

A Modular Approach to Construct Signature-Free BRB Algorithms under a Message Adversary

Timothé Albouy, Davide Frey, Michel Raynal, François Taïani

Univ Rennes, IRISA, CNRS, Inria, 35042 Rennes, France

15th December 2022



UNIVERSITÉ DE
RENNES 1

Inria



UMR IRISA

There is a World beyond Byzantium!




There is a World beyond Byzantium!



Byzantine Fault Tolerance (BFT) is not the end of the road

Process faults & Network faults

Byzantine faults




Model for **Process** faults:

- ▶ attackers
- ▶ crashes
- ▶ implementation errors
- ▶ ...

Process faults & Network faults


Byzantine faults



Model for **Process** faults:

- ▶ attackers
- ▶ crashes
- ▶ implementation errors
- ▶ ...

Message adversary (MA)




Model for **Network** faults:

- ▶ message losses
- ▶ disconnections
- ▶ interference
- ▶ ...

Process faults & Network faults


Byzantine faults



Model for **Process** faults:

- ▶ attackers
- ▶ crashes
- ▶ implementation errors
- ▶ ...

Message adversary (MA)





Model for **Network** faults:



- ▶ message losses
- ▶ disconnections
- ▶ interference
- ▶ ...

- ▶ **2 orthogonal fault models**
- ▶ What happens if we combine the 2?
 - ▶ Case of **Reliable Broadcast**



Roadmap

1. New fault model:  + 




Roadmap

1. New fault model:  + 
2. New reliable broadcast specification: **MBRB**

Roadmap

1. New fault model:  + 
2. New reliable broadcast specification: **MBRB**
3. New many-to-many primitive: *k2l-cast*

Roadmap

1. New fault model:  + 
2. New reliable broadcast specification: **MBRB**
3. New many-to-many primitive: *k2l-cast*
4. 2 MBRB implementations 


System Model

- ▶ **Asynchronous** network (arbitrary message delays)


System Model

- ▶ **Asynchronous** network (arbitrary message delays)
- ▶ n processes in the network: p_1, p_2, \dots, p_n

System Model


- ▶ **Asynchronous** network (arbitrary message delays)
- ▶ n processes in the network: p_1, p_2, \dots, p_n
- ▶ $t < n$ processes can be Byzantine 

System Model

- ▶ **Asynchronous** network (arbitrary message delays)
- ▶ n processes in the network: p_1, p_2, \dots, p_n
- ▶ $t < n$ processes can be Byzantine 
- ▶ $c \in [n-t..n]$ correct processes


System Model

- ▶ **Asynchronous** network (arbitrary message delays)
- ▶ n processes in the network: p_1, p_2, \dots, p_n
- ▶ $t < n$ processes can be Byzantine 🦋
- ▶ $c \in [n-t..n]$ correct processes
- ▶ $d < n$ messages can be deleted each sending step (**MA power**)



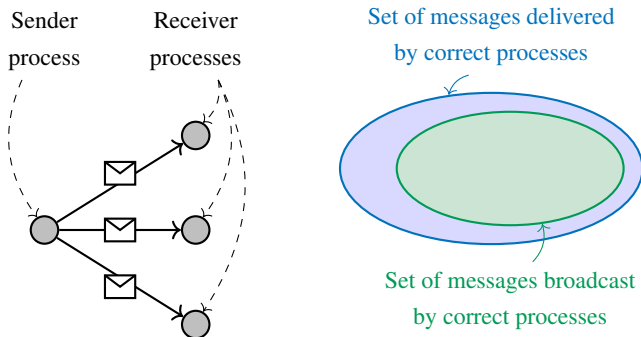
System Model

- ▶ **Asynchronous** network (arbitrary message delays)
- ▶ n processes in the network: p_1, p_2, \dots, p_n
- ▶ $t < n$ processes can be Byzantine 🐉
- ▶ $c \in [n-t..n]$ correct processes
- ▶ $d < n$ messages can be deleted each sending step (**MA power**)



- ▶ **Deterministic** algorithm

Reminder: Byzantine Reliable Broadcast (BRB) 🐼



Many practical application: consensus algorithms, voting systems, state-machine replication, payment systems, distributed memory...

