Título: **An Adaptive Replica Placement Approach for Distributed Key-Value Stores**

Data: **12/03/2019**

Horário: **13:30h**

Local: **Sala de Seminários - Bloco 952**

Resumo:

The use of distributed key-value stores (KVS) has experienced fast adoption by various typesof applications in recent years due to key advantages such as HTTP-based RESTful APIs, highavailability and elasticity. Due to great scalability characteristics, KVS systemscommonly useconsistent hashing as data placement mechanism. Although KVS systems offer many advantages,they were not designed to dynamically adapt to changing workloads which oftenincludedataaccess skew. Furthermore, the underlying physicalstorage nodes may be heterogeneous anddo not expose their performance capabilities to higher level data placement layers. In thispaper, we address those issues and propose an essential step towards a dynamic autonomous solution by leveraging deep reinforcement learning. We design a self-learning approach that incrementally changes the consistent hashing, improving the load balancing among storagenodes. Our approach is dynamic in the sense that is capable of avoiding hot spots, i.e. overloaded storage nodes when facing different workloads including uneven data popularity situations. Also,we design our solution to be pluggable. It assumes no previous

knowledge of the storage nodescapabilities, thus different KVS deployments may make use of it. Our experiments show that ourmethod performs well on changing workloads including data access skew aspects. In addition,we evaluate our strategy on scenarios when storage nodes heterogeneity changes. The resultsdemonstrate that our approach can adapt, building up on the knowledge about the storage node'sperformance it has already acquired.

Banca:

- Prof. Dr. Javam de Castro Machado (MDCC/UFC - Orientador)
- Prof. Dr. Leonardo Oliveira Moreira (UFC - Coorientador)
- Prof. Dr. João Paulo Pordeus Gomes (MDCC/UFC)
- Prof. Dr. Marcial Porto Fernandez (UECE)