


more

# Broadcast Algorithms

Gustav Cedersjö



# Outline

Part 1 Delivery order

Part 2 Byzantine processes

Module

# Reliable Broadcast

**Validity:** If a correct process  $p$  broadcasts a message  $m$ , then  $p$  eventually delivers  $m$ .

**No duplication:** No message is delivered more than once.

**No creation:** If a process delivers a message  $m$  with sender  $s$ , then  $m$  was previously broadcast by process  $s$ .

**Agreement:** If a message  $m$  is delivered by some correct process, then  $m$  is eventually delivered by every correct process.

old stuff...

Part I

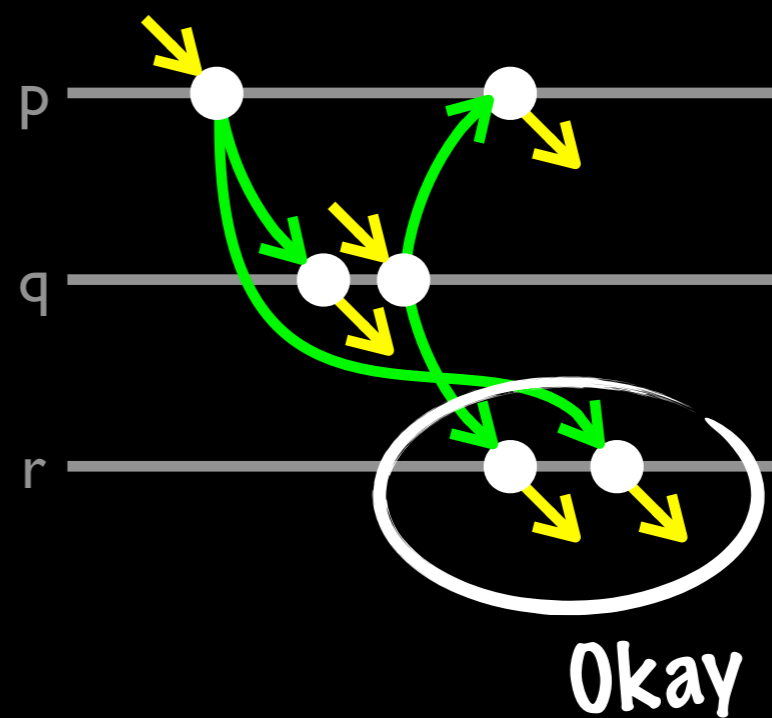

# Delivery Order

Module

# FIFO Reliable Broadcast

Validity, no duplication, no creation, and agreement:  
Reliable Broadcast

**FIFO delivery:** If some process broadcasts message  $m_1$  before it broadcasts  $m_2$ , then no correct process delivers  $m_2$  unless it has already delivered  $m_1$ .



# Broadcast with Sequence Number

- Use a Reliable Broadcast module
- Add sequence numbers to each message
- Keep track of next expected sequence number for each process
- Hold delivery until previous messages are delivered

# Broadcast with Sequence Number

**upon event**  $\langle frb, Init \rangle$  **do**

$lsn := 0;$

$pending := \emptyset;$

$next := [1]^N;$

**upon event**  $\langle frb, Broadcast \mid m \rangle$  **do**

$lsn := lsn + 1;$

**trigger**  $\langle rb, Broadcast \mid [DATA, self, m, lsn] \rangle;$

**upon event**  $\langle rb, Deliver \mid p, [DATA, s, m, sn] \rangle$  **do**


$pending := pending \cup \{(s, m, sn)\};$

**while exists**  $(s, m', sn') \in pending$  such that  $sn' = next[s]$  **do**

$next[s] := next[s] + 1;$

$pending := pending \setminus \{(s, m', sn')\};$

**trigger**  $\langle frb, Deliver \mid s, m' \rangle;$





# Broadcast with Sequence Number

**upon event**  $\langle frb, Init \rangle$  **do**

$lsn := 0;$

$pending := \emptyset;$

$next := [1]^N;$

**upon event**  $\langle frb, Broadcast \mid m \rangle$  **do**

$lsn := lsn + 1;$

**trigger**  $\langle rb, Broadcast \mid [DATA, self, m, lsn] \rangle;$

**upon event**  $\langle rb, Deliver \mid p, [DATA, s, m, sn] \rangle$  **do**

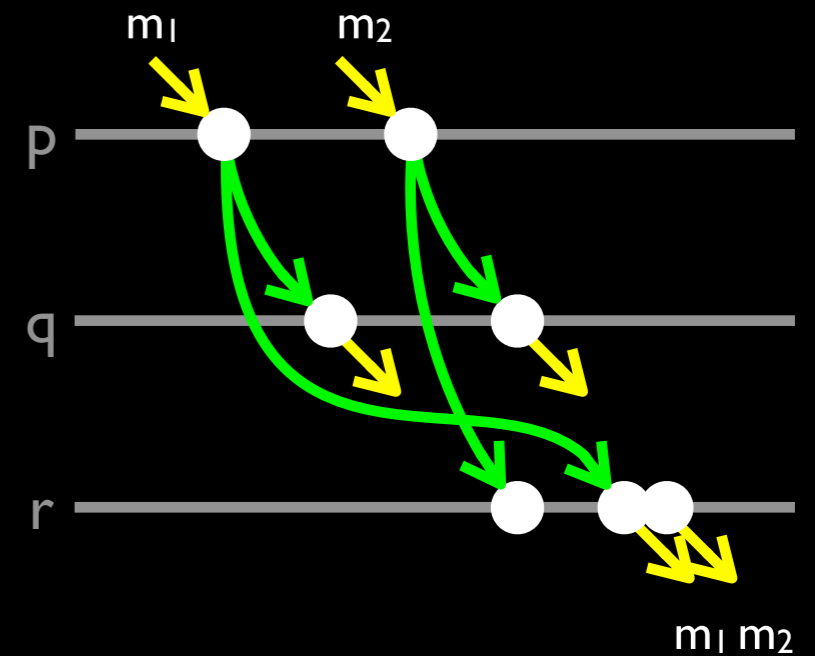
$pending := pending \cup \{(s, m, sn)\};$

**while exists**  $(s, m', sn') \in pending$  such that  $sn' = next[s]$  **do**

$next[s] := next[s] + 1;$

$pending := pending \setminus \{(s, m', sn')\};$

**trigger**  $\langle frb, Deliver \mid s, m' \rangle;$





Module

# Causal Reliable Broadcast





Validity, no duplication, no creation, and agreement:  
Reliable Broadcast





**Causal order:** For any message  $m_1$  that potentially caused a message  $m_2$ , no process delivers  $m_2$  unless it has already delivered  $m_1$ .






# No-Waiting Causal Broadcast


- Use a Reliable Broadcast module
- Send the history of all received messages together with each message
- When receiving a message, deliver all undelivered messages in its history before delivering the new message itself







yes [hi hej]



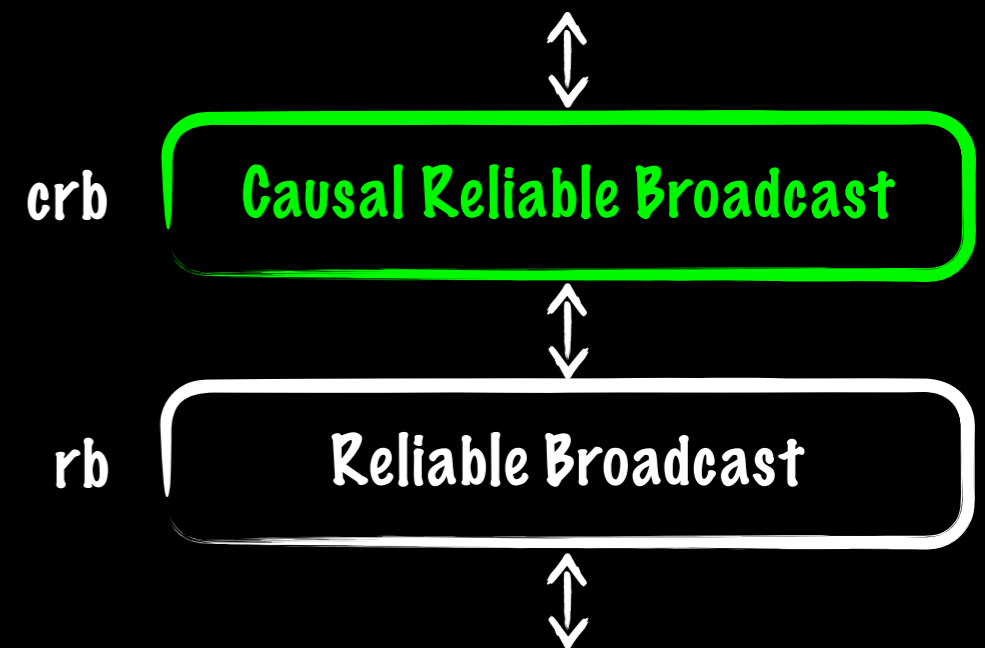
## Fail-Silent Algorithm

# No-Waiting Causal Broadcast

```
upon event  $\langle crb, \text{Init} \rangle$  do
     $delivered := \emptyset$ ;  $past := []$ ;

upon event  $\langle crb, \text{Broadcast} \mid m \rangle$  do
    trigger  $\langle rb, \text{Broadcast} \mid [\text{DATA}, past, m] \rangle$ ;
    append( $past, (self, m)$ );

upon event  $\langle rb, \text{Deliver} \mid p, [\text{DATA}, mpast, m] \rangle$  do
    if  $m \notin delivered$  then
        forall  $(s, n) \in mpast$  do // by the order in the list
            if  $n \notin delivered$  then
                trigger  $\langle crb, \text{Deliver} \mid s, n \rangle$ ;
                 $delivered := delivered \cup \{n\}$ ;
                if  $(s, n) \notin past$  then append( $past, (s, n)$ );
            trigger  $\langle crb, \text{Deliver} \mid p, m \rangle$ ;
             $delivered := delivered \cup \{m\}$ ;
            if  $(p, m) \notin past$  then append( $past, (p, m)$ );
```



