

Time-stamp: < 10 Jun 2002 at 12:39:50 by charpov on *berlioz.cs.unh.edu* >

Atomic Commitment Protocol with Simple Broadcast primitive (ACP-SB)

From:

Sape Mullender , editor. Distributed Systems.

Chapter 6: Non-Blocking Atomic Commitment, by Ö. Babaoglu and S. Toueg.

1993.

Synchronous communication has been replaced with (implicit) asynchronous communication. Failures are detected “magically” instead of relying on timeouts.

This version of the protocol uses a “simple broadcast” where a broadcast is simply a series of messages sent, possibly interrupted by a failure. Consequently, this algorithm is “non terminating” and property *AC5* does not hold.

CONSTANTS

<i>participants</i> ,	set of participants
<i>yes, no</i> ,	vote
<i>undecided, commit, abort</i> ,	decision
<i>waiting</i> ,	coordinator state wrt a participant
<i>notsent</i>	broadcast state wrt a participant

VARIABLES

<i>participant</i> ,	participants (N)
<i>coordinator</i>	coordinator (1)

$$TypeInvParticipant \triangleq participant \in [$$

$$participants \rightarrow [$$

$$vote : \{yes, no\},$$

$$alive : BOOLEAN ,$$

$$decision : \{undecided, commit, abort\},$$

$$faulty : BOOLEAN ,$$

$$voteSent : BOOLEAN$$

$$]$$

$$]$$

$$TypeInvCoordinator \triangleq coordinator \in [$$

$$request : [participants \rightarrow BOOLEAN],$$

$$vote : [participants \rightarrow \{waiting, yes, no\}],$$

$$broadcast : [participants \rightarrow \{commit, abort, notsent\}],$$

$$decision : \{commit, abort, undecided\},$$

$$alive : BOOLEAN ,$$

$$faulty : BOOLEAN$$

$$]$$

$$TypeInv \triangleq TypeInvParticipant \wedge TypeInvCoordinator$$

Initially:

All the participants:

have a yes/no vote

are alive and not faulty

have not sent in their votes yet

are undecided about final decision

The coordinator:

has not sent vote requests yet

has not received votes from any participant

is alive and not faulty

has not sent broadcast messages to any participant

is undecided about final decision

$$\text{InitParticipant} \triangleq \text{participant} \in [$$

$$\quad \text{participants} \rightarrow [$$

$$\quad \quad \text{vote} \quad : \{\text{yes}, \text{no}\},$$

$$\quad \quad \text{alive} \quad : \{\text{TRUE}\},$$

$$\quad \quad \text{decision} : \{\text{undecided}\},$$

$$\quad \quad \text{faulty} \quad : \{\text{FALSE}\},$$

$$\quad \quad \text{voteSent} : \{\text{FALSE}\}$$

$$\quad]$$

$$]$$

$$\text{InitCoordinator} \triangleq \text{coordinator} \in [$$

$$\quad \text{request} \quad : [\text{participants} \rightarrow \{\text{FALSE}\}],$$

$$\quad \text{vote} \quad \quad : [\text{participants} \rightarrow \{\text{waiting}\}],$$

$$\quad \text{alive} \quad \quad : \{\text{TRUE}\},$$

$$\quad \text{broadcast} : [\text{participants} \rightarrow \{\text{notsent}\}],$$

$$\quad \text{decision} \quad : \{\text{undecided}\},$$

$$\quad \text{faulty} \quad \quad : \{\text{FALSE}\}$$

$$]$$

$$\text{Init} \triangleq \text{InitParticipant} \wedge \text{InitCoordinator}$$

COORDINATOR STATEMENTS

```
request(i):
IF
  coordinator is alive
  request for vote has not been sent to participant i
THEN
  request for vote is sent to participant i
```

