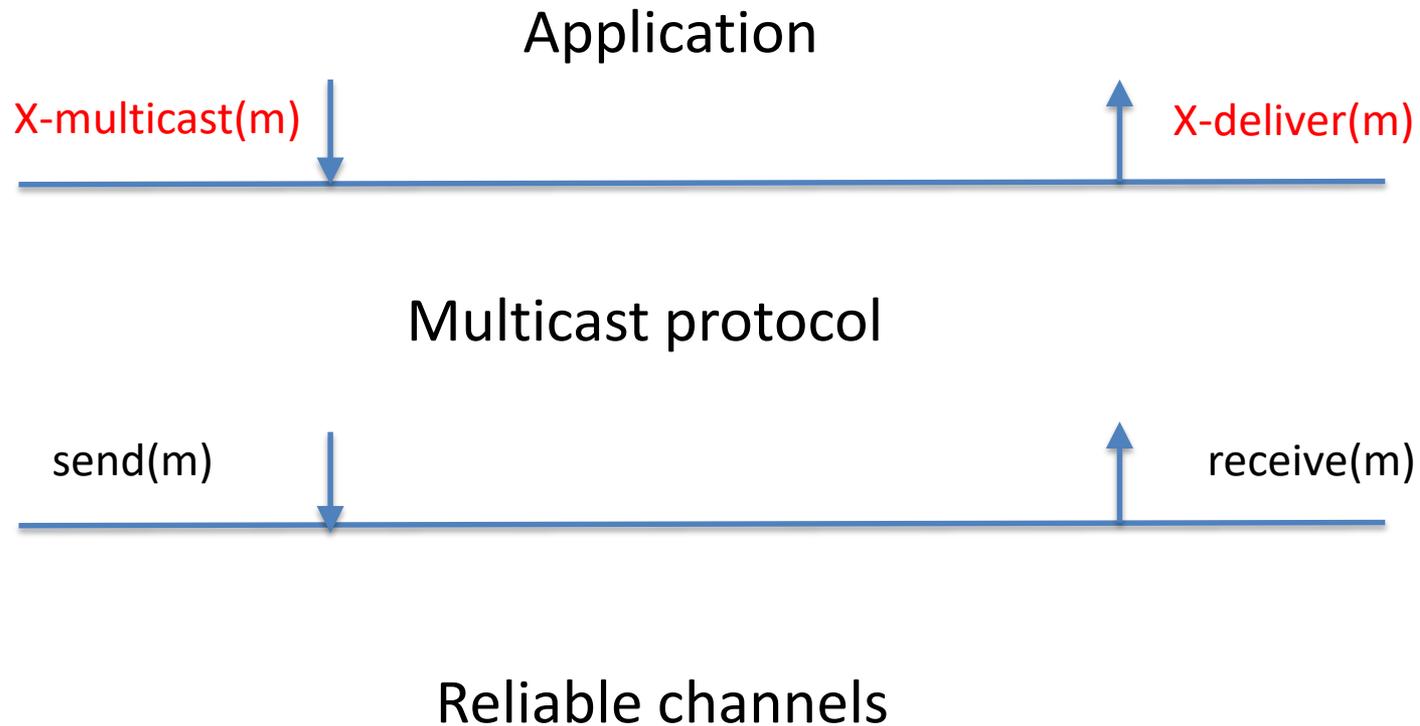


Multicast protocols (protocolos de radiado)

