

EXTENDS *TLC*

This is a preliminary version of a formal specification of the Web Services Atomic Transaction protocol, described in the document by Cabrera et al. We are specifying the safety property of the protocol (what is allowed to happen), not its liveness property (what must eventually happen).

The protocol involves an initiator, a transaction coordinator (*TC*), and a set of participants. The *TC* exchanges messages with the participants. For convenience, we assume that the initiator and *TC* are actually executed on the same processor, so they can be considered to be a single process.

The protocol allows messages to be lost, duplicated, or received out of order. A process is therefore free to resend a message at any time. An implementation will resend a message if it times out without receiving a reply. Such resending is not described explicitly in the specification. Instead, we represent the communication infrastructure by a set *msgs* of all messages that have ever been sent. Since resending a sent message does not change the specification's state, it is allowed by our specification.

An action that, in an implementation, would be enabled by the receipt of certain messages is here enabled by the existence of those messages in *msgs*. Loss of a message is represented by simply not executing that action, even though it is enabled. This works because we are specifying only safety properties, so there is no requirement that an enabled action is ever executed.

Since messages are never removed from *msgs*, receipt of the same message twice is allowed. This can happen in an implementation either because the communication infrastructure delivers a duplicate copy, or because the sender mistakenly believed the original copy had been lost and resent the message. In most cases, the specification says that such duplicate copies are ignored because the response has already been sent, so *msgs* already contains the response. However, in an implementation, receipt of a duplicate message may indicate that the sender resent the message because it never received the response. Hence, an implementation would resend the response. The specification sometimes asserts that such responses are in *msgs*, indicating that they would probably be resent in an implementation.

The CONSTANT and VARIABLES sections define constant parameters and variables.

CONSTANT *Participant* The set of all participants.

VARIABLES *iState*, The state of the initiator.
tcData, The data maintained by the coordinator.
pData, *pData*[*p*] is the data maintained by participant *p*.
msgs The set of all sent messages.

Message \triangleq

The set of all possible messages. A message sent from the *TC* to a participant has a *dest* field indicating its destination. A message from a participant to the *TC* has a *src* field indicating its sender. A participant's "Register" message also has a *reg* field indicating if it's registering as volatile or durable.

[*type* : { "RegisterResponse", "Prepare", "Commit", "Rollback" },
dest : *Participant*]
 \cup [*type* : { "Prepared", "ReadOnly", "Committed", "Aborted" },
src : *Participant*]
 \cup [*type* : { "Register" },
reg : { "volatile", "durable" },
src : *Participant*]

$TypeOK \triangleq$

The type-correctness invariant, indicating the possible values that can be assumed by the variables.

$\wedge iState \in \{\text{"active"}, \text{"committed"}, \text{"aborted"}, \text{"completing"}\}$

Because we are assuming that the initiator and the *TC* are the same process, there is no need for the initiator to have a separate “aborting” state. Once the decision to abort has been made by the *TC*, the then initiator knows that the transaction is aborted.

For convenience, we ignore the action in which the initiator registers the transaction with the *TC*. Instead, we assume that the initiator and *TC* begin in the “active” state. We also assume that the initiator never forgets the outcome. (In an implementation, the initiator can forget about the transaction as soon as it knows the outcome. It is forgetting by the *TC* that’s tricky.)

$\wedge tcData \in$

$[st : \{\text{"active"}, \text{"preparingVolatile"}, \text{"preparingDurable"},$
 $\text{"aborting"}, \text{"committing"}\},$
 $reg : [Participant \rightarrow \{\text{"unregistered"}, \text{"volatile"}, \text{"durable"},$
 $\text{"prepared"}, \text{"readOnly"}, \text{"committed"}\}]]$

$\cup [st : \{\text{"ended"}\},$
 $res : \{\text{"committed"}, \text{"aborted"}\}]$

tcData has an *st* component that indicates the *TC*’s state. While the *TC* is performing the transaction, *tcData* also has a *reg* field such that *tcData.reg[p]* indicates the *TC*’s knowledge of the state of participant *p*. The *TC* enters the “ended” state when it forgets about the transaction. For convenience in understanding the protocol, when *tcData* = “ended”, we let *tcData* have a *res* component that indicates the outcome (whether the transaction committed or aborted). No *TC* actions depend on the value of *tcData.res*.