BRB SAFETY:

- ▶ **Validity:** If a correct process p_i delivers m from a correct process p_j , then p_j broadcast m
- ▶ **No-duplication:** A correct process p_i delivers at most one app-message from a process p_j
- ▶ **No-duplicity:** No two correct processes deliver different app-messages from a process p_i

BRB LIVENESS:

- ▶ **Local-delivery:** If a correct process p_i broadcasts m , then at least one correct process p_j delivers m from p_i
- ▶ **Global-delivery:** If a correct process p_i delivers m , then **all c** correct processes deliver m

MBRB SAFETY:

- ▶ **Validity:** If a correct process p_i delivers m from a correct process p_j , then p_j broadcast m
- ▶ **No-duplication:** A correct process p_i delivers at most one app-message from a process p_j
- ▶ **No-duplicity:** No two correct processes deliver different app-messages from a process p_i


MBRB LIVENESS:

- ▶ **Local-delivery:** If a correct process p_i broadcasts m , then at least one correct process p_j delivers m from p_i
- ▶ **Global-delivery:** If a correct process p_i delivers m , then at least $\ell_{MBRB} \leq c - d$ correct processes deliver m

As the MA can totally isolate d processes


Albouy, Frey, Raynal and Taïani (SSS 2021):

Byzantine-tolerant reliable broadcast in the presence of silent churn

- ▶ First (**signature-based**) MBRB implementation 
 - ▶ assumes $n > 3t + 2d$ (optimal)
 - ▶ guarantees $\ell_{MBRB} = c - d$ (optimal)
- ▶ Theorem: $n > 3t + 2d$ **tight** for MBRB
 - ▶ **optimal** implementation

Albouy, Frey, Raynal and Taïani (SSS 2021):

Byzantine-tolerant reliable broadcast in the presence of silent churn


- ▶ First (**signature-based**) MBRB implementation 
 - ▶ assumes $n > 3t + 2d$ (optimal)
 - ▶ guarantees $\ell_{MBRB} = c - d$ (optimal)
- ▶ Theorem: $n > 3t + 2d$ **tight** for MBRB
 - ▶ **optimal** implementation

Question: Is signature-free MBRB possible?




Intuition

How BRB algorithms work (e.g. Bracha 1987) ~~0~~



Intuition

How BRB algorithms work (e.g. Bracha 1987) ~~0~~




Quorum Construction Pattern

Let $q_f, q_d \in \mathbb{N}^*$, with $1 \leq q_f \leq q_d \leq n$


Quorum Construction Pattern

Let $q_f, q_d \in \mathbb{N}^*$, with $1 \leq q_f \leq q_d \leq n$



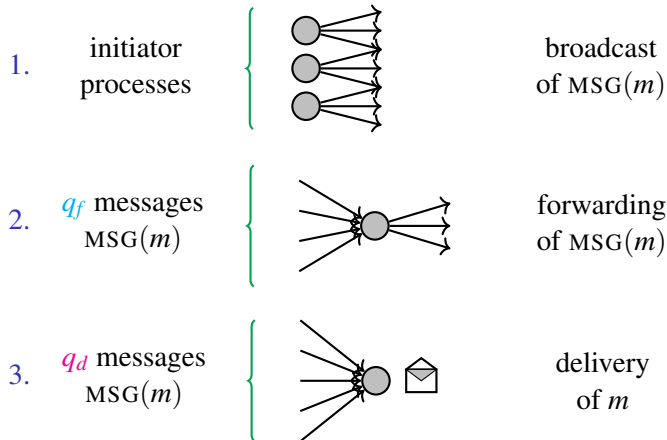
Quorum Construction Pattern

Let $q_f, q_d \in \mathbb{N}^*$, with $1 \leq q_f \leq q_d \leq n$



Quorum Construction Pattern


Let $q_f, q_d \in \mathbb{N}^*$, with $1 \leq q_f \leq q_d \leq n$



$k2\ell$ -Cast Abstraction


$k2\ell$ -cast (for k -to- ℓ -cast): many-to-many protocol

(k correct processes $k2\ell$ -cast) \rightarrow (ℓ correct processes $k2\ell$ -deliver)



Guarantee	Meaning	Best value
k	(k $k2\ell$ -casts) \rightarrow (one $k2\ell$ -delivery)	Low
ℓ	(one $k2\ell$ -delivery) \rightarrow (ℓ $k2\ell$ -deliveries)	High

