

Master Project Proposal: Modular and Adaptive Key-Value Storage Systems

Advisor: João Carlos Antunes Leitão
NOVA Laboratory for Computer Science and Informatics (NOVA LINCS)
and
Departamento de Informática
Faculdade de Ciências e Tecnologia
Universidade NOVA de Lisboa

First Semester, 2018/2019

1 Context & Motivation

Data management, including data replication, is a central aspect in the design of distributed systems and information systems, ranging from web applications to Internet of Things applications. Considering the increasing scale of these systems, which is highly motivated by increasing user bases and larger and richer data sources, the systems commonly used for storing and manage data have to evolve as to naturally support emergent new applications and requirements associated with this.

In this thesis we plan to bootstrap this process, by considering a novel design and implementation of the very popular key-value store design, that was made popular by systems such as Dynamo [4] and Cassandra [3]. In particular, our concern is on making these system more flexible and easier to modify, by exploring a modular design and implementation.

2 Thesis Goals

The goal is to design, implement, and benchmark a modular version of the popular open-source Cassandra no-SQL key value store. This version will offer a design that will allow to easily modify, potentially at run-time, the replication protocol being employed (including the consistency guarantees offered by that protocol), and the replication scheme employed for data objects present in the data storage system. Optionally we also plan to develop solutions to allow the co-existence of multiple replication protocols at the same time.

In summary this thesis aims at:

- Proposal a novel design of a modular key-value store based on Cassandra.

- Propose a set of simple replication protocols, offering different guarantees, that can be plugged into the proposal design of the modular key value store.

- Implement and document a prototype of all developed solutions.

- Conduct an experimental evaluation of the prototypes by comparing with the baseline Cassandra system (this can be potentially conducted by exploiting the FMKe [5] benchmark tool developed in the Computer Systems Group of the NOVA LINCS Laboratory).

3 Notes & Observations

The thesis work will be advised by João Leitão (jc.leitao@fct.unl.pt). Please send an e-mail if you are interested in this master thesis proposal.

This thesis project will be pursued in close collaboration with PhD Student Pedro Fouto and MSc student Gonalo Tomas.

This master thesis proposal is related with the NG-STORAGE (New Generation of data STORAge And manaGement systEMs) financed by the Fundaao para a Ciencia e Tecnologia.

Other relevant bibliography include the ChainReaction system [1] and the C3 system [2].

References

- [1] S. Almeida, J. a. Leitao, and L. Rodrigues. Chainreaction: A causal+ consistent datastore based on chain replication. In *Proceedings of the 8th ACM European Conference on Computer Systems*, EuroSys '13, pages 85–98, New York, NY, USA, 2013. ACM.
- [2] P. Fouto, J. Leitao, and N. Preguia. Consistencia causal em sistemas geo-distribuídos com replicaao parcial. In *Proceedings of the 9th Portuguese Symposium on Informatics*, Coimbra, Portugal, Sept. 2018.
- [3] A. Lakshman and P. Malik. Cassandra: A decentralized structured storage system. *SIGOPS Oper. Syst. Rev.*, 44(2):35–40, Apr. 2010.
- [4] S. Sivasubramanian. Amazon dynamodb: A seamlessly scalable non-relational database service. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, SIGMOD '12, pages 729–730, New York, NY, USA, 2012. ACM.
- [5] G. Tomas, P. Zeller, V. Balesgas, D. Akkoorath, A. Bieniusa, J. a. Leitao, and N. Preguia. Fmke: A real-world benchmark for key-value data stores. In *Proceedings of the 3rd International Workshop on Principles and Practice of Consistency for Distributed Data*, PaPoC '17, pages 7:1–7:4, New York, NY, USA, 2017. ACM.