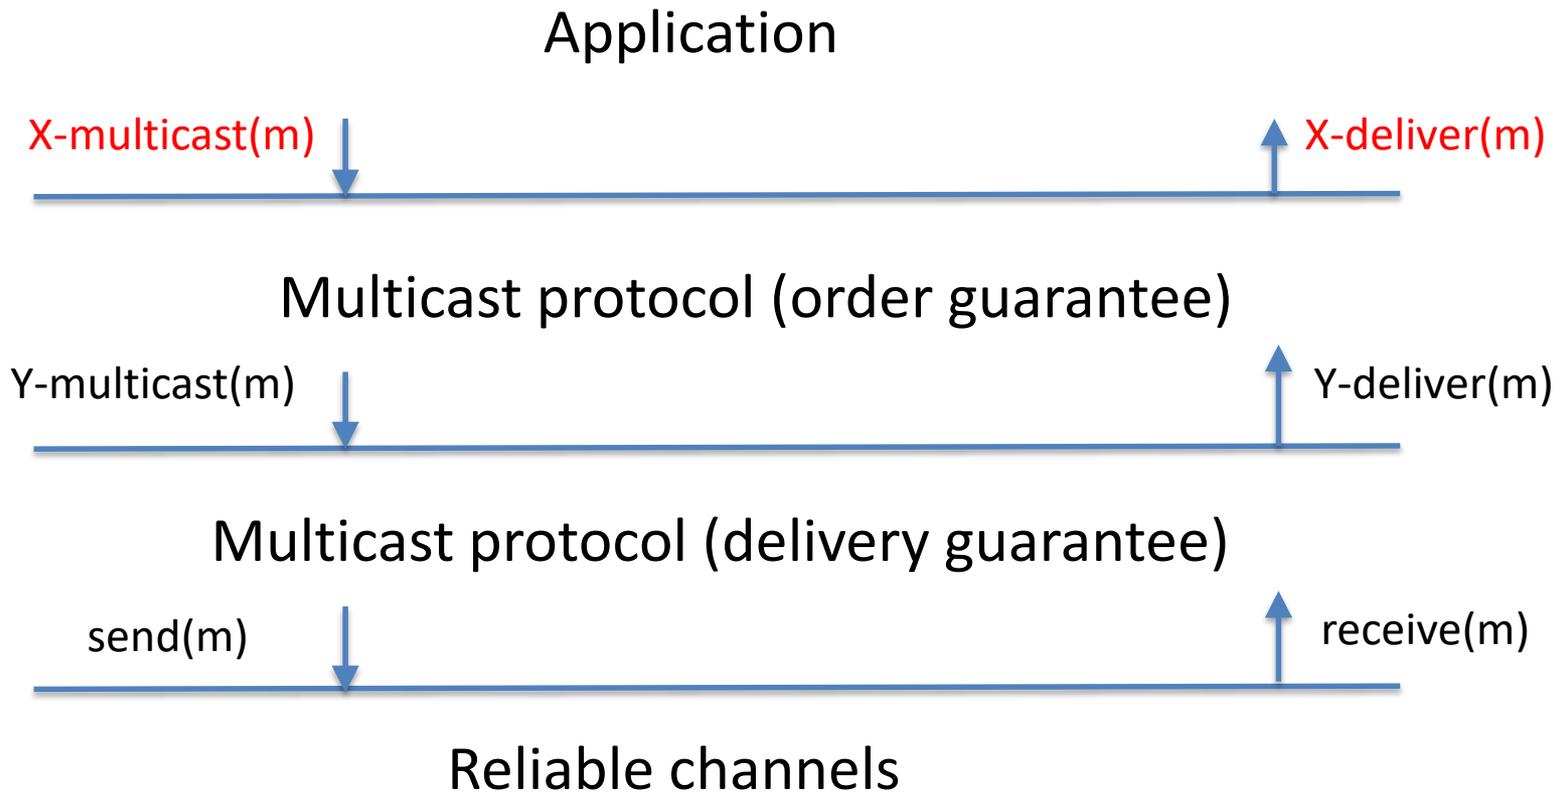
Ernesto Jiménez

Prometeo

Multicast Protocol one layer architecture



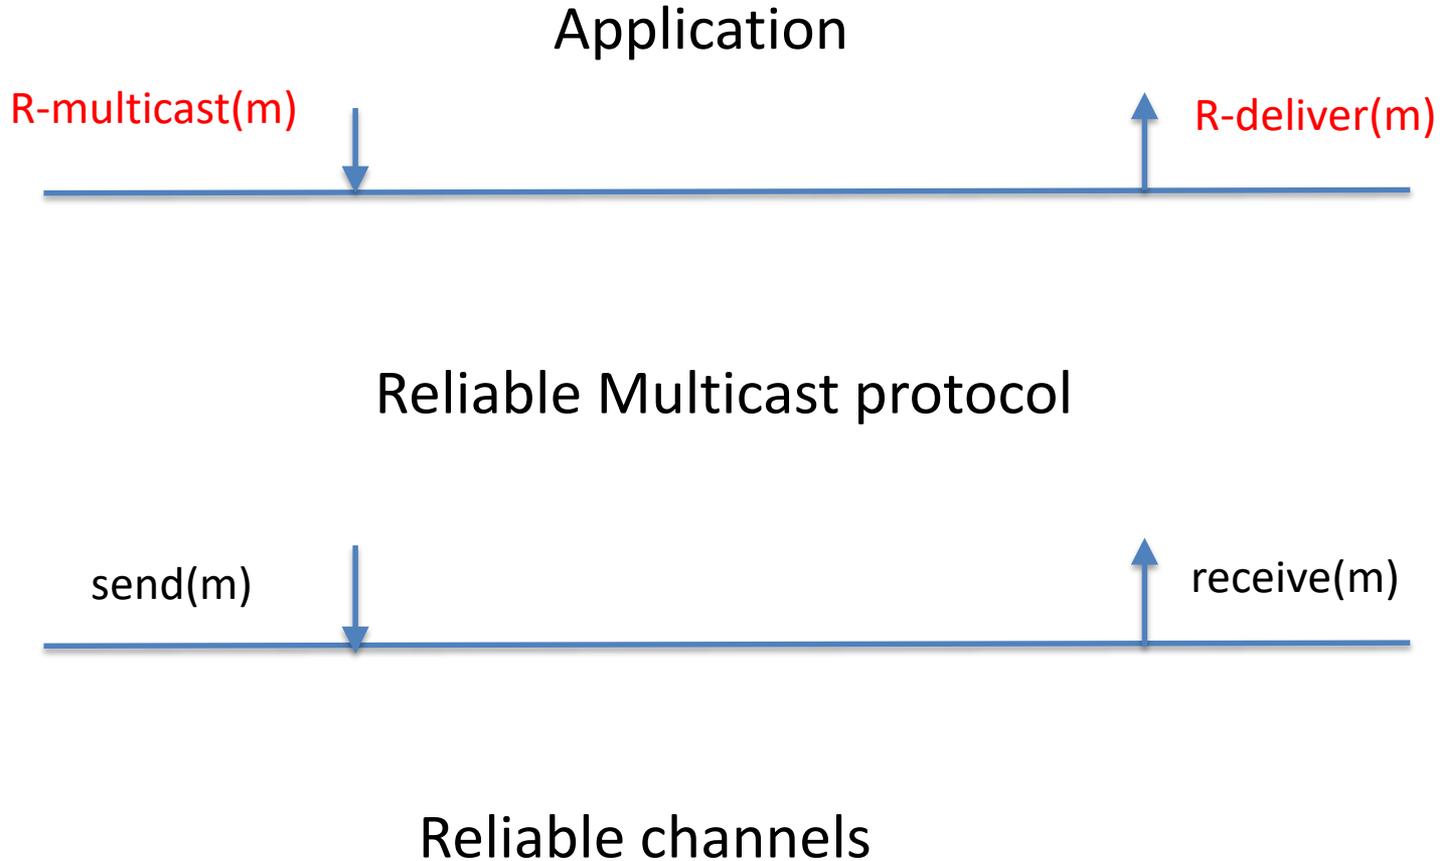
Multicast Protocol two layer architecture



