

Self-stabilizing Systems in Spite of High Dynamics

Karine Altisen¹ Stéphane Devismes² Anaïs Durand³
Colette Johnen⁴ Franck Petit⁵

¹ VERIMAG, Grenoble

² MIS, Amiens

³ LIMOS, Clermont Ferrand

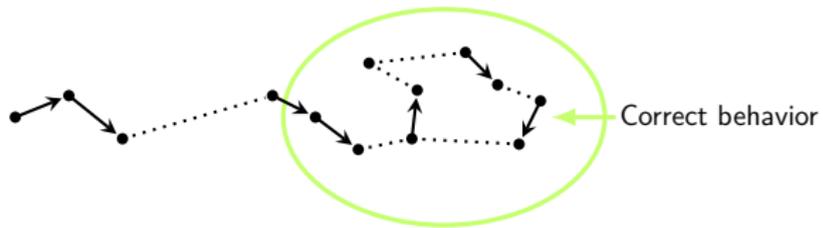
⁴ LaBRI, Bordeaux

⁵ LIP6, Paris

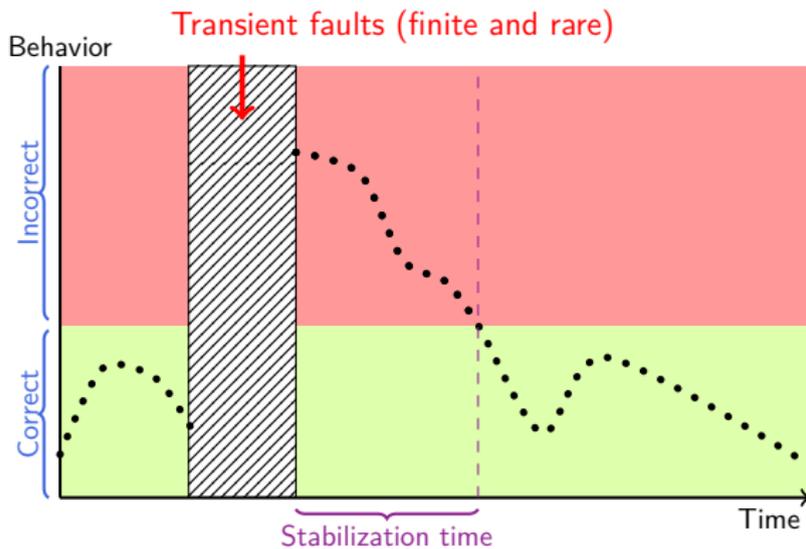
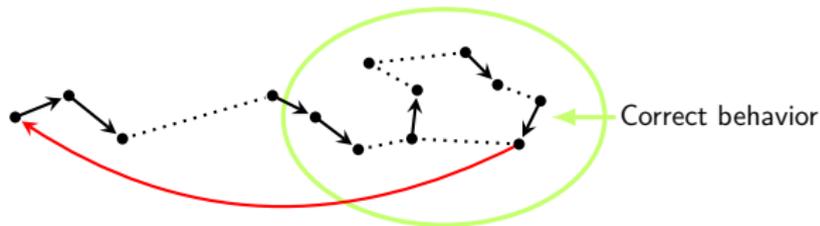
ANR TEMPOGRAL, November 2022



Self-stabilization, [Dijkstra, ACM Com., 74]



Self-stabilization, [Dijkstra, ACM Com., 74]



General approach for tolerating **transient faults**

■ Transient faults :

- ▶ Self-stabilization

■ Permanent failures :

- ▶ process crashes \rightsquigarrow **fault-tolerant self-stabilization** [Beauquier, Kekkonen-Moneta, 97]
- ▶ Byzantine failures \rightsquigarrow **strict stabilization** [Nesterenko, Arora, 02]

■ Intermittent failures :

- ▶ message loss, duplication, and/or reordering [Delaët *et al.*, 06]

... in static networks

Self-stabilization vs. Dynamic Networks

- If **topological changes** are :

