

Inside Virtual Synchrony

Prepared by:

Steven Dake

7/12/05

Definitions

- Group Messaging
 - Sending messages from 1 sender to many receivers.
- Processor
 - The entity responsible for executing group messaging and membership protocols.
- Regular Configuration Change
 - An event containing a unique view identifier and list of processors contained within the configuration. Denoted as RCn in examples.
- Transitional Configuratoin
 - An event containing a unique view identifier and list of processors transitioning from the old regular configuration to the new regular configuration. Denoted as TCn in exmaples.

Virtual Synchrony Property #1 – self delivery

- self delivery – A message sent by a processor is delivered to that processor.

Example:

Processor P1 sends message M1

M1 is self-delivered to P1

Virtual Synchrony Property #2 – AGREED ordering

- agreed ordering – all processors agree upon delivery order of messages.

Example:

P1: M1 M2 M3 M4

P2: M1 M2 M3 M4

P3: M1 M2 M3 M4

Can't happen:

P1: M1 M2 M3 M4

P2: M1 M2 M4 M3

P3: M1 M2 M3 M4

Virtual Synchrony Property #3 – SAFE ordering

- SAFE ordering – agreed ordering extended such that a message may not be delivered until every processor within the configuration has a copy.

P1 sends M1 M2 M3.

P2 and P3 only receive M1 and M2.

P1 P2 P3 deliver M1 M2.

P2 and P3 recover M3.

M3 may now be delivered in safe order.

Virtual Synchrony Property #4 – Virtual Synchrony

- virtual synchrony – messages are delivered in agreed order and configuration changes are delivered in agreed order relative to messages.

Example:

P1: M1 M2 M3 M4 CT1 CR1

P2: M1 M2 M3 M4 CT1 CR1

Can't happen:

P1: M1 M2 M3 M4 CT1 CR1

P2: M1 M2 M3 CT1 CR1

The Ring Protocol

- Sequence Number
- Retransmit List
- flow control count
- group arut

ORF Token

Processor #1

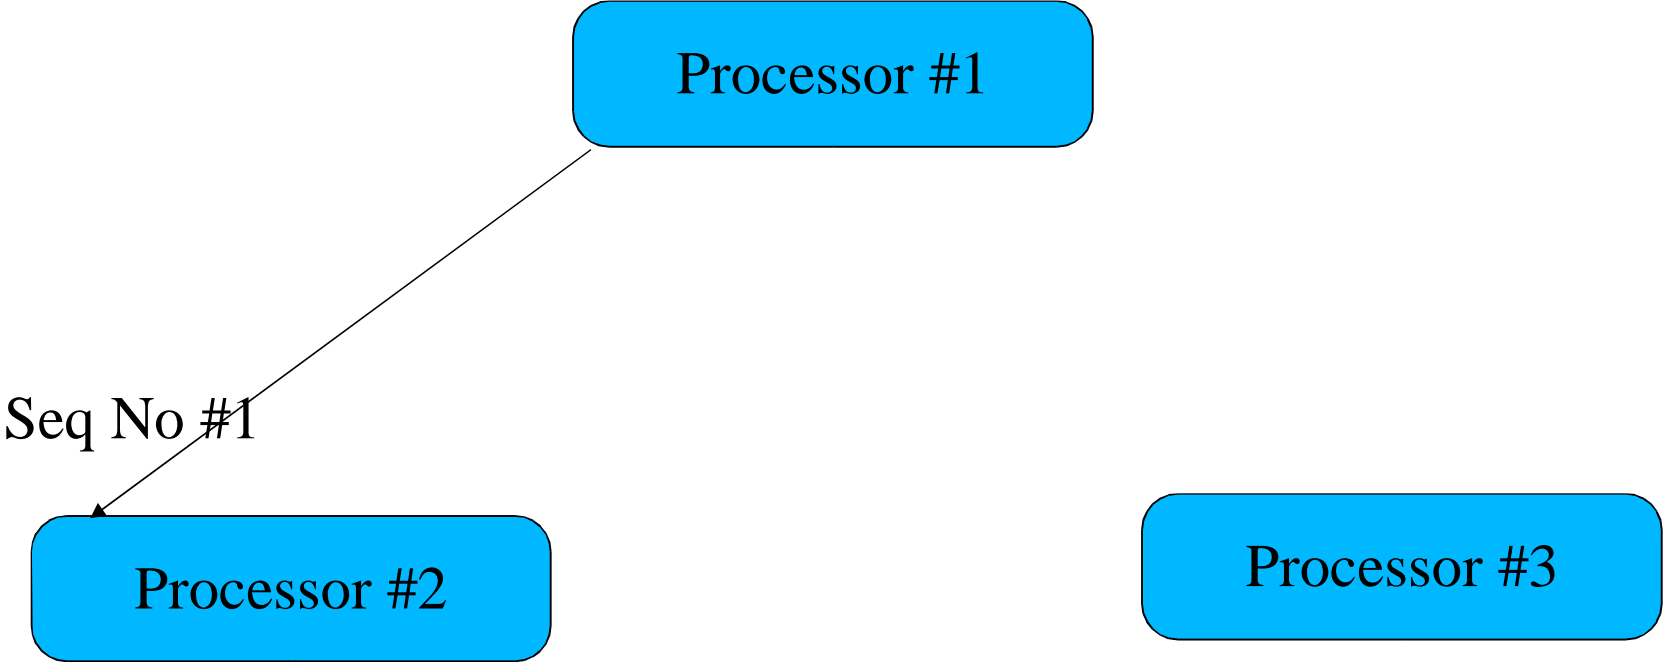
Processor #2

Processor #3

The Ring Protocol

- Sequence Number
- Retransmit List