$$\text{request}(i) \triangleq \wedge \text{coordinator.alive}$$

$$\wedge \neg \text{coordinator.request}[i]$$

$$\wedge \text{coordinator}' = [\text{coordinator EXCEPT !.request} =$$

$$\quad [@ \text{ EXCEPT } ![i] = \text{TRUE}]$$

$$]$$

$$\wedge \text{UNCHANGED} \langle \text{participant} \rangle$$

```
getVote(i):
IF
  coordinator is alive
  coordinator is still undecided
  coordinator has sent request for votes to all participants
  coordinator is waiting to receive a vote from participant i
  participant i has sent the vote message
THEN
  the coordinator can record the vote of participant i
```

$$\text{getVote}(i) \triangleq \wedge \text{coordinator.alive}$$

$$\wedge \text{coordinator.decision} = \text{undecided}$$

$$\wedge \forall j \in \text{participants} : \text{coordinator.request}[j]$$

$$\wedge \text{coordinator.vote}[i] = \text{waiting}$$

$$\wedge \text{participant}[i].\text{voteSent}$$

$$\wedge \text{coordinator}' = [\text{coordinator EXCEPT !.vote} =$$

$$\quad [@ \text{ EXCEPT } ![i] = \text{participant}[i].\text{vote}]$$

$$]$$

$$\wedge \text{UNCHANGED} \langle \text{participant} \rangle$$

```
detectFault(i):
IF
  coordinator is alive
```

```

coordinator is still undecided
coordinator has sent request for votes to all participants
coordinator is waiting for vote from participant  $i$ 
participant  $i$  has died without sending its vote
THEN
coordinator times out on participant  $i$  and decides to abort

```

$$\begin{aligned}
detectFault(i) \triangleq & \wedge coordinator.alive \\
& \wedge coordinator.decision = undecided \\
& \wedge \forall j \in participants : coordinator.request[j] \\
& \wedge coordinator.vote[i] = waiting \\
& \wedge \neg participant[i].alive \\
& \wedge \neg participant[i].voteSent \\
& \wedge coordinator' = [coordinator \text{ EXCEPT } !.decision = abort] \\
& \wedge \text{UNCHANGED } \langle participant \rangle
\end{aligned}$$

```

makeDecision:
IF
coordinator is alive
coordinator is undecided
coordinator has received votes from all participants
THEN
IF
all votes are yes
THEN
coordinator decides commit
ELSE
coordinator decides abort

```

$$\begin{aligned}
makeDecision \triangleq & \wedge coordinator.alive \\
& \wedge coordinator.decision = undecided \\
& \wedge \forall j \in participants : coordinator.vote[j] \in \{yes, no\} \\
& \wedge \forall j \in participants : coordinator.vote[j] = yes \\
& \quad \wedge coordinator' = [coordinator \text{ EXCEPT } !.decision = commit] \\
& \quad \vee \exists j \in participants : coordinator.vote[j] = no \\
& \quad \quad \wedge coordinator' = [coordinator \text{ EXCEPT } !.decision = abort] \\
& \wedge \text{UNCHANGED } \langle participant \rangle
\end{aligned}$$

```

coordBroadcast( $i$ ) (simple broadcast):
IF
coordinator is alive
coordinator has made a decision
coordinator has not sent the decision to participant  $i$ 
THEN
coordinator sends its decision to participant  $i$ 

```

$$\begin{aligned}
coordBroadcast(i) \triangleq & \wedge coordinator.alive \\
& \wedge coordinator.decision \neq undecided \\
& \wedge coordinator.broadcast[i] = notsent \\
& \wedge coordinator' = [coordinator \text{ EXCEPT } !.broadcast = \\
& \quad [@ \text{ EXCEPT } ![i] = coordinator.decision \\
& \quad] \\
& \wedge \text{UNCHANGED } \langle participant \rangle
\end{aligned}$$

```

coordDie:
IF
coordinator is alive
THEN
coordinator dies
coordinator is now faulty

```

$$\begin{aligned} \text{coordDie} &\triangleq \wedge \text{coordinator.alive} \\ &\wedge \text{coordinator}' = [\text{coordinator EXCEPT !.alive = FALSE, !.faulty = TRUE}] \\ &\wedge \text{UNCHANGED } \langle \text{participant} \rangle \end{aligned}$$

