
MODULE *DistributedCommit*

This is a high-level TLA+ specification of a distributed commit protocol. A set of nodes individually decide whether they wish to commit or abort a transaction. Second, the transaction will be committed iff every node wishes to commit.

CONSTANT *Node* set of participating nodes

VARIABLE *nState* state of each participant

Possible states of nodes.

$NState \triangleq \{ \text{"preparing"}, \text{"readyCommit"}, \text{"readyAbort"}, \text{"committed"}, \text{"aborted"} \}$

Initially, every node is preparing the transaction.

$Init \triangleq nState = [n \in Node \mapsto \text{"preparing"}]$

A node decides whether it wishes to commit or abort.

$Decide(n) \triangleq$
 $\wedge nState[n] = \text{"preparing"}$
 $\wedge \vee nState' = [nState \text{ EXCEPT } ![n] = \text{"readyCommit"}]$
 $\vee nState' = [nState \text{ EXCEPT } ![n] = \text{"readyAbort"}]$

A node may commit only if all nodes wish to do so.

$Commit(n) \triangleq$
 $\wedge \forall q \in Node : nState[q] \in \{ \text{"readyCommit"}, \text{"committed"} \}$
 $\wedge nState' = [nState \text{ EXCEPT } ![n] = \text{"committed"}]$

A node aborts if some node requests an abort.

$Abort(n) \triangleq$
 $\wedge \exists q \in Node : nState[q] \in \{ \text{"readyAbort"}, \text{"aborted"} \}$
 $\wedge nState' = [nState \text{ EXCEPT } ![n] = \text{"aborted"}]$

The next-state relation is the disjunction of the above actions.

$Next \triangleq \exists n \in Node : Decide(n) \vee Commit(n) \vee Abort(n)$

Overall specification.

$Spec \triangleq Init \wedge \square [Next]_{nState} \wedge WF_{nState}(Next)$

Check type correctness.

$TypeOK \triangleq nState \in [Node \rightarrow NState]$

Main safety property: nodes may commit only if all nodes agree.

$Safety \triangleq$
 $\forall n \in Node : nState[n] = \text{"committed"} \Rightarrow$
 $\forall q \in Node : nState[q] \in \{ \text{"readyCommit"}, \text{"committed"} \}$

Non-property: no node may ever commit.

$NeverCommit \triangleq \forall n \in Node : nState[n] \neq \text{"committed"}$

Liveness property: nodes will eventually commit or abort. (Requires fairness condition on next-state relation.)

$Liveness \triangleq \forall n \in Node : \diamond(nState[n] \in \{\text{"committed"}, \text{"aborted"}\})$