- ▶ **rare**
- ▶ eventually **locally detected**

↪ self-stabilization is a good approach

- Variants **dedicated** to dynamic networks, e.g.:

- ▶ **Superstabilization**: quick convergence after few topological changes from a legitimate configuration [Dolev and Herman, *Chicago Journal of Theoretical Computer Science*, 1997]
- ▶ **Robust convergence**: quick convergence to a safe configuration, then convergence to a legitimate configuration [Johnen and Nguyen, OPODIS'06, Kakugawa and Masuzawa, IPDPS'06]

- **But**, highly **ineffective** when the frequency of topological changes increases

Self-stabilization in *Highly Dynamic Networks* ?

where **topological changes** are not :

{ **transient**
an anomaly but { **intermittent**
inherent

To tolerate both **transient faults** and **high dynamics**

Case Study: Leader Election

Negative result: Even if the network is **always connected over the time**, **silent self-stabilization is (almost always) impossible!**

[Braud-Santoni *et al.*, IJNC, 2016]

Silence: converges within finite time to a configuration from which the values of the communication registers used by the algorithm remain fixed.

- Finding **conditions** under which **self-stabilizing leader election** can be achieved.

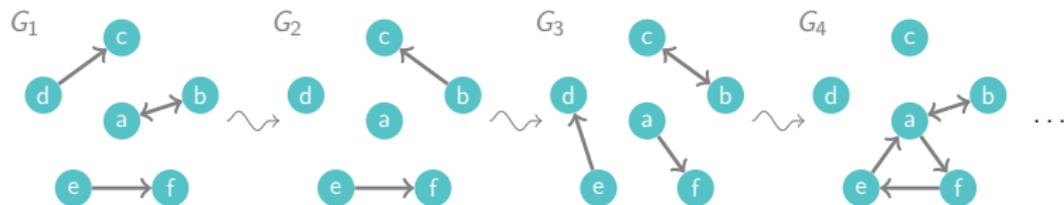
We look for (non-silent) self-stabilizing algorithm for
general classes of dynamic networks

(e.g., we do not enforce the network to be in a particular topology at a given time)

- Finding the **limits** where self-stabilizing leader election becomes impossible?
- Studying **lower bounds** on the convergence time

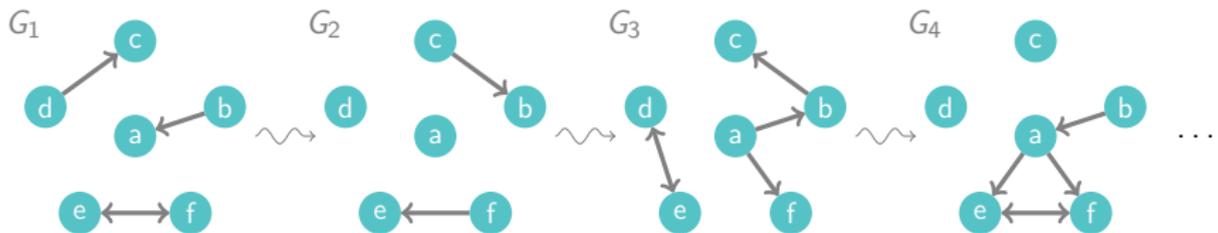
Context

Dynamics modeled with a **Dynamic Graph on V** , V set of nodes.

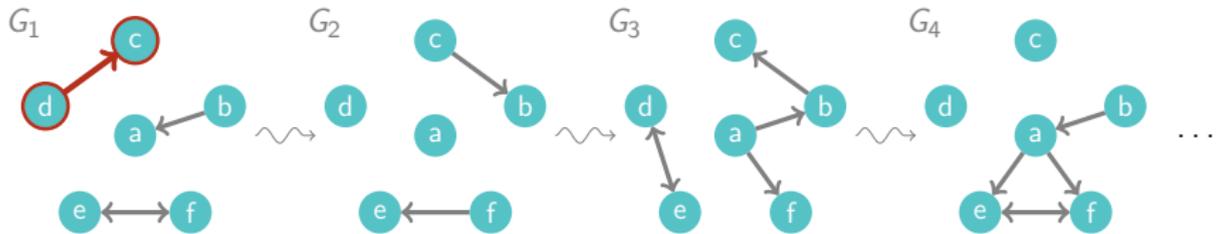


[Xuan *et al.*, 03], [Casteigts *et al.*, 13]

Can **d** transmit information to **a** ?

