

Transactions

D 3.1a First report on formal models for
transactions over structured overlay networks
M12

Task description

T3.1 Formal models for transactions over a structured overlay network (ZIB:3, UCL:3, FT R&D:2)

This task will investigate how to do transactions over a structured overlay network. These will likely be different from classical transactions. One of the important research questions is to resolve the tension between the distributed system and the needs of the application. We will design different compromises and see which are appropriate for applications.

Atomic Commit – Using Paxos

- Atomic Commit built on Consensus
 - Many messages, but small
 - Tolerates F failures in case of $2F + 1$ acceptors
 - Non-blocking
 - Only one additional communication round in case of a failure
- Use Participants as acceptors:
 - Each participants must know all the others
 - Garbage collection more difficult
- Use replicated light-weight Transaction Managers
 - Replicated according to the replication scheme on a transactionID

Transaction Models

- What to do next:
- Analyze:
 - 2PL with Paxos Commit
 - Log-based (Versioning) Concurrency Control with Paxos Commit
- What happens in case of a failure?
 - Transaction Manager fails
 - Transaction Participant fails
 - Paxos Commit can cope with failure of TM and TP up to a certain number

Next Steps

- Usage scenarios
 - More emphasis on response time, prevention of need for roll-back?
 - Contact Manolis Koubarakis
 - Classic database update scenario: RDF database stored in a DHT (in context of a system called Atlas)
- Test an implementation of the transaction models on PlanetLab
- Analyze the transaction models regarding response time, messages ...
- Transactions with replication
 - Better performance through coupling of both mechanisms