Reliable Channel properties

- Validity: If a correct process $\text{send}(m)$, this message m must be eventually $\text{receive}(m)$ in all correct processes.
- Integrity: If a process $\text{receive}(m)$, then this message m has been previously $\text{send}(m)$, and a process $\text{receive}(m)$ at most once.

Reliable Multicast protocol architecture



Reliable Multicast protocol properties

- Validity: if a correct process R-multicast a message m , m must eventually R-deliver it in all correct processes.
- Integrity: if a process R-deliver(m), then this message m has been R-multicast by some process, and each process must R-deliver m at most once.
- Agreement: if a correct process R-deliver a message m , it must eventually R-deliver m in all correct processes.

Reliable multicast protocol

init:

received_i ← ∅; //set of messages received by process p_i

when R-multicast (m) is executed by p_i's application:

send (m) to all processes.

when receive (m) is executed by p_i's system:

if (m ∉ received_i) **then**

received_i ← received_i ∪ {m};

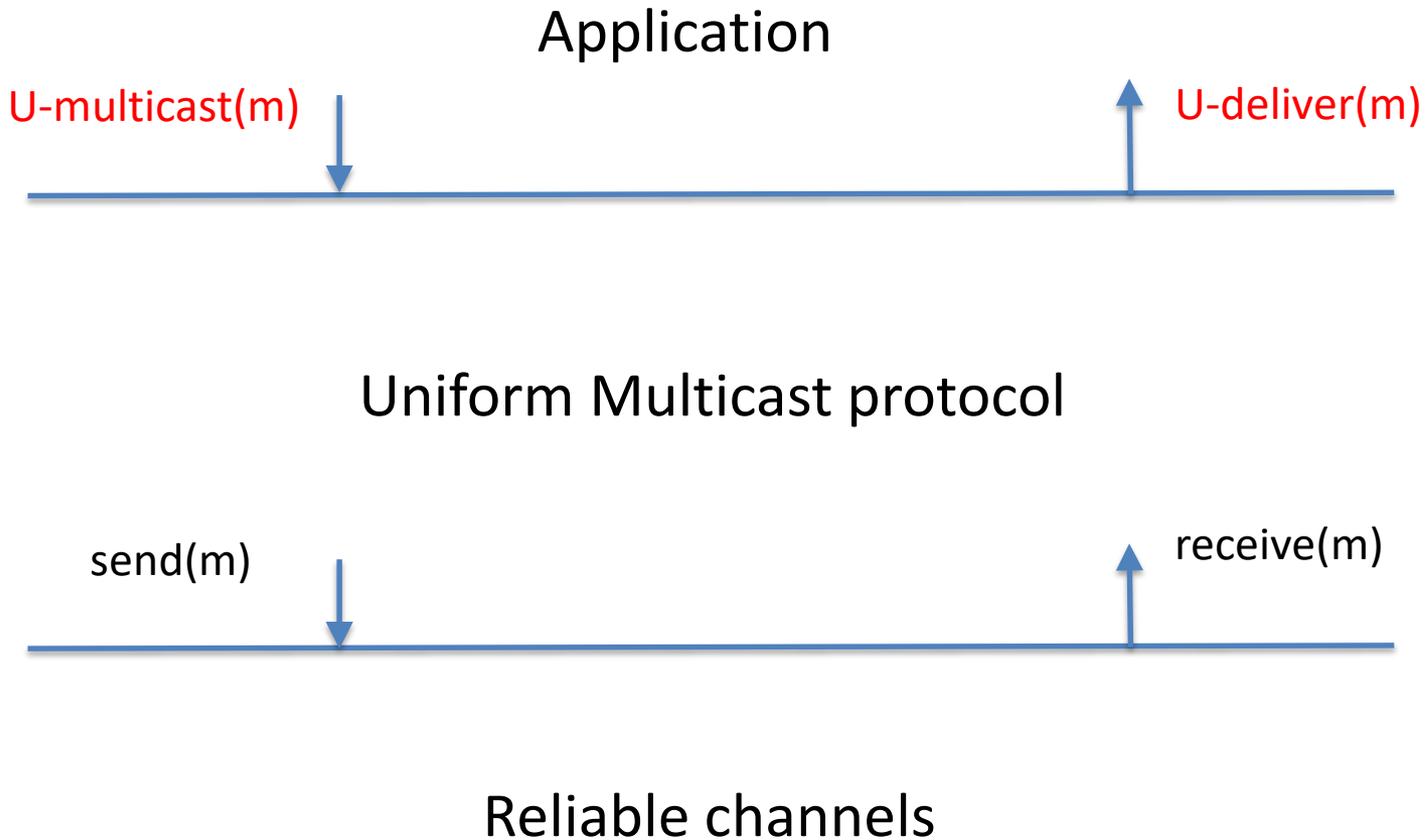
R-deliver(m);

send (m) to all processes

end if.