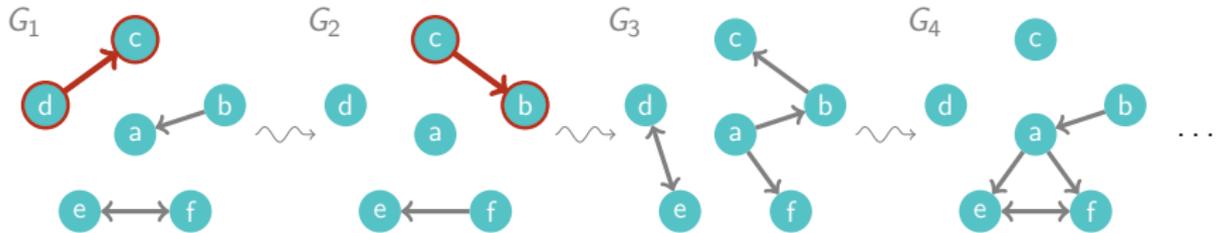


Can **d** transmit information to **a** ?



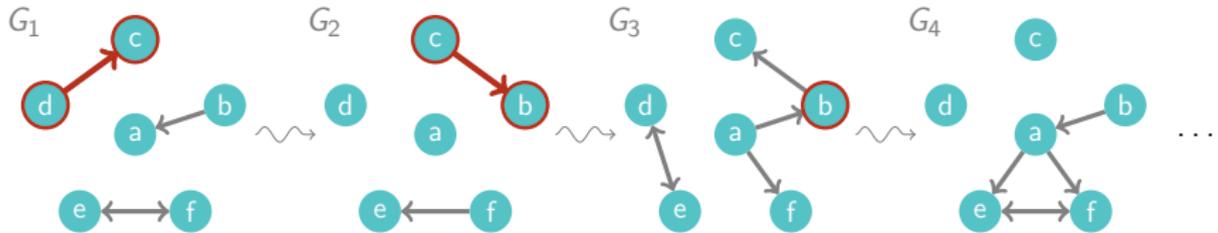
$1, (d, c)$

Can **d** transmit information to **a** ?



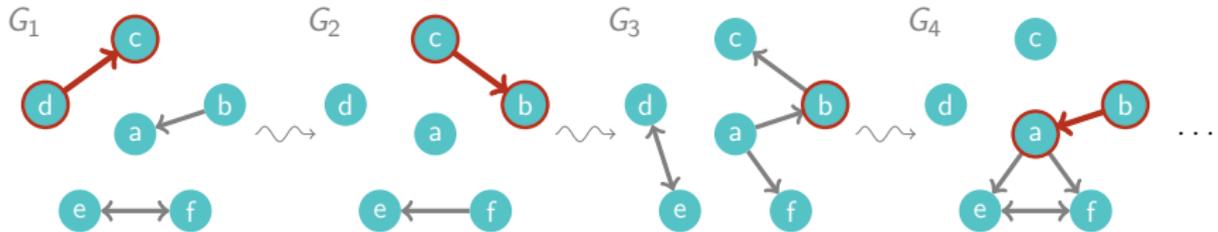
1, (d, c) ; 2, (c, b)

Can **d** transmit information to **a** ?



1, (d, c) ; 2, (c, b)

Can **d** transmit information to **a** ?



1, (d, c) ; 2, (c, b) ; 4, (b, a) = Journey from **d** to **a** of length 4.

