $n - 1$ IDs different from id can be inserted after the insertion of id . Hence, id cannot be removed from $members(p)$. \square

By definition of class $\mathcal{TC}^Q(\Delta)$, for every pair of processes p and q , there exists $t \geq \Delta$ such that $\hat{d}_{q, o_{\mathcal{T}}+t-1}(p) \leq \Delta$. We denote by $t(q, p)$ the minimum value t that satisfies the above property, namely $t(q, p)$ represents the first date after $\Delta + o_{\mathcal{T}} - 1$ (i.e., after Δ rounds) from which q can broadcast an information to p in no more than Δ rounds.

LEMMA 5.4. $\forall p, q \in V$, by the end of Round $t(q, p) + \Delta$, $id(q)$ is in $members(p)$ forever.

PROOF. Let $q \in V$. Remark, first, that $id(q) \in members(q) \wedge members(q)[q] = 0$ by the end of Round 1, by definition of Algorithm 2, see Line 7.

Let $p \in V$. If $q = p$ then using the remark above and since $t(q, p) + \Delta \geq 1$, we are done. We now assume $q \neq p$. As $\hat{d}_{q, o_{\mathcal{T}}+t(q, p)-1}(p) \leq \Delta$, there exists a journey $\mathcal{J} = \{(e_1, t_1), \dots, (e_k, t_k)\}$ such that $t_1 > o_{\mathcal{T}} + t(q, p) - 1$, $t_k = t(q, p) + \hat{d}_{q, o_{\mathcal{T}}+t(q, p)-1}(p) + o_{\mathcal{T}} - 1 \leq t(q, p) + \Delta + o_{\mathcal{T}} - 1$ and for every $i \in \{1, \dots, k\}$, $e_i = (p_{i-1}, p_i)$ with $p_0 = q$ and $p_k = p$. To simplify the notations, let $\tau_i = t_i - o_{\mathcal{T}} + 1$ for every i in $\{1, \dots, k\}$ such that the edge $e_i = (p_{i-1}, p_i)$ is present during Round τ_i . We have $\tau_1 > t(q, p)$, $\tau_k \leq t(q, p) + \Delta$, and $\tau_i - \tau_{i-1} < \Delta$.

We prove (by induction on i) that for all $i \in \{1, \dots, k\}$, (1) $id(q)$ is forever in $members(p_i)$ by the end of Round τ_i and (2) $members(p_i)[q] \leq \tau_i - \tau_{i-1} + 1$ at the end of Round τ_i .

Base case: for $i = 1$, the edge (q, p_1) exists at Round τ_1 . Using the first remark in the proof, at the beginning of Round τ_1 , since $\tau_1 > t(q, p) \geq \Delta \geq 1$, we have $id(q) \in members(q) \wedge members(q)[q] = 0$. Hence, at Round τ_1 , q sends $\langle id(q), 0 \rangle$ in its message to p_1 . Following the algorithm, p_1 inserts $id(q)$ in $members(p_1)$ during Round $\tau_1 > \Delta$. So, $id(q)$ is forever in $members(p_1)$ by the end of Round τ_1 ; see Lemma 5.3. Still following the algorithm, $members(p_1)[q] = 1$ at the end of Round τ_1 .

Induction Step: Let $i > 1$. We assume the result holds for $i - 1$: $id(q)$ is forever in $members(p_{i-1})$ by the end of Round τ_{i-1} and $members(p_{i-1})[q] \leq \tau_{i-1} - \tau_{i-2} + 1$ at the end of Round τ_{i-1} . Hence, at the beginning of Round τ_i (and so, at the end of Round $\tau_i - 1$), we have: $id(q)$ in $members(p_{i-1})$ and as the timestamps are at most incremented by one at the end of each round, $members(p_{i-1})[q] \leq \tau_{i-1} - \tau_{i-2} + 1 + \tau_i - 1 - \tau_{i-1} = \tau_i - \tau_{i-1} < \Delta$.