Process p_i pseudocode

Uniform Multicast Protocol architecture



Uniform Multicast protocol properties

- Validity: if a correct process U-multicast a message m , m must eventually U-deliver it in all correct processes.
- Integrity: if a process U-deliver(m), then this message m has been U-multicast by some process, and a process must U-deliver m at most once.
- Uniform Agreement: if a process (correct or not) U-deliver a message m , it must eventually U-deliver m in all correct processes.

Uniform multicast protocol

init:

$\text{next_seqnum}_i \leftarrow 1$; $\text{received}_i \leftarrow \emptyset$; set of messages received by process p_i

when U-multicast (m) is executed by p_i 's application:

$\text{m.sender} = i$; $\text{m.seqnum} \leftarrow \text{next_seqnum}_i$;

send (m) to all processes; $\text{next_seqnum}_i \leftarrow \text{next_seqnum}_i + 1$;

when receive (m) is executed by p_i 's system:

if ($m \notin \text{received}_i$) **then**

$\text{received}_i \leftarrow \text{received}_i \cup \{m\}$; $\text{pending}_i \leftarrow \text{pending}_i \cup \{m\}$;

send(ACK, i) to m.sender;

send (m) to all processes

end if.

when receive (ACK, id_j) is executed by p_i 's system:

// set of ACK senders

let $\text{ACK-IDs}(m) = \{\text{id}(p) : (\text{ACK}, \text{id}_j) \text{ has been received and } (\text{ACK}, \text{id}_j).\text{sender} = p \text{ and } \text{m.sender} = \text{id}_j\}$;

while ($\exists m \in (\text{pending}_i) : \text{m.sender} = \text{id}_j$) **and**

$(|\text{ACK-IDs}(m) \cup \text{m.sender} \cup \{i\}| > \lceil \lceil n \rceil / 2)$ **do** // A majority of processes received m

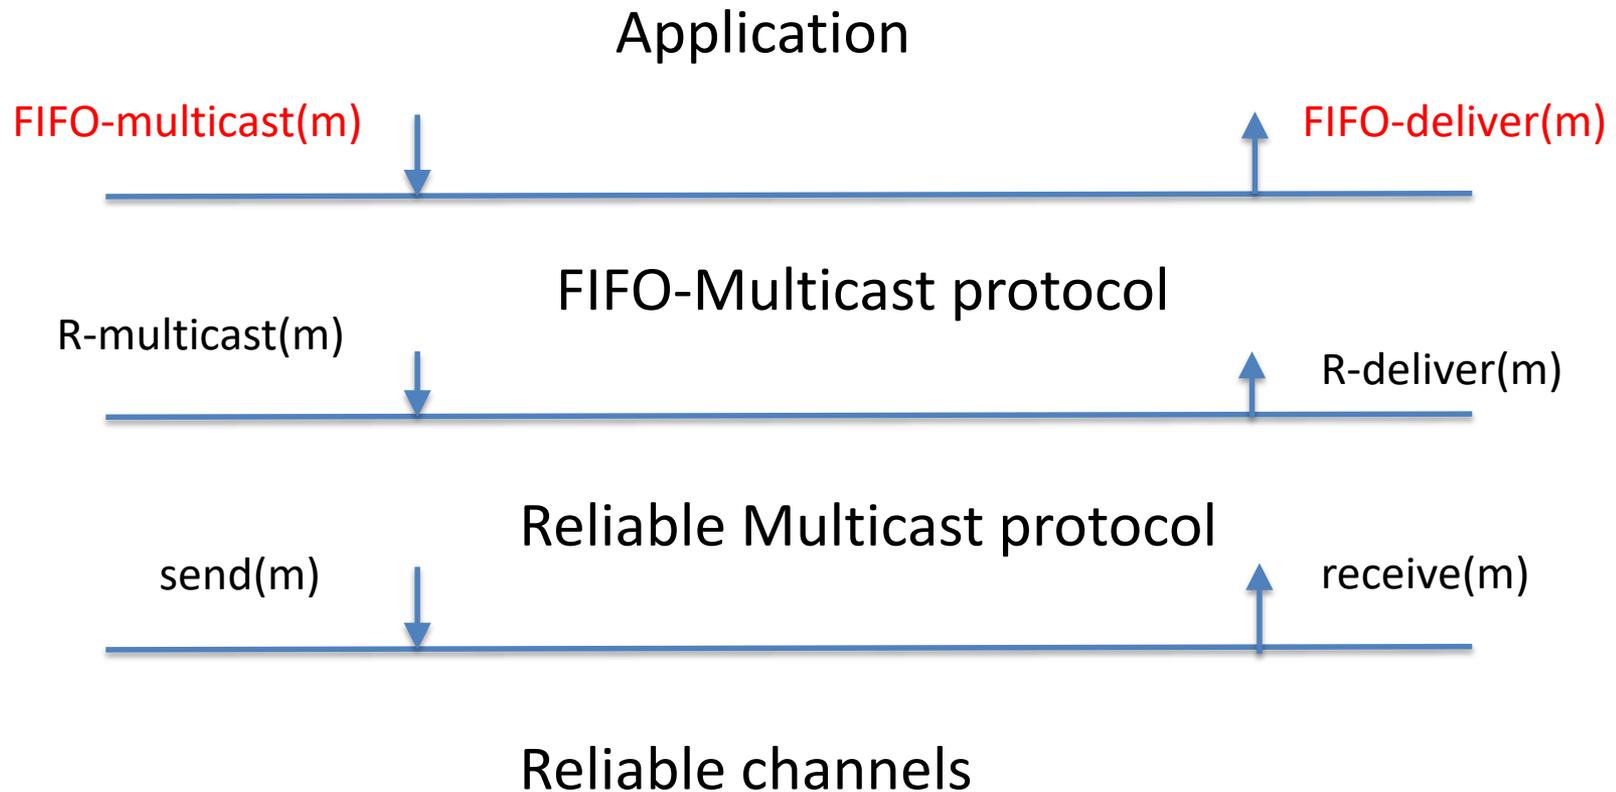
U-deliver(m);