## Fail-Silent Algorithm

# No-Waiting Causal Broadcast

**upon event**  $\langle crb, Init \rangle$  **do**

$delivered := \emptyset;$   $past := [];$

**upon event**  $\langle crb, Broadcast \mid m \rangle$  **do**

**trigger**  $\langle rb, Broadcast \mid [DATA, past, m] \rangle;$

$append(past, (self, m));$

**upon event**  $\langle rb, Deliver \mid p, [DATA, mpast, m] \rangle$  **do**

**if**  $m \notin delivered$  **then**

**forall**  $(s, n) \in mpast$  **do** // by the order in the list

**if**  $n \notin delivered$  **then**

**trigger**  $\langle crb, Deliver \mid s, n \rangle;$

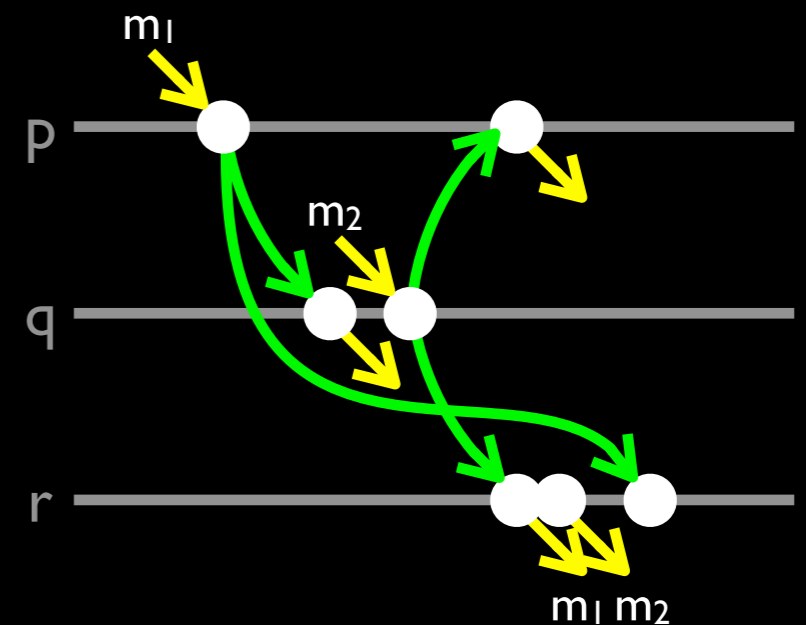
$delivered := delivered \cup \{n\};$

**if**  $(s, n) \notin past$  **then**  $append(past, (s, n));$

**trigger**  $\langle crb, Deliver \mid p, m \rangle;$

$delivered := delivered \cup \{m\};$

**if**  $(p, m) \notin past$  **then**  $append(past, (p, m));$







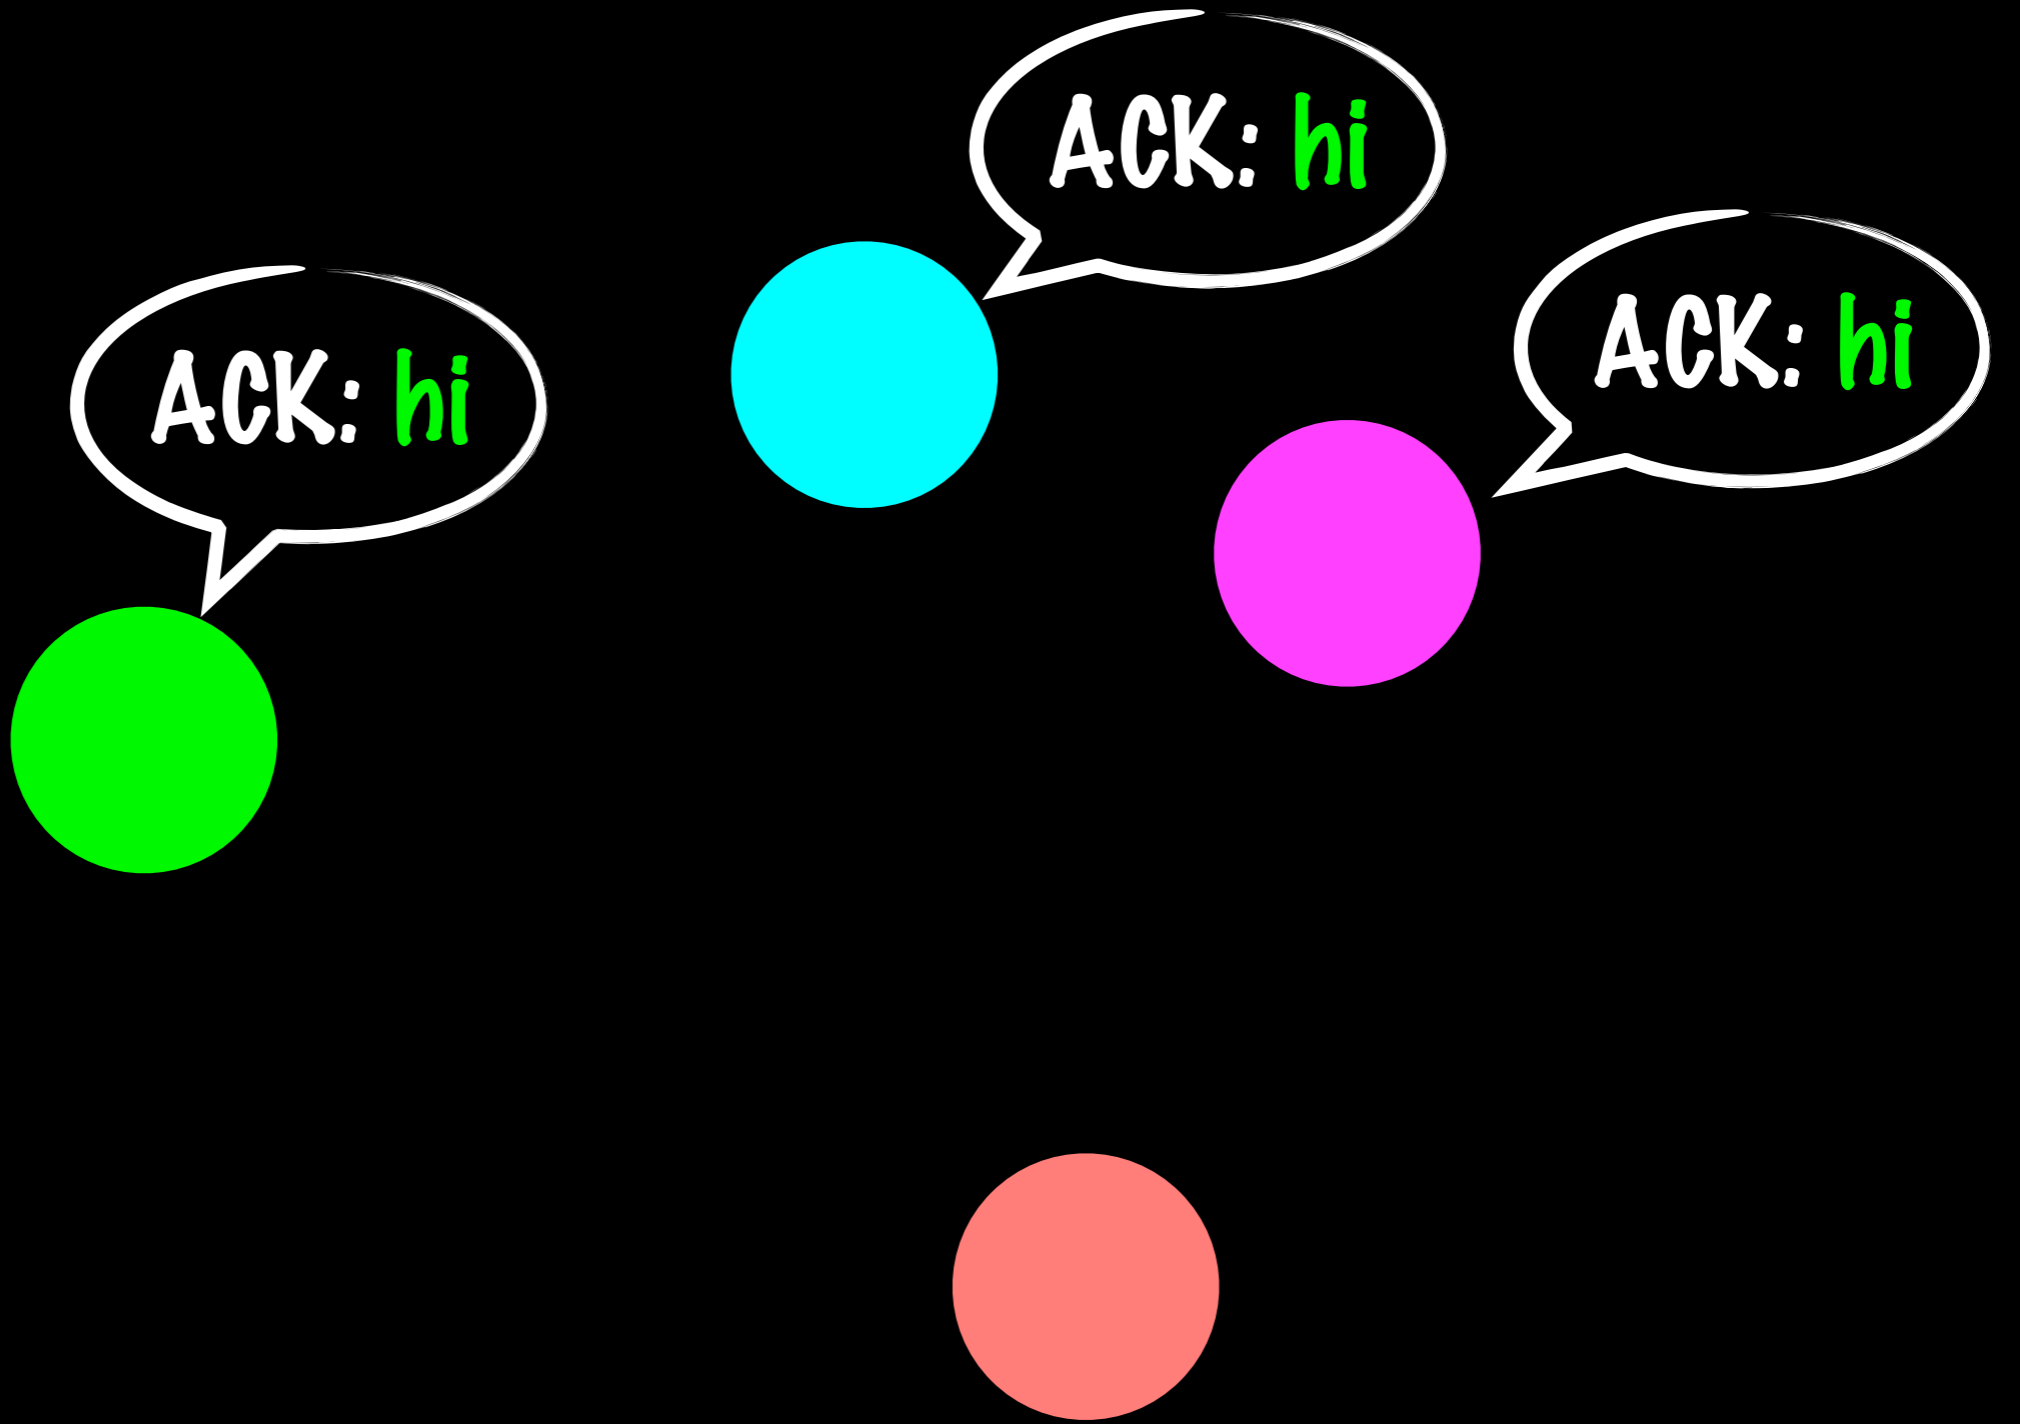






# Garbage-Collection of Causal Past

- Use a Reliable Broadcast module
- Use a Perfect Failure Detector module
- Send message history with each message
- Broadcast an *ack* for each delivered message.
- Remove a message from history when all correct processes have *ack*'ed a message

DATA: hi [1]










DATA: hej [hi]

ACK: hej




ACK: hej






ACK: hej


ACK: hej





ACK: hi




DATA: *yes* [1]



ACK: yes




ACK: yes



ACK: yes

ACK: yes



# Garbage-Collection of Causal Past

**upon event**  $\langle crb, Init \rangle$  **do**

$delivered := \emptyset;$

$past := [];$

$correct := \Pi;$

**forall**  $m$  **do**  $ack[m] := \emptyset;$

**upon event**  $\langle P, Crash \mid p \rangle$  **do**

$correct := correct \setminus \{p\};$

**upon exists**  $m \in delivered$  such that  $self \notin ack[m]$  **do**

$ack[m] := ack[m] \cup \{self\};$

**trigger**  $\langle rb, Broadcast \mid [ACK, m] \rangle;$

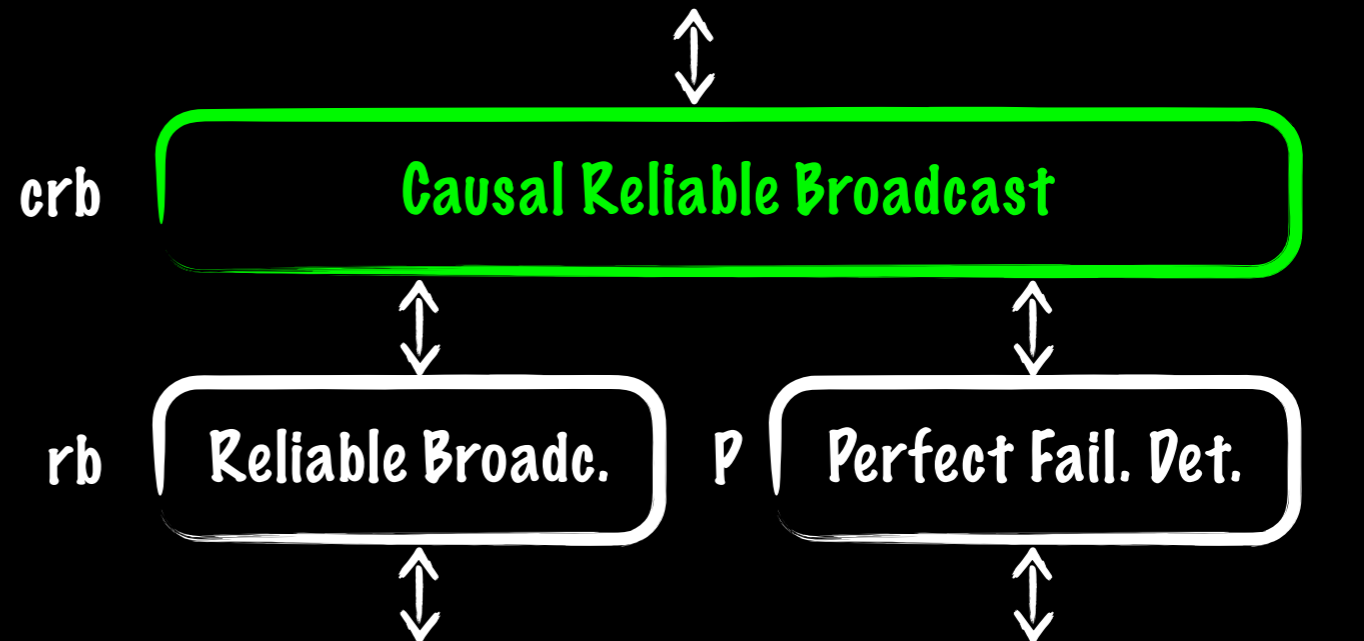
**upon event**  $\langle rb, Deliver \mid p, [ACK, m] \rangle$  **do**