Note: The document “Web Services Atomic Transaction *Commit*” lists a “Preparing” and a “PreparedSuccess” state. We have split the “Preparing” state into “preparingVolatile”, entered initially, and “preparingDurable”, entered after all volatile participants have prepared. We have eliminated the “PreparedSuccess” state because that is an internal state, not visible to other processes. We model the protocol by having the *TC* go directly from the “preparingDurable” state into either the “Committing” or “Aborting” state. (An implementation is free to split any of our specification’s states into substates.)

$\wedge pData \in [Participant \rightarrow [st : \{\text{"unregistered"}, \text{"prepared"}\}]$

\cup

$[st : \{\text{"registering"}, \text{"active"}, \text{"preparing"}\},$
 $reg : \{\text{"volatile"}, \text{"durable"}\}]$

\cup

$[st : \{\text{"ended"}\},$
 $res : \{\text{"?"}, \text{"committed"}, \text{"aborted"}\}]]$

pData[p] is the data maintained by participant *p*. It contains an *st* field indicating the participant’s state. When in the “registering” or “active” state, there is also a *reg* field indicating if the participant is volatile or durable. When a registered participant forgets about the transaction, it enters the “ended” state.

To help us understand the protocol, when *pData[p].st* = “ended”, there is a field *pData[p].res* that indicates the participant’s knowledge of the outcome when it forgot about the transaction. The value “?” indicates that the participant was read-only, so it didn’t learn the outcome.

Note that the document “Web Services Atomic Transaction *Commit*” does not distinguish between the “unregistered” and “registering” state, which it lumps with the “ended” state into a single “None” state. The document also has “Prepared” and “*PreparedSuccess*” states that we have combined into the single “prepared” state. Those states are not visible externally (that is, by the *TC*) as different states. (An implementation is free to split any of our specification’s states into substates.) We might have lumped the “preparing” and “prepared” states together as well, but it seems convenient to keep them separated because of the interaction of preparing and registration of durable participants.

We have also eliminated the “committing” and “aborting” states, having the participant immediately forget the transaction by going to the “ended” state. In a similar way, we could have eliminated the “aborting” state of the *TC*, but didn’t for no good reason. Perhaps we will in the next version.

$\wedge \text{msgs} \subseteq \text{Message}$

msgs equals a set of messages.

Consistency \triangleq

A predicate that implies that the protocol is not in an inconsistent final or finishing state—that is, where one process thinks the protocol committed and another thinks it has aborted. There are two separate conjuncts, one asserting what’s true if the initiator has reached the “committed” state, and the other asserting what’s true if a participant has reached the “committed” state. These two conjuncts are not logically independent, but we have not eliminated the redundancy in order to make it clear what is being asserted. The invariance of this predicate is the correctness property that we check.

$$\begin{aligned} & \wedge (iState = \text{“committed”}) \\ & \Rightarrow \vee \wedge tcData.st = \text{“ended”} \\ & \quad \wedge tcData.res = \text{“committed”} \\ & \quad \wedge \forall p \in Participant : \\ & \quad \quad \vee pData[p].st = \text{“unregistered”} \\ & \quad \quad \vee \wedge pData[p].st = \text{“ended”} \\ & \quad \quad \quad \wedge pData[p].res \in \{?, \text{“committed”}\} \\ & \vee \wedge tcData.st = \text{“committing”} \\ & \quad \wedge \forall p \in Participant : \\ & \quad \quad \vee pData[p].st \in \{\text{“unregistered”, “prepared”}\} \\ & \quad \quad \vee \wedge pData[p].st = \text{“ended”} \\ & \quad \quad \quad \wedge pData[p].res \in \{?, \text{“committed”}\} \\ & \wedge \forall p \in Participant : \\ & \quad \wedge pData[p].st = \text{“ended”} \\ & \quad \wedge pData[p].res = \text{“committed”} \\ & \Rightarrow \wedge iState = \text{“committed”} \\ & \quad \wedge \vee \wedge tcData.st = \text{“ended”} \\ & \quad \quad \wedge tcData.res = \text{“committed”} \\ & \quad \quad \wedge iState = \text{“committed”} \\ & \quad \vee tcData.st = \text{“committing”} \\ & \quad \wedge \forall pp \in Participant : \\ & \quad \quad \vee pData[pp].st \in \{\text{“unregistered”, “prepared”}\} \\ & \quad \quad \vee \wedge pData[pp].st = \text{“ended”} \\ & \quad \quad \quad \wedge pData[pp].res \in \{?, \text{“committed”}\} \end{aligned}$$