$\text{pending}_i \leftarrow \text{pending}_i \setminus \{m\}$

end while.

Process p_i pseudocode

FIFO Reliable Multicast Protocol architecture



FIFO-Reliable Multicast protocol properties

- Validity: if a correct process FIFO-multicast a message m , m must be FIFO-deliver in all correct processes.
- Integrity: if a process FIFO-deliver(m), then m has been FIFO-multicast by some process, and a process must FIFO-deliver m at most once.
- Agreement: if a correct process FIFO-deliver a message m , it must be FIFO-deliver in all correct processes.
- FIFO message delivery. If a process p_i FIFO-multicast(m) and then FIFO-multicast(m'), then no process FIFO-deliver(m') unless it has previously FIFO-deliver(m).

FIFO reliable multicast protocol

init:

```
seqnumi ← 1; //sequence number of process pi  
∀k, next_seqnumi[k] ← 1; //next message to delivered by process pi  
msg_seti ← ∅; //set of messages waiting to be delivered by process pi
```

when FIFO-multicast (m) is executed by p_i's application:

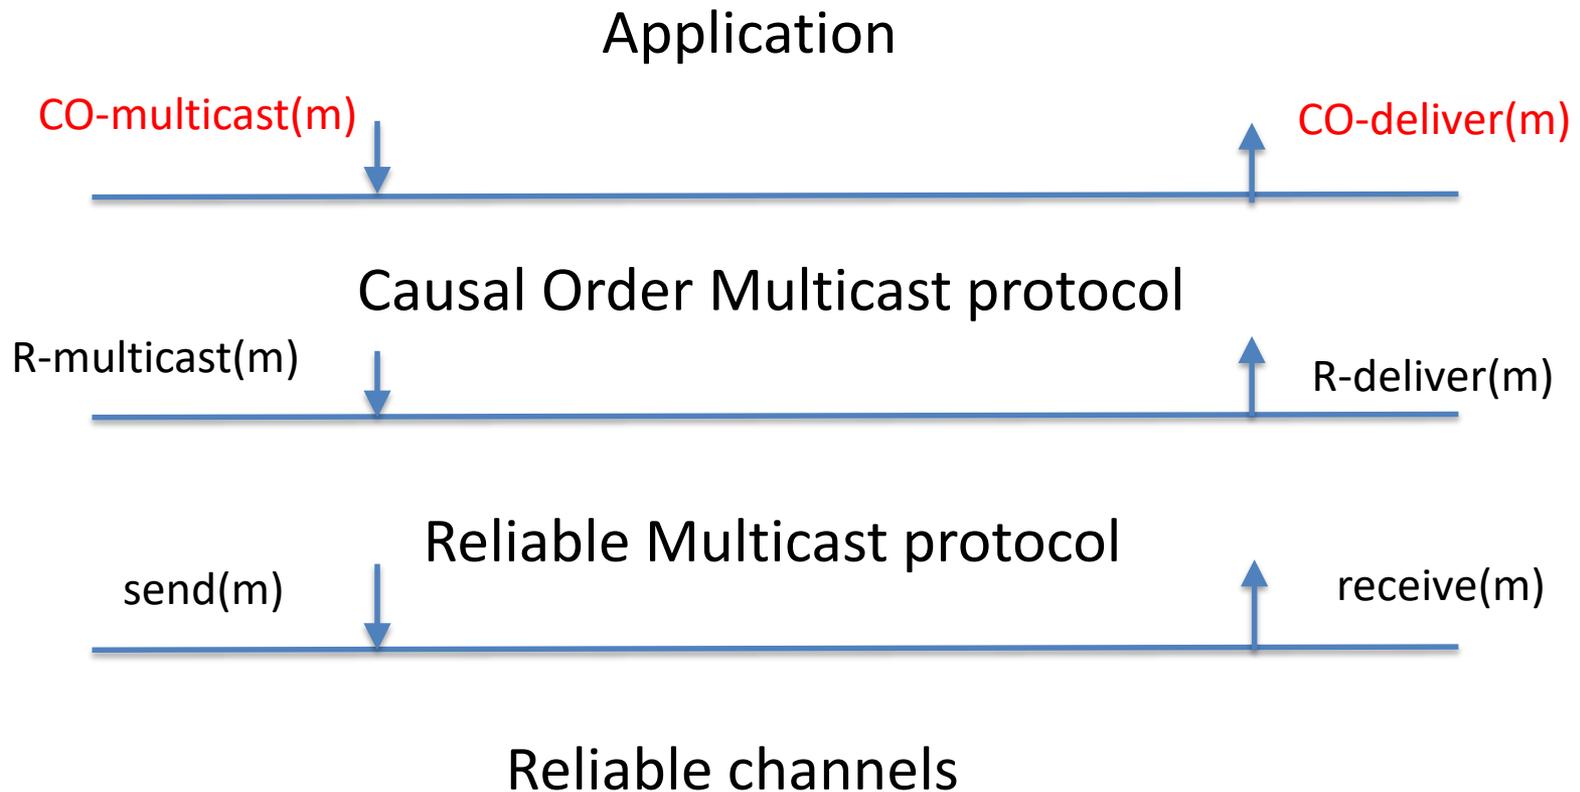
```
m.sender ← i;  
m.seqnum ← seqnumi;  
R-multicast (m) ;  
seqnumi ← seqnumi + 1.
```

