SCIENCESPRINGDAY

FCT FACULDADE DE CIÊNCIAS E TECNOLOGIA UNIVERSIDADE NOVA DE LISBOA

CITI - Center of Informatics and Information Technology

Consistent data in the Cloud

Computing Systems / CR4 Team

Valter Balegas

Advised by Nuno Preguiça and the participation of Sérgio Duarte, Rodrigo Rodrigues

I am currently a first year Ph.D student with keen interest on Geo-Replication and No-SQL databases

Objectives

- *Cloud infra-structures* support an increasing range of services and businesses.
  - E-Commerce platforms, Social networks, media hosting...
- Clients require low latency and high availability.
  - P**oor quality** of **services** leads to **loss of revenue**.
- Current *geo-replicated* cloud infrastructures reduce latency and improve availability, at the cost of increasing inconsistency of data replicas.
- The aim of our work is to provide geo-replicated databases that assure **data invariants** while preserving **low-latency**.

Methodology

- Survey applications for identifying relevant invariants.
- Define local invariants that assure the preservation of global invariants.
- Preserve invariant through access **reservations**.
  - Define reservations model for different data types;
  - Example: Replicated bank account with **always positive balance:**
    - Each replica holds a **reservation** for a **portion** of the balance;
    - Concurrent withdraws do not exceed the reserved portion;
    - Global invariant preserved by the local restriction.

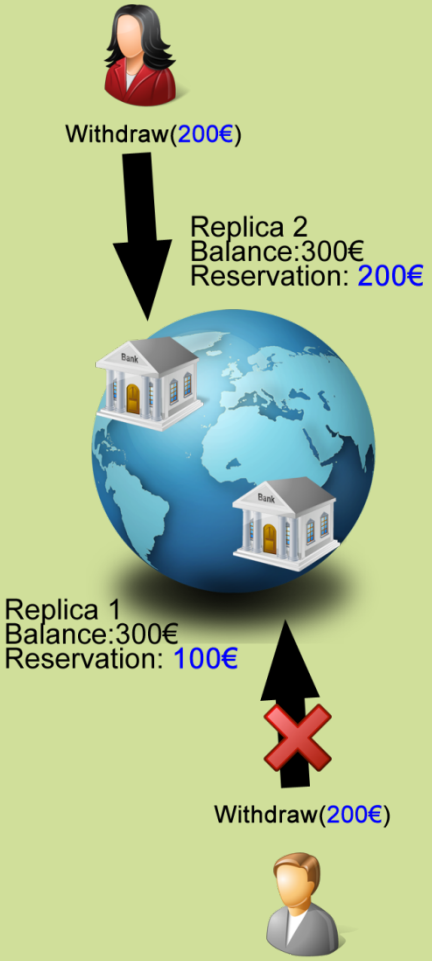

Expected Results

- Definition of relevant data reservations.
- Algorithms to manage reservation assignment.

- System preserving application invariants built on top of an weakly consistent database.
  - How far can we reduce the consistency requirements?
    - Causal consistency, eventual consistency?
- Extend the model to support weakly connected devices.

Funding:

FCT Fundação para a Ciência e a Tecnologia
MINISTÉRIO DA EDUCAÇÃO E CIÊNCIA

PEst-OE/EEI/UI0527/2011
SFRH/BD/87540/2012
PTDC/EIA-EIA/108963/2008