

Self-optimizing transactional data grids for elastic cloud environments*

Paolo Romano

*Coordinator of the Cloud-TM Project
INESC-ID, Lisbon, Portugal*

Abstract

For several decades, relational databases have represented the indisputable reference solution for transactional data management. Over the last years, however, we have witnessed the proliferation of a new generation of in-memory, transactional data platforms, often referred to as NoSQL data grids.

By relying on a simpler data model (key/value vs relational), lightweight application interfaces (embedded vs. JDBC/ODBC connections) and efficient mechanisms to achieve data durability (in-memory replication vs disk-based logging), NoSQL data grids are designed from the ground up to maximize the scalability of applications deployed on commodity, shared-nothing, distributed infrastructures, such as those typically offered by IaaS cloud providers.

On the other hand, the inherently dynamic nature of elastic cloud computing environments raises the issue of how to ensuring the optimal efficiency of NoSQL data grids in face of fluctuations of the applications' workloads and of the scale of the platform over which these data management platforms are deployed. This is an extremely relevant problem given that, in the usage-based model of the cloud, contenting with static configurations that achieve suboptimal (or not always optimal) performance means having to acquire a larger than required amount of resources to achieve predetermined QoS level, and thus incurring in higher operational costs. Further, manually identifying the optimal configuration of the many (often tightly interdependent) parameters of these platforms, in presence of rapidly fluctuating workloads, is an expensive, challenging and error prone task.

In this talk I will present some of the recent results achieved in the area of self-optimization of NoSQL data grids within the context of Cloud-TM. Cloud-

*This work has been partially supported by the project Cloud-TM (co-financed by the European Commission through the contract no. 257784), the FCT project ARISTOS (PTDC/EIA-EIA/102496/2008) and by FCT (INESC-ID multiannual funding) through the PIDDAC Program Funds.

TM¹ is an FP7 project whose goal is the development of a scalable transactional data platform that will abate the development and administration costs of cloud applications by:

- offering a simple and intuitive programming model, inspired by recent advances in the area of (Distributed) Transactional Memories, that will spare developers from dealing with low-level, error prone mechanisms, such as inter-process synchronization, data distribution, persistence and fault-tolerance;
- automating the provisioning of resources from the Cloud Computing infrastructure based on user defined Quality of Service/operational costs criteria;
- transparently maximizing scalability and efficiency (i.e. the costs/benefits ratio in the Cloud Computing usage-based pricing model) via pervasive self-tuning schemes operating across the various components of the Cloud-TM platform.

I will start by providing an overview of the Cloud-TM project, and then focus on some case studies concerning the self-optimization of two essential (and tightly related) components of NoSQL data grids, namely the replication manager and the group communication system. In particular, I will present techniques for dynamically adapting, based on the workload characteristics and on the number of nodes used by the data grid:

- the distributed algorithms used to ensure consistency across the data grid;
- the configuration of parameters having a deep impact on the performance of fundamental group communication abstractions, such as atomic broadcast or consensus.

I will conclude by highlighting some open research questions in the area of self-optimizing cloud data grids, and by outlining our future work in the context of the Cloud-TM project.

¹ <http://www.cloudtm.eu>