when R-deliver (m) is executed by p_i's system:

```
let j = m.sender;  
msg_seti ← msg_seti ∪ {m};  
while (∃m' ∈ msg_seti : (m'.sender = j) and (next_seqnumi[j] = m'.seqnum)) do  
    FIFO-deliver(m');  
    msg_seti ← msg_seti \ {m'};  
    next_seqnumi[j] ← next_seqnumi[j] + 1  
end while.
```

Process p_i pseudocode

Causal Order Reliable Protocol architecture



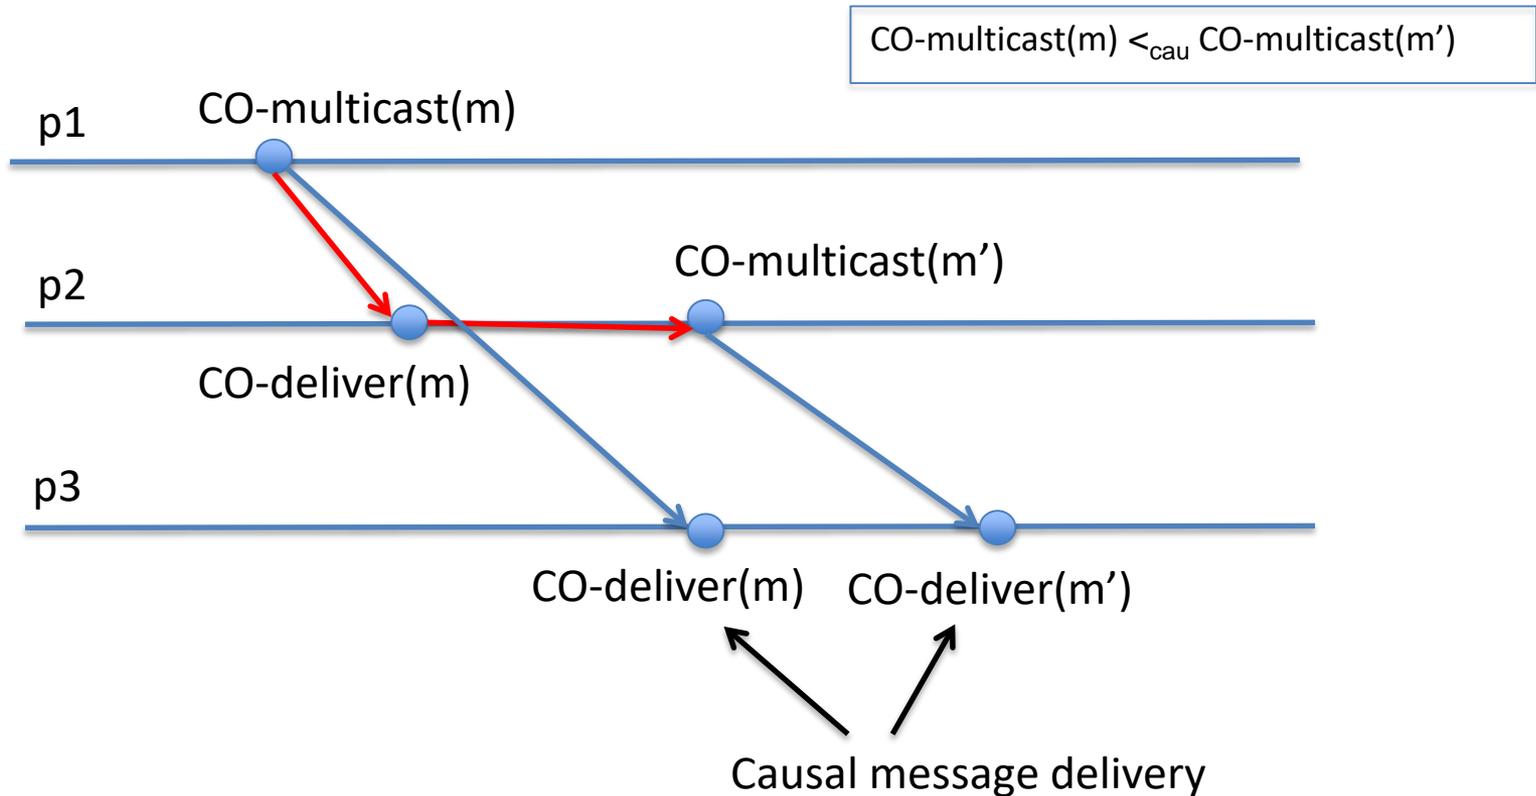
Causal Reliable Multicast protocol properties