PARTICIPANT STATEMENTS

sendVote(i):

IF
 participant is alive
 participant has received a request for vote
 THEN
 participant sends vote

$$\begin{aligned} \text{sendVote}(i) &\triangleq \wedge \text{participant}[i].\text{alive} \\ &\wedge \text{coordinator.request}[i] \\ &\wedge \text{participant}' = [\text{participant EXCEPT ![i] =} \\ &\quad [@ \text{ EXCEPT !.voteSent = TRUE} \\ &\quad] \\ &\wedge \text{UNCHANGED } \langle \text{coordinator} \rangle \end{aligned}$$

abortOnVote(i):

IF
 participant is alive
 participant is undecided
 participant has sent its vote to the coordinator
 participant's vote is no
 THEN
 participant decides (unilaterally) to abort

$$\begin{aligned} \text{abortOnVote}(i) &\triangleq \wedge \text{participant}[i].\text{alive} \\ &\wedge \text{participant}[i].\text{decision} = \text{undecided} \\ &\wedge \text{participant}[i].\text{voteSent} \\ &\wedge \text{participant}[i].\text{vote} = \text{no} \\ &\wedge \text{participant}' = [\text{participant EXCEPT ![i] =} \\ &\quad [@ \text{ EXCEPT !.decision} = \text{abort} \\ &\quad] \\ &\wedge \text{UNCHANGED } \langle \text{coordinator} \rangle \end{aligned}$$

abortOnTimeoutRequest(i):

IF
 participant is alive
 participant is still undecided
 coordinator has died without sending request for vote
 THEN
 participant decides (unilaterally) to abort

$$\begin{aligned} \text{abortOnTimeoutRequest}(i) &\triangleq \wedge \text{participant}[i].\text{alive} \\ &\wedge \text{participant}[i].\text{decision} = \text{undecided} \\ &\wedge \neg \text{coordinator.alive} \\ &\wedge \neg \text{coordinator.request}[i] \\ &\wedge \text{participant}' = [\text{participant EXCEPT ![i] =} \\ &\quad [@ \text{ EXCEPT !.decision} = \text{abort} \\ &\quad] \\ &\wedge \text{UNCHANGED } \langle \text{coordinator} \rangle \end{aligned}$$

decide(i):

IF
 participant is alive
 participant is undecided

participant has recieved decision from the coordinator
 THEN
 participant decides according to decision from coordinator

$$\begin{aligned}
 decide(i) \triangleq & \wedge participant[i].alive \\
 & \wedge participant[i].decision = undecided \\
 & \wedge coordinator.broadcast[i] \neq notsent \\
 & \wedge participant' = [participant \text{ EXCEPT } ![i] = \\
 & \quad [@ \text{ EXCEPT } !.decision = coordinator.broadcast[i]] \\
 & \quad] \\
 & \wedge \text{UNCHANGED } \langle coordinator \rangle
 \end{aligned}$$

parDie(i):
 IF
 participant is alive
 THEN
 participant dies and is now faulty

$$\begin{aligned}
 parDie(i) \triangleq & \wedge participant[i].alive \\
 & \wedge participant' = [participant \text{ EXCEPT } ![i] = \\
 & \quad [@ \text{ EXCEPT } !.alive = \text{FALSE}, !.faulty = \text{TRUE}] \\
 & \quad] \\
 & \wedge \text{UNCHANGED } \langle coordinator \rangle
 \end{aligned}$$

FOR N PARTICIPANTS

$$parProg(i) \triangleq sendVote(i) \vee abortOnVote(i) \vee abortOnTimeoutRequest(i) \vee decide(i)$$

$$parProgN \triangleq \exists i \in participants : parDie(i) \vee parProg(i)$$

$$coordProgA(i) \triangleq request(i) \vee getVote(i) \vee detectFault(i) \vee coordBroadcast(i)$$

$$coordProgB \triangleq makeDecision \vee \exists i \in participants : coordProgA(i)$$

$$coordProgN \triangleq coordDie \vee coordProgB$$

$$progN \triangleq parProgN \vee coordProgN$$

Death transitions are left outside of fairness

$$\begin{aligned}
 fairness \triangleq & \wedge \forall i \in participants : \text{WF}_{\langle coordinator, participant \rangle}(parProg(i)) \\
 & \wedge \text{WF}_{\langle coordinator, participant \rangle}(coordProgB)
 \end{aligned}$$

$$Spec \triangleq Init \wedge \square [progN]_{\langle coordinator, participant \rangle} \wedge fairness$$

