

Failure Detection and Recovery for Doubly Distributed Transactions


For Parallel and Distributed Computing



Jay Lofstead, Jai Dayal (student), Ivo Jimenez (student), Carlos Maltzahn
 gflfst@sandia.gov, jdayal3@gatech.edu, ivo@cs.ucsc.edu, carlosm@soe.ucsc.edu



Example Architecture




Messaging Structure

Gather/Broadcast in a tree structure

- Sub-coordinators remove duplicates
- Singleton sub-transactions must be defined early
- Global sub-transactions defined at any time

Subordinate Failure



Addressed Scenarios

Data Movement

- Simulation writes to staging
- Analysis reads, processes, writes to different staging area
- Visualization reads data, generates images, writes to disk, deletes intermediate data


System Reconfiguration

- Deploy resources, wait for proper setup, update registry
- Hide resources, quiesce service, start new successfully, update registry, advertise


Failure Detection/Recovery

Timeouts for message send/receive


- Account for multi-level
- Manage independent processing
- Maintain global state



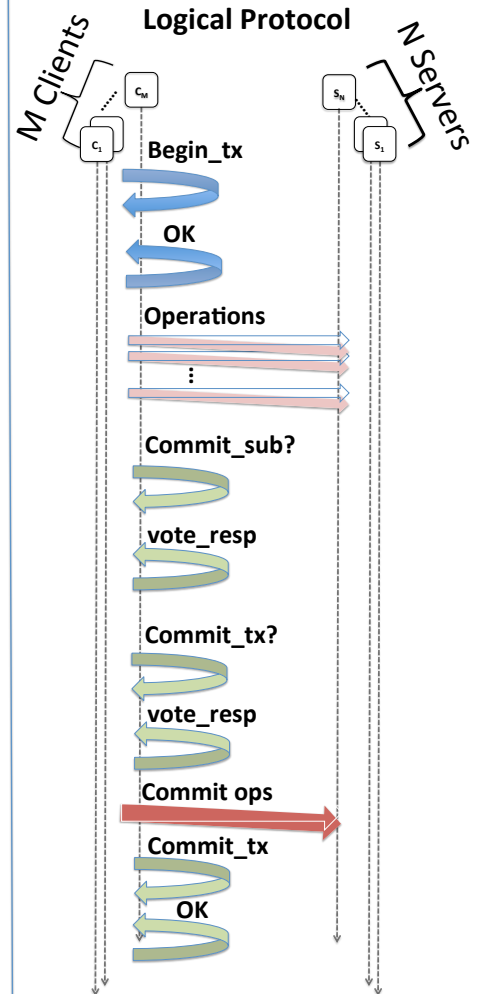
5 seconds timeout




10 seconds timeout



10 seconds timeout




Transactions Performance



Time spent executing
 Transaction protocol is negligible

Sub-Coordinator Failure



Coordinator Failure