Synchronous Rounds:



[Charron-Bost and Moran, STACS'18 / Barjon et al., CJ, 2019]

Execution in \mathcal{G} : infinite sequence of configurations $\gamma_0, \gamma_1, \dots$ such that

- γ_0 is arbitrary
- $\forall i > 0$, γ_i is obtained from γ_{i-1} as follows:
 -  Every node p sends a message consisting of all or a part of its local state in γ_{i-1} ,
 -  p receives all messages sent by nodes in $\mathcal{N}(p)^i$, and
 -  p computes its state in γ_i .

$\forall i > 0$,

- $\mathcal{N}(p)^i$ set of nodes that are neighbor of p during the round i .
- γ_{i-1} is the configuration **at the beginning of Round i**
- γ_i is the configuration **at the end of Round i**

Considered Classes of TVGs

Class $\mathcal{TC}^{\mathcal{B}}(\Delta)$ with $\Delta \in \mathbb{N}^*$ (Bounded Temporal Diameter): Every node can always reach any other node through a journey of length $\leq \Delta$, *i.e.*, **the temporal diameter is bounded by Δ** , [Gómez-Calzado *et al.*, Euro-Par'15]

Class $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ with $\Delta \in \mathbb{N}^*$ (Quasi Bounded Temporal Diameter): Every node can infinitely often reach any other node through a journey of length $\leq \Delta$. [Altisen *et al.*, ICDCN'21]

Class $\mathcal{TC}^{\mathcal{R}}$ (Recurrent Temporal Connectivity): every node can infinitely often reach any other node through a journey. [Casteigts *et al.*, IJPEDES, 2012]

$$\mathcal{TC}^{\mathcal{B}}(\Delta) \subseteq \mathcal{TC}^{\mathcal{Q}}(\Delta) \subseteq \mathcal{TC}^{\mathcal{R}}$$

Self-stabilization in an Highly Dynamic Context

Adaptation of the definition in the book [Self-Stabilization, Dolev, 2000]

Let V be a set of nodes.

An algorithm \mathcal{A} is **self-stabilizing** for the specification $SP(V)$ on a class \mathcal{C} of DGs if there exists a subset of configurations $\mathcal{L}(V)$ of \mathcal{A} , called **legitimate configurations**, such that:

- 1 for every $\mathcal{G}(V) \in \mathcal{C}$ and every execution of \mathcal{A} in $\mathcal{G}(V)$ contains a legitimate configuration $\gamma' \in \mathcal{L}(V)$ (**Convergence**), and
- 2 for every $\mathcal{G}(V) \in \mathcal{C}$, every legitimate configuration $\gamma \in \mathcal{L}(V)$, $SP(V)$ holds on the execution of \mathcal{A} in $\mathcal{G}(V)$ starting from γ (**Correctness**).

n identified nodes: $\forall p \in V$, $id(p)$ is the unique identifier of p

Let $IDSET$ be the definition domain of identifiers ($|IDSET| \gg n$)

$\forall v \in IDSET$,

- v is a **real ID** if $\exists p \in V$, $id(p) = v$
- v is a **fake ID** otherwise

$lid(p)$ contains the identifier of the leader computed by p

Initially, $lid(p)$ may contain a fake ID

GOAL: converge to a configuration from which all lid variables constantly designates **the same real ID**

Convergence Time of a Leader Election algorithm Unboundable in $\mathcal{TC}^Q(\Delta)$

$\mathcal{TC}^Q(\Delta)$ – every node can infinitely often reach any other node through a journey of length $\leq \Delta$

For $i \in \mathbb{N}^*$ and $\mathcal{G}(V) \in \mathcal{TC}^Q(\Delta)$,
let $\mathcal{G}_i(V)$ be the dynamic graph $G_{isolated}(V)^i \mathcal{G}(V)$.

$\mathcal{G}_i(V) \in \mathcal{TC}^Q(\Delta)$

In $G_{isolated}(V)$ no node has a neighbor.