$ack[m] := ack[m] \cup \{p\};$

**upon**  $correct \subseteq ack[m]$  **do**

**forall**  $(s', m') \in past$  such that  $m' = m$  **do**

$remove(past, (s', m));$



**upon event**  $\langle crb, Broadcast \mid m \rangle$  **do**

// same as before

**upon event**  $\langle rb, Deliver \mid p, [DATA, mp, m] \rangle$  **do**

// same as before

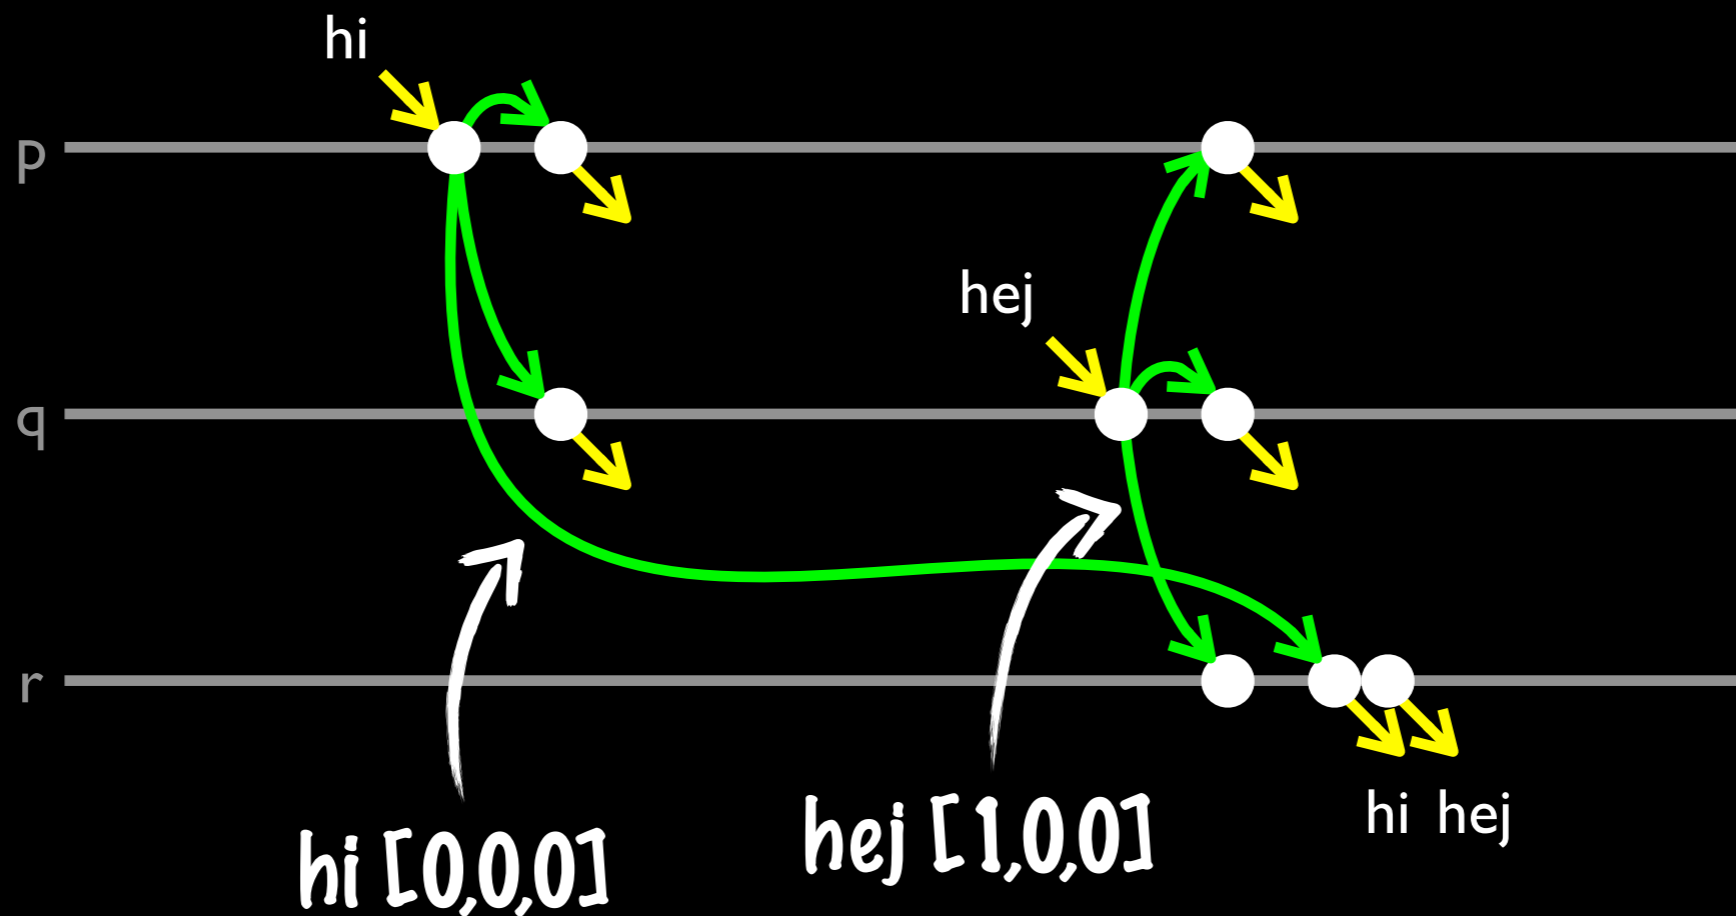


# Waiting Causal Broadcast

- Use a Reliable Broadcast module
- Record the number of delivered messages from each process in a vector  $V$
- Send the vector  $V$  with each message
- Delay the delivery of a message  $m$  with vector  $W$  until  $W \leq V \Leftrightarrow \forall i W[i] \leq V[i]$

Fail-Silent Algorithm

# Waiting Causal Broadcast



## Fail-Silent Algorithm

# Waiting Causal Broadcast

**upon event**  $\langle crb, Init \rangle$  **do**

$V := [0]^N;$

$lsn := 0;$

$pending := \emptyset;$

**upon event**  $\langle crb, Broadcast \mid m \rangle$  **do**

$W := V;$

$W[rank(self)] := lsn;$

$lsn := lsn + 1;$

**trigger**  $\langle rb, Broadcast \mid [DATA, W, m] \rangle;$

**upon event**  $\langle rb, Deliver \mid p, [DATA, W, m] \rangle$  **do**


$pending := pending \cup \{(p, W, m)\};$

**while exists**  $(p', W', m') \in pending$  such that  $W' \leq V$  **do**

$pending := pending \setminus \{(p', W', m')\};$

$V[rank(p')] := V[rank(p')] + 1;$

**trigger**  $\langle crb, Deliver \mid p', m' \rangle;$



Part 2

# Byzantine Processes

# Byzantine Consistent Broadcast

**Validity:** If a correct process  $p$  broadcasts a message  $m$ , then every correct process eventually delivers  $m$ .

**No duplication:** Every correct process delivers at most one message.

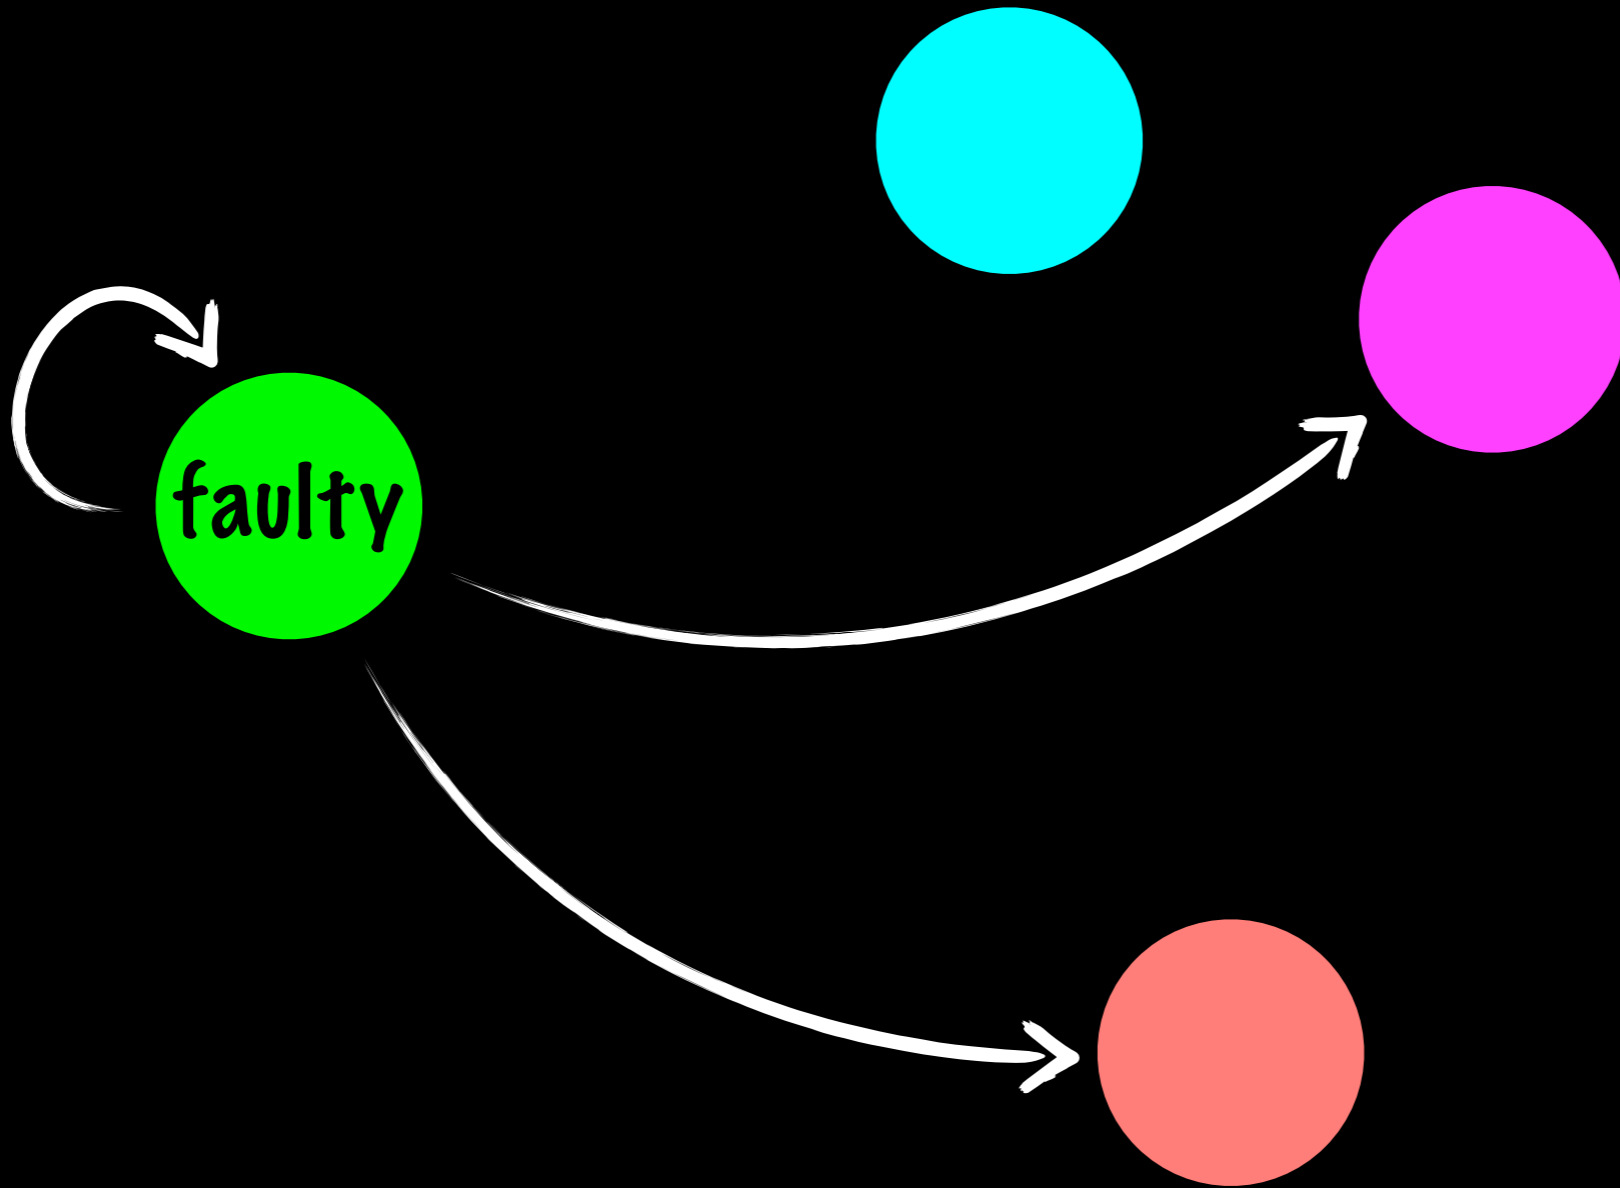
**Integrity:** If some correct process delivers a message  $m$  with sender  $p$  and process  $p$  is correct, then  $m$  was previously broadcast by  $p$ .