During Round τ_i , the edge $e_i = (p_{i-1}, p_i)$ is present and p_{i-1} sends in its message to p_i a pair $\langle id(q), t_q \rangle$ such that $t_q \leq \tau_i - \tau_{i-1}$ since $\tau_i - \tau_{i-1} < \Delta$. As p_i receives it, it inserts $id(q)$ in $members(p_i)$ in Round τ_i . Since $\tau_i > \tau_{i-1} > \Delta$, Lemma 5.3 ensures that $id(q)$ remains forever in $members(p_i)$ by the end of Round τ_i . Moreover, following the algorithm, at the end of Round τ_i , we have $members(p_i)[q] \leq \tau_i - \tau_{i-1} + 1$.

With $i = k$, $id(q)$ is forever in $members(p)$ by the end Round $\tau_k \leq t(q, p) + \Delta$. \square

Let V be a set of processes. We define a *legitimate* configuration of Algorithm 2 for V as any configuration of Algorithm 2 for V where for every process p , we have $lid(p) = id(\ell)$ and $\{id : \langle id, _ \rangle \in members(p)\} = \{id(q) : q \in V\}$. Remark that the set of legitimate

configurations of Algorithm 2 for V is closed in $\mathcal{TC}^Q(\Delta)$. Indeed, by definition of the algorithm, no message containing a fake ID can be sent from such a configuration. Hence, the set $members(p)$ of every process p remains constant and $\min\{id : \langle id, _ \rangle \in members(p)\} = id(\ell)$ forever. Hence:

LEMMA 5.5. *Any execution of Algorithm 2 that starts from a legitimate configuration in an arbitrary TVG satisfies SP_{LE} .*

The next lemma is a direct consequence of Corollary 5.2 and Lemma 5.4.

LEMMA 5.6. $\exists t \geq \Delta$ such that the configuration reached at the end of Round $t + \Delta$ is legitimate.

By Lemmas 5.5-5.6, follows.

THEOREM 5.7. *Algorithm 2 is a self-stabilizing leader election algorithm for $\mathcal{TC}^Q(\Delta)$. It requires $O(n(B + \log \Delta))$ bits per process and messages of size $O(n(B + \log \Delta))$ bits.*

Speculative Stabilization. Stabilization time cannot be bounded in $\mathcal{TC}^Q(\Delta)$, except for trivial specifications. Indeed, even if there exist infinitely many journeys of length bounded by Δ between any pair of distinct processes and Δ is known by all processes, the time between any two consecutive such journeys is unbounded, by definition of $\mathcal{TC}^Q(\Delta)$. Consequently, we cannot bound the time necessary to route any piece of information from some process p to another process q , making the stabilization time unbounded in any case. We now show that Algorithm 2 is speculative in the sense that its stabilization time cannot be bounded in $\mathcal{TC}^Q(\Delta)$, but in a more favorable case, actually in $\mathcal{TC}^B(\Delta) \subseteq \mathcal{TC}^Q(\Delta)$, its stabilization time is at most 2Δ rounds. To see this, remark that the previous proof holds for $\mathcal{TC}^B(\Delta)$ since $\mathcal{TC}^B(\Delta) \subseteq \mathcal{TC}^Q(\Delta)$. Yet, for every processes p and q , $t(q, p)$ is exactly Δ in this class of TVGs. Hence in the proof of Lemma 5.6, we have $T = \max\{t(q, p) : q, p \in V\} = \Delta$; ensuring that in Class $\mathcal{TC}^B(\Delta)$, the configuration reached at the end of Round 2Δ is legitimate.