$k2\ell$ -Cast Implementation

Encapsulates *Quorum Construction Pattern* 

Parameter	Meaning
q_f	Forwarding threshold
q_d	Delivery threshold

Theorem: This implementation guarantees




$$\blacktriangleright k = \left\lfloor \frac{c(q_f-1)}{c-d-q_d+q_f} \right\rfloor + 1$$

$$\blacktriangleright \ell = c - \left\lfloor \frac{cd}{c-q_d+1} \right\rfloor$$

k2l-Cast in Action: Signature-Free MBRB

k2l-cast allows

- ▶ ~~✗~~-tolerant & simpler BRB ~~0~~
- ▶ systematic approach to analyze quorums
 - ▶ **side-effect:** optimized algorithms

Examples:

- ▶ **Bracha's** BRB (1987)
- ▶ **Imbs-Raynal's** BRB (2016)

Reconstructing Bracha's BRB (1/2)

```

operation brb_broadcast( $m, sn$ ) is
(1) broadcast(INIT( $m, sn$ )).

when INIT( $m, sn$ ) is received from  $p_j$  do
(2) if (ECHO( $-, sn, j$ ) not yet broadcast)
(3) then broadcast(ECHO( $m, sn, j$ ))
(4) end if.

when ECHO( $m, sn, j$ ) is received do
(5) if (ECHO( $m, sn, j$ ) is received from strictly more than  $\frac{n+t}{2}$  processes) then
(6) if (ECHO( $-, sn, j$ ) not yet broadcast)
(7) then broadcast(ECHO( $m, sn, j$ ))
(8) end if;
(9) if (READY( $-, sn, j$ ) not yet broadcast)
(10) then broadcast(READY( $m, sn, j$ ))
(11) end if
(12) end if.

when READY( $m, sn, j$ ) is received do
(13) if (READY( $m, sn, j$ ) is received from at least  $t + 1$  processes
       $\wedge$  READY( $-, sn, j$ ) not yet broadcast)
(14) then broadcast(READY( $m, sn, j$ ))
(15) end if;
(16) if (READY( $m, sn, j$ ) is received from at least  $2t + 1$  processes
       $\wedge$  ( $-, sn, j$ ) not already brb-delivered)
(17) then brb_deliver( $m, sn, j$ )
(18) end if.
    
```

Original version



BRB 
18 lines

$k2\ell$ -cast

init: $obj_e \leftarrow K2LCast(q_d = \lfloor \frac{n+t}{2} \rfloor + 1, q_f = t+1);$
 $obj_r \leftarrow K2LCast(q_u = 2t+d+1, q_l = t+1).$

(1) **operation** mbrb_broadcast(m, sn) **is** broadcast(INIT(m, sn)).
(2) **when** INIT(m, sn) **is received from** p_j **do** $obj_e.k2\ell_cast$ (ECHO(m), (sn, j)).
(3) **when** (ECHO(m), (sn, j)) **is** $obj_e.k2\ell_delivered$ **do** $obj_e.k2\ell_cast$ (READY(m), (sn, j)).
(4) **when** (READY(m), (sn, j)) **is** $obj_e.k2\ell_delivered$ **do** mbrb_deliver(m, sn, j).

Reconstructed version

MRRB  
4 lines

Reconstructing Bracha's BRB (1/2)

```

operation brb_broadcast( $m, sn$ ) is
(1) broadcast(INIT( $m, sn$ )).

when INIT( $m, sn$ ) is received from  $p_i$  do
(2) if (ECHO( $-, sn, j$ ) not yet broadcast)
(3) then broadcast(ECHO( $m, sn, j$ ))
(4) end if.

when ECHO( $m, sn, j$ ) is received do
(5) if (ECHO( $m, sn, j$ ) is received from strictly more than  $\frac{n+1}{2}$  processes) then
(6) if (ECHO( $-, sn, j$ ) not yet broadcast)
(7) then broadcast(ECHO( $m, sn, j$ ))
(8) end if;
(9) if (READY( $-, sn, j$ ) not yet broadcast)
(10) then broadcast(READY( $m, sn, j$ ))
(11) end if
(12) end if.

when READY( $m, sn, j$ ) is received do
(13) if (READY( $m, sn, j$ ) is received from at least  $t + 1$  processes
      & READY( $-, sn, i$ ) not yet broadcast)
(14) then brb_deliver( $m, sn, j$ )
(15) end if.
    
```

Original version

BRB 

18 lines

Takeaway: Better & simpler algorithms

$k2\ell$ -cast

init: $obj_e \leftarrow K2LCast(q_d = \lfloor \frac{n+1}{2} \rfloor + 1, q_f = t+1)$;
 $obj_r \leftarrow K2LCast(q_d = 2t+d+1, q_f = t+1)$.