**Consistency:** If some correct process delivers a message  $m$  and another correct process delivers a message  $m'$ , then  $m = m'$ .


# Authenticated Echo Broadcast

- Use Authenticated Point to Point Links
- When you receive the first message  $m$  from sender  $s$ , send  $m$  as ECHO to all other processes
- If you get more than  $(N + f) / 2$  ECHO with the same message  $m$ , then deliver  $m$

**SEND**




ECHO





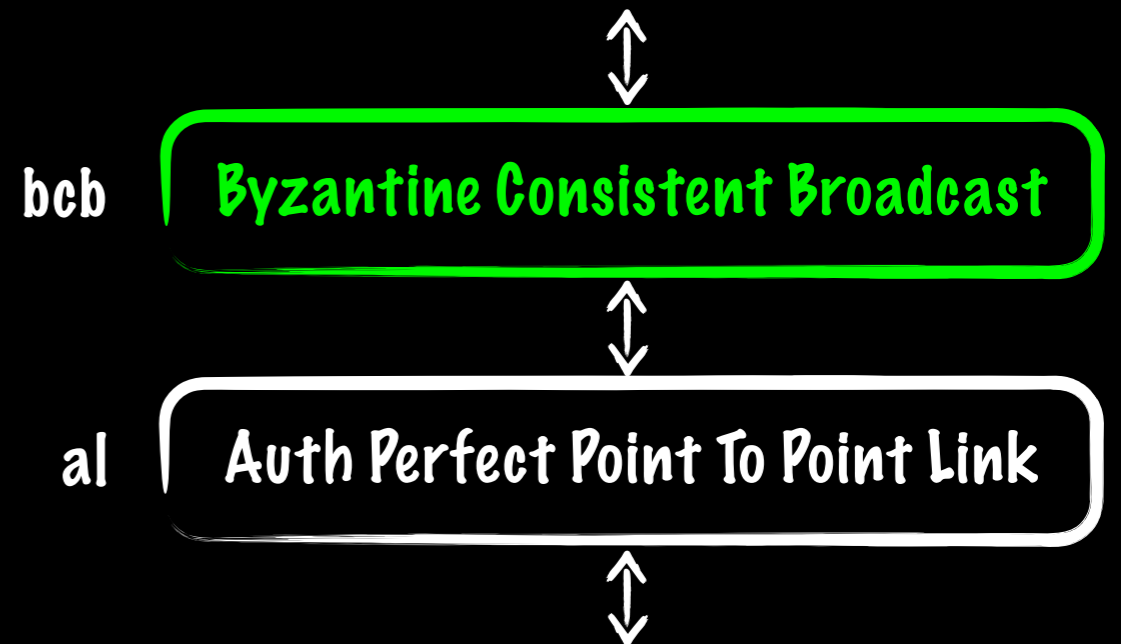
DELIVERY



## Fail-Arbitrary Algorithm

# Authenticated Echo Broadcast

```
upon event  $\langle bcb, Init \rangle$  do  
     $sentecho := FALSE;$   
     $delivered := FALSE;$   
     $echos := [\perp]^N;$   
  
upon event  $\langle bcb, Broadcast \mid m \rangle$  do  
    forall  $q \in \Pi$  do  
        trigger  $\langle al, Send \mid q, [SEND, m] \rangle;$   
  
upon event  $\langle al, Deliver \mid p, [SEND, m] \rangle$  such that  $p = s$   
    and  $sentecho = FALSE$  do  
         $sentecho := TRUE;$   
        forall  $q \in \Pi$  do  
            trigger  $\langle al, Send \mid q, [ECHO, m] \rangle;$   
  
upon event  $\langle al, Deliver \mid p, [ECHO, m] \rangle$  do  
    if  $echos[p] = \perp$  then  
         $echos[p] := m;$   
  
upon exists  $m \neq \perp$  such that  $\#(\{p \in \Pi \mid echos[p] = m\}) > (N+f)/2$   
    and  $delivered = FALSE$  do  
         $delivered := TRUE;$   
        trigger  $\langle bcb, Deliver \mid s, m \rangle;$ 
```



Module

# Byzantine Reliable Broadcast

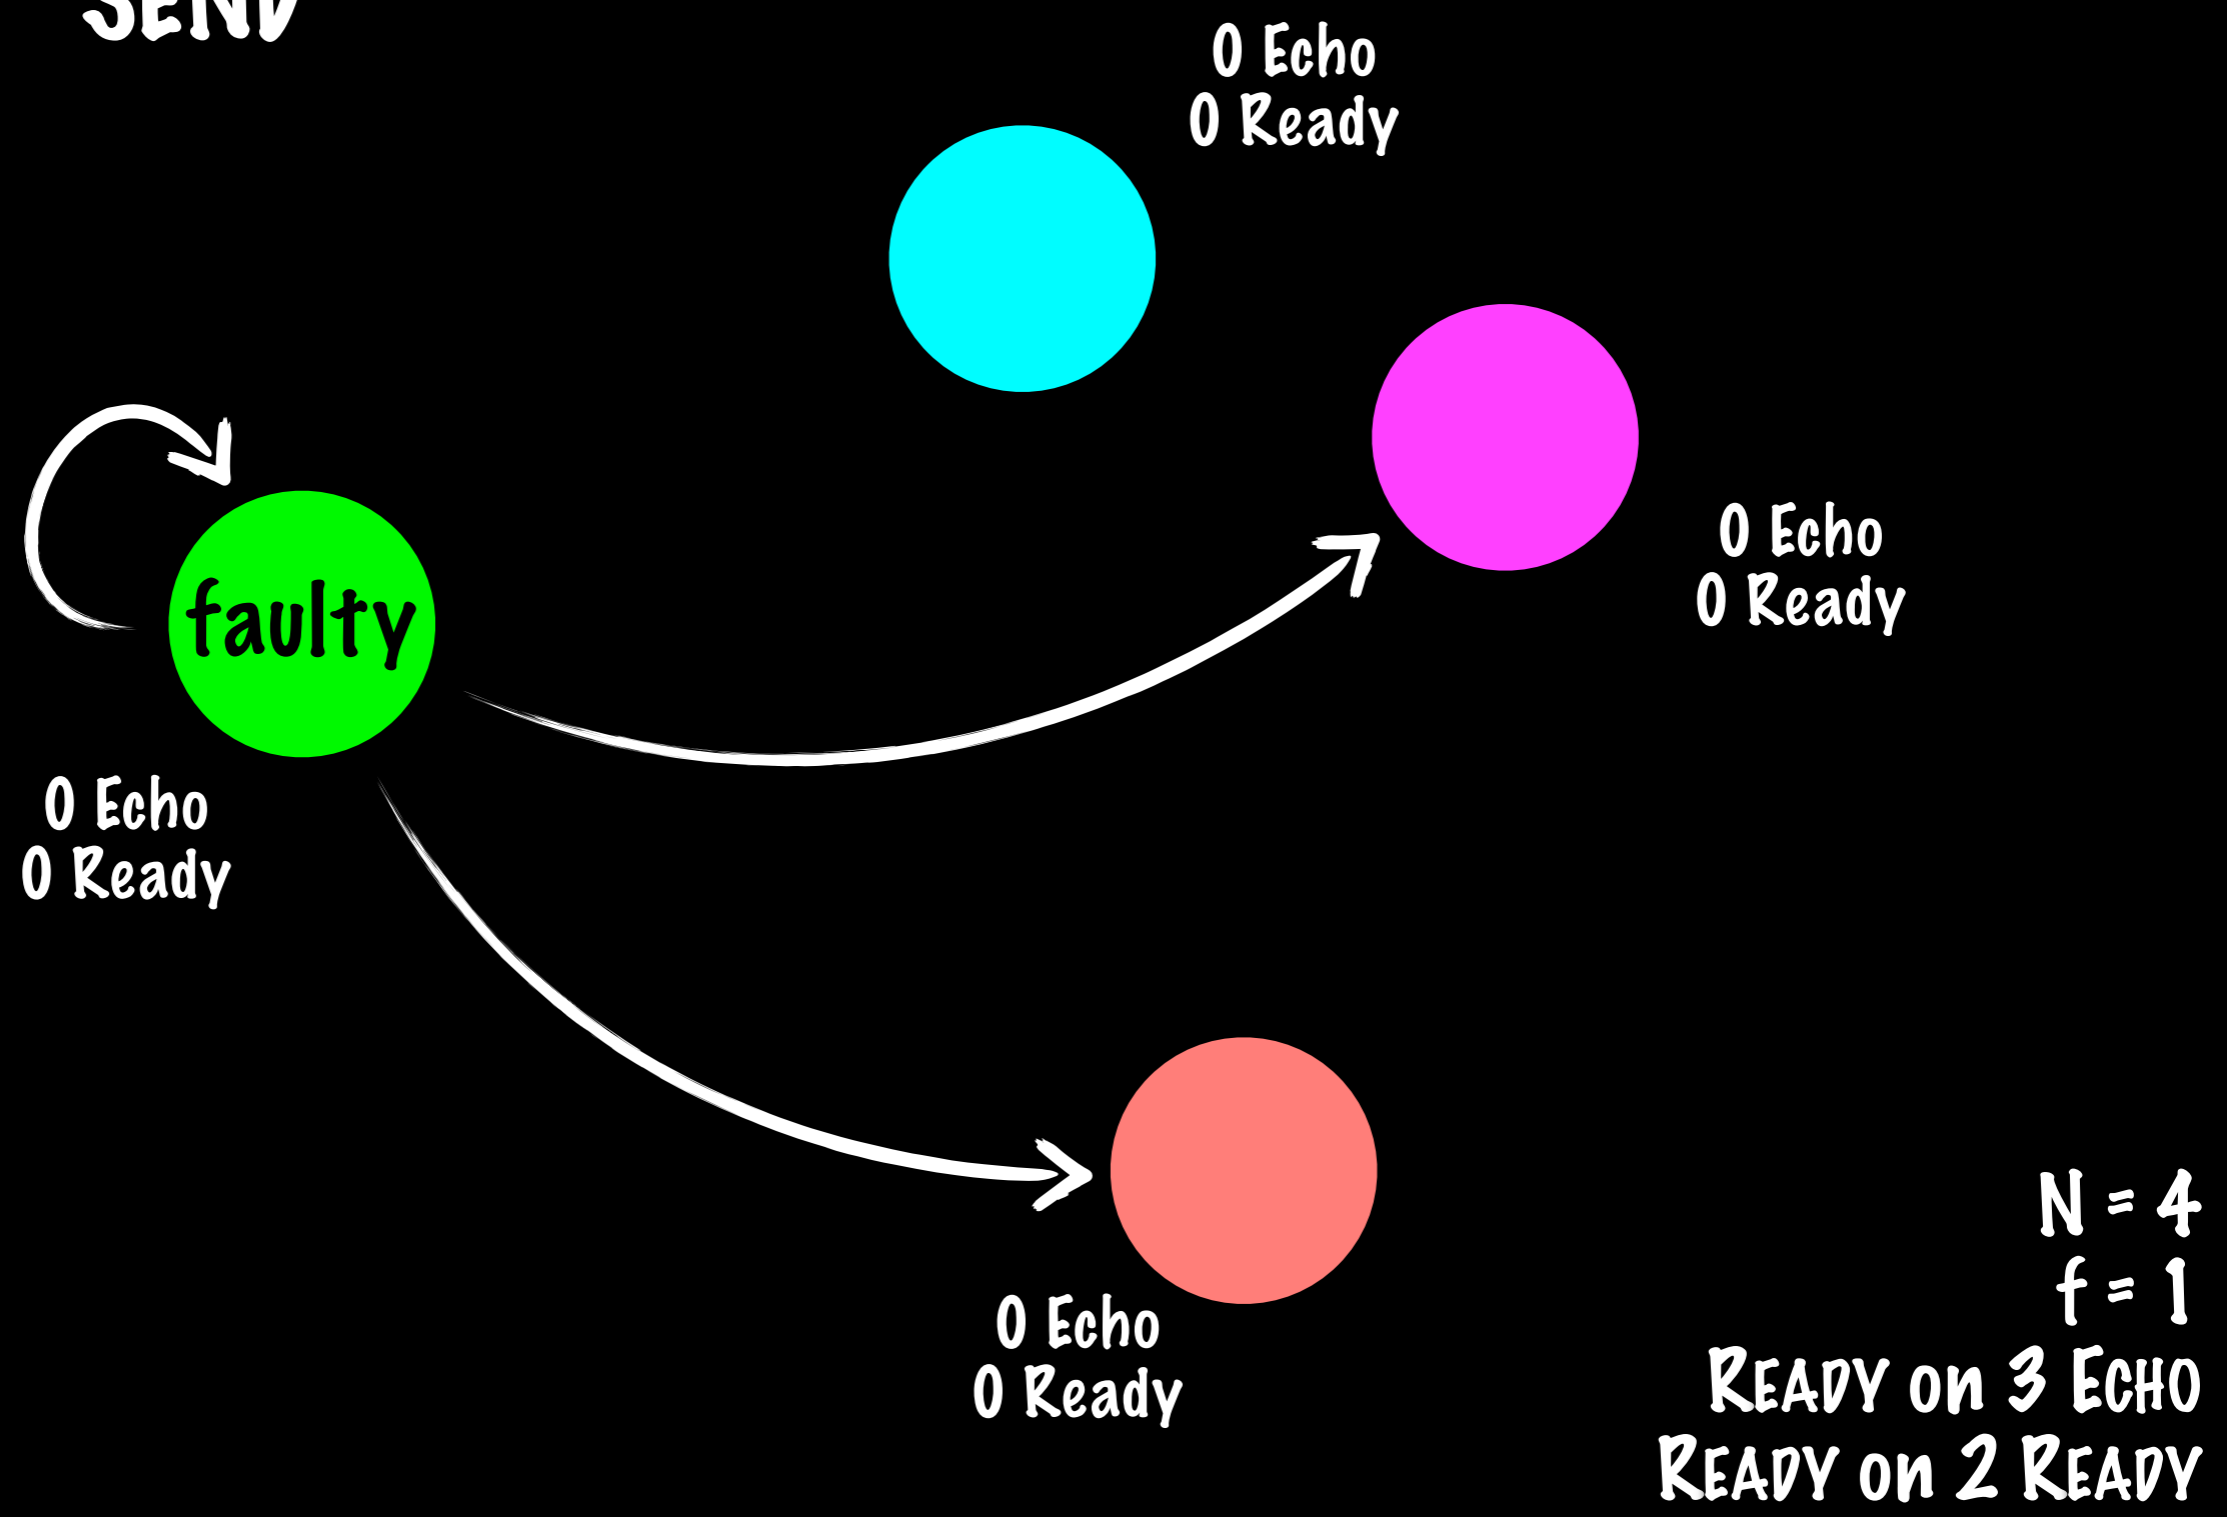
**Validity, No duplication, Integrity and Consistency:** Byzantine Consistent Broadcast