$Init \triangleq$

The initial predicate.

$$\begin{aligned} \wedge iState &= \text{"active"} \\ \wedge tcData &= [st \mapsto \text{"active"}, \\ &\quad reg \mapsto [p \in Participant \mapsto \text{"unregistered"}]] \\ \wedge pData &= [p \in Participant \mapsto [st \mapsto \text{"unregistered"}]] \\ \wedge msgs &= \{\} \end{aligned}$$

THE ACTIONS

The next-state action is the disjunction of the four actions $TCInternal$, $TCRcvMsg$, $PInternal$, and $PRcvMsg$. The major part of the specification consists of the definitions of these four actions, which follow.

$TCInternal \triangleq$

This action describes the actions of the initiator and the internal actions of the TC —that is, the actions of the initiator prompted by timeouts or by a spontaneous action of the initiator. It also describes actions enabled by the TC having received enough messages. (Those actions could be described as occurring when the the TC receives the last message needed to enable it, but it's more convenient to let it be done as a separate internal action.)

$\wedge \vee$ The initiator decides to commit the transaction. It sets its state to "completing". At the same time, the TC sets its state to "preparingVolatile" and sends "Prepare" messages to every participant that has registered as volatile.

$$\begin{aligned} \wedge iState &= \text{"active"} \\ \wedge \forall p \in Participant : \\ &\quad (pData[p].st = \text{"registering"}) \Rightarrow \\ &\quad \wedge pData[p].reg = \text{"durable"} \\ &\quad \wedge \exists pp \in Participant : \\ &\quad \quad \wedge pData[pp].st \in \{\text{"active"}, \text{"preparing"}\} \\ &\quad \quad \wedge pData[pp].reg = \text{"volatile"} \end{aligned}$$

The only participants that may be registering are durable ones that are being installed by a volatile participant that is not yet prepared. It is up to the application to ensure that this condition is met.

$$\begin{aligned} \wedge iState' &= \text{"completing"} \\ \wedge tcData' &= [tcData \text{ EXCEPT } !.st = \text{"preparingVolatile"}] \\ \wedge msgs' &= msgs \cup [type : \{\text{"Prepare"}\}, \\ &\quad dest : \{p \in Participant : \\ &\quad \quad tcData.reg[p] = \text{"volatile"}\}] \end{aligned}$$

\vee Either the initiator decides to abort the transaction while in its "active" state, or else the TC decides to abort it while in a "preparing" state—presumably because of some timeout. The initiator's state is set to "aborted", the TC state is set to "aborting", and "Rollback" messages are sent to every registered participant from which the TC did not already receive a "ReadOnly" message.

$$\begin{aligned} \wedge \vee iState &= \text{"active"} \\ &\quad \vee tcData.st \in \{\text{"preparingVolatile"}, \text{"preparingDurable"}\} \\ \wedge iState' &= \text{"aborted"} \\ \wedge tcData' &= [tcData \text{ EXCEPT } !.st = \text{"aborting"}] \end{aligned}$$

$$\wedge \text{msgs}' = \text{msgs} \cup [\text{type} : \{\text{"Rollback"}\}, \\ \text{dest} : \{p \in \text{Participant} : \\ \text{tcData.reg}[p] \notin \{\text{"unregistered"}, \\ \text{"readOnly"}\}\}]$$

- ∨ The *TC* ends the volatile prepare and begins the durable prepare. It does this when it has received "Prepared" or "ReadOnly" messages from every participant that registered as volatile, and it sends "Prepare" message to every participant that registered as durable.