- (1) operation mbrb_broadcast(m, sn) is broadcast(INIT(m, sn)).
- (2) when INIT(m, sn) is received from p_i do $obj_e.k2\ell_cast$ (ECHO(m), (sn, j)).
- (3) when (ECHO(m), (sn, j)) is $obj_e.k2\ell_delivered$ do $obj_r.k2\ell_cast$ (READY(m), (sn, i)).
- (4) when (READY(m), (sn, j)) is $obj_r.k2\ell_delivered$ do mbrb_deliver(m, sn, j).

Reconstructed version

MRB  

4 lines

Reconstructing Bracha's BRB (2/2)

When $d = 0$ (no MA), 2 versions have same

- ▶ t -resilience: $n > 3t$
- ▶ round complexity: ≤ 3 rounds
- ▶ message complexity: $\leq n + 2n^2$ messages

Reconstructing Bracha's BRB (2/2)

When $d = 0$ (no MA), 2 versions have same

- ▶ t -resilience: $n > 3t$
- ▶ round complexity: ≤ 3 rounds
- ▶ message complexity: $\leq n + 2n^2$ messages

However, new version uses smaller quorums (ECHO phase):

Threshold	Original version	New version	
Forwarding q_f	$\lfloor \frac{n+t}{2} \rfloor + 1$	$t + 1$	BETTER
Delivery q_d	$\lfloor \frac{n+t}{2} \rfloor + 1$	$\lfloor \frac{n+t}{2} \rfloor + 1$	

Needs less messages to progress \rightarrow **terminates earlier**

Reconstructing Imbs-Raynal's BRB

When $d = 0$ (no MA), 2 versions have same

- ▶ t -resilience: $n > 5t$
- ▶ round complexity: ≤ 2 rounds
- ▶ message complexity: $\leq n + n^2$ messages

Reconstructing Imbs-Raynal's BRB

When $d = 0$ (no MA), 2 versions have same

- ▶ t -resilience: $n > 5t$
- ▶ round complexity: ≤ 2 rounds
- ▶ message complexity: $\leq n + n^2$ messages



However, new version uses smaller quorums:

Threshold	Original version	New version
Forwarding q_f	$n - 2t$	$\lfloor \frac{n+t}{2} \rfloor + 1$
Delivery q_d	$n - t$	$\lfloor \frac{n+3t}{2} \rfloor + 1$



BETTER

Needs less messages to progress \rightarrow **terminates earlier**





Conclusion

- ▶ New computing model:  + 





Conclusion

- ▶ New computing model:  + 
- ▶ New modular many-to-many primitive: ***k2ℓ-cast***
 - ▶ Better quorum engineering

Conclusion

- ▶ New computing model:  + 
- ▶ New modular many-to-many primitive: ***k2ℓ-cast***
 - ▶ Better quorum engineering
- ▶ Applications to quorum-based algorithms:
 - ▶ BRB → -tolerant BRB 
 - ▶ Self-stabilizing systems?

Conclusion

- ▶ New computing model:  + 
- ▶ New modular many-to-many primitive: *k2ℓ-cast*
 - ▶ Better quorum engineering
- ▶ Applications to quorum-based algorithms:
 - ▶ BRB → -tolerant BRB 
 - ▶ Self-stabilizing systems?

Thank you for your attention!