**Totally:** If some message is delivered by any correct process, every correct processes eventually delivers a message

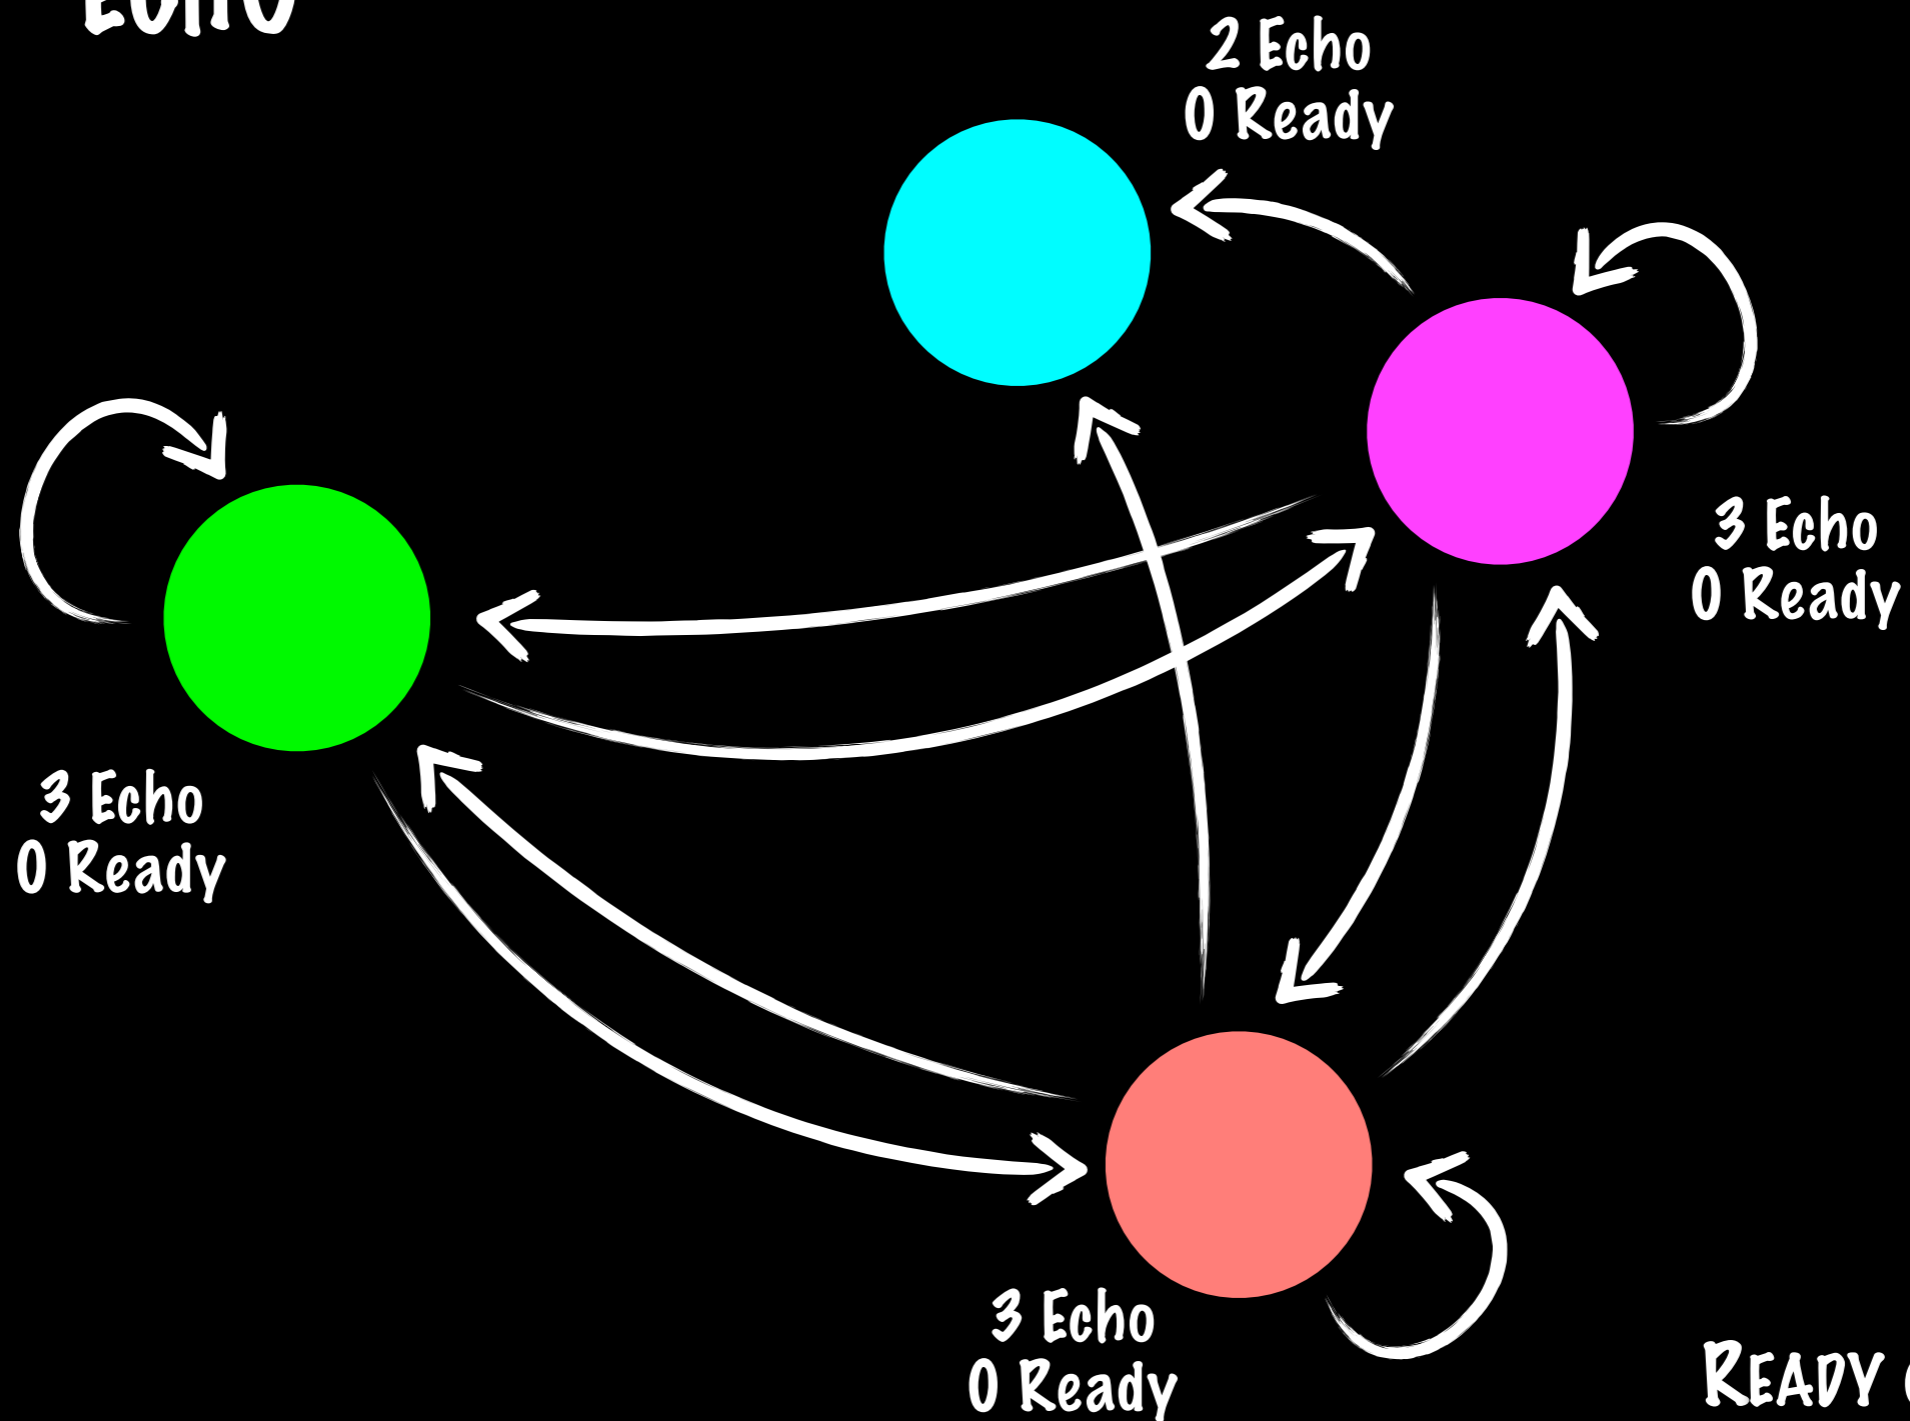
# Authenticated Double-Echo Broadcast

- Use Authenticated Point to Point Links
- When you receive the first message  $m$  from sender  $s$ , send  $m$  as ECHO to all other processes
- If you get more than  $(N + f) / 2$  ECHO or more than  $f$  READY with the same message  $m$ , then send  $m$  as READY to all other processes
- If you get more than  $2f$  READY with the same message  $m$ , then deliver  $m$

