Optimal deployment of microservices

PROJECT DESCRIPTION

This proposal relates to the research grant funded by MUR for the PRIN 2022 project titled "FREEDA: Failure-Resilient, Energy-aware, and Explainable Deployment of Microservice-based Applications over Cloud-IoT Infrastructures."

The research project's focus is on developing customizable constraint reasoning techniques for deploying Microservices Architectures (MSA) on Cloud-IoT infrastructure. These techniques take into account trade-offs among costs and various deployment requirements, such as hardware, software, network quality, security, failure resilience, and energy sustainability.

Specifically, the objective is to create a novel constraint-based model for optimal MSA deployment using state-of-the-art constraint solving technologies (e.g., CP, SAT/SMT, or MIP - possibly combined via a portfolio approach). Moreover, the FREEDA framework aims to provide explanations for the returned deployments.

The activity plan for the project is outlined below:

ACTIVITY PLAN

State-of-the-art analysis and use-case(s) definition:

Review and analyze scientific literature on constraint solving, microservice architectures, and Cloud-IoT infrastructure. Define one or more use cases for MSA deployment, possibly selecting from existing benchmarks/testbeds or ongoing collaborations with industrial partners.

Model definition

In this stage, the holistic deployment problem is formally abstracted, encompassing infrastructure parameters, variables, domains, as well as hard and soft deployment prerequisites and goals. Various models might be formulated, allowing for the tailored adjustment of the balance between conflicting deployment requirements.

Model implementation

This phases implements the model defined in the previous stage in one or more *solver-independent* constraint modeling language(s) (e.g., MiniZinc or SMT-LIB). In this way, different constraint solvers can be used to solve the model, possibly via a portfolio approach. The framework also aims to explain why/how a given deployment trades-off the MSA deployment requirements.

Model evaluation

Once the framework is defined and integrated with the other components of the FREEDA toolchain, we evaluate its performance on the previously defined use case(s) to assess the viability of the proposed approach. The main output of this workstream will be a realistic use case illustrating the (re)configuration of a MSA deployment over a testbed Cloud-IoT infrastructure.