- flow control count
- group arut

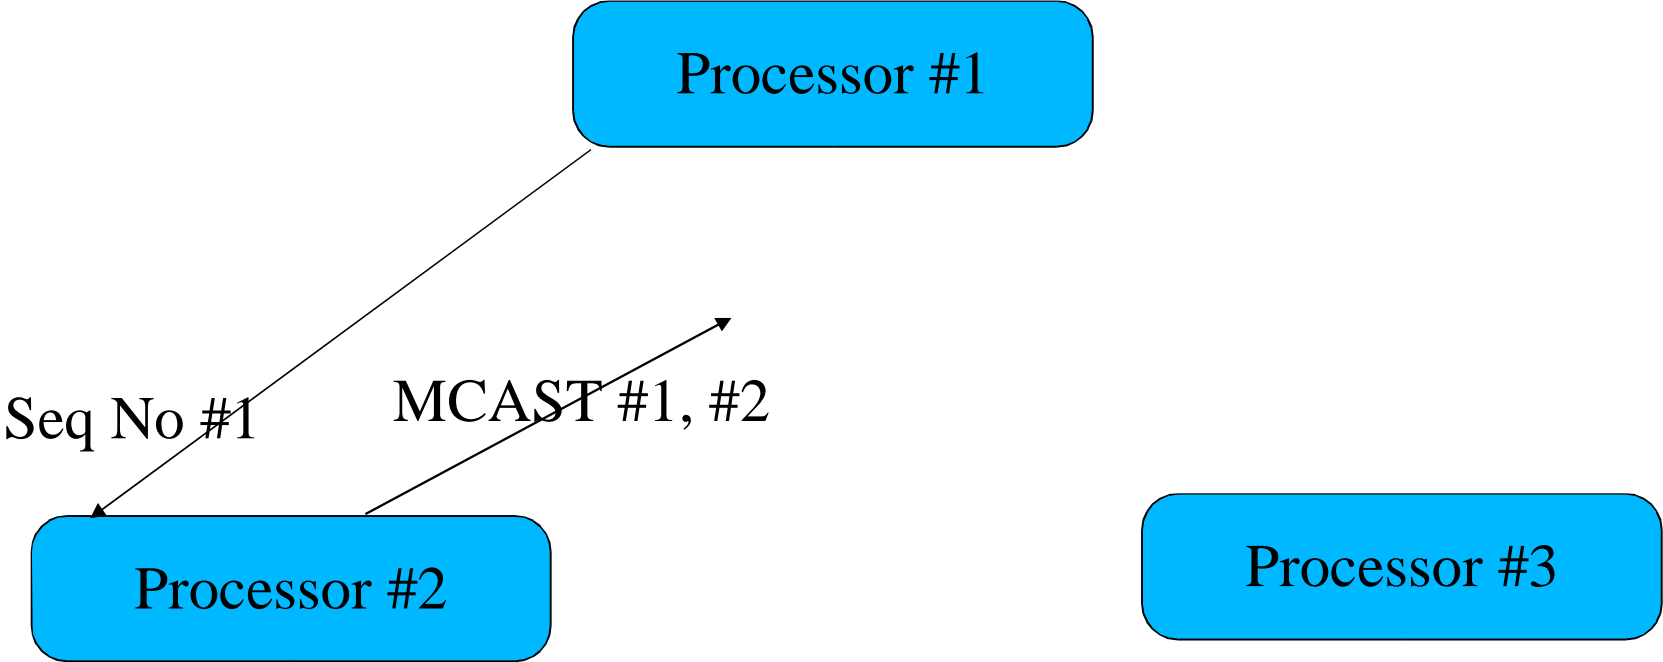
ORF Token



The Ring Protocol

- Sequence Number
- Retransmit List
- flow control count
- group arut

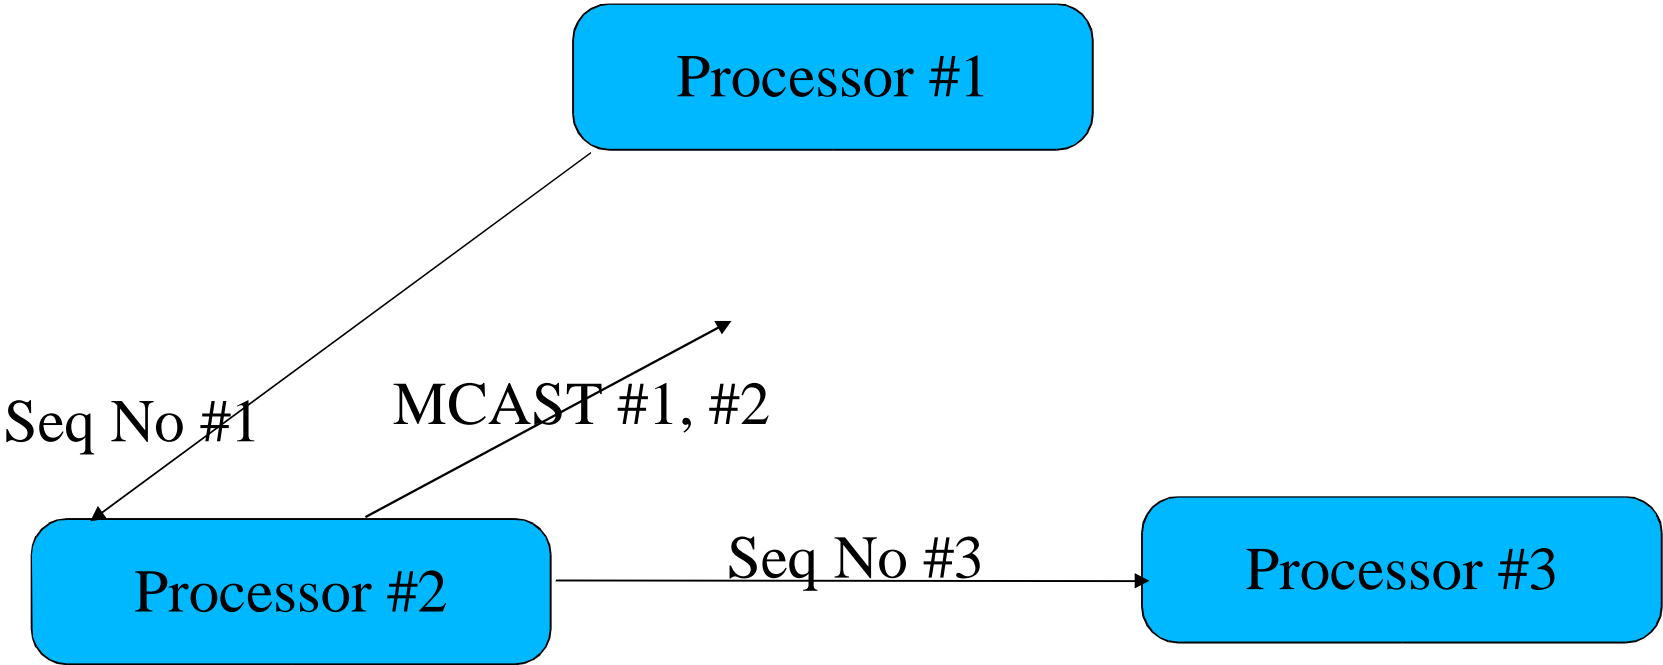
ORF Token



The Ring Protocol

- Sequence Number
- Retransmit List