**SEND**

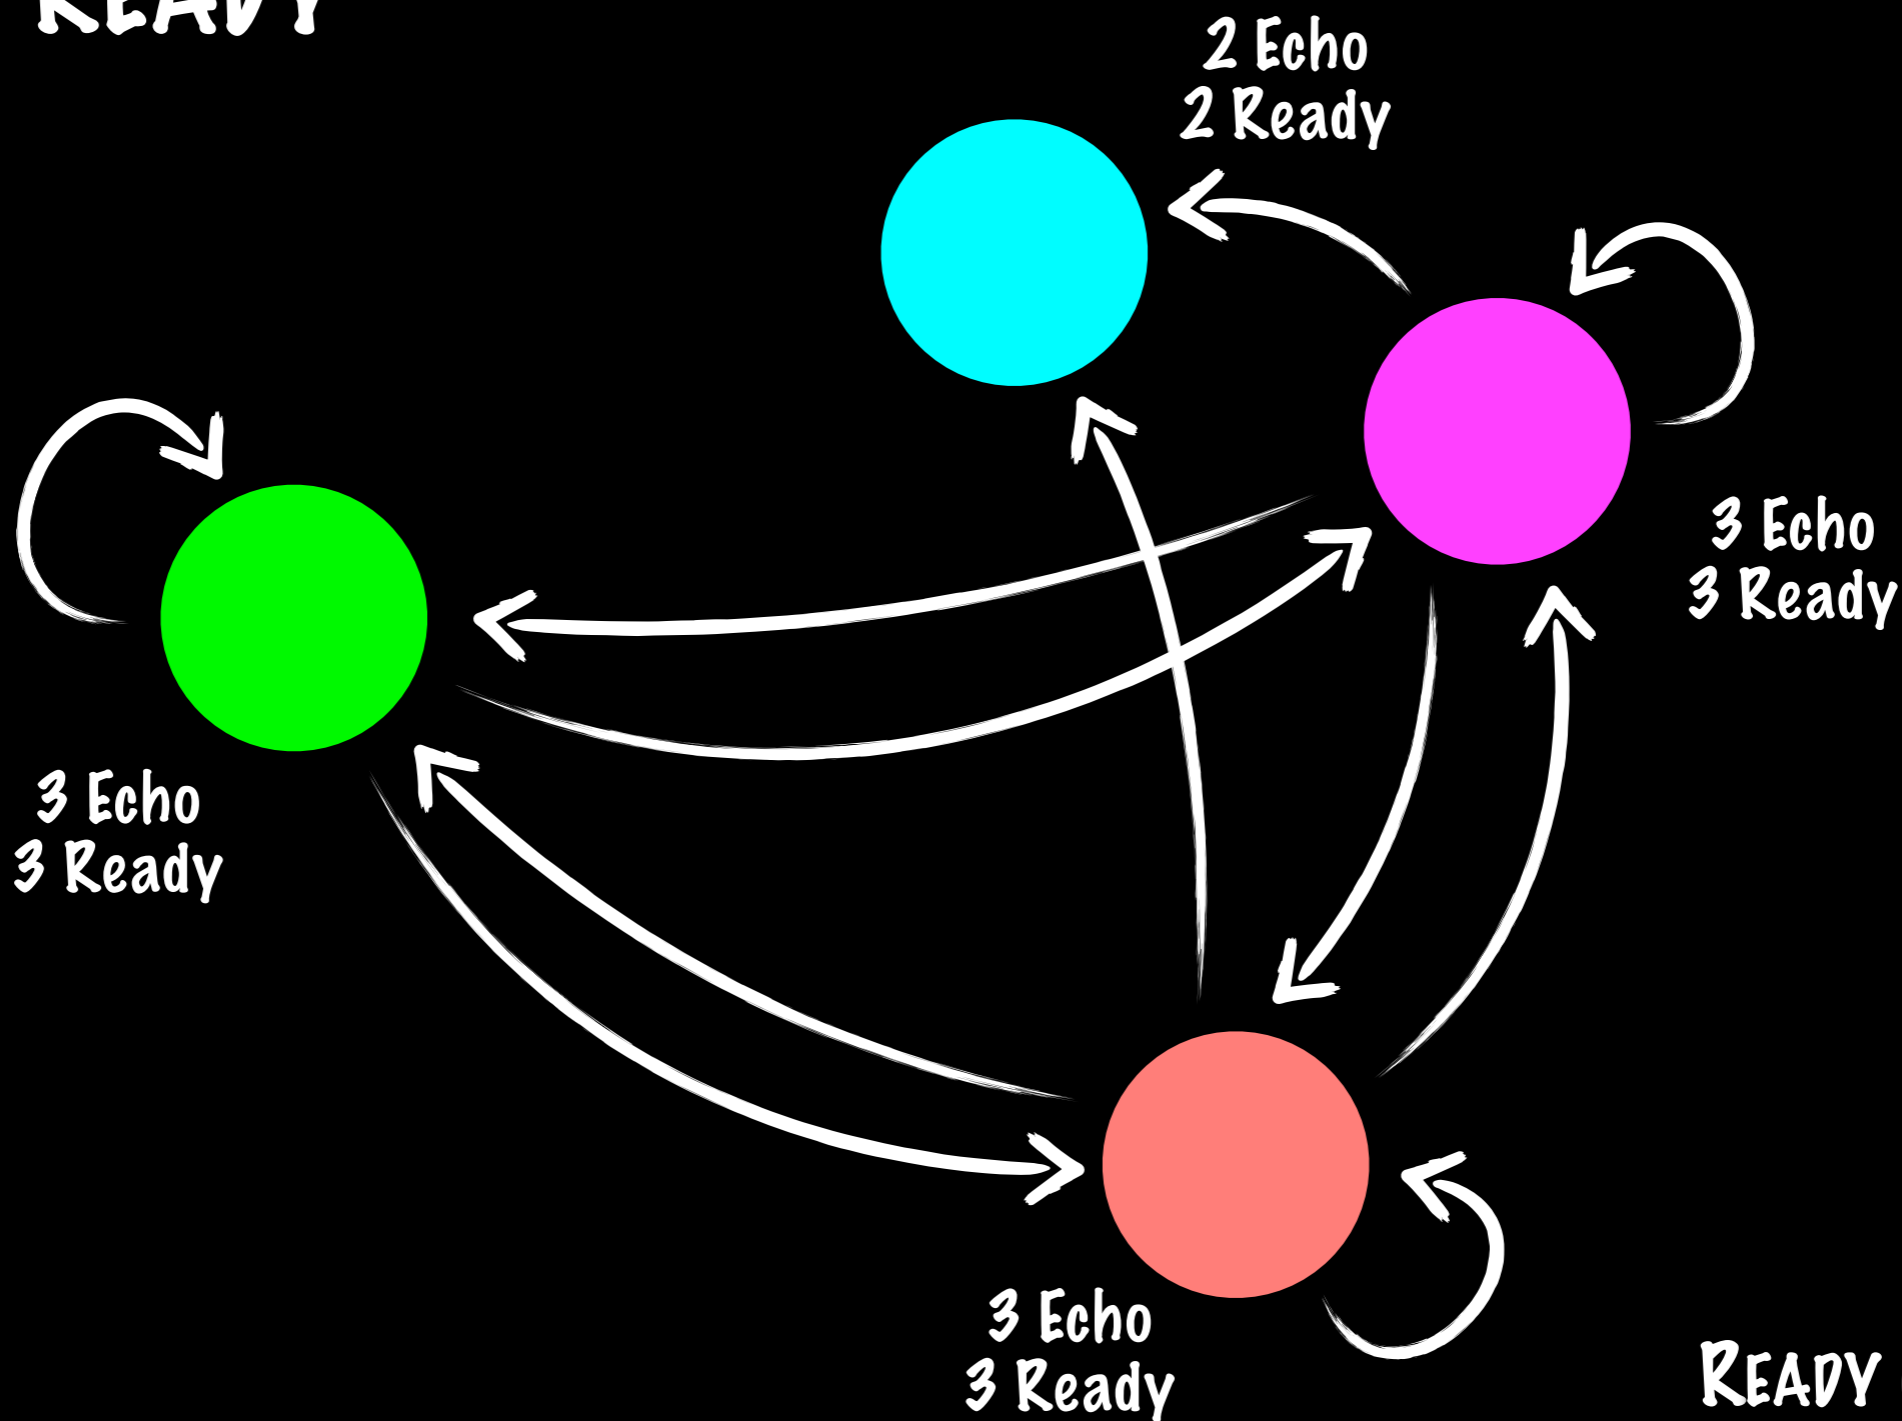


# ECHO



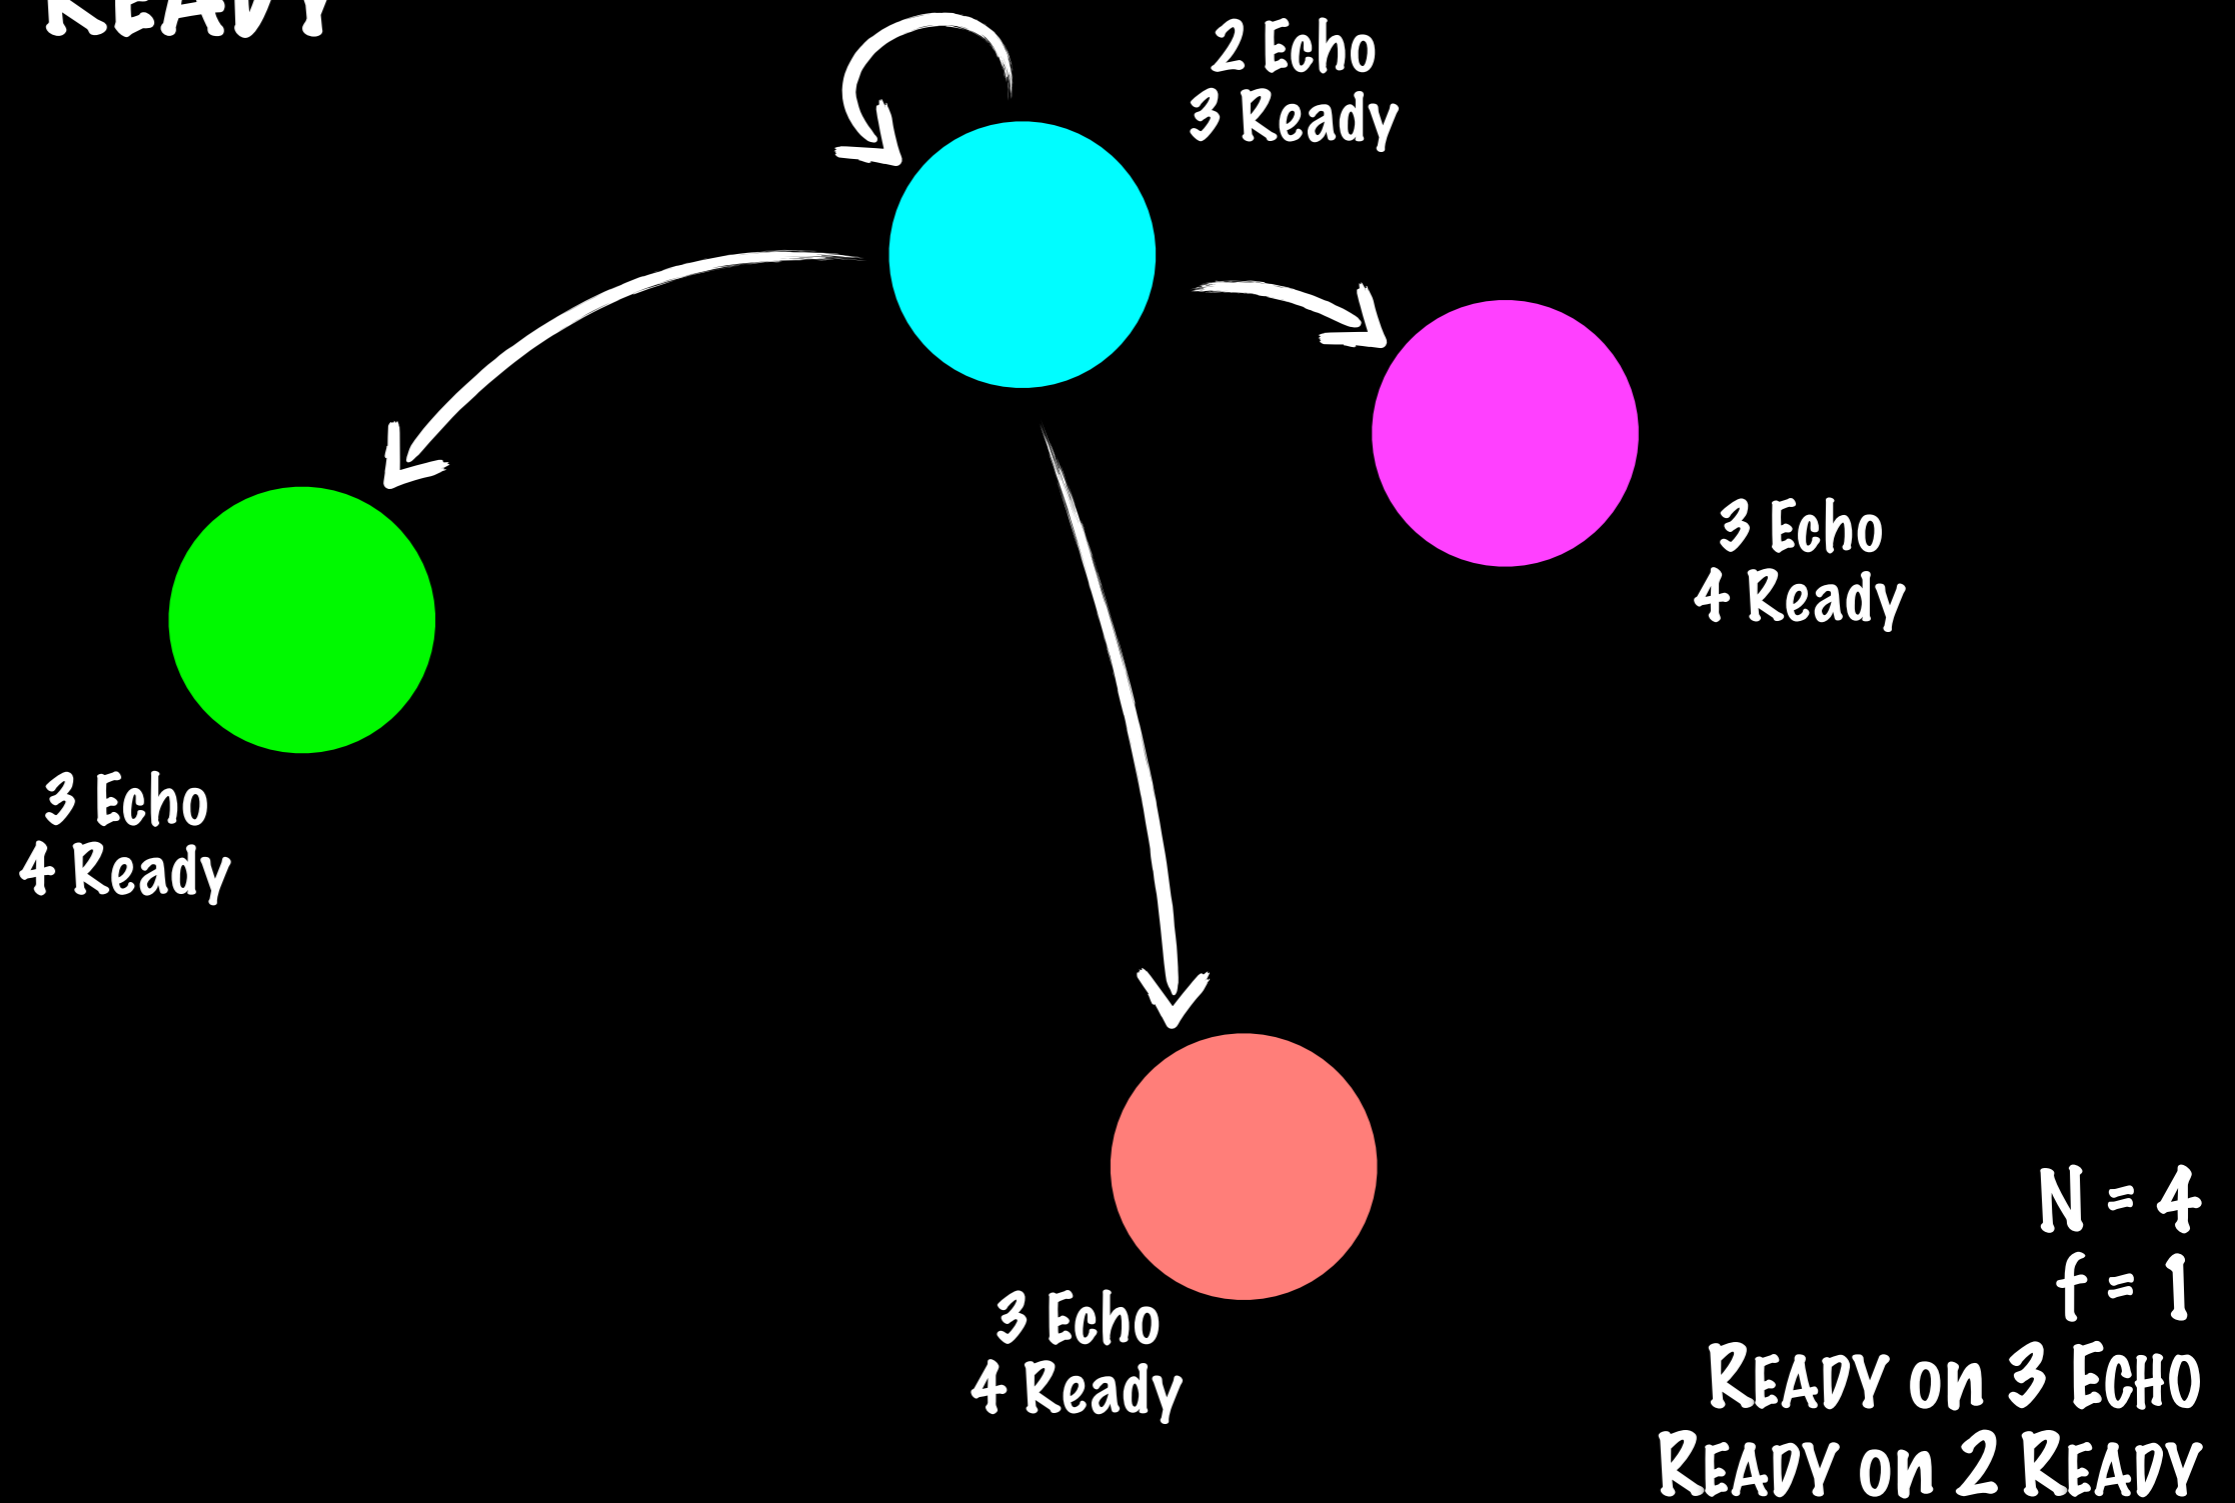
$N = 4$   
 $f = 1$   