$$\wedge \text{tcData.st} = \text{"preparingVolatile"} \\ \wedge \forall p \in \text{Participant} : \text{tcData.reg}[p] \neq \text{"volatile"} \\ \wedge \text{tcData}' = [\text{tcData} \text{ EXCEPT } !.st = \text{"preparingDurable"}] \\ \wedge \text{msgs}' = \text{msgs} \cup [\text{type} : \{\text{"Prepare"}\}, \\ \text{dest} : \{p \in \text{Participant} : \\ \text{tcData.reg}[p] = \text{"durable"}\}]$$

∧ UNCHANGED *iState*

- ∨ The *TC* finishes the durable prepare and decides to commit the transaction. It can do this when it has received a "Prepared" or "ReadOnly" message from every durable participant. It sets its state to "committing", notifies the initiator that the transaction has committed, and sends "Commit" messages to every participant that replied with a "Prepared" message (instead of a "ReadOnly" message).

$$\wedge \text{tcData.st} = \text{"preparingDurable"} \\ \wedge \forall p \in \text{Participant} : \text{tcData.reg}[p] \neq \text{"durable"} \\ \wedge \text{tcData}' = [\text{tcData} \text{ EXCEPT } !.st = \text{"committing"}] \\ \wedge \text{msgs}' = \text{msgs} \cup [\text{type} : \{\text{"Commit"}\}, \\ \text{dest} : \{p \in \text{Participant} : \\ \text{tcData.reg}[p] = \text{"prepared"}\}]$$

∧ *iState*' = "committed"

- ∨ The action by which the *TC* forgets the transaction, entering the "ended" state. It can do this if it is in the "aborting" state, or if it is in the "committing" state and has received a "ReadOnly" or "Committed" message from every registered participant.

$$\wedge \vee \text{tcData.st} = \text{"aborting"} \\ \vee \wedge \text{tcData.st} = \text{"committing"} \\ \wedge \forall p \in \text{Participant} : \\ \text{tcData.reg}[p] \in \{\text{"unregistered"}, \text{"readOnly"}, \text{"committed"}\} \\ \wedge \text{tcData}' = [st \mapsto \text{"ended"}, \\ res \mapsto \text{IF } \text{tcData.st} = \text{"aborting"} \text{ THEN } \text{"aborted"} \\ \text{ELSE } \text{"committed"}]$$

∧ UNCHANGED $\langle iState, \text{msgs} \rangle$

∧ UNCHANGED *pData*

The participants' states are left unchanged.

The following action definition uses the TLA+ construct

CASE $B1 \rightarrow P1 \square \dots \square Bn \rightarrow Pn$

This construct is used here (and in most places) when the state predicates B_i , which are called the *guards*, are mutually disjoint—that is, no two of them can be true in the same state. When the guards are mutually disjoint, the CASE expression equals P_i if guard B_i is true. If none of the guards are true, then the value of the expression is undefined. If *TLC* ever evaluates such an undefined expression, it reports an error. Thus, this CASE statement is equivalent to the formula

$$(B_1 \wedge P_1) \vee \dots \vee (B_n \wedge P_n)$$

except that when none of the guards are true, the formula equals FALSE while the value of the CASE statement is undefined.

In this following action, the guards of each CASE formula describe all the possible cases in which the *TC* can receive a particular message.

$TCRcvMsg \triangleq$

The action in which the *TC* receives a message from a participant.

$\exists m \in msgs :$

m is the message being received.

LET $Reply(tp) \triangleq msgs' = msgs \cup \{[type \mapsto tp, dest \mapsto m.src]\}$

Locally defines $Reply(tp)$ to be the action of sending a message of type tp to the sender of message m .

$HaveSent(tp) \triangleq \exists mm \in msgs : (mm.type = tp) \wedge (mm.dest = m.src)$

Locally defines $HaveSent(tp)$ to be true iff the *TC* has sent a message of type tp to the sender of m .

IN $\vee m$ is a “Register” message.

$\wedge m.type = \text{“Register”}$

\wedge CASE The normal case, in which the *TC* state is either “active”, or else this is a durable participant registering while the *TC* is performing the volatile prepare.

