

Internship Proposal:

Applying Graceful Degrading Approach to Self-Stabilization

Advisors: Swan Dubois and Franck Petit

Laboratory: LIP6, project-team Inria DELYS, Campus Jussieu, 4 place Jussieu, 75005 Paris, France

Mail: <firstname>.<lastname>@lip6.fr

Dates: February to July 2020 (6 months)

Funding: Regular internship gratification by ANR project ESTATE

Context

Self-stabilization is a desirable property for modern distributed systems. Indeed, a self-stabilizing (distributed) system [4] is guaranteed to converge to its intended behavior in finite time regardless of the initial state of the processors and initial messages in the links. It is a versatile technique to design distributed algorithms that withstand transient faults. Moreover, it requires no initialization phase, which is particularly suitable for the development of large-scale systems.

Adaptive algorithms [5, 3, 1, 2] are a promising approach for distributed computing. Indeed, such algorithms adapt themselves to their environment to provide a best-effort solution. There exists mainly two adaptive approaches. *Speculation* [5, 3] provides protocols that adapt their performance to their environment (the “easier” is the environment, the better is the performance). On the other hand, *graceful degradation* [1, 2] leads to protocols adapting the quality of their solution to the environment (the “harder” is the environment, the lower is the quality). Those two approaches are intrinsically orthogonal. The first approach is motivated by lower bounds results on complexity while the second is motivated by impossibility results.

Previous results [3] show that speculation may be applied to self-stabilization and allows the design of self-stabilizing algorithms whose stabilization time depends on the environment at the time of their execution. For now, graceful degradation was not studied in the scope of self-stabilization.

Internship Assignment

The main goal of this internship is to address the question whether the graceful degrading approach may be applied to self-stabilization. The goal is to circumvent impossibility results of these area by providing self-stabilizing algorithms that autonomously adapt their behavior to their environment to provide a best-effort solution.

The scientific agenda is mainly threefold:

- Study of the state-of-the-art related to the internship ;
- Definition of a new variant of self-stabilization integrating the graceful degrading approach ;

- Application of this new variant to a classical problem of the literature at the choice of the intern. This includes the definition of pertinent weaker variants of this problem, the design and validation of an algorithm satisfying the proposed definition and the writing of some impossibility proofs.

This internship subject is purely algorithmic and requires some appetite to read and write formal definitions, algorithms, and proofs.

References

- [1] M. Biely, P. Robinson, U. Schmid, M. Schwarz, and K. Winkler. Gracefully degrading consensus and k -set agreement in directed dynamic networks. *TCS*, 726:41–77, 2018.
- [2] M. Bournat, S. Dubois, and F. Petit. Gracefully degrading gathering in dynamic rings. In *SSS*, 2018.
- [3] S. Dubois and R. Guerraoui. Introducing speculation in self-stabilization: an application to mutual exclusion. In *PODC*, pages 290–298, 2013.
- [4] E. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11):643–644, 1974.
- [5] R. Guerraoui. Foundations of speculative distributed computing. In *DISC*, pages 204–205, 2010.