READY on 3 ECHO  
READY on 2 READY

READY



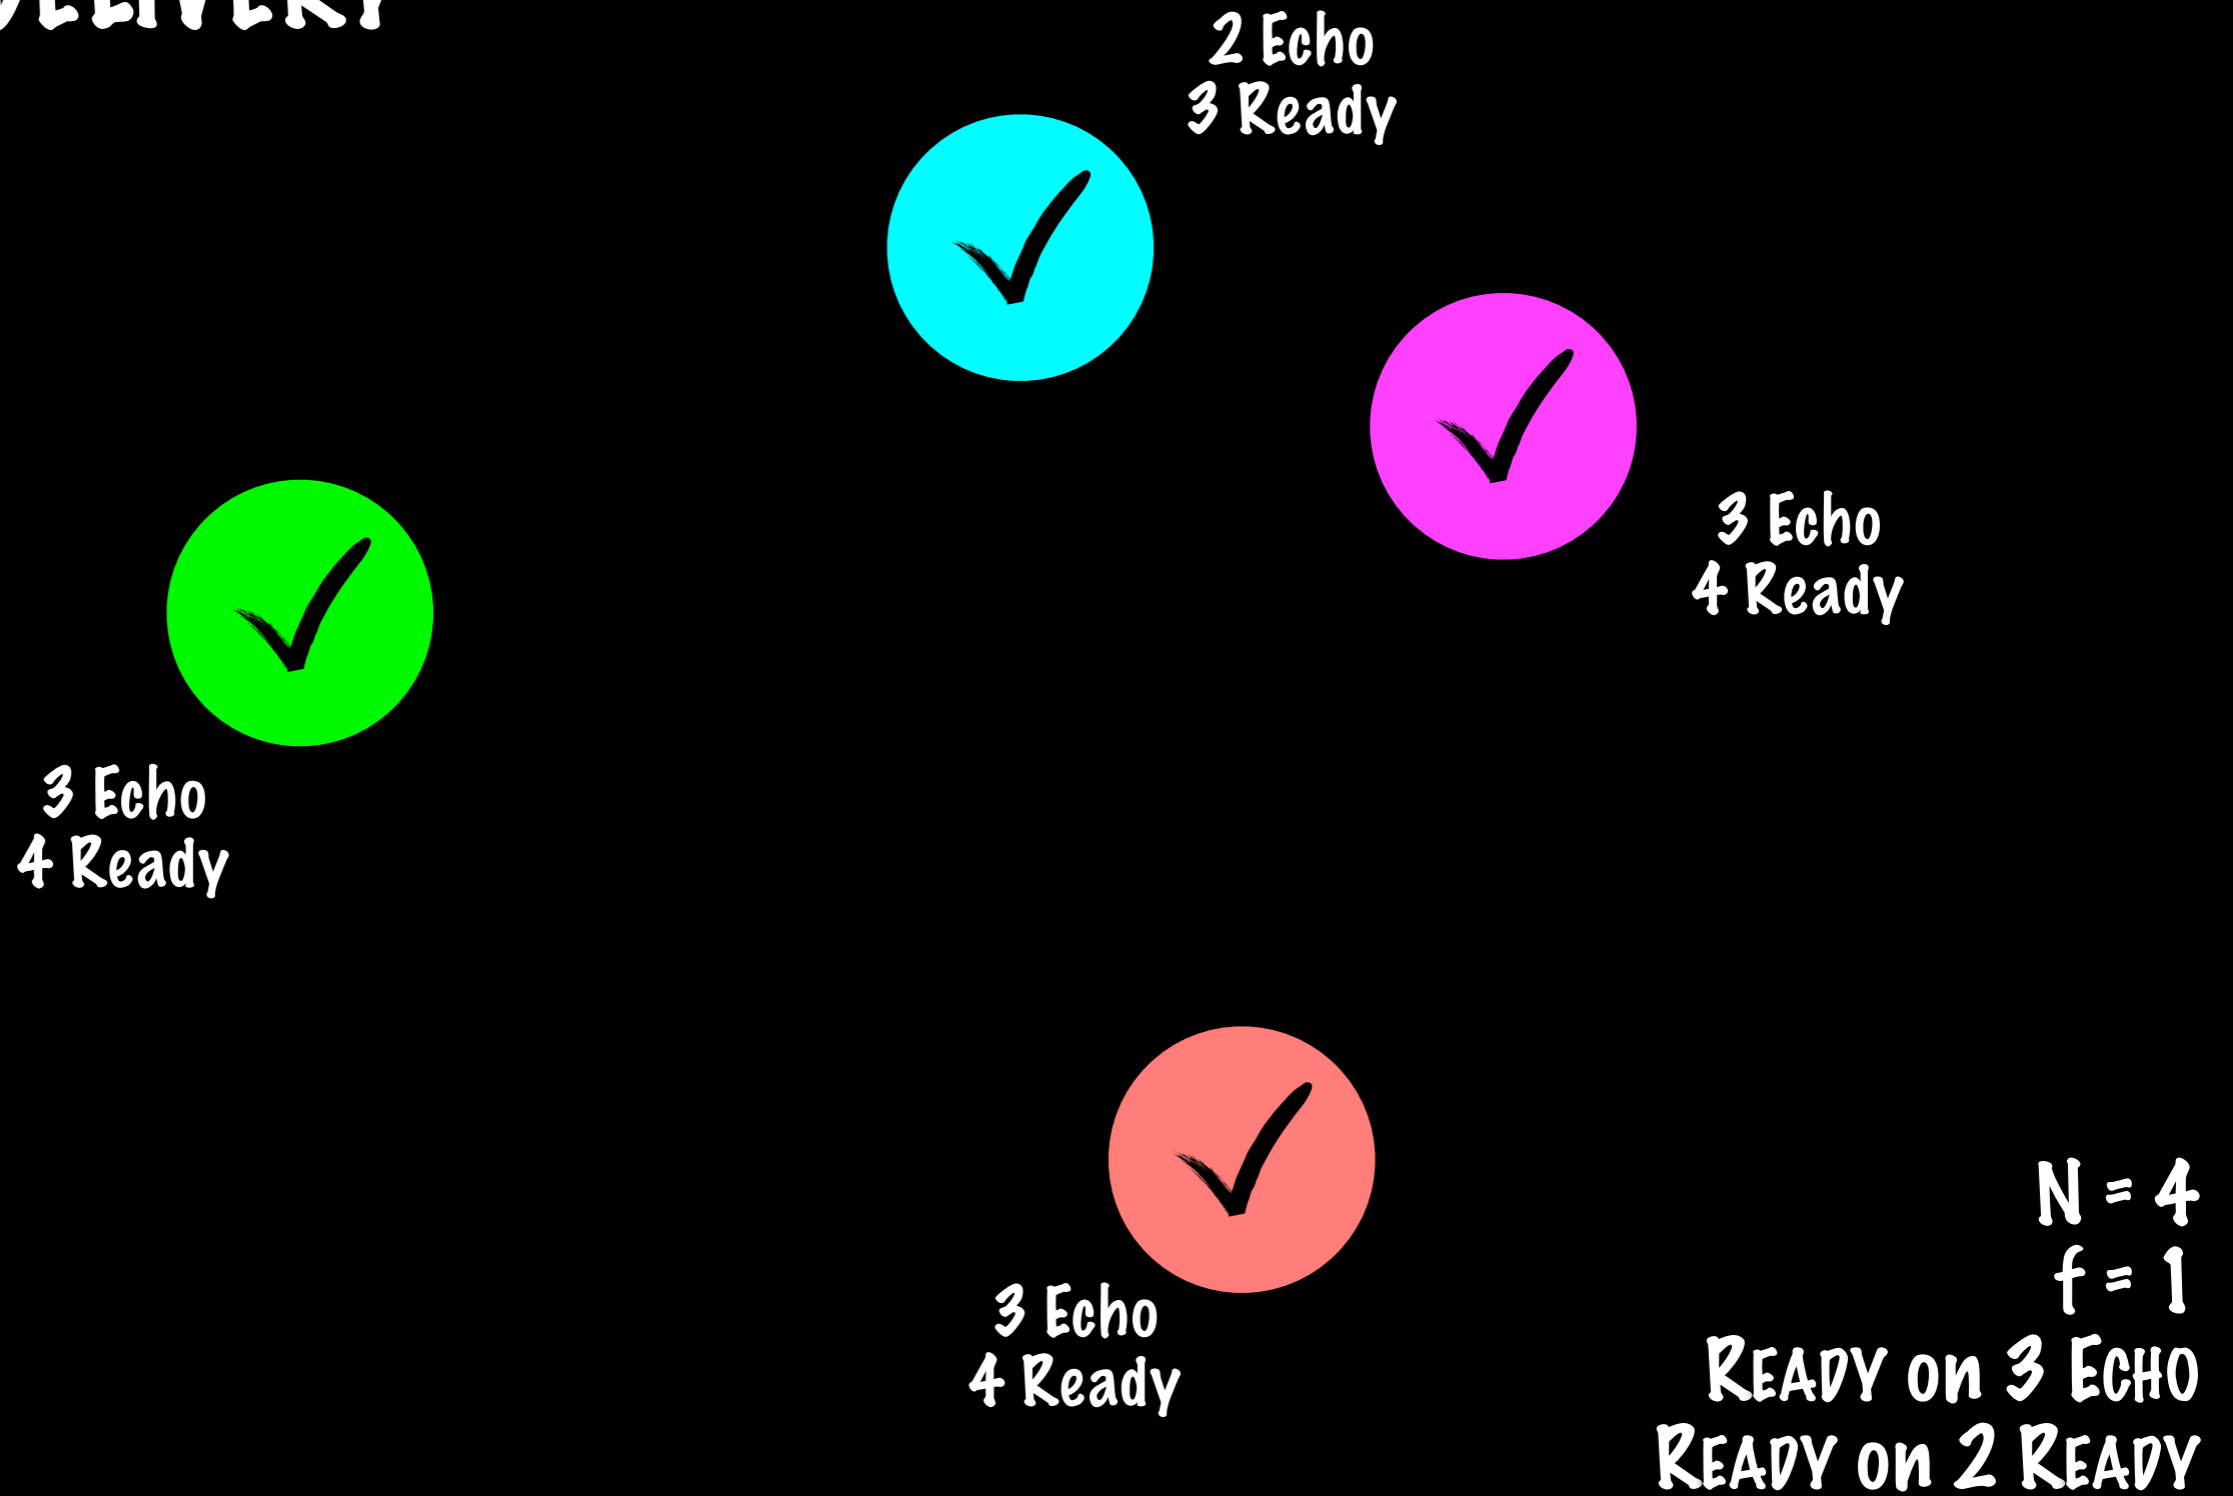
$N = 4$   
 $f = 1$   
READY on 3 ECHO  
READY on 2 READY