$\vee tcData.st = \text{“active”}$

$\vee \wedge tcData.st = \text{“preparingVolatile”}$

$\wedge m.reg = \text{“durable”}$

\rightarrow The *TC* sends a “RegisterResponse” reply to the sender, and sets the appropriate $tcData.reg$ component (the one corresponding to the sender) to the contents of the reg field of m .

$\wedge Reply(\text{“RegisterResponse”})$

$\wedge tcData' = [tcData \text{ EXCEPT } !.reg[m.src] = m.reg]$

\square If the *TC* is in a “preparing” or “committing” state or has forgotten the transaction, then m is a duplicate message to which the *TC* has already responded and so is ignored.

$\wedge \vee \wedge tcData.st = \text{“preparingVolatile”}$

$\wedge m.reg = \text{“volatile”}$

$\vee tcData.st \in \{\text{“preparingDurable”}, \text{“committing”}\}$

$\wedge HaveSent(\text{“RegisterResponse”})$

\rightarrow

UNCHANGED $\langle tcData, msgs \rangle$

□ If the *TC* is in the “aborting” state, then if the sender is not already registered, then the decision to abort was made before the sender could register, in which case a “Rollback” message is sent. (Is this correct?) Otherwise, this is a duplicate message to which the *TC* has already responded, and it has already sent a “Rollback” message to the sender (unless the participant responded “ReadOnly” to a “Prepare” message).

$\wedge tcData.st = \text{“aborting”}$
 $\wedge \vee tcData.reg[m.src] \in \{\text{“unregistered”, “readOnly”}\}$
 $\vee \wedge tcData.reg[m.src] \in \{\text{“volatile”, “durable”, “prepared”}\}$
 $\wedge HaveSent(\text{“Rollback”})$

→

\wedge IF $tcData.reg[m.src] = \text{“unregistered”}$
 THEN $Reply(\text{“Rollback”})$
 ELSE UNCHANGED $msgs$

\wedge UNCHANGED $tcData$

□ If the *TC* is in the “ended” state, then either the transaction has been aborted before the sender had a chance to register, or else the “Register” message is old and, if it was committed, then the sender has already forgotten the transaction and hence will ignore any message it receives. Therefore, it’s safe to send a “Rollback” message.

$\wedge tcData.st = \text{“ended”}$
 $\wedge \vee tcData.res = \text{“aborted”}$
 $\vee pData[m.src].st = \text{“ended”}$

→

$\wedge Reply(\text{“Rollback”})$

\wedge UNCHANGED $tcData$

\wedge UNCHANGED $\langle iState, pData \rangle$

The initiator’s state and the participants’ data are not changed.

∨ m is a “Prepared” message.

$\wedge m.type = \text{“Prepared”}$

\wedge CASE The normal case, in which the *TC* has sent a “Prepare” message and is waiting for the sender’s reply.

$\vee \wedge tcData.st = \text{“preparingVolatile”}$

$\wedge tcData.reg[m.src] = \text{“volatile”}$

$\vee \wedge tcData.st = \text{“preparingDurable”}$

$\wedge tcData.reg[m.src] = \text{“durable”}$

→

$\wedge tcData' = [tcData \text{ EXCEPT } !.reg[m.src] = \text{“prepared”}]$

\wedge UNCHANGED $msgs$

□ The *TC* has forgotten the transaction.

$tcData.st = \text{“ended”}$

→

The *TC* sends a “Rollback” reply to the sender. The transaction could either have aborted or committed. However, if the transaction has committed, then the sender will know that it has and will ignore the “Rollback” message.

$\wedge \text{Reply}(\text{“Rollback”})$
 $\wedge \text{UNCHANGED } tcData$

□

If the *TC* is either in a “preparing” state, or in the “aborting” or “committing” state, then it has already received and acted on a copy of *m*, so it does nothing.

