

1 Introduction

1.1 Model

1.1.1 Model Properties

The system consists of n asynchronous processes communicating via reliable point-to-point message passing.

Each process has a unique, unforgeable identifier and knows the identifiers of all other processes.

Up to $f < n$ processes may crash (fail-stop).

The network is reliable : if a correct process sends a message to another correct process, it is eventually delivered.

Messages are uniquely identifiable : two messages sent by distinct processes or at different rounds are distinguishable

2 messages sent by the same process in two different rounds are different

Property 1 (Message Uniqueness) *If two messages are sent by different processes, or by the same process in different rounds, then the messages are distinct.*

Formally :

$$\forall p_1, p_2, \forall r_1, r_2, \forall m_1, m_2, \left(\begin{array}{l} \text{send}(p_1, r_1, m_1) \wedge \text{send}(p_2, r_2, m_2) \\ \wedge (p_1 \neq p_2 \vee r_1 \neq r_2) \end{array} \right) \Rightarrow m_1 \neq m_2$$

1.1.2 Reliable Broadcast Properties

Property 2 Integrity *Every message received was previously sent.*

Formally :

$$\forall p_i : bc\text{-recv}_i(m) \Rightarrow \exists p_j : bc\text{-send}_j(m)$$

Property 3 No Duplicates *No message is received more than once at any single processor.*

Formally :

$$\forall m, \forall p_i : bc\text{-recv}_i(m) \text{ occurs at most once}$$

Property 4 Validity *All messages broadcast by a correct process are eventually received by all non faulty processors.*

Formally :

$$\forall m, \forall p_i : \text{correct}(p_i) \wedge bc\text{-send}_i(m) \Rightarrow \forall p_j : \text{correct}(p_j) \Rightarrow bc\text{-recv}_j(m)$$

1.1.3 AtomicBroadcast Properties

Property 5 AB Totally ordered $\forall m_1, m_2, \forall p_i, p_j : ab\text{-recv}_{p_i}(m_1) < ab\text{-recv}_{p_i}(m_2) \Rightarrow ab\text{-recv}_{p_j}(m_1) < ab\text{-recv}_{p_j}(m_2)$

1.1.4 DenyList Properties

Let Π_M be the set of processes authorized to issue APPEND operations, and Π_V the set of processes authorized to issue PROVE operations.

Let S be the set of valid values that may be appended. Let Seq be the linearization of operations recorded in the DenyList.

Property 6 APPEND Validity *An operation APPEND(x) is valid iff : the issuing process $p \in \Pi_M$, and the value $x \in S$*

Property 7 PROVE Validity An operation $PROVE(x)$ is valid iff : the issuing process $p \in \Pi_V$, and there exists no $APPEND(x)$ that appears earlier in Seq .

Property 8 PROGRESS If an $APPEND(x)$ is invoked by a correct process, then all correct processes will eventually be unable to $PROVE(x)$.

Property 9 READ Validity $READ()$ return a list of tuples who is a random permutation of all valids $PROVE()$ associated to the identity of the emitter process.

1.2 Algorithms

We consider a set of processes communicating asynchronously over reliable point-to-point channels. Each process maintains the following local or shared variables :

- **received** : the set of messages that have been received via the reliable broadcast primitive but not yet ordered.
- **delivered** : the set of messages that have been ordered.
- **prop[r][j]** : the proposal set announced by process j at round r . It contains a set of messages that process j claims to have received but not yet delivered.
- **winner^r** : the set of processes that have issued a valid $PROVE$ for round r , as observed through the registry.
- **RB-cast**(PROP, S, r, j) : a reliable broadcast invocation that disseminates the proposal S from process j for round r .
- **RB-delivered**(PROP, S, r, j) : the handler invoked upon reception of a **RB-cast**, which stores the received proposal S into $prop[r][j]$.
- **READ()** : returns the current view of all valid operations stored in the DenyList registry.
- **ordered**(S) : returns a deterministic total order over a set S of messages.

1.3 proof

Theorem 1 (Integrity) If a message m is delivered by any process, then it was previously broadcast by some process via the **AB-broadcast** primitive.

Proof 1 Let j be a process such that **AB-deliver_j**(m) occurs.