**READY**





# DELIVERY



# Byzantine Consistent Broadcast Channel

**Validity:** If a correct process  $p$  broadcasts a message  $m$ , then every correct process eventually delivers  $m$ .

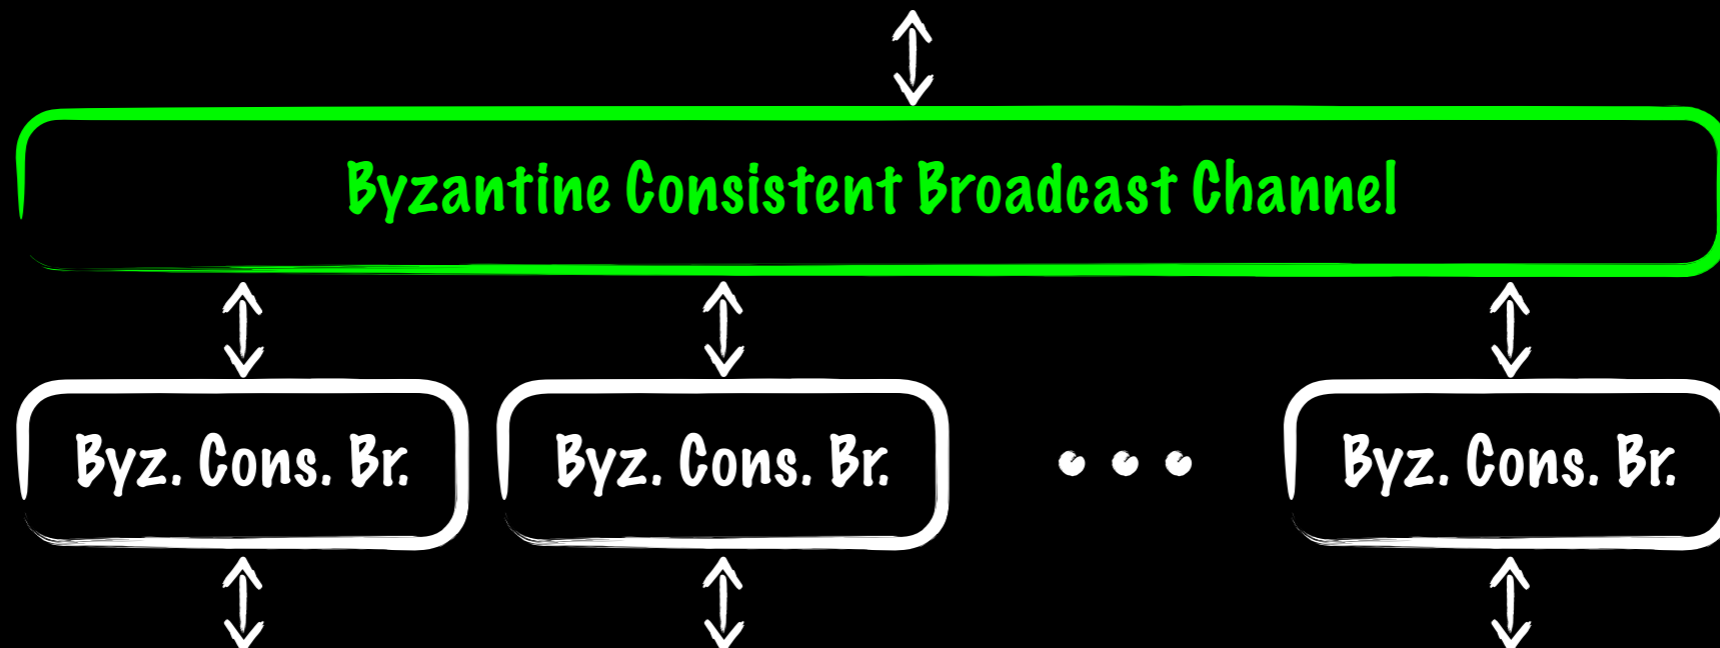
**No duplication:** For every process  $p$  and label  $l$ , every correct process delivers at most one message with label  $l$  and sender  $p$ .

**Integrity:** If some correct process delivers a message  $m$  with sender  $p$  and process  $p$  is correct, then  $m$  was previously broadcast by  $p$ .

**Consistency:** If some correct process delivers a message  $m$  with label  $l$  and sender  $s$ , and another correct process delivers a message  $m'$  with label  $l$  and sender  $s$ , then  $m = m'$ .

# Byzantine Consistent Channel

- Create an instance of a **Byzantine Consistent Broadcast** for each process.
- On delivery of a message from sender  $s$ , create a new instance of a **Byzantine Consistent Broadcast** for process  $s$ .



Module

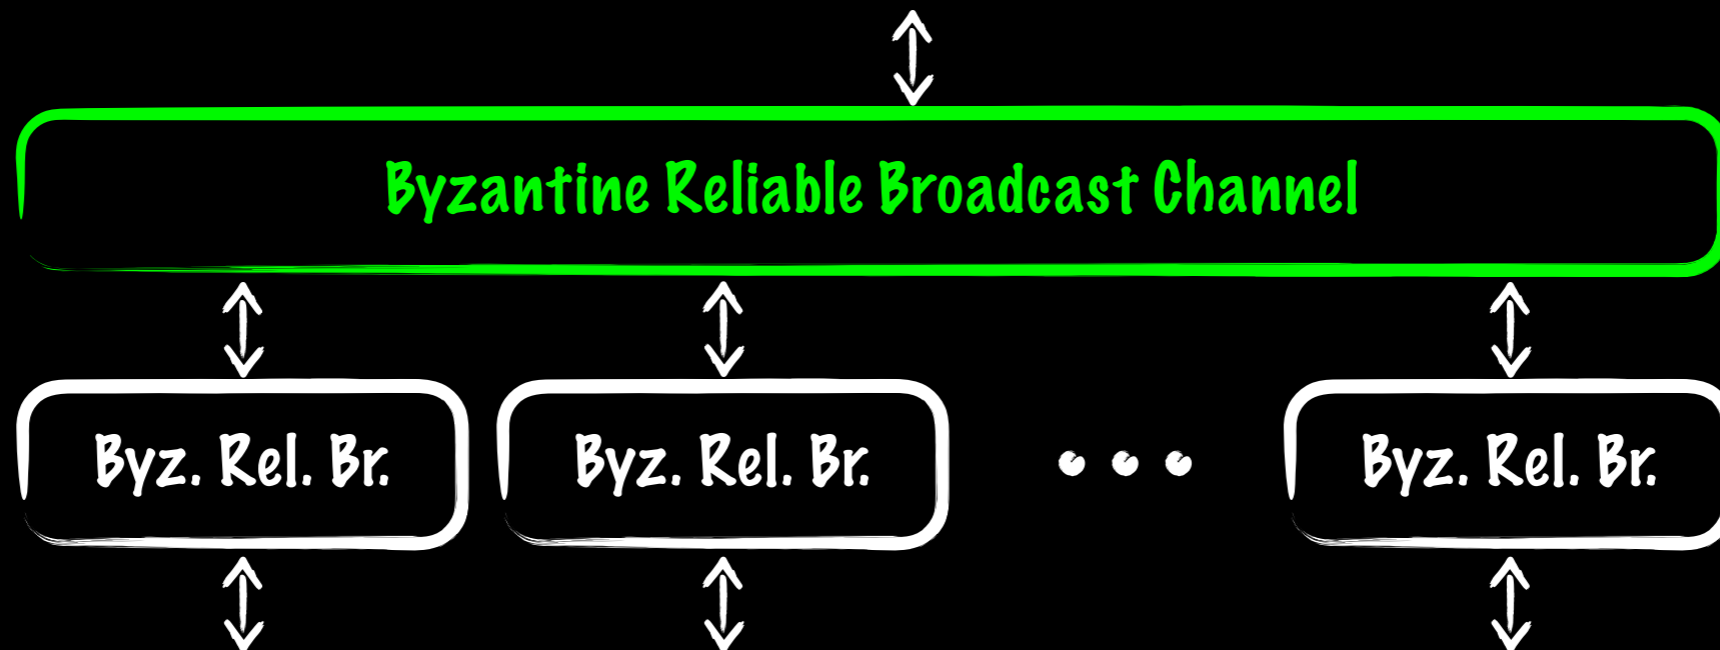
# Byzantine Reliable Broadcast Channel

**Validity, No duplication, Integrity and Consistency:** Byzantine Consistent Broadcast Channel.

**Agreement:** If some correct process delivers a message  $m$  with label  $l$  and sender  $s$ , then every correct process eventually delivers message  $m$  with label  $l$  and sender  $s$ .

# Byzantine Reliable Channel

- Create an instance of a **Byzantine Reliable Broadcast** for each process.
- On delivery of a message from sender  $s$ , create a new instance of a **Byzantine Reliable Broadcast** for process  $s$ .



LUNCH