$\vee \wedge tcData.st \in \{\text{“preparingVolatile”}, \text{“preparingDurable”}\}$
 $\wedge tcData.reg[m.src] = \text{“prepared”}$
 $\vee \wedge tcData.st = \text{“aborting”}$
 $\wedge \text{HaveSent}(\text{“Rollback”})$
 $\vee \wedge tcData.st = \text{“committing”}$
 $\wedge \text{HaveSent}(\text{“Commit”})$

→

$\text{UNCHANGED } \langle tcData, msgs \rangle$

$\wedge \text{UNCHANGED } \langle iState, pData \rangle$

The initiator’s state is not changed.

∨ *m* is a “ReadOnly” message.

$\wedge m.type = \text{“ReadOnly”}$

∧ CASE The normal case, in which the *TC* has sent a “Prepare” message and is waiting for the sender’s reply.

$\vee \wedge tcData.st = \text{“preparingVolatile”}$
 $\wedge tcData.reg[m.src] = \text{“volatile”}$
 $\vee \wedge tcData.st = \text{“preparingDurable”}$
 $\wedge tcData.reg[m.src] = \text{“durable”}$

→

The *TC* sets its *tcData.reg* component corresponding to the sender to “readOnly”.

$\wedge tcData' = [tcData \text{ EXCEPT } !.reg[m.src] = \text{“readOnly”}]$
 $\wedge \text{UNCHANGED } msgs$

□

If the *TC* has forgotten the transaction, then either *m* is a duplicate message, or else it was decided to abort the transaction before the *TC* received the response to the “Prepare” message it sent to the sender of *m*. In either case, the message is ignored.

$tcData.st = \text{“ended”}$

→

$\text{UNCHANGED } \langle tcData, msgs \rangle$

□

In the following cases, *m* is a duplicate of a message that the *TC* has already received and it is ignored.

$$\begin{aligned}
& \vee \wedge tcData.st \in \{ \text{"preparingVolatile"}, \text{"preparingDurable"} \} \\
& \quad \wedge tcData.reg[m.src] = \text{"readOnly"} \\
& \vee \wedge tcData.st = \text{"aborting"} \\
& \quad \wedge \vee tcData.reg[m.src] = \text{"readOnly"} \\
& \quad \quad \vee \wedge tcData.reg[m.src] \in \{ \text{"volatile"}, \text{"durable"} \} \\
& \quad \quad \wedge HaveSent(\text{"Rollback"}) \\
& \vee tcData.st = \text{"committing"} \\
\rightarrow & \\
& \text{UNCHANGED } \langle tcData, msgs \rangle
\end{aligned}$$

\wedge UNCHANGED $\langle iState, pData \rangle$

The initiator's state and the participants' data are not changed.

\vee m is an "Aborted" message.

$\wedge m.type = \text{"Aborted"}$

\wedge CASE The normal case, in which the TC receives the message when it is "active" or in a "preparing" state and the sender has not replied to a "Prepare" message. In this case, the transaction is aborted, the TC sends "Rollback" messages to all registered participants from which it has not already received a "ReadOnly" message, and the initiator is notified that the transaction has been aborted.

$$\begin{aligned}
& \wedge tcData.st \in \{ \text{"active"}, \text{"preparingVolatile"}, \text{"preparingDurable"} \} \\
& \wedge tcData.reg[m.src] \in \{ \text{"unregistered"}, \text{"volatile"}, \text{"durable"} \}
\end{aligned}$$

\rightarrow

$\wedge iState' = \text{"aborted"}$

$\wedge tcData' = [tcData \text{ EXCEPT } !.st = \text{"aborting"}]$

$\wedge msgs' = msgs \cup$

$[type : \{ \text{"Rollback"} \},$

$dest : \{ p \in Participant :$

$tcData.reg[p] \notin$

$\{ \text{"unregistered"}, \text{"readOnly"} \} \}$

11 Dec 2003, changed "ReadOnly" \rightarrow "readOnly"

typo found by Colin Campbell

\wedge UNCHANGED $pData$

\square If the TC is already in the "aborting" state or it has forgotten an aborted transaction, then the TC has already sent a "Rollback" message to the sender (perhaps because m is a duplicate of a message the TC already received). If the TC has forgotten a committed transaction, then this "Aborted" message was sent because the sender received an obsolete "Prepare" message after it had forgotten the transaction. In either case, m is ignored.