THEOREM 5.8. *The stabilization time of Algorithm 2 in Class $\mathcal{TC}^B(\Delta)$ is at most 2Δ rounds.*

6 CLASS \mathcal{TC}^R WITH n KNOWN

Overview of Algorithm 3. Similarly to Algorithm 2, each process p uses a variable $members(p)$ to collect IDs. However, this time, $members(p)$ is a map that can contain up to n IDs, each of them being associated with a timestamp (we denote by $members(p)[id]$ the timestamp associated to the identifier id belonging to $members(p)$).

At each round i , p sends the content of $members(p)$ (Line 2). Then, p updates $members(p)$ by calling function $insert$ on each received pair $\langle id, t \rangle$ such that $id \neq id(p)$ (Lines 4-5). The function $insert$ works as follows: if id already appears in $members(p)$, then the associated timestamp is updated by keeping the smallest value (Line i1). Otherwise, p tries to insert $\langle id, t \rangle$ in the map. Actually, $\langle id, t \rangle$ is inserted in the map if the map is not full (Line i2) or t is smaller than the greatest timestamp tM in the map (Lines i3-i5). In this latter case, $\langle id, t \rangle$ overwrites any value having this timestamp in $members(p)$ (Line i5). This overwriting mechanism allows to eventually remove all fake IDs from $members(p)$, since their timestamps will regularly increase. After $members(p)$ has been updated,

all timestamps of $members(p)$ are incremented (Line 6) and then, $\langle id(0), 0 \rangle$ is systematically inserted in the map (Line 7).

Actually, Algorithm 3 guarantees two main properties. First, at the beginning of any round i , any timestamp associated to a fake ID is greater than or equal to $i - 1$; see Lemma 6.1. Second, by definition of \mathcal{TC}^R , at any point in time, every process can reach all the others through a journey. The key property is then to show that if some broadcast initiated by process p reaches a process q at Round i , then the value of the timestamp in the message is small enough to ensure the insertion of $id(p)$ into $members(q)$; see Lemma 6.2. These two properties ensure that eventually $members(p)$ exactly contains all IDs of the network. Now, at the end of each round, p updates its leader variable with the smallest ID in $members(p)$ (Line 7). Hence, the process of lowest ID, ℓ , is eventually elected.

Algorithm 3: Self-stabilizing leader election for \mathcal{TC}^R , for each process p .

Inputs:

$n \in \mathbb{N}$: number of processes
 $id(p) \in IDSET$: ID of p

Local Variables:

$members(p)$: map of size at most n
 contains pairs $\langle id, t \rangle \in IDSET \times \mathbb{N}$
 $lid(p) \in IDSET$: ID of the leader

Macros:

$max(p)$:

m1: $\text{if } |members(p)| < n \text{ then return } \perp$
 m2: $\text{else return } \langle id, t \rangle \in members(p) \text{ with maximum timestamp } t$

$insert(p, \langle id, t \rangle)$:

i1: $\text{if } \langle id, _ \rangle \in members(p) \text{ then}$
 $\quad members(p)[id] := \min(t, members(p)[id])$
 i2: $\text{else if } max(p) = \perp \text{ then add } \langle id, t \rangle \text{ in } members(p)$
 i3: else
 i4: $\quad \langle idM, tM \rangle := max(p)$
 i5: $\quad \text{if } t < tM \text{ then remove } \langle idM, tM \rangle \text{ from } members(p); \text{ add}$
 $\quad \quad \langle id, t \rangle \text{ in } members(p)$

1: **Repeat Forever**

2: $\text{SEND}(\langle members(p) \rangle)$
 3: $\text{mailbox} := \text{RECEIVE}()$
 4: **forall pair** $\langle id, t \rangle$ **in a message of mailbox do**
 5: $\quad \text{if } id \neq id(p) \text{ then } insert(p, \langle id, t \rangle)$
 6: **forall id** : $\langle id, _ \rangle \in members(p)$ **do** $members(p)[id] ++$
 7: $\quad insert(p, \langle id(p), 0 \rangle); lid(p) := \min\{id : \langle id, _ \rangle \in members(p)\}$

Self-stabilization. The lemma below and its proof are identical to those of Lemma 5.1.