- Validity: if a correct process CO-multicast a message m , m must be delivered in all correct processes.
- Integrity: if a process CO-deliver(m), then this message m has been CO-multicast by some process, and a process must CO-deliver m at most once.
- Agreement: If a correct process CO-deliver a message m , this message m must be CO-deliver in all correct processes.
- Causal message delivery ($<_{cau}$): if CO-multicast(m) causally precedes CO-multicast(m'), then no process CO-deliver(m') unless it has previously CO-deliver(m).

Causal Reliable Multicast protocol properties (cont.)

- Causal Order ($\text{CO-multicast}(m) <_{\text{cau}} \text{CO-multicast}(m')$): if some of the following two cases happens:
 1. A process p_k invokes $\text{CO-multicast}(m)$ before this same process p_k invokes $\text{CO-multicast}(m')$.
 2. A process p_k invokes $\text{CO-multicast}(m)$, and a process p_s invokes $\text{CO-deliver}(m)$ before this same process p_s invokes $\text{CO-multicast}(m')$.
- Causal message delivery ($<_{\text{cau}}$): every process must preserve causal order $\text{CO-multicast}(m) <_{\text{cau}} \text{CO-multicast}(m')$, for all messages m and m' , when it executes CO-deliver .

Causal Reliable Multicast protocol properties (cont.)



Causal Order reliable multicast protocol

init:

$\forall k, \text{next_seqnum}_i[k] \leftarrow 1;$ // seq. number of next messages to be applied by p_i
 $\text{msg_set}_i \leftarrow \emptyset;$ // set of messages waiting to be delivered

when CO-multicast (m) is executed by p_i 's application:

$m.\text{sender} \leftarrow i;$ $m.\text{seqnum} \leftarrow \text{next_seqnum}_i;$

R-multicast (m);

CO-deliver (m); // immediate delivery of p_i 's message

$\text{next_seqnum}_i[i] \leftarrow \text{next_seqnum}_i[i] + 1.$

when R-deliver (m) is executed by p_i 's system:

let $j = m.\text{sender};$

if ($i \neq j$) then

$\text{msg_set}_i \leftarrow \text{msg_set}_i \cup \{m\};$

while ($\exists m' \in \text{msg_set}_i : \text{CausalCondition}(m')$) **do**

let $j = m'.\text{sender};$

CO-deliver(m');

$\text{next_seqnum}_i[j] \leftarrow m'.\text{seqnum}[j] + 1;$

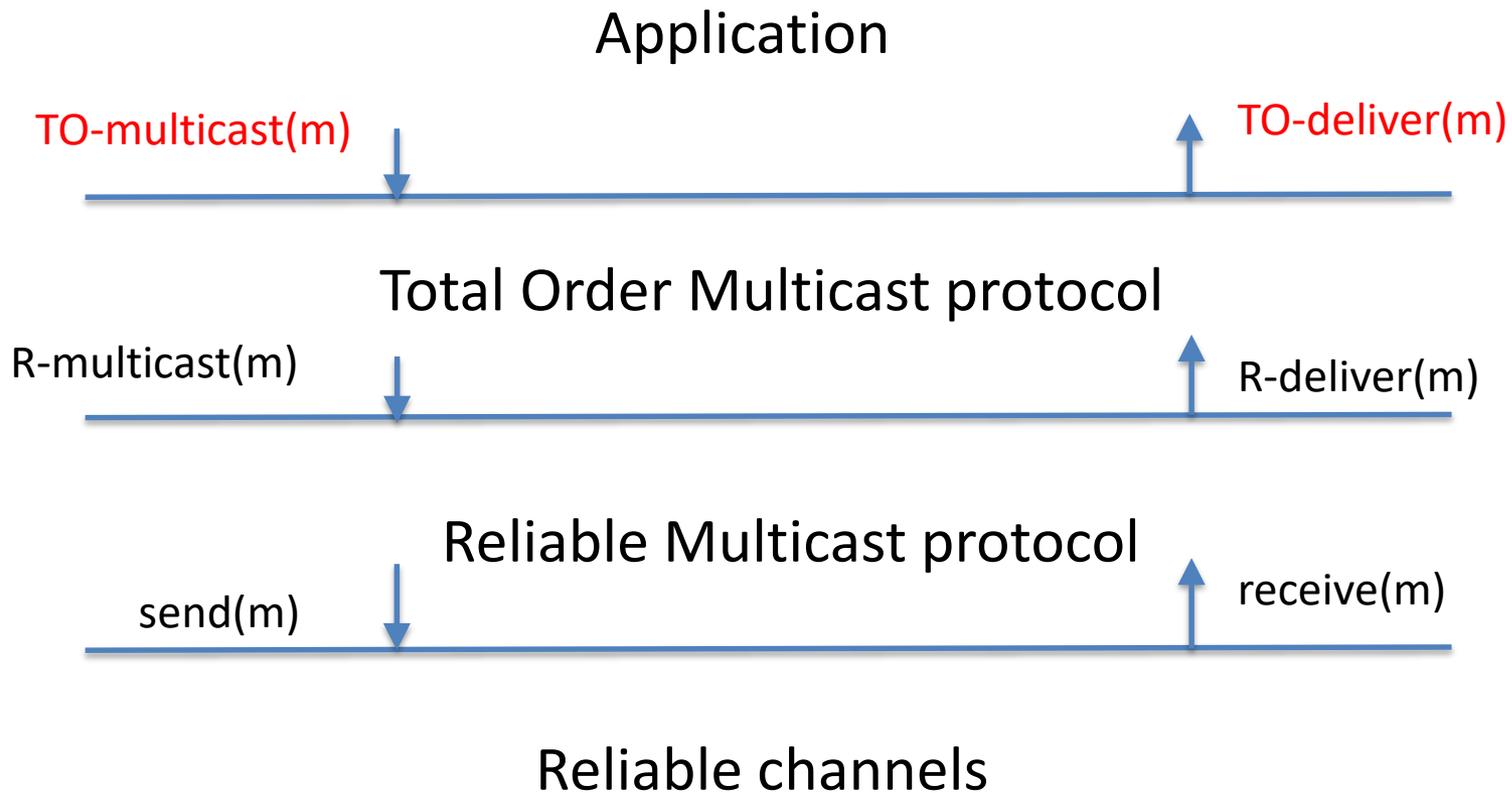
$\text{msg_set}_i \leftarrow \text{msg_set}_i \setminus \{m'\}$