$\wedge \vee tcData.st = \text{"aborting"}$

$\vee \wedge tcData.st = \text{"ended"}$

$\wedge \vee \wedge tcData.res = \text{"aborted"}$

$\wedge HaveSent(\text{"Rollback"})$

$\vee \wedge tcData.res = \text{"committed"}$

$$\begin{aligned} & \wedge pData[m.src].st = \text{"ended"} \\ & \wedge pData[m.src].res = \text{"committed"} \\ \rightarrow & \\ & \text{UNCHANGED } \langle tcData, pData, iState, msgs \rangle \\ \square & \text{ If the } TC \text{ is in the "committing" or "ended" state, or it is in a preparing} \\ & \text{ state and the sender has already responded to the "Prepared" message, then} \\ & \text{ the message is ignored. (It could have been sent in response to a duplicate} \\ & \text{ "Prepared" message after the participant had reached the "ended" state.)} \\ & \wedge \vee tcData.st \in \{\text{"committing"}, \text{"ended"}\} \\ & \quad \vee \wedge tcData.st \in \{\text{"active"}, \text{"preparingVolatile"}, \\ & \quad \quad \quad \text{"preparingDurable"}\} \\ & \quad \wedge tcData.reg[m.src] \in \{\text{"prepared"}, \text{"readOnly"}, \\ & \quad \quad \quad \text{"committed"}\} \\ \rightarrow & \\ & \text{UNCHANGED } \langle tcData, pData, iState, msgs \rangle \\ \vee & \text{ } m \text{ is a "Committed" message.} \\ & \wedge m.type = \text{"Committed"} \\ & \wedge \text{CASE } \text{The normal case, in which the } TC \text{ is in the "committing" state. In this} \\ & \quad \text{ case, it sets the element of } tcData.reg \text{ corresponding to the sender to} \\ & \quad \text{ "committed".} \\ & \quad tcData.st = \text{"committing"} \\ \rightarrow & \\ & \quad tcData' = [tcData \text{ EXCEPT } !.reg[m.src] = \text{"committed"}] \\ \square & \text{ If the } TC \text{ has forgotten the transaction, then the transaction has been} \\ & \text{ committed and } m \text{ is ignored.} \\ & \wedge tcData.st = \text{"ended"} \\ & \wedge tcData.res = \text{"committed"} \\ \rightarrow & \\ & \quad \text{UNCHANGED } tcData \\ & \wedge \text{UNCHANGED } \langle iState, pData, msgs \rangle \\ & \quad \text{The initiator's state and the participants' data are unchanged, and no messages} \\ & \quad \text{are sent.} \end{aligned}$$

$P_{Internal} \triangleq$

This action describes the internal actions of the participants—actions that occur “spontaneously”, either prompted by timeouts or, as in the case of the register action, by some communication external to the protocol.

$\exists p \in Participant :$

p is the participant performing the action.

LET $SendMsg(tp) \triangleq msgs' = msgs \cup \{[type \mapsto tp, src \mapsto p]\}$
 $SendRegisterMsg(rg) \triangleq msgs' = msgs \cup \{[type \mapsto \text{"Register"},$
 $\quad \quad \quad src \mapsto p,$
 $\quad \quad \quad reg \mapsto rg]\}$

Locally defined action expressions. $SendMsg(tp)$ sends a message of type tp from participant p to the TC . $SendRegisterMsg(rg)$ sends a $Register$ rg message, where rg is either “durable” or “volatile”.