- flow control count
- group arut

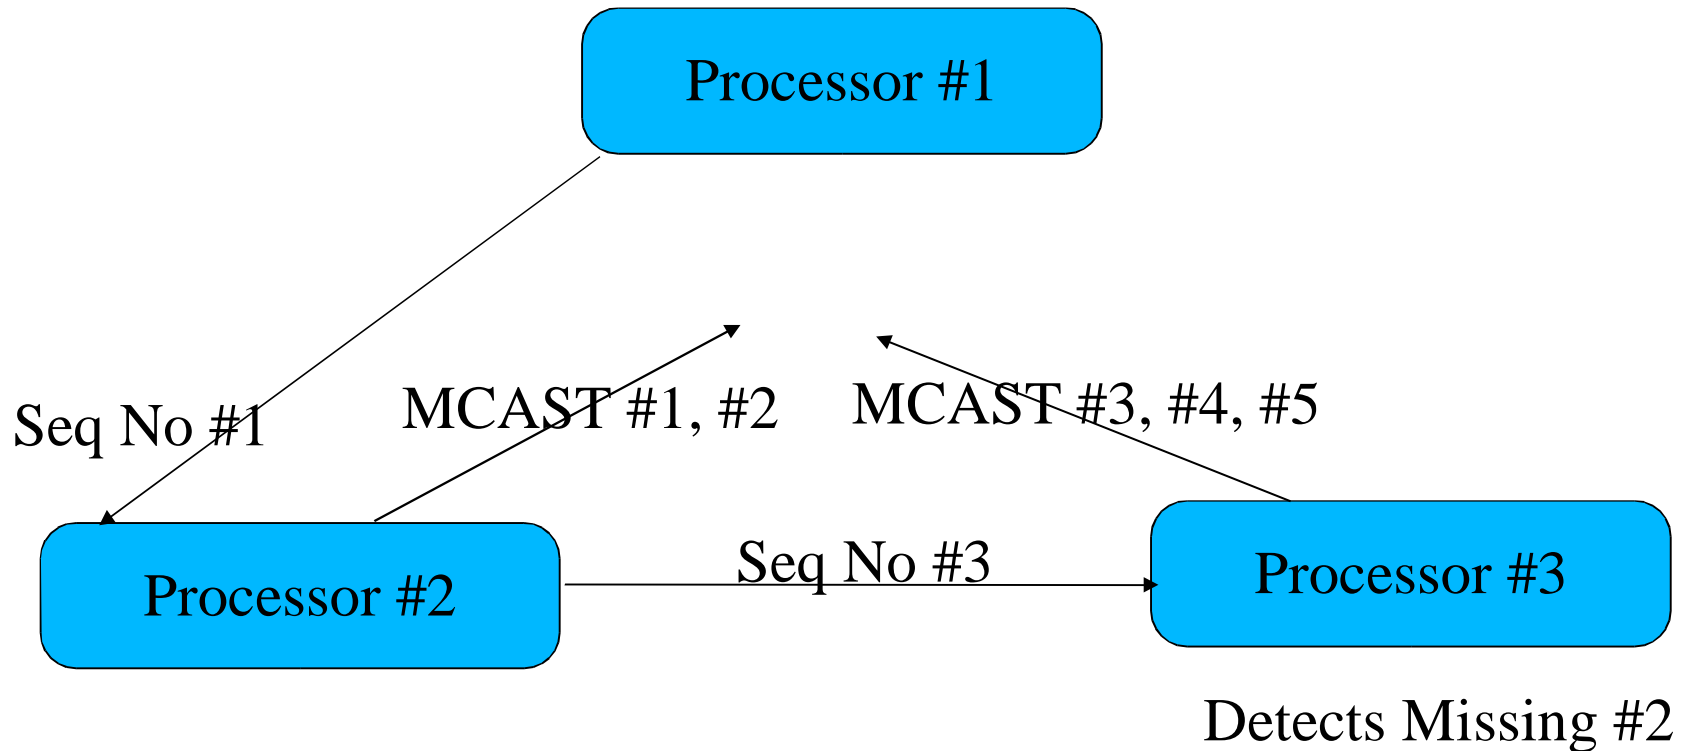
ORF Token



- Sequence Number
- Retransmit List
- flow control count
- group arut

ORF Token

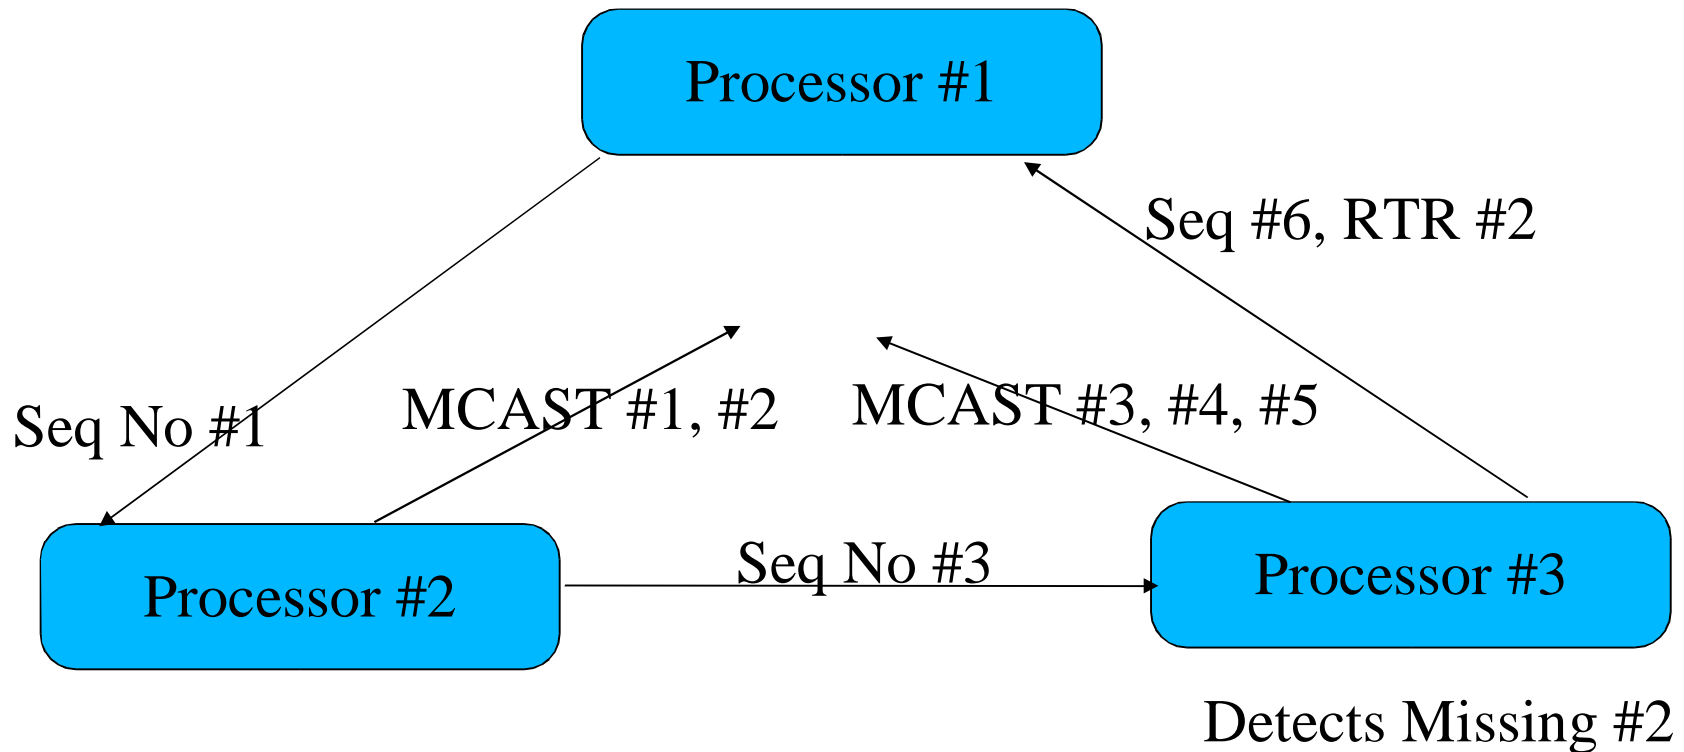
The Ring Protocol



The Ring Protocol

- Sequence Number
- Retransmit List (RTR)
