

Abstract

We show how to upgrade a Reliable Broadcast (RB) primitive to Atomic Reliable Broadcast (ARB) by leveraging a synchronous DenyList (DL) object. In a purely asynchronous message-passing model with crashes, ARB is impossible without additional power. The DL supplies this power by enabling round closing and agreement on a set of "+winners" for each round. We present the algorithm, its safety arguments, and discuss liveness and complexity under the assumed synchrony of DL.

Keywords Atomic broadcast, total order broadcast, reliable broadcast, consensus, synchrony, shared object, linearizability.

1 Introduction

Atomic Reliable Broadcast (ARB)—a.k.a. total order broadcast—ensures that all processes deliver the same sequence of messages. In asynchronous message-passing systems with crashes, implementing ARB is impossible without additional assumptions, as it enables consensus. We assume a synchronous DenyList (DL) object and demonstrate how to combine DL with an asynchronous RB to realize ARB.

2 Model

We consider a static set of n processes with known identities, communicating by reliable point-to-point channels, in a complete graph. Messages are uniquely identifiable.

Synchrony. The network is asynchronous. Processes may crash; at most f crashes occur.

Communication. Processes can exchange through a Reliable Broadcast (RB) primitive (defined below) which's invoked with the functions $\text{RB-cast}(m)$ and $\text{RB-received}(m)$. There exists a shared object called DenyList (DL) (defined below) that is interfaced with the functions $\text{APPEND}(x)$, $\text{PROVE}(x)$ and $\text{READ}()$.

Notation. Let Π be the finite set of process identifiers and let $n \triangleq |\Pi|$. Two authorization subsets are $\Pi_M \subseteq \Pi$ (processes allowed to issue APPEND) and $\Pi_V \subseteq \Pi$ (processes allowed to issue PROVE). Indices $i, j \in \Pi$ refer to processes, and p_i denotes the process with identifier i . Let \mathcal{M} denote the universe of uniquely identifiable messages, with $m \in \mathcal{M}$. Let $\mathcal{R} \subseteq \mathbb{N}$ be the set of round identifiers; we write $r \in \mathcal{R}$ for a round. We use the precedence relation \prec for the DL linearization: $x \prec y$ means that operation x appears strictly before y in the linearized history of DL. For any finite set $A \subseteq \mathcal{M}$, $\text{ordered}(A)$ returns a deterministic total order over A (e.g., lexicographic order on $(\text{senderId}, \text{messageId})$ or on message hashes). For any round $r \in \mathcal{R}$, define $\text{Winners}_r \triangleq \{j \in \Pi \mid (j, \text{PROVE}(r)) \prec \text{APPEND}(r)\}$, i.e., the set of processes whose $\text{PROVE}(r)$ appears before the first $\text{APPEND}(r)$ in the DL linearization.

3 Primitives

3.1 Reliable Broadcast (RB)

RB provides the following properties in the model.

- **Integrity:** Every message received was previously sent. $\forall p_i : \text{RB-received}_i(m) \Rightarrow \exists p_j : \text{RB-cast}_j(m)$.
- **No-duplicates:** No message is received more than once at any process.
- **Validity:** If a correct process broadcasts m , every correct process eventually receives m .

3.2 DenyList (DL)

The DL is a *shared, append-only* object that records attestations about opaque application-level tokens. It exposes the following operations:

- **APPEND(x)**
- **PROVE(x):** issue an attestation for token x ; this operation is *valid* (return true) only if no APPEND(x) occurs earlier in the DL linearization. Otherwise, it is invalid (return false).
- **READ():** return a (permutation of the) valid operations observed so far; subsequent reads are monotone (contain supersets of previously observed valid operations).

Validity. APPEND(x) is valid iff the issuer is authorized (in Π_M) and x belongs to the application-defined domain S . PROVE(x) is valid iff the issuer is authorized (in Π_V) and there is no earlier APPEND(x) in the DL linearization.

Progress. If a correct process invokes APPEND(x), then eventually all correct processes will be unable to issue a valid PROVE(x), and READ at all correct processes will (eventually) reflect that APPEND(x) has been recorded.