Algorithm 1 Atomic Broadcast with DenyList

```
1  $proves \leftarrow \emptyset$ 
1  $received \leftarrow \emptyset$ 
1  $delivered \leftarrow \emptyset$ 
1  $r_1 \leftarrow 0$ 

1 AB-Broadcast $_j(m)$ 
2   RB-Broadcast $_j(m)$ 

3 RB-delivered $_j(m)$ 
4    $received \leftarrow received \cup \{m\}$ 
5   repeat until  $received \setminus delivered \neq \emptyset$ 
6      $S \leftarrow received \setminus delivered$ 
7      $proves \leftarrow \text{READ}()$ 
8      $r_2 \leftarrow \max\{r : j, (j, \text{PROVE}(r)) \in proves\} + 1$ 
9     RB-cast(PROP,  $S, r_2, j$ )
10    PROVE( $r_2$ )
11    for  $r \in [r_1 + 1, \dots, r_2]$  do
12      APPEND( $r$ )
13       $proves \leftarrow \text{READ}()$ 
14       $winner^r \leftarrow \{j : (j, \text{PROVE}(r)) \in proves\}$ 
15      wait  $\forall j \in winner^r, prop[r][j] \neq \perp$ 
16       $T \leftarrow \bigcup_{j \in winner^r} prop[r][j] \setminus delivered$ 
17      for each  $m \in \text{ordered}(T)$ 
18         $delivered \leftarrow delivered \cup \{m\}$ 
19        AB-deliver $_j(m)$ 
20       $r_1 \leftarrow r_2$ 

21 RB-delivered $_j(\text{PROP}, S, r_1, j_1)$ 
22    $prop[r_1][j_1] \leftarrow S$ 
```

$AB\text{-deliver}_j(m)$ (line 18)
 $\Rightarrow m \in \text{ordered}(T)$, with $T = \bigcup_{j' \in \text{winner}^r} \text{prop}[r][j'] \setminus \text{delivered}$ (lines 16-17)
 $\Rightarrow \exists j_0, r_0 : m \in \text{prop}[r_0][j_0]$ (line 16)
 $\Rightarrow \text{prop}[r_0][j_0] = S$, with $RB\text{-delivered}_j(\text{PROP}, S, r_0, j_0)$ (line 22)
 $\Rightarrow S$ was sent in $RB\text{-cast}(\text{PROP}, S, r_0, j_0)$ (line 9)
 $\Rightarrow S = \text{received}_{j_0} \setminus \text{delivered}_{j_0}$ (line 6)
 $\Rightarrow m' \in \text{received}_{j_0}$ where m' broadcast by j_0 (line 4)
 \Rightarrow **if** $m = m'$
 $\Rightarrow RB\text{-Broadcast}_{j_0}(m)$ occurred (line 3)
 $\Rightarrow AB\text{-Broadcast}_{j_0}(m)$ occurred (line 1) \square
else : $m \in \text{received}_{j_0} \setminus \text{delivered}_{j_0}$
 $\Rightarrow m \in \text{received}_{j_0}$ (line 4)
 $\Rightarrow RB\text{-delivered}_{j_0}(m)$ occurred (line 3)
 $\Rightarrow \exists j_1 : RB\text{-Broadcast}_{j_1}(m)$ occurred (line 2)
 $\Rightarrow AB\text{-Broadcast}_{j_1}(m)$ occurred (line 1) \square

Therefore, every delivered message m must originate from some call to $AB\text{-Broadcast}$.

Theorem 2 (No Duplication) No message is delivered more than once by any process.

Proof 2 Assume by contradiction that a process j delivers the same message m more than once, i.e.,

$AB\text{-deliver}_j(m)$ occurs at least twice.

$AB\text{-deliver}_j(m)$ occurs (line 19)
 $\Rightarrow m \in \text{ordered}(T)$, where $T = \bigcup_{j' \in \text{winner}^r} \text{prop}[r][j'] \setminus \text{delivered}$ (lines 16-17)
 $\Rightarrow m \notin \text{delivered}$ at that time

However :

$\text{delivered} \leftarrow \text{delivered} \cup \{m\}$ (line 18)
 $\Rightarrow m \in \text{delivered}$ permanently
 \Rightarrow In any future round, $m \notin T'$ since $T' = \bigcup_{j' \in \text{winner}^r} \text{prop}[r'][j'] \setminus \text{delivered}$
 $\Rightarrow m$ will not be delivered again
 \Rightarrow Contradiction.

