Multi-armed Bandits for Self-distributing Stateful Services across Networking Infrastructures

Frederico Meletti Rappa
Instituto de Computação
Universidade Estadual de Campinas
Campinas, SP, Brazil

Roberto Rodrigues-Filho
Departamento de Computação
Universidade Federal de Santa Catarina
Araranguá, SC, Brazil
roberto.filho@ufsc.br

Alison R. Panisson
Departamento de Computação
Universidade Federal de Santa Catarina
Araranguá, SC, Brazil
alison.panisson@ufsc.br

Leandro Soriano Marcolino
School of Computing and Communications
Lancaster University
Lancaster, UK
l.marcolino@lancaster.ac.uk

Luiz F. Bittencourt
Instituto de Computação
Universidade Estadual de Campinas
Campinas, SP, Brazil
bit@ic.unicamp.br

***Abstract*—The investigation of stateful service mobility across networking infrastructures is becoming increasingly important as applications require stateful services capable of migrating from centralized cloud data centers to edge computing infrastructures. State-of-the-art approaches propose either machine learning solutions for *stateless service* placement or stateful service mobility using *static and inflexible* state management strategies. We believe these approaches fall short of addressing the full length of the stateful service mobility problem. In this paper, we revisit an emerging concept named self-distributing systems, where a local executing application manages to detach some of its constituent (often stateful) components and place them in remote machines as a solution for stateful service mobility. In previous work, a machine learning approach to support self-distributing systems has not been thoroughly investigated. We model the distribution of stateful components across networking infrastructures as a multi-armed bandits problem and use the UCB1 algorithm to solve it as a first attempt at a flexible solution for stateful service mobility. We conclude the paper by discussing the main challenges and opportunities in this area.**

***Index Terms*—stateful service mobility, edge-cloud infrastructures, reinforcement learning, self-distributing systems**

I. Introduction

Applications deployed over the edge-cloud continuum [1] are often required to move their services across the infrastructure to exploit the trade-off between resource availability and network latency. Therefore, in order to explore code mobility throughout networking infrastructures without adding the complexity of properly managing services' state consistency, the development of stateless service architectures, such as Function-as-a-service and microservices, became popular.

Designing a stateless service-based system is a convenient way to explore the underlying adaptive platforms that currently support modern applications. Cloud and edge-based infrastructures are supported by containers, container-orchestrators (e.g., Mesos[1], Kubernetes[2]), and softwarized networks that allow adaptation of the underlying infrastructure and enable mobility of stateless services wrapped inside containers.

However, avoiding state when developing services is not always possible. Many applications require stateful services to properly function. Nowadays, there is an increase in demand for service mobility capable of migrating from cloud-based platforms to edge computing infrastructures deployed in close proximity to end-user devices. To tackle such issues, some papers have looked into the concept of stateful Function-as-a-service [2], [3], but these solutions often employ static mechanisms to deal with state when moving services across platforms and do not employ any machine learning solutions for service mobility. Moreover, there are many papers in the literature that apply machine learning for service mobility, but they often target stateless services [4], [5].

Self-distributing systems concept [6], on the other hand, enables the flexible distribution of stateful components executing on a local container to other containers executing across infrastructures (e.g., edge to cloud and vice-versa), choosing a state management strategy that better fits the demands of the application. Previous work that employs such a concept, however, either does not explore machine learning in the process of distributing components (using a brute-force online strategy instead), or only targets stateless components. Thus, we currently lack a study of a machine learning approach for the problem of distributing stateful components across networking infrastructures.

In this context, this paper presents the following contributions: *i)* defining the problem of autonomously learning where to place stateful components over distributed platforms, *ii)* showing preliminary results that support the potential of applying a completely autonomous solution for autonomous placement of stateful components, and *iii)* identifying challenges and opportunities in this research field.

The remainder of this paper is organized as follows: Sec. II surveys the most relevant related work; Sec. III revisits the

[1]https://mesos.apache.org/
[2]https://kubernetes.io/