LEMMA 6.1. *Let f be a fake ID. For every $i \geq 1$, at the beginning of Round i , the following holds: $\forall p \in V$, if f is in $members(p)$, then $members(p)[f] \geq i - 1$.*

The following lemma can be proven by induction on i .

LEMMA 6.2. *For every $i \geq 1$, at the end of Round i , the following property holds: $\forall p, q \in V$, if $\hat{d}_{p, o_T}(q) \leq i - 1$, then $id(p)$ is in $members(q)$ and $members(q)[p] \leq i - 1$.*

Let V be a set of processes. We define a *legitimate* configuration of Algorithm 3 for V as any configuration of Algorithm 3 for V where $lid(p) = id(\ell)$ and $\{id : \langle id, _ \rangle \in members(p)\} = \{id(q) : q \in V\}$, for every process p . By definition of the algorithm, no message containing a fake ID can be sent from such a configuration. So, from

any legitimate configuration, the set $members(p)$ of every process p is constant and $\min\{id : \langle id, _ \rangle \in members(p)\} = id(\ell)$ forever. Thus, the set of legitimate configurations of Algorithm 3 for V is closed in \mathcal{TC}^R and we have:

LEMMA 6.3. *Any execution of Algorithm 3 that starts from a legitimate configuration in an arbitrary TVG satisfies SP_{LE} .*

THEOREM 6.4. *Algorithm 3 is a self-stabilizing leader election algorithm for \mathcal{TC}^R .*

PROOF. Let $p \in V$. By definition of \mathcal{TC}^R , $\forall q \in V, \exists \mathcal{J} \in \mathcal{J}(p, q)$ such that $departure(\mathcal{J}) > o_{\mathcal{T}}$. The temporal length of \mathcal{J} is finite. Thus, $\exists \delta(p) \in \mathbb{N}$ such that $\forall q \in V, d_{p, o_{\mathcal{T}}}(q) \leq \delta(p)$. Thus, at the end of Round $\delta(p) + 1$, $\forall q \in V, id(p)$ is in $members(q)$ by Lemma 6.2. Since $members(q)$ contains at most n entries, after $\max_{p \in V} \delta(p) + 1$ rounds, $members(q)$ contains the ID of every process and no fake ID. So q chooses $id(\ell)$ as leader at the end of Round $\max_{p \in V} \delta(p) + 1$. Hence, the system is in a legitimate configuration at the end of this round and, by Lemma 6.3, we are done. \square

Speculative Stabilization. Similarly to $\mathcal{TC}^Q(\Delta)$, stabilization time cannot be bounded in \mathcal{TC}^R , except for trivial specifications ($n.b.$, $\mathcal{TC}^Q(\Delta) \subseteq \mathcal{TC}^R$). We now show that Algorithm 3 is speculative in the sense that we cannot bound its stabilization time in \mathcal{TC}^R , but in a more favorable case, precisely in $\mathcal{TC}^B(\Delta) \subseteq \mathcal{TC}^R$, its stabilization time is at most $\Delta + 1$ rounds, despite Δ is unknown. The proof of the theorem below is the same as the one of Theorem 6.4 but as we consider a TVG in $\mathcal{TC}^B(\Delta)$, for every $p \in V, \delta(p) \leq \Delta$. Hence the system reaches a legitimate configuration at the end of Round $\max_{p \in V} \delta(p) + 1 = \Delta + 1$.