- flow control count
- group arut

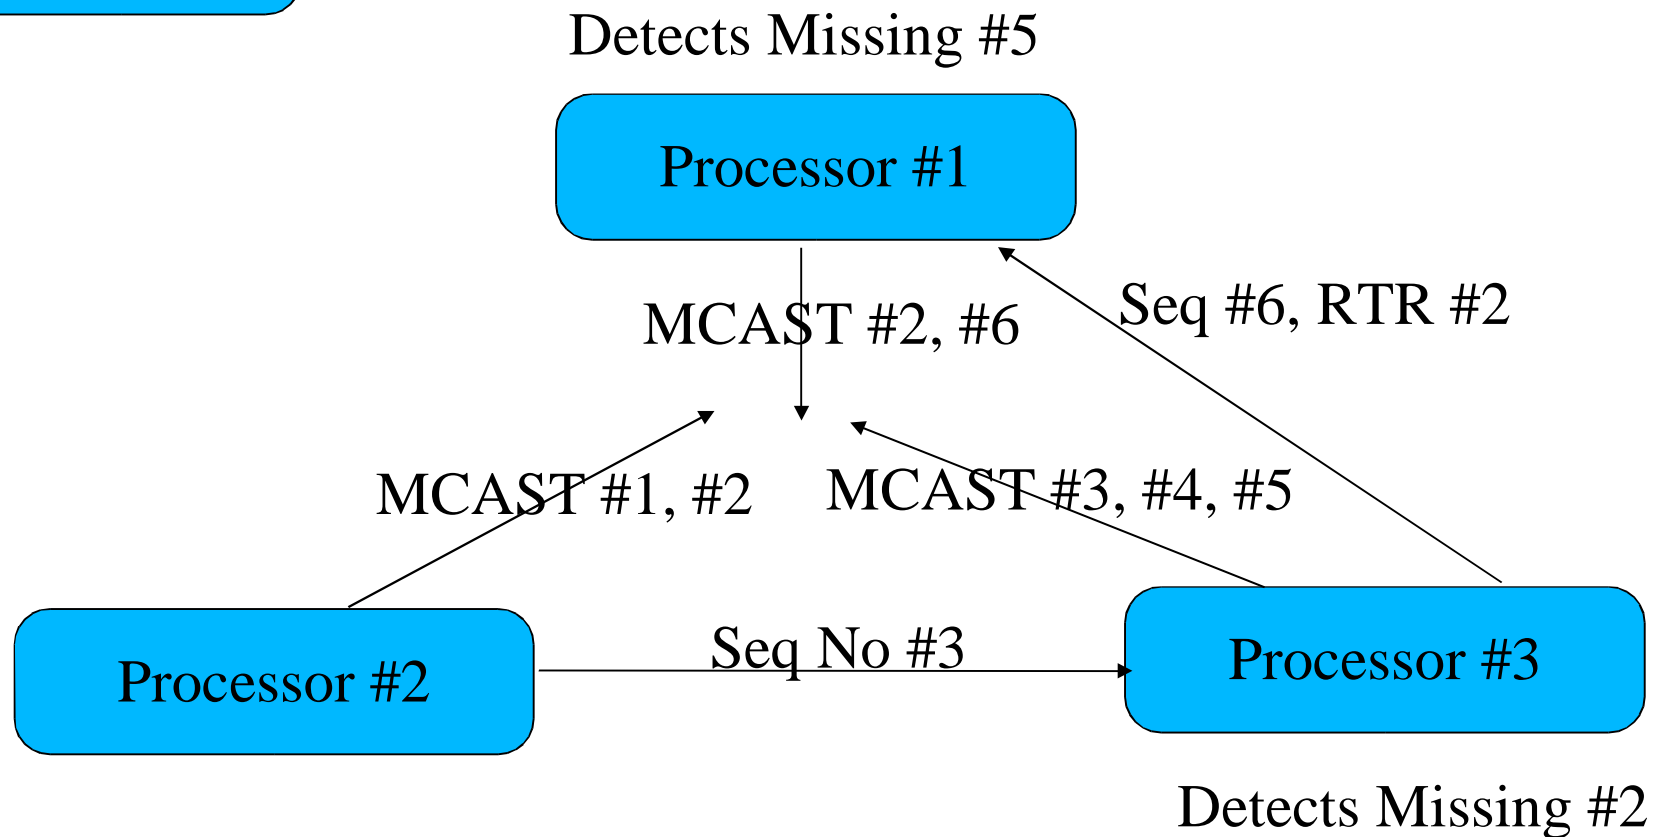
ORF Token



The Ring Protocol

- Sequence Number
- Retransmit List (RTR)
- flow control count
- group arut

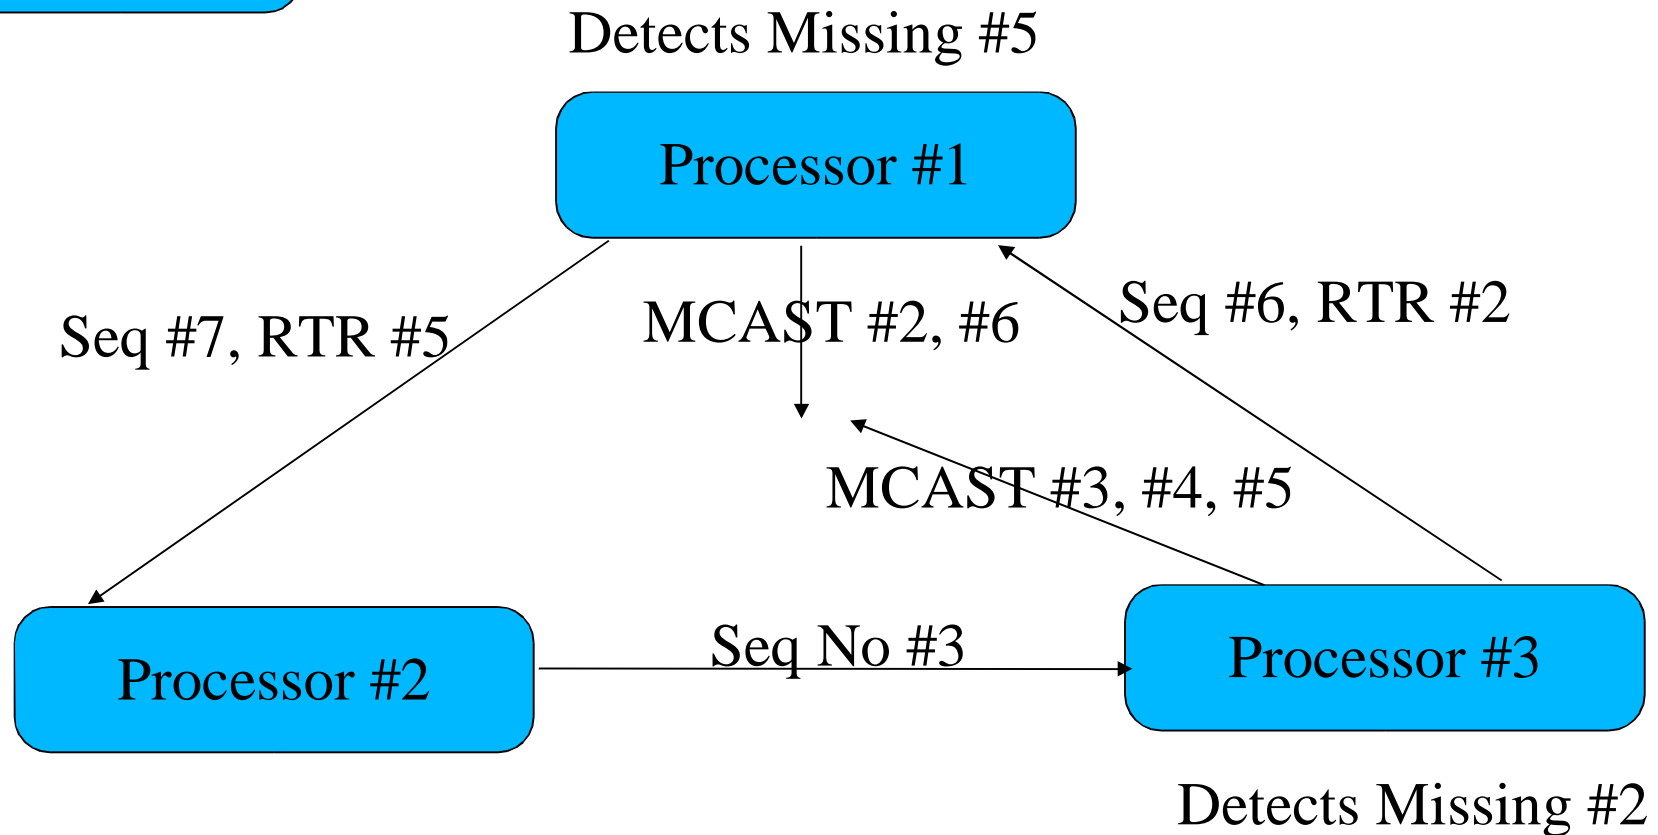
ORF Token



The Ring Protocol

- Sequence Number
- Retransmit List (RTR)
