

MODULE *TwoPhaseCommit*

This is a TLA+ specification of the two-phase commit protocol used for distributed data bases. It models only one instance of the protocol, *i.e.* for a single transaction.

CONSTANT *Node* set of nodes other than the coordinator

VARIABLES

cState, the state of the coordinator
nState, the state of the non-coordinator nodes
committed, nodes that the coordinator knows are *OK* for committing
msgs messages sent during the protocol

vars \triangleq $\langle cState, nState, committed, msgs \rangle$

Possible states of coordinator.

CState \triangleq { "preparing", "committed", "aborted" }

Possible states of non-coordinator participants.

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

Messages sent during the protocol.

Message \triangleq
 node informs coordinator about its decision
 $[kind : \{ "commit", "abort" \}, node : Node] \cup$
 coordinator tells nodes whether to commit or abort
 $[kind : \{ "doCommit", "doAbort" \}]$

commit(*n*) $\triangleq [kind \mapsto "commit", node \mapsto n]$

abort(*n*) $\triangleq [kind \mapsto "abort", node \mapsto n]$

doCommit $\triangleq [kind \mapsto "doCommit"]$

doAbort $\triangleq [kind \mapsto "doAbort"]$

The following predicate specifies what values the variables can take during an execution of the protocol.

TypeOK \triangleq
 $\wedge cState \in CState$
 $\wedge nState \in [Node \rightarrow NState]$
 $\wedge committed \subseteq Node$
 $\wedge msgs \subseteq Message$

The initial state of the protocol.

Init \triangleq
 $\wedge cState = "preparing"$
 $\wedge nState = [n \in Node \mapsto "preparing"]$
 $\wedge committed = \{ \}$
 $\wedge msgs = \{ \}$

The following action formulas describe the possible transitions of the nodes.

A participant decides and informs the coordinator of its decision.

$$\begin{aligned}
Decide(n) &\triangleq \\
&\wedge nState[n] = \text{"preparing"} \\
&\wedge \vee \wedge nState' = [nState \text{ EXCEPT } ![n] = \text{"readyCommit"}] \\
&\quad \wedge msgs' = msgs \cup \{commit(n)\} \\
&\quad \vee \wedge nState' = [nState \text{ EXCEPT } ![n] = \text{"readyAbort"}] \\
&\quad \wedge msgs' = msgs \cup \{abort(n)\} \\
&\wedge \text{UNCHANGED } \langle cState, committed \rangle
\end{aligned}$$

The coordinator receives a new commit decision for some participant. If all participants wish to commit, it sends an order to commit.

$$\begin{aligned}
RcvCommit(n) &\triangleq \\
&\wedge n \notin committed \wedge commit(n) \in msgs \\
&\wedge committed' = committed \cup \{n\} \\
&\wedge \text{IF } committed' = Node \\
&\quad \text{THEN } \wedge cState' = \text{"committed"} \\
&\quad \quad \wedge msgs' = msgs \cup \{doCommit\} \\
&\quad \text{ELSE UNCHANGED } \langle cState, msgs \rangle \\
&\wedge nState' = nState
\end{aligned}$$

The coordinator receives an abort decision and sends an order to abort.

$$\begin{aligned}
RcvAbort(n) &\triangleq \\
&\wedge abort(n) \in msgs \\
&\wedge cState' = \text{"aborted"} \\
&\wedge msgs' = msgs \cup \{doAbort\} \\
&\wedge \text{UNCHANGED } \langle nState, committed \rangle
\end{aligned}$$

A participant receives a commit or abort order from the coordinator.

$$\begin{aligned}
Execute(n) &\triangleq \\
&\wedge \vee doCommit \in msgs \wedge nState' = [nState \text{ EXCEPT } ![n] = \text{"committed"}] \\
&\quad \vee doAbort \in msgs \wedge nState' = [nState \text{ EXCEPT } ![n] = \text{"aborted"}] \\
&\wedge \text{UNCHANGED } \langle cState, committed, msgs \rangle
\end{aligned}$$

The overall next-state relation is the disjunction of the action formulas defined previously.

$$\begin{aligned}
Next &\triangleq \\
&\exists n \in Node : Decide(n) \vee RcvCommit(n) \vee RcvAbort(n) \vee Execute(n)
\end{aligned}$$

$$Spec \triangleq Init \wedge \Box [Next]_{vars} \wedge WF_{vars}(Next)$$

$$\begin{aligned}
NoConflictingOrders &\triangleq \\
&doCommit \in msgs \Rightarrow \\
&\quad \wedge \neg(doAbort \in msgs) \\
&\quad \wedge \forall q \in Node : nState[q] \in \{\text{"readyCommit"}, \text{"committed"}\}
\end{aligned}$$

Correctness properties.

The coordinator never sends both a *doCommit* and a *doAbort* message.

$$CommitOrAbort \triangleq \neg(doCommit \in msgs \wedge doAbort \in msgs)$$

The coordinator may commit only if all participants wish to commit and no participant wishes to abort.

$$\begin{aligned} AbortWins &\triangleq \\ &doCommit \in msgs \Rightarrow \\ &\forall n \in Node : \\ &\quad \wedge commit(n) \in msgs \wedge nState[n] \in \{ \text{"readyCommit"}, \text{"committed"} \} \\ &\quad \wedge abort(n) \notin msgs \end{aligned}$$

$$\begin{aligned} Terminal &\triangleq \\ &\vee cState = \text{"aborted"} \wedge \forall n \in Node : nState[n] = \text{"aborted"} \\ &\vee cState = \text{"committed"} \wedge \forall n \in Node : nState[n] = \text{"committed"} \end{aligned}$$

The protocol may only terminate in a terminal state.

$$CorrectTermination \triangleq (\neg ENABLED \langle Next \rangle_{vars}) \Rightarrow Terminal$$

$$Liveness \triangleq \diamond Terminal$$

Two-phase commitment implements distributed commitment.

$$DC \triangleq \text{INSTANCE } DistributedCommit$$

$$\text{THEOREM } Spec \Rightarrow DC!Spec$$