THEOREM 6.5. *The stabilization time of Algorithm 3 in $\mathcal{TC}^B(\Delta)$ is at most $\Delta + 1$ rounds.*

7 CONCLUSION

We have addressed self-stabilization in highly dynamic identified message passing systems by proposing self-stabilizing leader election algorithms for three major classes of time-varying graphs: $\mathcal{TC}^B(\Delta)$, $\mathcal{TC}^Q(\Delta)$, and \mathcal{TC}^R . It is worth noting that the impossibility result of Braud-Santoni *et al.* [9] applies to the class of *always connected over the time TVGs* of n processes which is actually included and so stronger than \mathcal{TC}^R , as well as $\mathcal{TC}^B(\Delta)$ and $\mathcal{TC}^Q(\Delta)$, for every $n \geq 0$ and $\Delta \geq n - 1$. Precisely, this result forbids the existence of *silent* self-stabilizing solutions for many non-trivial static problems. Actually, silent self-stabilization additionally requires all processes to eventually keep their local state constant [21]. Now, leader election is a static problem. We have chosen to avoid this issue by proposing non-silent, *a.k.a.*, *talkative* [7] solutions, *i.e.*, in our algorithms, a small part of the local state of each process (namely, the timestamps) is modified infinitely often. Beyond extending our results to other particular problems, further research could focus on studying expressiveness of self-stabilization in TVGs. To that goal, broadcast problems should be investigated, again in very general TVG classes. Indeed, coupled with our leader election solutions, they should allow to build generic transformers, following, for example, the approaches proposed in [15, 28].