Therefore, no message can be delivered more than once by the same process. \square

Theorem 3 (Validity) If a correct process invokes $AB\text{-Broadcast}_j(m)$, then all correct processes eventually deliver m .

Proof 3 Let j be a correct process such that $AB\text{-Broadcast}_j(m)$ occurs (line 5).

$$\begin{aligned}
& AB\text{-Broadcast}_j(m) && \text{(line 1)} \\
\Rightarrow & RB\text{-Broadcast}_j(m) \text{ occurs} && \text{(line 2)} \\
\Rightarrow & \forall j_0 : RB\text{-delivered}_{j_0}(m) && \text{(line 3)} \\
\Rightarrow & m \in received_{j_0} && \text{(line 4)} \\
\Rightarrow & \mathbf{if} \ m \in delivered_{j_0} && \\
& \quad \Rightarrow delivered_{j_0} \leftarrow textitdelivered_{j_0} \cup \{m\} && \text{(line 18)} \\
& \quad \Rightarrow AB\text{-delivered}_{j_0}(m) && \text{(line 19)} \quad \square \\
& \mathbf{else} \ m \notin delivered_{j_0} : && \\
& \quad \Rightarrow m \in S_{j_0} \text{ since } S_{j_0} = received_{j_0} \setminus delivered_{j_0} && \text{(line 6)} \\
& \quad \Rightarrow \exists r : RB\text{-cast}_{j_0}(textitPROP, S_{j_0}, r, j_0) && \text{(line 9)} \\
& \quad \Rightarrow \forall j_1 : RB\text{-Deliver}_{j_1}(PROP, S_{j_0}, r, j_0) \text{ occurs} && \text{(line 21)} \\
& \quad \Rightarrow prop[r][j_0] = S_{j_0} && \text{(line 22)} \\
& \quad \Rightarrow \exists j_2 \in j_0 : PROVE_{j_2}(r) \text{ is valid} && \text{(line 10)} \\
& \quad \Rightarrow j_2 \in textitwinner^r && \text{(line 14)} \\
& \quad \Rightarrow T_{j_0} \ni prop[r][j_2] \setminus delivered_{j_0} && \text{(line 16)} \\
& \quad \Rightarrow T_{j_0} \ni S_{j_2} \setminus delivered_{j_0} \ni m && \text{(line 16)} \\
& \quad \Rightarrow AB\text{-deliver}_{j_0}(m) && \text{(line 19)} \quad \square
\end{aligned}$$

Theorem 4 (Total Order) If two correct processes deliver two messages m_1 and m_2 , then they deliver them in the same order.

Proof 4

$$\begin{aligned}
& \forall j_0 : AB\text{-Deliver}_{j_0}(m_0) \wedge AB\text{-Deliver}_{j_0}(m_1) && \text{(line 19)} \\
\Rightarrow & \exists r_0, r_1 : m_0 \in ordered(T^{r_0}) \wedge m_1 \in ordered(T^{r_1}) && \text{(line 17)} \\
\Rightarrow & T^{r_0} = \bigcup_{j' \in winner^{r_0}} prop[r_0][j'] \setminus delivered \wedge \\
& T^{r_1} = \bigcup_{j' \in winner^{r_1}} prop[r_1][j'] \setminus delivered && \text{(line 16)} \\
\Rightarrow & \mathbf{if} \ r_0 = r_1 && \\
& \quad \Rightarrow T^{r_0} = T^{r_1} && \\
& \quad \Rightarrow m_0, m_1 \in ordered(T^{r_0}) \text{ since } ordered \text{ is deterministic} && \\
& \quad \Rightarrow \mathbf{if} \ m_0 < m_1 : && \\
& \quad \quad \Rightarrow AB\text{-Deliver}_{j_0}(m_0) < AB\text{-Deliver}_{j_0}(m_1) && \square \\
& \mathbf{else if} \ r_0 < r_1 && \\
& \quad \Rightarrow \forall m \in T^{r_0}, \forall m' \in T^{r_1} : AB\text{-Deliver}(m) < AB\text{-Deliver}(m') && \square
\end{aligned}$$

Therefore, for all correct processes, messages are delivered in the same total order.