Termination. Every operation invoked by a correct process eventually returns.

Interface and Semantics. The DL provides a single global linearization of operations consistent with each process's program order. READ is prefix-monotone; concurrent updates become visible to all correct processes within bounded time (by synchrony). Duplicate requests may be idempotently coalesced by the implementation.

4 Target Abstraction: Atomic Reliable Broadcast (ARB)

Processes export AB-broadcast(m) and AB-deliver(m). ARB requires total order:

$$\forall m_1, m_2, \forall p_i, p_j : \text{AB-deliver}_i(m_1) < \text{AB-deliver}_i(m_2) \Rightarrow \text{AB-deliver}_j(m_1) < \text{AB-deliver}_j(m_2),$$

plus Integrity/No-duplicates/Validity (inherited from RB and the construction).

5 Algorithm

Each process p_i maintains:

- **received** — set of RB-received messages not yet delivered;
- **delivered** — set of messages already delivered;

- $\text{prop}[r][j]$ — proposal set announced by process p_j for round r (possibly \perp locally);
- Local view of DL via $\text{READ}()$.
- next — lowest round index not yet delivered.

5.1 Handlers and Procedures

Algorithm A RB handler (at process p_i)

```

A1 on RB-received( $m, S, r, j$ ) do
A2 received  $\leftarrow$  received  $\cup \{m\}$ 
A3  $\text{prop}[r][j] \leftarrow S$  ▷ Record sender  $j$ 's proposal  $S$  for round  $r$ 

```

An AB-broadcast(m) chooses the next open round from the DL view, proposes all pending messages together with the new m , disseminates this proposal via RB, then executes $\text{PROVE}(r)$ followed by $\text{APPEND}(r)$ to freeze the winners of the round. The loop polls DL until (i) some winner's proposal includes m in a *closed* round and (ii) all winners' proposals for closed rounds are known locally, ensuring eventual inclusion and delivery.

Algorithm B AB-broadcast(m) (at process p_i)

```

B1 on AB-broadcast( $m, S, r, j$ ) do
B2  $P \leftarrow \text{READ}()$  ▷ Fetch latest DL state to learn recent PROVE operations
B3  $r \leftarrow \max\{r' : \exists j, (j, \text{PROVE}(r')) \in P\}$  ▷ Pick next open round: one past the most recent proved round
B4  $S \leftarrow (\text{received} \setminus \text{delivered}) \cup \{m\}$  ▷ Propose all pending messages plus the new  $m$ 
B5  $\text{RB-cast}(m, S, r, \text{self})$  ▷ Asynchronously disseminate proposal via RB
B6  $\text{PROVE}(r)$  ▷ Attest participation in round  $r$  while it is still open
B7  $\text{APPEND}(r)$  ▷ Close round  $r$ ; freezes  $\text{Winners}_r$  in the DL linearization
B8 while ( $\nexists j : \exists r, (j, \text{PROVE}(r)) \in P \wedge m \in \text{prop}[r][j]$ ) do ▷ Wait until  $m$  is included in some closed round and all needed proposals are known
B9    $P \leftarrow \text{READ}()$  ▷ Refresh local view of DL
B10    $\text{RB-cast}(m, S, r, \text{self})$ 
B11    $\text{PROVE}(r)$ 
B12    $\text{APPEND}(r)$ 
B13 end while