REFERENCES

- [1] Karine Altisen, Stéphane Devismes, Swan Dubois, and Franck Petit. 2019. *Introduction to Distributed Self-Stabilizing Algorithms*. Vol. 8. Morgan & Claypool Publishers. 1–165 pages.
- [2] Karine Altisen, Stéphane Devismes, Anaïs Durand, and Franck Petit. 2019. Gradual stabilization. *JPDC* 123 (2019), 26–45.
- [3] Dana Angluin. 1980. Local and Global Properties in Networks of Processors. In *Proceedings of STOC'80*. 82–93.
- [4] Baruch Awerbuch, Boaz Patt-Shamir, George Varghese, and Shlomi Dolev. 1994. Self-Stabilization by Local Checking and Global Reset (Extended Abstract). In *Distributed Algorithms, 8th International Workshop, WDAG*. 326–339.
- [5] Matthieu Barjon, Arnaud Casteigts, Serge Chaumette, Colette Johnen, and Yessin M. Neggaz. 2019. Maintaining a Distributed Spanning Forest in Highly Dynamic Networks. *Comput. J.* 62, 2 (2019), 231–246.
- [6] Joffroy Beauquier and Synnöve Kekkonen-Moneta. 1997. On FTSS-Solvable Distributed Problems. In *PODC*. 290.
- [7] Lélia Blin and Sébastien Tixeuil. 2018. Compact deterministic self-stabilizing leader election on a ring: the exponential advantage of being talkative. *Dist. Comput.* 31, 2 (2018), 139–166.
- [8] Marjorie Bournat, Ajoy K. Datta, and Swan Dubois. 2019. Self-stabilizing robots in highly dynamic environments. *Theor. Comput. Sci.* 772 (2019), 88–110.
- [9] Nicolas Braud-Santoni, Swan Dubois, Mohamed-Hamza Kaouachi, and Franck Petit. 2016. The Next 700 Impossibility Results in Time-Varying Graphs. *IJNC* 6, 1 (2016), 27–41.
- [10] Binh-Minh Bui-Xuan, Afonso Ferreira, and Aubin Jarry. 2003. Computing shortest, fastest, and foremost journeys in dynamic networks. *IJFCS* 14, 02 (2003), 267–285.
- [11] Shukai Cai, Taisuke Izumi, and Koichi Wada. 2012. How to Prove Impossibility Under Global Fairness: On Space Complexity of Self-Stabilizing Leader Election on a Population Protocol Model. *Theory Comput. Syst.* 50, 3 (2012), 433–445.
- [12] Arnaud Casteigts. 2018. *A Journey through Dynamic Networks (with Excursions)*. HDR. Université de Bordeaux.
- [13] Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. 2012. Time-varying graphs and dynamic networks. *Inter. J. of Parall., Emergent and Dist. Systems* 27, 5 (2012), 387–408.
- [14] Bernadette Charron-Bost and Shlomo Moran. 2018. The Firing Squad Problem Revisited. In *STACS*. 20:1–20:14.
- [15] Alain Cournier, Ajoy Kumar Datta, Stéphane Devismes, Franck Petit, and Vincent Villain. 2016. The expressive power of snap-stabilization. *Theor. Comput. Sci.* 626 (2016), 40–66.
- [16] Ajoy K. Datta and Lawrence L. Larmore. 2018. Self-Stabilizing Leader Election in Dynamic Networks. *Theory Comput. Syst.* 62, 5 (2018), 977–1047.
- [17] Sylvie Delaët, Bertrand Ducourthial, and Sébastien Tixeuil. 2006. Self-stabilization with r-operators revisited. *J. of Aerospace Comp., Inf., and Comm.* 3, 10 (2006), 498–514.
- [18] Edsger W. Dijkstra. 1974. Self-stabilizing Systems in Spite of Distributed Control. *Commun. ACM* 17, 11 (1974), 643–644.
- [19] Shlomi Dolev. 1998. Optimal Time Self-Stabilization in Uniform Dynamic Systems. *Parallel Process. Lett.* 8, 1 (1998), 7–18.
- [20] Shlomi Dolev. 2000. *Self-Stabilization*. MIT Press.
- [21] Shlomi Dolev, Mohamed G. Gouda, and Marco Schneider. 1999. Memory Requirements for Silent Stabilization. *Acta Inf.* 36, 6 (1999), 447–462.
- [22] Shlomi Dolev, Ariel Hanemann, Elad Michael Schiller, and Shantanu Sharma. 2012. Self-stabilizing End-to-End Communication in (Bounded Capacity, Omitting, Duplicating and non-FIFO) Dynamic Networks - (Extended Abstract). In *SSS 2012*, Vol. 7596. Springer, 133–147.
- [23] Shlomi Dolev and Ted Herman. 1995. Superstabilizing Protocols for Dynamic Distributed Systems. *Chicago Journal of Theoretical Computer Science* (1995).
- [24] Shlomi Dolev, Amos Israeli, and Shlomo Moran. 1993. Self-Stabilization of Dynamic Systems Assuming Only Read/Write Atomicity. *Dist. Comput.* 7, 1 (1993), 3–16.
- [25] Shlomi Dolev, Amos Israeli, and Shlomo Moran. 1997. Resource Bounds for Self-Stabilizing Message-Driven Protocols. *SIAM J. Comput.* 26, 1 (1997), 273–290.
- [26] Swan Dubois and Rachid Guerraoui. 2013. Introducing speculation in self-stabilization: an application to mutual exclusion. In *PODC*. 290–298.
- [27] Carlos Gómez-Calzado, Arnaud Casteigts, Alberto Lafuente, and Mikel Larrea. 2015. A Connectivity Model for Agreement in Dynamic Systems. In *Euro-Par*. 333–345.
- [28] Shmuel Katz and Kenneth J. Perry. 1993. Self-Stabilizing Extensions for Message-Passing Systems. *Dist. Comput.* 7, 1 (1993), 17–26.
- [29] Ramakrishna Kotla, Lorenzo Alvisi, Michael Dahlin, Allen Clement, and Edmund L. Wong. 2009. Zyzzyva: Speculative Byzantine fault tolerance. *ACM Trans. Comp. Syst.* 27 (2009), 1–39.
- [30] Mikhail Nesterenko and Anish Arora. 2002. Dining Philosophers that Tolerate Malicious Crashes. In *ICDCS*. 191–198.
- [31] George Varghese. 2000. Self-Stabilization by Counter Flushing. *SIAM J. Comput.* 30, 2 (2000), 486–510.