

“Adaptation of micro-services applications in the edge-cloud spectrum”

Ph.D. thesis proposal in the Cloud and Large-Scale Systems group (CLSC) @ UCLouvain, in the context of a 4-year collaborative project between UCLouvain and UNamur.

Advisor: **Pr. Etienne Rivière**, UCLouvain

Context

Emerging collaborative applications such as smart cyber-physical environments, augmented reality, or virtual spaces typically rely on cloud infrastructures to host applications backends, store shared state, and implement the logic of interactions between users.

Traditional cloud infrastructures formed of centralized data centers have enabled ease-of-programming, scalability, and cost-efficiency but are a poor fit for latency-sensitive applications due to the significant network latency between users and data centers.

An emerging model to solve this issue is the edge-cloud model, where large data centers are complemented by smaller ones distributed over the globe and operating close to the users.

Problem

The complexity of edge-cloud infrastructures makes the task of writing efficient application backends a tedious task for programmers. The manual management of application state, the sharing and coherency of this data, and the placement of application components over the edge-cloud spectrum requires to factor in a lot of complexity in the application development, departing from the ease-of-use and high-level abstractions enabled by a centralized cloud model. Moreover, it is nearly impossible for programmers to plan for the diversity of edge-cloud deployments or the need to dynamically migrate or share data between components deployed across different sites.

Approach

The goal of this Ph.D. is to study and propose novel methods for the automatic adaptation of applications in the edge-cloud spectrum.

A specific focus will be on the modern approach of structuring applications as a collection of microservices communicating through modern cloud middleware (event buses, service meshes such as *istio*, etc.) and using modern cloud storage (*e.g.*, CouchDB, redis, etc.) and orchestration (*e.g.*, Kubernetes) solutions.

The goal of the research will be to allow, at runtime, (1) the placement of the application components, copies of the different micro-services, across the different edge and cloud sites; (2) their configuration according to the load of the system and service-level-objectives as expressed by developers and operators (DevOps) alike; and (3) the migration and re-configuration of the services and of their data following changing needs, deployment conditions, and users' locations.

A distinctive feature of this work will be its integration of *software evolution* considerations. An application based on microservices and designed for the centralized cloud may not be directly amenable to the automated adaptation mechanisms that will be developed in this research. It is, therefore, necessary to guide the developer in adapting the application so that

it can benefit fully from this adaptation. This adaptation may require, for instance, revisiting schemas and organization of microservices' data, their mode of interaction (synchronous vs. asynchronous), or the level of consistency that is offered for shared, replicated data.

To this end, the research will contribute, in cooperation with other researchers working on the collaborative project, to the definition of joint models of the application, of its execution, and of its adaptation at runtime. These models and the collected data will drive a feedback loop for the developer and allow her/him to evolve the application towards more amenability to edge-cloud deployment and automated adaptation.

Profile and skills

This thesis is particularly suited for a candidate with interests in Cloud Computing, Operating Systems, Distributed Systems, and Databases. The research will focus on a *systems* approach, meaning that the conceptual contributions must be validated by the development and evaluation of proof-of-concept prototypes. These prototypes will target and integrate with state-of-the-art systems such as the istio service mesh and the Envoy proxy, NoSQL databases such as redis and CouchDB, and cloud management and observability solutions such as Kubernetes and Prometheus.

Proficiency with a low-level systems language such as C or Rust is required. Knowledge of modern cloud languages such as Go and Javascript (node.js), or a strong will to learn them, is also required. The researcher should have a good command of GNU/Linux systems. Experience with distributed systems or cluster computing is a plus.

In addition to the technical skills, the candidate is expected to be able to work in collaboration with other junior and senior researchers active in fields such as software engineering, database engineering, networking, and security. A team player attitude is necessary. Doing a Ph.D. in computer science and particularly in computer systems is a very rewarding but also a very challenging experience due to the need to build prototypes and to run evaluation campaigns on real hardware. The candidate should, therefore, be ready to invest significant conceptual and technical effort towards the realization of the Ph.D.

English will be the primary working language (i.e., speaking French is not a requirement).

Environment

The Ph.D. position is funded for a total of 4 years (competitive salary + research expenses). The Ph.D. candidate will be hosted in the Cloud and Large-Scale Computing group at UCLouvain's ICTEAM (Institute for Information and Communication Technologies, Electronics and Applied Mathematics). She or he will interact with other members of the group as well as members of ICTEAM's Security and Performance of Networked Systems group and of the Research Center on Information Systems Engineering (PReCISE) at the University of Namur.

ICTEAM is a top-level research institute with excellent facilities. The candidate will have access to several distributed clusters with high-end, recent servers and edge resources to test the contributions. UCLouvain is Belgium's largest French-speaking university and is ranked in the first 100 institutions for computer science and information systems by the QS World University Ranking. The position is in Louvain-la-Neuve, a modern city 30 kms south of Brussels with excellent quality of life and connectivity to Belgium, Europe, and the world.

Application

Interested students should send an application file (PDF only, in English) with the following information to both of the promoter:

- An up-to-date CV;
- Link to a copy of the publication(s), if any, and to Masters' thesis;
- If applicable, a link to a GitHub (or similar) profile with examples of personal software contributions;
- Masters' degree transcripts;
- The name and emails of up to three reference persons able to provide an assessment of the application. Reference letters sent by the applicant her/himself will be ignored;
- A short cover email detailing your motivation for pursuing a Ph.D. and your interest in this specific topic.

Application deadline: ASAP -- applications will be screened as they arrive. Feel free to send a declaration of intent or request for clarification before sending a formal application.

Target starting date: September 1st, 2022 (flexible)

Contact: etienne.riviere@uclouvain.be