```

Algorithm C AB-deliver() at process p_i

```
C1 on AB-deliver() do                                ▷ Called when the process wants to receive the next message
C2  $P \leftarrow \text{READ}()$                                 ▷ Fetch latest DL state to learn recent PROVE operations
C3 if  $\nexists j : (j, \text{PROVE}(\text{next})) \in P$  then          ▷ No process has proved round  $\text{next}$ 
C4   return  $\perp$                                          ▷ No closed round ready for delivery
C5 end if
C6  $\text{APPEND}(\text{next})$                                     ▷ Close round  $\text{next}$  if not already closed
C7 if  $\text{PROVE}(\text{next}) == \text{FALSE}$  then                    ▷ Process closed rounds strictly in order
C8    $P \leftarrow \text{READ}()$                                 ▷ Refresh local view of DL
C9    $\text{Winners}_{\text{next}} \leftarrow \{j : (j, \text{PROVE}(\text{next})) \in P\}$     ▷ Frozen winners of round  $\text{next}$ 
C10  if  $\forall j \in \text{Winners}_{\text{next}} : \text{prop}[\text{next}][j] \neq \perp$  then
C11     $M_{\text{next}} \leftarrow \bigcup_{j \in \text{Winners}_{\text{next}}} \text{prop}[\text{next}][j]$     ▷ Round- $\text{next}$  message set
C12    for all  $m \in \text{ordered}(M_{\text{next}})$  do                ▷ Deterministic per-round order
C13      if  $m \notin \text{delivered}$  then
C14         $\text{delivered} \leftarrow \text{delivered} \cup \{m\}$ 
C15        return  $m$                                      ▷ Deliver exactly one message per call
C16      end if
C17    end for
C18     $\text{next} \leftarrow \text{next} + 1$                             ▷ This round fully processed; advance
C19    continue                                         ▷ Try the next closed round (if any)
C20  else
C21    return  $\perp$                                          ▷ Some winners' proposals not yet known via RB
C22  end if
C23 end if
C24 return  $\perp$                                          ▷ No closed round ready for delivery
```

6 Correctness

Definition 1 (Closed round). Given a DL linearization H , a round $r \in \mathcal{R}$ is *closed* in H iff H contains an operation $\text{APPEND}(r)$. Equivalently, there exists a time after which every $\text{PROVE}(r)$ is invalid in H .

Lemma 1 (Stable round closure). *If a round r is closed, then there exists a unique linearization point t_0 of $\text{APPEND}(r)$ in the DL, and from that point on, no $\text{PROVE}(r)$ can be valid. Once closed, a round never becomes open again.*

Lemma 2 (Across rounds). *Rounds are delivered in strictly increasing r . Concatenating the per-round sequences yields a single global total order. Equivalently, $\forall r_1, r_2$ such that r_1, r_2 are closed and $r_1 < r_2$, all round r such that $r_1 < r < r_2$ are also closed.*

Lemma 3 (Well-defined winners). *In any execution, whenever some process computes Winners_r , its current DL-view contains $\text{APPEND}(r)$; hence Winners_r is computed only for closed round.*

Lemma 4 (View-Invariant Winners). *For any closed round r , there exists a unique set*

$$\text{Winners}_r = \{j : (j, \text{PROVE}(r)) \prec \text{APPEND}^*(r)\}$$

with $\text{APPEND}^(r)$ being the earliest $\text{APPEND}(r)$ in the DL linearization.*

Proof. We admit that some $\text{APPEND}(r)$ occurs; if $\text{APPEND}(r)$ occurs then it exists an operation $\text{APPEND}^*(r)$ that is the earliest in the linearization.

Due to the validity property of DL, a $\text{PROVE}(r)$ is valid iff $\text{PROVE}(r) \prec \text{APPEND}^*(r)$. Thus,

$$\text{Winners}_r \triangleq \{j : (j, \text{PROVE}(r)) \prec \text{APPEND}^*(r)\}$$

Let's consider two correct processes p_i and p_j and two READ responses P_i and P_j such that

$$\text{APPEND}^*(r) \prec (i, \text{PROVE}(r)) \prec P_i \text{APPEND}^*(r) \prec (j, \text{PROVE}(r)) \prec P_j$$

By the DL interface, READ returns a permutation of the valid operations observed so far, and hence every $\text{PROVE}(r)$ that precedes $\text{APPEND}^*(r)$ is valid while any $\text{PROVE}(r)$ that follows $\text{APPEND}^*(r)$ is invalid. We ensure that

$$\{j : (j, \text{PROVE}(r)) \in P_i\} = \{j : (j, \text{PROVE}(r)) \in P_j\} = \text{Winners}_r.$$

□

Lemma 5 (No empty winners). *For any round r closed, $\text{Winners}_r \neq \emptyset$.*

Proof. Assume some correct process invokes $\text{APPEND}(r)$. Hence an $\text{APPEND}(r)$ occurs, and there exists an operation $\text{APPEND}^*(r)$ that is earliest in the DL linearization.