Along $G_{isolated}(V)^i$, no information is exchanged: **leader election impossible!**

In $\mathcal{G}_i(V)$ the leader election requires at least $i + 1$ rounds. \square

Since $\mathcal{TC}^Q(\Delta) \subset \mathcal{TC}^R$, the result also holds for \mathcal{TC}^R .

Speculation [Kotla *et al.*, ACM Trans. Comput. Syst., 2009]:

- the system satisfies its requirements for all executions,
- but also exhibits significantly better performances in a subset of executions.

Idea: worst possible scenarios are often rare in practice.

When the stabilization time of an algorithm cannot be bounded in a class (i.e. $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ and $\mathcal{TC}^{\mathcal{R}}$), we exhibit a bound of this algorithm in the class $\mathcal{TC}^{\mathcal{B}}(\Delta)$ (i.e. every node can always reach any other node through a journey of length $\leq \Delta$).

Overview of our algorithms

n identified nodes: $\forall p \in V$, $id(p)$ is the unique identifier of p

Let $IDSET$ be the definition domain of identifiers ($|IDSET| \gg n$)

$\forall v \in IDSET$,

- v is a **real ID** if $\exists p \in V$, $id(p) = v$
- v is a **fake ID** otherwise

Every node p computes the identifier of the leader in $lid(p)$

Initially, the value of $lid(p)$ may be a fake ID

Strategy:

- 1 First, **eliminate all fake IDs**, and then
- 2 Compute in all output variables the **minimum real ID**, noted $id(\ell)$.

Positive Results

| | Algorithm 3 | Algorithm 2 | Algorithm 1 |
|-----------|------------------------------|--------------------------------------|--------------------------------------|
| Knowledge | n | n, Δ | Δ |
| Class | $\mathcal{TC}^{\mathcal{R}}$ | $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ | $\mathcal{TC}^{\mathcal{B}}(\Delta)$ |
| Time | not bounded | not bounded | 3Δ |
| Memory | not bounded | $O(n \log(n + \Delta))$ | $O(\log(n + \Delta))$ |

n : number of nodes

Positive Results

| | Algorithm 3 | | Algorithm 2 | | Algorithm 1 |
|-----------|------------------------------|--------------------------------------|--------------------------------------|--------------------------------------|--------------------------------------|
| Knowledge | n | | n, Δ | | Δ |
| Class | $\mathcal{TC}^{\mathcal{R}}$ | $\mathcal{TC}^{\mathcal{B}}(\Delta)$ | $\mathcal{TC}^{\mathcal{Q}}(\Delta)$ | $\mathcal{TC}^{\mathcal{B}}(\Delta)$ | $\mathcal{TC}^{\mathcal{B}}(\Delta)$ |
| Time | not bounded | $\Delta + 1$ | not bounded | 2Δ | 3Δ |
| Memory | not bounded | | $O(n \log(n + \Delta))$ | | $O(\log(n + \Delta))$ |

Speculation

n : number of nodes

Results :

- **Definition** of self-stabilization in highly dynamic networks
- Focus on **self-stabilizing leader election**
 - ▶ **Negative Results:**
 - in $\mathcal{TC}^B(\Delta)$ if n is not known then the legitimate configurations set cannot be closed
 - in \mathcal{TC}^R or $\mathcal{TC}^Q(\Delta)$, the leader election is impossible if n is not known
 - ▶ **Positive Results:**
 - we have circumvented the impossibility result of [Braud-Santoni *et al.*, IJNC, 2016] by considering non-silent solutions algorithms for \mathcal{TC}^R , $\mathcal{TC}^Q(\Delta)$, $\mathcal{TC}^B(\Delta)$

Future work: Other TVG classes

Thank you for your attention

Joint Work:



- Karine Altisen, Stéphane Devismes, and Anaïs Durand, Colette Johnen, and Franck Petit. [Self-stabilizing Systems in Spite of High Dynamics](#). *ICDCN'21*.
- full version of the paper in HAL :
[<https://hal.archives-ouvertes.fr/hal-02376832>]

Funding: ANR Estate & Descartes