

# Internship Proposal: Performance Analysis of Replication Schemes in Computer Systems

Supervisor: Jonatha Anselmi (INRIA)  
<http://polaris.imag.fr/jonatha.anselmi>

February 5, 2020

**Keywords:** Markov chains, Stability, Load balancing

**Lab:** Laboratoire d'Informatique de Grenoble (LIG), <http://www.liglab.fr> (head: Eric Gaussier)

**Team:** POLARIS, <https://team.inria.fr/polaris> (head: Arnaud Legrand)

## Scientific and Technologic Context

Nowadays, it is crucial that computer systems respond to interactive users in a few milliseconds. In practice, user requests, or jobs in the following, are subject to a number of factors that increase the variability of response times. These include unfortunate disk seek times, run-time contention phenomena among CPU cores, caches, memory bandwidth, and network bandwidth. As a result, some jobs take significantly longer than expected to complete while keeping blocked resources that could be used by other concurrent interactive jobs. This has brought researchers to propose the adoption of *redundant* jobs, used, e.g., in Google's big table services [1]. The principle consists in replicating a job multiple times and use the results from whichever replica responds first. Several replication schemes based on this principle have been recently proposed in the literature. The current general picture is that mathematical properties such as steady-state stability are not understood when a stochastic and dynamic setting is considered, unless restrictive assumptions are taken.

## Goal of the Internship

The objective is to study a new class of replication algorithms and investigate mathematical properties such as Lyapunov stability and asymptotic optimality. Asymptotic optimality refers to the ideal situation where latency approaches zero in the limit where the system size grows to infinity [2]. The first main contribution of the intern will consist in investigating these properties numerically. Towards this purpose, he/she will extend an existing code that simulates the dynamics of jobs inside the network. Then, an analytical investigation will follow up where dynamics will be modelled by means of Markov processes in continuous time.

## Requirements

The intern will have a background in applied probability and computer programming.

## Location, Timings and Contact

The intern will be hosted in the POLARIS team. The POLARIS team is a joint team between Inria and LIG (Grenoble Computer Science Laboratory) and is located in Grenoble University main campus (<https://batiment.imag.fr>).

The duration of the internship is from three up to six months, in 2020.

The intern will receive an allowance to cover accommodation costs.

For more information, please contact [jonatha.anselmi@inria.fr](mailto:jonatha.anselmi@inria.fr).

## References

- [1] J. Dean and L. A. Barroso. “*The Tail at Scale.*” Commun. ACM 56, 2 (Feb. 2013), 74–80.
- [2] J. Anselmi and F. Dufour. “*Power-of-d-choices with memory: Fluid limit and optimality*”. Mathematics of Operations Research, to appear.