Persistent and Replicated Software Transactional Memory*

António Sousa
Computer Science and Technology Center (CCTC) / U. Minho
Campus de Gualtar, 4710-057 Braga PORTUGAL
email: als@di.uminho.pt, phone: +351 253604477, fax: +351 253604471

MAP-i Thesis Proposal

1 Context

The hardware industry's commitment to the use of multi-core processors as the only practical way to improve computing power for the new generation of computers, brought concurrent programming, finally, into the realms of mainstream programming. Yet, almost all modern programming languages lack adequate abstractions for concurrent programming.

Software Transactional Memories (STMs) are emerging as a powerful paradigm to develop concurrent applications. By relishing the programmer from the burden of managing locks or other low-level concurrency control mechanisms, the reliability of the code is increased and the software development time significantly shortened.

FénixEDU, is a middleware university management system based on the STM technology, in production at the Instituto Superior Técnico (IST) of the Technical University of Lisbon. FénixEDU already augments the basic STM model with persistence, to provide ACID properties to web application functionalities, and replication, because due to the high load of the system the application needs to be deployed in more than one cluster server. However, current solutions to tackle these two important aspects suffer several limitations, including:

- The interface with the persistence susbsystem (a relational database) often requires the use of an excessive amount of memory, increasing garbage collecting and degrading the system performance.
- It is only able to interface a single datastore, which introduces a single point of failure.
- Resolution of conflicts between STMs running on different nodes of the cluster requires nodes to obtain exclusive access to the datastore during the commit phase, limiting the concurrency of the system.

*Pastramy research project (PTDC/EIA/72405/2006).

These limitations are not surprising given that, techniques to persist and replicate STMs are still an open research problem.

Database replication is an effective technique to improve the reliability of a system. In a scenario with a high share of update operations, increasing the number of replicas also increases the I/O requirements. Database replication results in an increase in the I/O load which may not be handled without reformulating the current I/O sub-system.

By ensuring that every piece of information is reliably stored but only by some of the replicas it should be possible enlarge the life time of the I/O sub-system when compared to a full replication scenario.

2 Objectives

The goal of this project is to advance the state of art on techniques to persist and replicate STMs:

- To design and implement a collection of partial replication strategies tailored to build reliable and distributed STMs, in order to increase the scalability and availability of STMs based systems.
- To leverage previous experience with GCS-based replication in the context of the GORDA project for relational database models, in order to develop novel partial replication algorithms, adapted to the unique characteristics of STMs.