Let p_k be the issuer of $\text{APPEND}^*(r)$. By the algorithm (AB-broadcast), any process that issues $\text{APPEND}(r)$ (B6) must have previously invoked $\text{PROVE}(r)$ (B5) in its program order. Because the DL is linearizable and preserves per-process program order in the linearization, we obtain

$$(k, \text{PROVE}(r)) \prec \text{APPEND}^*(r).$$

Consequently, $(k, \text{PROVE}(r))$ is a *valid* attestation for round r (no prior $\text{APPEND}(r)$ exists before $\text{APPEND}^*(r)$ by definition of $\text{APPEND}^*(r)$). By the definition of Winners_r ,

$$k \in \text{Winners}_r.$$

Therefore $\text{Winners}_r \neq \emptyset$, i.e., $|\text{Winners}_r| > 0$. □

Lemma 6 (Proposal convergence). *For any closed round r , after all RB messages from Winners_r are received, every correct process computes the same $M_r = \bigcup_{j \in \text{Winners}_r} \text{prop}[r][j]$.*

Proof. Fix a round r such that some $\text{APPEND}(r)$ occurs; hence r is *closed*. By Lemma 3, all correct processes that observe $\text{READ}()$ after $\text{APPEND}^*(r)$ compute the same winner set Winners_r .

For any $j \in \text{Winners}_r$, in process p_j 's program order we have, for round r :

$$\text{RB-cast}(_, S_j, r, j) \text{ (B4) before } \text{PROVE}(r) \text{ (B5) before } \text{APPEND}(r) \text{ (B6)}.$$

Since $\text{Winners}_r = \{j : (j, \text{PROVE}(r)) \prec \text{APPEND}^*(r)\}$, each winner j executed (B4) before $\text{APPEND}^*(r)$, thereby broadcasting a RB message for (r, j) with some (uniquely determined) payload S_j .

The round chosen at (B2) is the next open round. After (B6) closes r , any later invocation of AB-broadcast by p_j picks a strictly larger round. Thus p_j executes at most one $\text{RB-cast}(_, \cdot, r, j)$, so the set S_j attached to (r, j) is unique.

Now consider any two correct processes p_a and p_b that have (via RB) received all winners' messages for round r . By the RB *Integrity* and *No-duplicates* properties, every delivery of p_j 's RB message for (r, j) carries the same S_j , and the handler sets

$$\text{prop}[r][j] \leftarrow S_j \text{ (A3)}$$

at both p_a and p_b . Therefore,

$$M_r^{(a)} \triangleq \bigcup_{j \in \text{Winners}_r} \text{prop}^{(a)}[r][j] = \bigcup_{j \in \text{Winners}_r} S_j = \bigcup_{j \in \text{Winners}_r} \text{prop}^{(b)}[r][j] \triangleq M_r^{(b)}.$$

Hence, after all RB messages from Winners_r have been received, every correct process computes the same M_r . \square

Lemma 7 (Per-round determinism). *Given Lemma 6, all correct processes iterate $\text{ordered}(M_r)$ in the same order; hence the delivery order inside round r is identical at all correct processes.*

Corollary 8 (No duplicates). *Each message is delivered at most once since membership in delivered is tested before delivery.*

Lemma 9 (Inclusion). *If some process invokes $\text{AB-broadcast}(m)$, then there exist a (closed) round r and a process $j \in \text{Winners}_r$ such that j invokes*

$$\text{RB-cast}(_, S, r, j) \quad \text{with} \quad m \in S.$$

Equivalently, every $\text{AB-broadcast}(m)$ is included in the proposal of some winner of some closed round.

Theorem 10 (ARB). *Under the assumed DL synchrony and RB reliability, the algorithm implements Atomic Reliable Broadcast.*

References