



Título: PopRing: A Popularity-aware Replica Placement for Reducing Latency on Distributed Key-Value Stores

Data: 03/09/2018 Horário: 10:00h Local: Sala de Seminários – Bloco 952

Resumo:

Distributed key-value stores (KVS) are a well-established approach for cloud data-intensive applications, but they were not designed to consider workloads with data access skew, mainly caused by popular data. In this work, we analyze the problem of replica placement on KVS for workloads with data access skew. We formally define our problem as a multi-objective optimization problem because not only load imbalance cost, but replica maintenance and reconfiguration costs affect system performance as well. To solve the replica placement problem, we present the PopRing replica placement component based on Genetic algorithms to find new replica placements efficiently. Next, we extend PopRing framework with a hyper-parameter optimization component based on Bayesian optimization in order to efficiently find the proper importance of load imbalance, replica maintenance, and reconfiguration objectives according to the system latency. To validate our PopRing engine in practice, we implemented a full prototype of PopRing to generate new replica placement schemes in the format of the DHT interface of the popular object store OpenStack-Swift. Then, in our lab environment, we deployed a distributed cluster of the OpenStack-Swift, a benchmark node of the COSBench and the PopRing prototype to run some experiments. From results evaluation, we verified that our solution was able to reduce system latency for different levels of data

access skew without human intervention by auto-tuning its parameters to find a proper replica placement scheme to a given scenario.

Banca:

- Prof. Dr. José Neuman de Souza (MDCC/UFC - Orientador)
- Prof. Dr. Javam de Castro Machado (MDCC/UFC - Coorientador)
- Prof. Dr. Leonardo Oliveira Moreira (UFC)
- Prof. Dr. Joaquim Celestino Junior (UECE)