IN \vee p registers as a volatile participant. It can do this only if it is *unregistered* and the initiator is in the “active” state.

$$\begin{aligned} &\wedge pData[p].st = \text{“unregistered”} \\ &\wedge iState = \text{“active”} \\ &\wedge pData' = [pData \text{ EXCEPT } ![p] = [st \mapsto \text{“registering”}, \\ &\hspace{10em} reg \mapsto \text{“volatile”}]] \\ &\wedge SendRegisterMsg(\text{“volatile”}) \\ &\wedge \text{UNCHANGED } \langle iState, tcData \rangle \end{aligned}$$

\vee p registers as a durable participant. It can do this only if it is *unregistered* and, if either the initiator is in the “active” state, or there is some volatile participant that is willing to wait for p to register before preparing. Since we don’t model “willingness to wait”, we allow the participant to register as long as there is some volatile participant that can wait for it.

$$\begin{aligned} &\wedge pData[p].st = \text{“unregistered”} \\ &\wedge \vee iState = \text{“active”} \\ &\quad \vee \exists pp \in Participant : \\ &\quad \quad \wedge pData[pp].st \in \{ \text{“active”}, \text{“preparing”} \} \\ &\quad \quad \wedge pData[pp].reg = \text{“volatile”} \\ &\wedge pData' = [pData \text{ EXCEPT } ![p] = [st \mapsto \text{“registering”}, \\ &\hspace{10em} reg \mapsto \text{“durable”}]] \\ &\wedge SendRegisterMsg(\text{“durable”}) \\ &\wedge \text{UNCHANGED } \langle iState, tcData \rangle \end{aligned}$$

\vee p spontaneously aborts and forgets about the transaction. We do not allow it to abort in the “registering” state. If we allowed this, then we could wind up with a situation in which a participant aborted before it registered, and the transaction committed anyway. In practice, there will have to be some way for a participant to give up when it hasn’t received a *RegisterResponse* message. However, to do this, it must learn from the initiator or a volatile participant that the transaction aborted so it can forget about it. Since we are not modeling this kind of inter-participant communication, we do not model this procedure.

$$\begin{aligned} &\wedge pData[p].st \in \{ \text{“active”}, \text{“preparing”} \} \\ &\wedge pData' = [pData \text{ EXCEPT } ![p] = [st \mapsto \text{“ended”}, \\ &\hspace{10em} res \mapsto \text{“aborted”}]] \\ &\wedge SendMsg(\text{“Aborted”}) \\ &\wedge \text{UNCHANGED } \langle iState, tcData \rangle \end{aligned}$$

\vee p either prepares or becomes read-only. If p is volatile, then it cannot do this if there is a durable participant that is in the “registering” state, and there is no other volatile participant to wait for it to register.

$$\begin{aligned} &\wedge pData[p].st = \text{“preparing”} \\ &\wedge \vee pData[p].reg = \text{“durable”} \\ &\quad \vee \neg \exists dp \in Participant : \wedge pData[dp].st = \text{“registering”} \\ &\quad \quad \wedge pData[dp].reg = \text{“durable”} \end{aligned}$$

$$\wedge pData' =$$

$$[pData \text{ EXCEPT } ![m.dest] = [st \mapsto \text{"ended"},$$

$$res \mapsto \text{"aborted"}]]$$

$$\wedge Reply(\text{"Aborted"})$$

- If the participant has already finished, then this is ignored. It is possible for this message to arrive even though the participant has ended by committing the transaction. In this case, the "Rollback" message was generated by a duplicate "Register" message arriving at the *TC* after it had forgotten the transaction.

$$pData[m.dest].st = \text{"ended"}$$

$$\rightarrow$$

$$\text{UNCHANGED } \langle pData, msgs \rangle$$

$$\wedge \text{UNCHANGED } \langle iState, tcData \rangle$$

$$Next \triangleq TCInternal \vee TCRcvMsg \vee PInternal \vee PRcvMsg$$

The specification's next-state action.

$$vars \triangleq \langle iState, tcData, pData, msgs \rangle$$

The tuple of all variables.

$$Spec \triangleq Init \wedge \square[Next]_{vars}$$

The complete spec of the two-phase *Commit* protocol.

THEOREM $Spec \Rightarrow \square(\text{TypeOK} \wedge \text{Consistency})$

This theorem asserts that the predicates *TypeOK* and *Consistency* are invariants of the specification. *TLC* checked this with 4 participants. It generated 10269919 states, 504306 of them were distinct. The longest non-repeating behavior had 45 states. It took *TLC* about 4-1/4 minutes on a 2-processor, 2.4GHz machine.

Last modified on *Thu Dec 11 13:27:41 PT 2003* by lamport