Revisiting Epsilon Serializabilty to improve the Database State Machine*

A. Correia Jr.
Universidade do Minho

December 10, 2004

Recently, a large body of research has been exploiting group communication based techniques to improve the dependability and performance of synchronously replicated database systems. In particular, the Database State Machine (DBSM) allows a transaction to be executed at any site and postpone the interaction among distributed concurrent transactions, which can be seen as an optimistic execution. Upon receiving the commit request, it propagates relevant information of the transaction to all replicas. If conflicts arise among concurrent transactions, the order in which the transactions were delivered is used to decide which of them commit or abort.

Unfortunately, the optimistic execution of transactions combined with the strictness of the serializability consistency criterion adopted by the DBSM may lead to a considerable number of aborts. We have been investigating how to relax the consistency criteria of DBSM in a controlled manner according to the *Epsilon Serializability* (ESR) concepts and evaluate the direct benefits in terms of performance.

Basically, ESR relies on the assumption that some transactions may tolerate a certain degree of imprecision to improve the overall performance. Instead of weakening the database consistency, it stretches it in a bounded and consentaneous manner. It allows controlled inconsistencies using a framework that can be applied regardless of the application semantics. For instance, a transaction that retrieves a warehouse's amount of sales may accept a value that does not represent the amount in the last millisecond but an available value in the last couple of seconds.

To evaluate the benefits of our approach, we use a workload based on the TPC-C benchmark annotating the transactions with the allowed degree of inconsistency.

* Research funded by FCT, StrongRep project (FCT POSI / CHS / 41285 / 2001).