CORRECTNESS SPECIFICATION

This specification follows the original paper, except that AC3 is stronger: It forces participants to abort if one vote at least is NO (in the absence of failure).

The specification is split between safety and liveness.

SAFETY

All participants that decide reach the same decision

$$\begin{aligned}
 AC1 \triangleq & \square \forall i, j \in participants : \\
 & \quad \vee participant[i].decision \neq commit \\
 & \quad \vee participant[j].decision \neq abort
 \end{aligned}$$

If any participant decides commit, then all participants must have votes YES
 $AC2 \triangleq \Box ((\exists i \in participants : participant[i].decision = commit) \Rightarrow (\forall j \in participants : participant[j].vote = yes))$

If any participant decides abort, then:
 at least one participant voted NO, or
 at least one participant is faulty, or
 coordinator is faulty

$AC3_1 \triangleq \Box ((\exists i \in participants : participant[i].decision = abort) \Rightarrow \vee (\exists j \in participants : participant[j].vote = no) \vee (\exists j \in participants : participant[j].faulty) \vee coordinator.faulty)$

Each participant decides at most once

$AC4 \triangleq \Box \wedge (\forall i \in participants : participant[i].decision = commit \Rightarrow \Box (participant[i].decision = commit)) \wedge (\forall j \in participants : participant[j].decision = abort \Rightarrow \Box (participant[j].decision = abort))$

LIVENESS

(stronger for $AC3$ than in the original paper)

$AC3_2 \triangleq \Diamond \vee \forall i \in participants : participant[i].decision \in \{abort, commit\} \vee \exists j \in participants : participant[j].faulty \vee coordinator.faulty$

(SOME) INTERMEDIATE PROPERTIES USED IN PROOFS

$FaultyStable \triangleq \wedge \forall i \in participants : \Box (participant[i].faulty \Rightarrow \Box participant[i].faulty) \wedge \Box (coordinator.faulty \Rightarrow \Box coordinator.faulty)$

$VoteStable \triangleq \forall i \in participants : \vee \Box (participant[i].vote = yes) \vee \Box (participant[i].vote = no)$

$StrongerAC2 \triangleq \Box ((\exists i \in participants : participant[i].decision = commit) \Rightarrow \wedge (\forall j \in participants : participant[j].vote = yes) \wedge coordinator.decision = commit)$

$StrongerAC3_1 \triangleq \Box ((\exists i \in participants : participant[i].decision = abort) \Rightarrow \vee (\exists j \in participants : participant[j].vote = no) \vee \wedge \exists j \in participants : participant[j].faulty \wedge coordinator.decision = abort \vee \wedge coordinator.faulty \wedge coordinator.decision = undecided)$

($AC1$ follows from $StrongerAC2 \wedge StrongerAC3_1$)

$NoRecovery \triangleq \Box \wedge \forall i \in participants : participant[i].alive \equiv \neg participant[i].faulty \wedge coordinator.alive \equiv \neg coordinator.faulty$

(SOME) INVALID PROPERTIES

$DecisionReachedNoFault \triangleq (\forall i \in participants : participant[i].alive) \rightsquigarrow (\forall k \in participants : participant[k].decision \neq undecided)$

$AbortImpliesNoVote \triangleq \Box ((\exists i \in participants : participant[i].decision = abort) \Rightarrow (\exists j \in participants : participant[j].vote = no))$

The following is the termination property that this SB algorithm doesn't have

$AC5 \triangleq \diamond \forall i \in participants : \vee participant[i].decision \in \{abort, commit\}$
 $\vee participant[i].faulty$
