

Research internship (Master 2)

Time Simulator for the Declarative DevOps tool Ballet

Hélène Coullon
IMT Atlantique, LS2N, Inria, France

Jolan Philippe
IMT Atlantique, LS2N, Inria, France

2023-2024

Keywords : *DevOps, Reconfiguration, Service-oriented software, Distributed system* .

1 Context

In recent years, mainly because of the advent of distributed paradigms such as Cloud computing, service-oriented (SO) software architectures have become the norm and have evolved towards micro-services applications and systems, sometimes made of thousands of small components. At the same time, the DevOps profession has gradually taken shape. The DevOps concept consists of filling the gap between development and operation concerns, which very often conflicts. One targets very fast integration of new software features, while the other targets reliability and stability. DevOps engineers are in practice responsible for automating and accelerating every possible procedure between the development and the operation [3, 4].

One important aspect of this job is the process of deploying and configuring long-running services while avoiding their interruptions which is made easier through Infrastructure-as-Code (IaC) techniques [2]. In IaC, procedures are defined as well-structured and easy-to-read codes. In particular, declarative approaches have become widely used for their high abstraction level. DevOps users specify what is required, not the imperative program (or plan) to get it. This program is instead automatically inferred, which avoids errors.

When dealing with large and geographically distributed software systems deployed in environments such as multi-Cloud, cyber-physical systems, Edge/Fog computing, DevOps engineers may have a limited view of the whole system, for scale reasons and to avoid building and replicating a coherent picture of the full system's state. In practice, in this context, multiple DevOps teams work in a loosely coupled manner to handle different parts of systems, each team using its local instances of centralized DevOps declarative tools. However, a management operation in one part of the system may require modifications in other parts, thus requiring coordination between teams. Because existing declarative tools are designed in a centralized manner (one central entity that controls the overall procedure to apply), this coordination is often done manually between teams. Ballet stands out as one of the rare decentralized declarative IaC tools, with Muse [5] a decentralized extension of Pulumi.

Ballet is a research prototype that automatically processes the DevOps deployment, management, and update operations in a decentralized manner. It offers two key contributions : firstly a planner that infers the actions required by each Ballet instance to attain a given global objective (*e.g.*, reaching a specific system state, or executing a particular operation); secondly, an executor that is responsible for executing the local set of actions, while performing necessary communications with other Ballet's instances. In particular, Ballet focuses (among other objectives) on improving the speed of DevOps operations. In this internship, we want to design, develop, and verify a performance model for Ballet. In other words, we want to develop and validate a time simulator to predict the duration of any DevOps management operation by using Ballet. This contribution is very important to be able to optimize the set

of actions chosen by the planner, but also to avoid costly experiments on real infrastructures in future publications. Indeed, simulation is one important leverage to reduce the energy footprint of research in computer science.

2 Work

We expect the selected candidate to contribute to the state-of-the-art in decentralized reconfigurations and service-oriented software. The candidate has to create and validate a model that accurately predicts Ballet's executor performances, taking into account the network on which it runs.

The objectives of the internship are :

- Formalize and concretely define the performance model for Ballet's executor. The model can take inspiration in [1].
- Integrate this performance model within a network simulator (*e.g.*, SimGrid) to incorporate accurate communication times.
- Validate the model through an extensive set of experiments on Grid'5000¹, the biggest share network dedicated to research in Computer Science. The experiments will encompass a variety of synthetic and real-world application scenarios (some of which are already available).
- Submit a research paper detailing the findings and contributions of the work.

We expect the successful candidate to create reusable artifacts, that will be used in further development of decentralized reconfiguration engine.

Note that if satisfactory, the successful candidate will probably have an opportunity to start a Ph.D. thesis after the internship.

3 Expected skills

The following skills are expected from the successful candidate :

- a student in the last year of a Master's degree in Computer Science (or in the last year of an engineering school with a computer science option) ;
- knowledge and experience in distributed software systems, in particular, micro-services ;
- knowledge and experience on DevOps approaches such as Infrastructure-as-Code, orchestration etc. ;
- knowledge and experience in modeling ;
- knowledge of the Python programming language ;
- a good level of English to contribute to the writing of a research paper ;
- an ability to collaborate and communicate ;
- curiosity and an appetite for learning new things.

4 Additional information

Advisors

- [Hélène COULLON](#), IMT Atlantique & Inria team STACK, helene.coullon@imt-atlantique.fr
- [Jolan PHILIPPE](#), IMT Atlantique & Inria team STACK, jolan.philippe@imt-atlantique.fr

1. www.grid5000.fr

Duration 6 months

Salary legal amount of 4,05€ / hour, full time

Location IMT Atlantique, équipe Inria Stack, laboratoire LS2N à Nantes

Références

- [1] Maverick Chardet, Hélène Coullon, and Christian Pérez. Predictable Efficiency for Reconfiguration of Service-Oriented Systems with Concerto. In *CCGrid 2020 : 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing*, 2020. doi:10.1109/CCGrid49817.2020.00-59.
- [2] Maverick Chardet, Hélène Coullon, and Simon Robillard. Toward Safe and Efficient Reconfiguration with Concerto. *Science of Computer Programming*, 2021. doi:10.1016/j.scico.2020.102582.
- [3] Leonardo Leite, Carla Rocha, Fabio Kon, Dejan Milojicic, and Paulo Meirelles. A survey of devops concepts and challenges. *ACM Comput. Surv.*, 2019. doi:10.1145/3359981.
- [4] Kristian Nybom, Jens Smeds, and Ivan Porres. On the impact of mixing responsibilities between devs and ops. In *Agile Processes, in Software Engineering, and Extreme Programming*, 2016.
- [5] Daniel Sokolowski, Pascal Weisenburger, and Guido Salvaneschi. Automating serverless deployments for devops organizations. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2021.