Reconstructing Imbs-Raynal's BRB (Cont'd)

Original version (BRB)

```
operation brb_broadcast( $m, sn$ ) is  
(1) broadcast(INIT( $m, sn$ )).  
  
when INIT( $m, sn$ ) is received from  $p_j$  do  
(2) if (WITNESS( $-, sn, j$ ) not yet broadcast)  
(3)   then broadcast WITNESS( $m, sn, j$ )  
(4) end if.  
  
when WITNESS( $m, sn, j$ ) is received do  
(5) if (WITNESS( $m, sn, j$ ) is received from at least  $n - 2t$  processes  
     $\wedge$  WITNESS( $m, sn, i$ ) not yet broadcast)  
(6)   then broadcast(WITNESS( $m, sn, i$ ))  
(7) end if.  
(8) if (WITNESS( $m, sn, j$ ) is received from at least  $n - t$  processes  
     $\wedge$  ( $-, sn, j$ ) not already brb-delivered)  
(9)   then brb_deliver( $m, sn, j$ )  
(10) end if.
```

\downarrow $k2\ell$ -cast \downarrow

Reconstructed version (MBRB)

```
init:  $obj_w \leftarrow \text{K2LCast}(q_d = \lfloor \frac{n+3t}{2} \rfloor + 3d + 1, q_r = \lfloor \frac{n+t}{2} \rfloor + 1, \text{single} = \text{false})$ .  
(1) operation mbrb_broadcast( $m, sn$ ) is broadcast(INIT( $m, sn$ )).  
(2) when INIT( $m, sn$ ) is received from  $p_j$  do  $obj_w.k2\ell\_cast(\text{WITNESS}(m), (sn, j))$ .  
(3) when (WITNESS( $m$ ), ( $sn, j$ )) is  $obj_w.k2\ell\_delivered$  do mbrb_deliver( $m, sn, j$ ).
```

Signature-Based $k2\ell$ -Cast Implementation (1/2)

object SigBasedK2LCast(q_d) **is**

(1) **operation** $k2\ell_cast(m, id)$ **is**

(2) **if** $(-, id)$ not already signed by p_i **then**

(3) $sig_i \leftarrow$ signature of (m, id) by p_i ;

(4) $sigs_i \leftarrow$ {all valid signatures for (m, id) broadcast by p_i } \cup $\{sig_i\}$;

(5) broadcast(BUNDLE($m, id, sigs_i$));

(6) check_delivery();

(7) **end if.**

(8) **when** BUNDLE($m, id, sigs$) **is received do**

(9) **if** ($sigs$ contains valid signatures for (m, id) not already broadcast by p_i) **then**

(10) $sigs_i \leftarrow$ {all valid signatures for (m, id) broadcast by p_i }

\cup {all valid signatures for (m, id) in $sigs$ };

(11) broadcast(BUNDLE($m, id, sigs_i$));

(12) check_delivery();

(13) **end if.**

(14) **internal operation** check_delivery() **is**

(15) **if** (p_i broadcast at least q_d valid signatures for (m, id)

\wedge $(-, id)$ not already $k2\ell$ -delivered)

(16) **then** $k2\ell_deliver(m, id)$

(17) **end if.**

end object.

Signature-Based $k2\ell$ -Cast Implementation (2/2)

Parameter:

- ▶ q_d : nb of matching signatures that must be received in order to $k2\ell$ -deliver an app-message

No q_f (no forwarding) and no *single* (only a single app-message can be endorsed)

Guarantees:

- ▶ $k' = q_d - n + c$
- ▶ $\delta = q_d > \frac{n+t}{2}$
- ▶ $k = q_d$
- ▶ $\ell = c - d$ (optimal)

This algorithm can still be greatly improved (forwarding, better round and message complexity...)

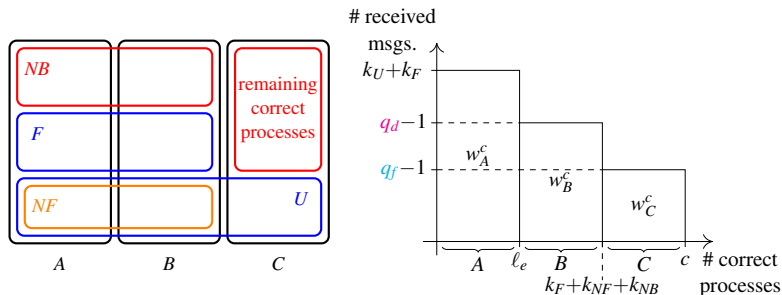
Liveness of Signature-Free $k2\ell$ -Cast (1/3)

Let us assume that k_f correct processes $k2\ell$ -cast (m, id)

We define 7 sets of correct processes:

- ▶ A : processes that reach the delivery threshold q_d ($|A| = \ell_e$)
- ▶ B : processes that reach the forwarding threshold q_f but not q_d
- ▶ C : remaining processes that do not reach q_f
- ▶ U : processes that broadcast $MSG(m, id)$ after a $k2\ell$ -cast
- ▶ F : processes of $A \cup B$ that forward $MSG(m, id)$
- ▶ NF : processes of $A \cup B$ that broadcast $MSG(m, id)$ after a $k2\ell$ -cast
- ▶ NB : processes of $A \cup B$ that never broadcast $MSG(m, id)$, because they have already broadcast some other $MSG(m' \neq m, id)$

Liveness of Signature-Free $k2\ell$ -Cast (2/3)



w_A^c, w_B^c, w_C^c : nb of $\text{MSG}(m, id)$ received by A, B and C

$$w_A^c \leq (k_U + k_F)\ell_e \quad (1)$$

$$w_B^c \leq (q_d - 1)(k_{NF} + k_{NB} + k_F - \ell_e) \quad (2)$$

$$w_C^c \leq (q_f - 1)(c - k_{NF} - k_{NB} - k_F) \quad (3)$$

Liveness of Signature-Free $k2\ell$ -Cast (3/3)

Given that the MA can suppress d copies of each $\text{MSG}(m, id)$ broadcast, we have

$$w_A^c + w_B^c + w_C^c \geq (k_U + k_F)(c - d) \quad (4)$$

By combining (1), (2), (3) and (4) we can obtain

$$\begin{aligned} \ell_e \times (k_U + k_F - q_d + 1) &\geq (k_U + k_F)(c - d - q_d + q_f) \\ &\quad - c(q_f - 1) - k_{NB}(q_d - q_f) \end{aligned} \quad (5)$$

Proving liveness: For each property to prove, plug the values obtained from the hypotheses in the above equation, and derive

Signature-Free $k2\ell$ -Cast Implementation (Cont'd)

Parameters:

- ▶ q_d : nb of $\text{MSG}(m, id)$ to receive to $k2\ell_deliver(m, id)$
- ▶ q_f : nb of $\text{MSG}(m, id)$ to receive to forward $\text{MSG}(m, id)$
- ▶ *single*: true if processes can send only one $\text{MSG}(-, id)$, false otherwise

Guarantees:

- ▶ $k' = q_f - n + c$
- ▶ $\delta = (q_f > \frac{n+t}{2}) \vee (\textit{single} \wedge q_d > \frac{n+t}{2})$
- ▶ $k = \lfloor \frac{c(q_f-1)}{c-d-q_d+q_f} \rfloor + 1$
- ▶ $\ell = c - \lfloor \frac{cd}{c-q_d+1} \rfloor$

k2l-Cast Variables

Variable	Meaning	Best value
k'	minimal nb of correct $k2l$ -casts if there is a correct $k2l$ -delivery	High
δ	true iff no-duplicity is guaranteed, false otherwise	true
k	minimal nb of correct processes that $k2l$ -cast a message	Low
l	minimal nb of correct processes that $k2l$ -deliver a message	High

k2l-Cast Safety

k2l-Validity:

If a correct process p_i *k2l*-delivers (m, id) , then $k' < c$ correct processes *k2l*-cast (m, id)

k2l-No-duplication:

A correct process p_i *k2l*-delivers at most one $(-, id)$

k2l-Conditional-no-duplicity:

If $\delta = \text{true}$, then no two correct processes *k2l*-deliver (m, id) and $(m' \neq m, id)$

$k2\ell$ -Cast Liveness

$k2\ell$ -Local-delivery:

If $k \leq c$ correct processes $k2\ell$ -cast (m, id) and no correct process $k2\ell$ -casts $(m' \neq m, id)$, then at least one correct process p_i $k2\ell$ -delivers (m, id)

$k2\ell$ -Weak-global-delivery:

If a correct process p_i $k2\ell$ -delivers (m, id) , then at least $\ell \leq c - d$ correct processes $k2\ell$ -deliver $(-, id)$

$k2\ell$ -Strong-global-delivery:

If a correct process p_i $k2\ell$ -delivers (m, id) and no correct process $k2\ell$ -casts $(m' \neq m, id)$, then at least $\ell \leq c - d$ correct processes $k2\ell$ -deliver (m, id)