end while

end if.

$\text{CausalCondition}(m') \equiv (\forall k : \text{next_seqnum}_i[k] \geq m'.\text{seqnum}[k])$

Total order reliable Protocol architecture



Total Order Reliable Multicast protocol properties

- Validity if a correct process TO-multicast a message m , this message m must be delivered in all correct processes.
- Integrity: if a process TO-deliver(m), then m has been TO-multicast by some process, and a process must TO-deliver m at most once.
- Agreement: if a correct process TO-deliver a message m , this message m must be TO-deliver in all correct processes.
- Total order message deliver. If a process TO-deliver(m) and then TO-deliver(m'), then no process TO-deliver(m') unless it has previously TO-deliver(m).

Total Order reliable multicast

init:

$\text{next_seqnum}_i \leftarrow 1$; //next message to delivered
 $\text{msg_set}_i \leftarrow \emptyset$; // set of messages waiting to be delivered

when TO-multicast (m) is executed by p_i 's appl.:
send (m) to p_s ; // sequencer

when R-deliver (m) is executed by p_i 's system:
 $\text{msg_set}_i \leftarrow \text{msg_set}_i \cup \{m\}$;
while ($\exists m' \in \text{msg_set}_i$:
 $\text{next_seqnum}_i = m'.\text{seqnum}$) **do**
 TO-deliver(m');
 $\text{msg_set}_i \leftarrow \text{msg_set}_i \setminus \{m'\}$;
 $\text{next_seqnum}_i \leftarrow \text{next_seqnum}_i + 1$;
end while.

Process p_i pseudocode

init:

$\text{seq_glob}_s \leftarrow 1$; // global sequence number

when receive (m) is executed by p_i 's system:
 $m.\text{seqnum} \leftarrow \text{seq_glob}_s$;
R-multicast (m);
 $\text{seq_glob}_s \leftarrow \text{seq_glob}_s + 1$;

Process sequencer (p_s) pseudocode

NOTE: if the process sequencer crashes, then alive processes must detect it reliably and elect a new sequencer.

FIFO+Total Order reliable multicast

init:

```
next_seqnumi ← 1; //next message to delivered  
msg_seti ← ∅; // set of messages waiting to be  
delivered
```

**when FIFO+TO-multicast (m) is executed by p_i's
appl.:**

```
m.sender ← i;  
m.seqnum ← next_seqnumi;  
send (m) to ps; // sequencer  
next_seqnumi ← next_seqnumi + 1;
```

when R-deliver (m) is executed by p_i's system:

```
msg_seti ← msg_seti ∪ {m};  
while (∃ m' ∈ msg_seti :  
    next_seqnumi = m'.seqnum) do  
    FIFO+TO-deliver(m');  
    msg_seti ← msg_seti \ {m'};  
    next_seqnumi ← next_seqnumi + 1;  
end while.
```

init:

```
seq_globs ← 1; // global sequence number  
∀j, next_FIFOs[j] ← 1;  
msg_sets ← ∅;
```

when receive (m) is executed by p_i's system:

```
msg_sets ← msg_sets ∪ {m};  
while ((∃ m ∈ msg_sets) and (∃ k = m.sender) :  
    next_FIFOs[k] = m.seqnum) do  
    m.seqnum ← seq_globs;  
    R-multicast (m);  
    seq_globs ← seq_globs + 1;  
    msg_seti ← msg_seti \ {m};  
    next_FIFOs[k] ← next_FIFOs[k] + 1;  
end while.
```

Process sequencer (p_s) pseudocode

NOTE: if the process sequencer crashes, then alive processes must detect it reliably and elect a new sequencer.