- flow control count
- group arut

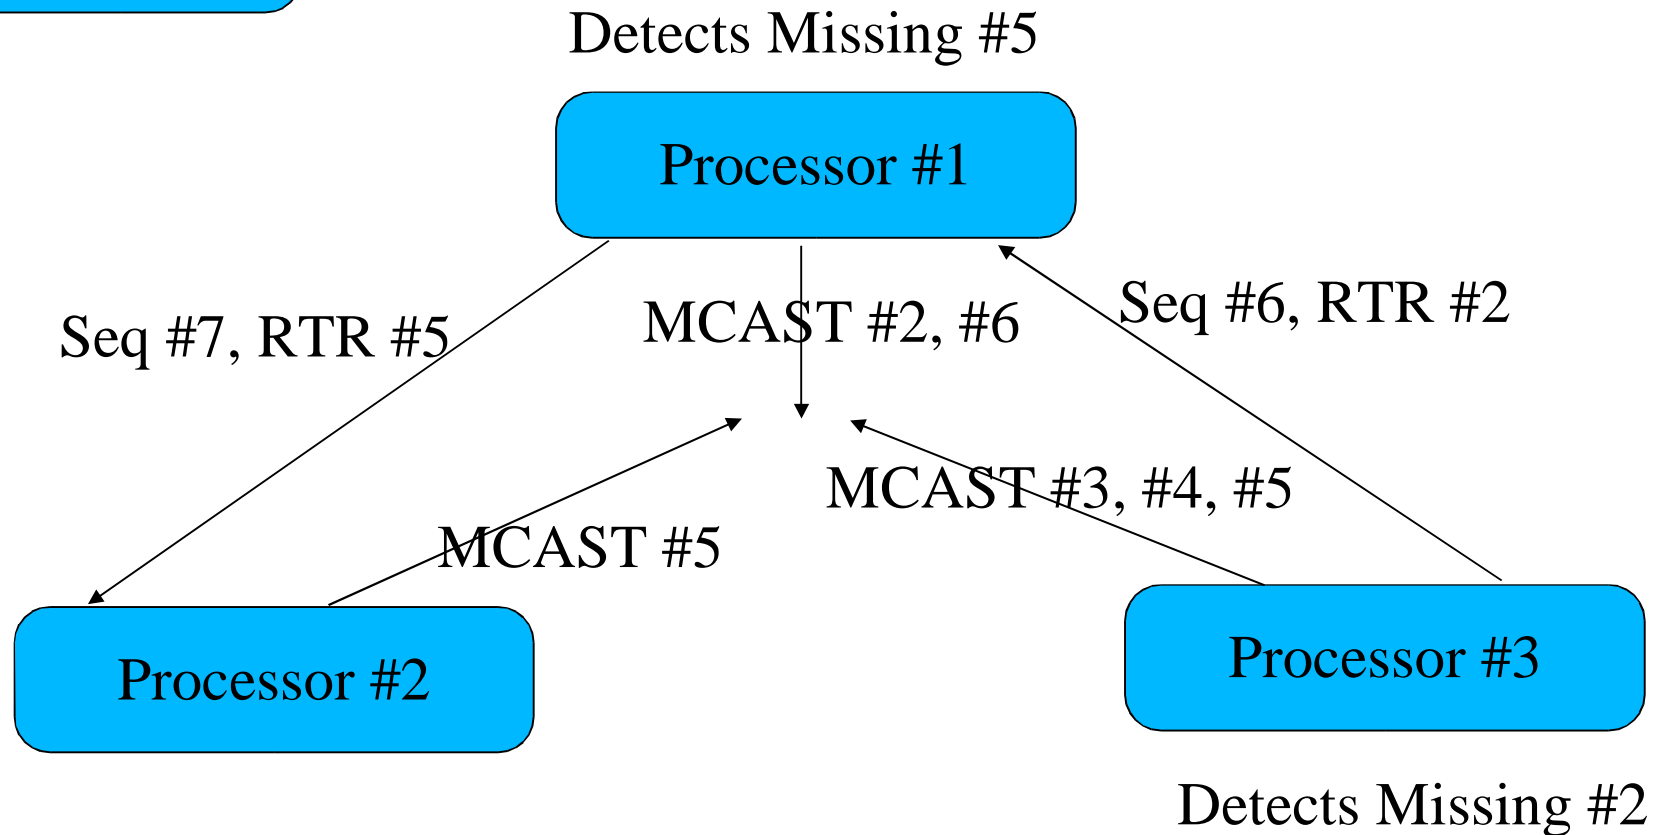
ORF Token



The Ring Protocol

- Sequence Number
- Retransmit List (RTR)
- flow control count
- group arut

ORF Token



Example Problem – Lock Service

Client / Server Approach

- One server contains list of locks.
- A lock request is sent to the server.
- The server processes the request.
- The server responds to the client.
- maximum 1700 locks per second – tied directly to ethernet access time.

Example Problem – Lock Server

Virtual Synchrony Approach

- List of all locks contained on all processors
- processor acquires lock by sending message requesting lock
- when message is self-delivered lock is acquired
- because all processors have replica of locks, no request/response is required
- maximum locks per second – depends on cpu speed but atleast 20,000 per second