

Time, Clocks, and the Ordering of Events in a Distributed System

Leslie Lamport

Presented by Mark Colbert

Overview

- Why Order?
- Partial Ordering
- Logical (Lamport) Clocks
- Mutual Exclusion
- Physical Clocks
- References

Why Order?

- Distributed Systems
 - Mutual Exclusion
 - Database Serialization

Partial Ordering

- usually we order by physical time
 - cannot perfectly synchronize clocks
- a “happened before” b
- what if b “happened before” a too?
 - concurrent


Logical (Lamport) Clocks

- use notion of partial ordering to find an arbitrary total ordering
- logical clocks do not require any basis in physical time, just a counter
- clock ticks represented by passing events
- processes send messages to one another with timestamps (current time for process)

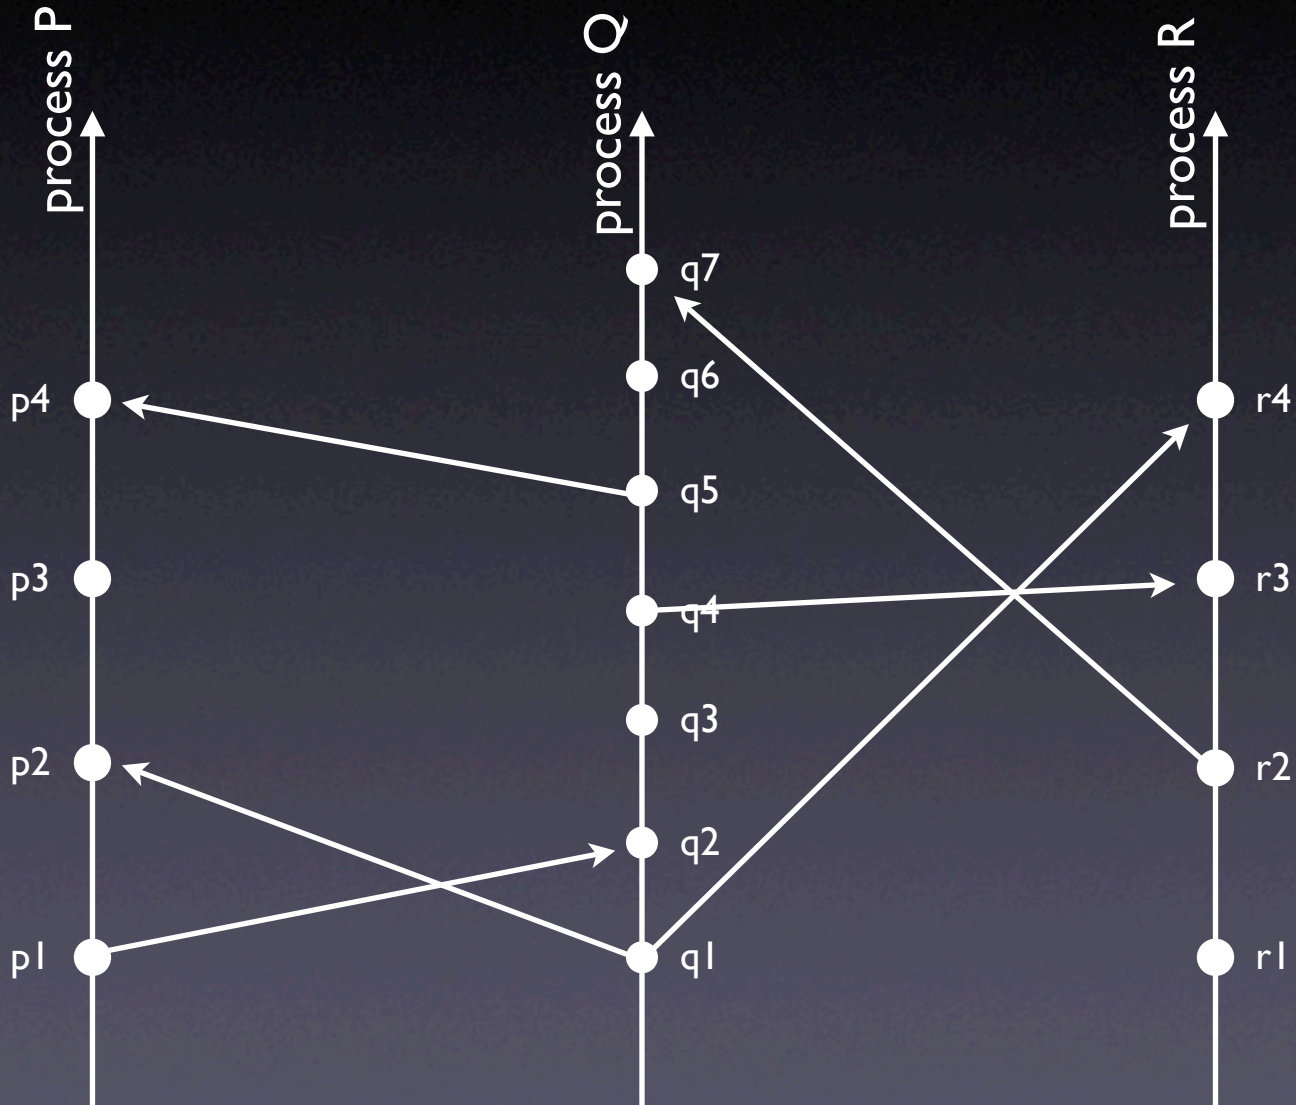
Logical (Lamport) Clocks

- rules for logical clocks
 - each process increments its counter (clock) after every event
 - if an event is a message, set the counter, c , to:
$$c = \max(\text{timestamp} + 1, c + 1)$$


Using Logical Clocks



Example



Example

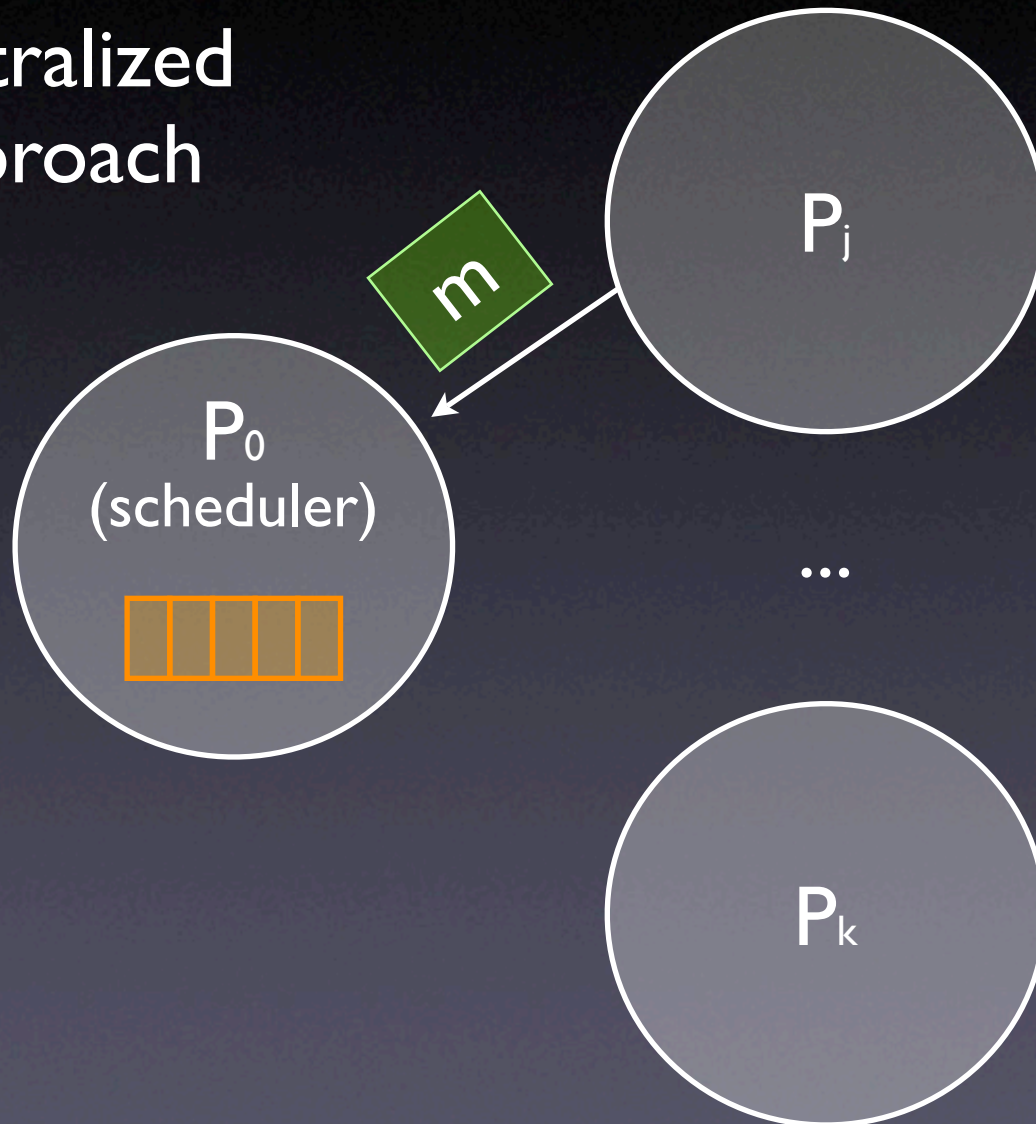


Mutual Exclusion


- Shared resource between processes
- Requirements
 - A process must release a resource before another can obtain
 - Different requests must be granted in the order they were requested
 - Every process which is granted the resource must release it

Mutual Exclusion

Centralized
Approach



m
request message

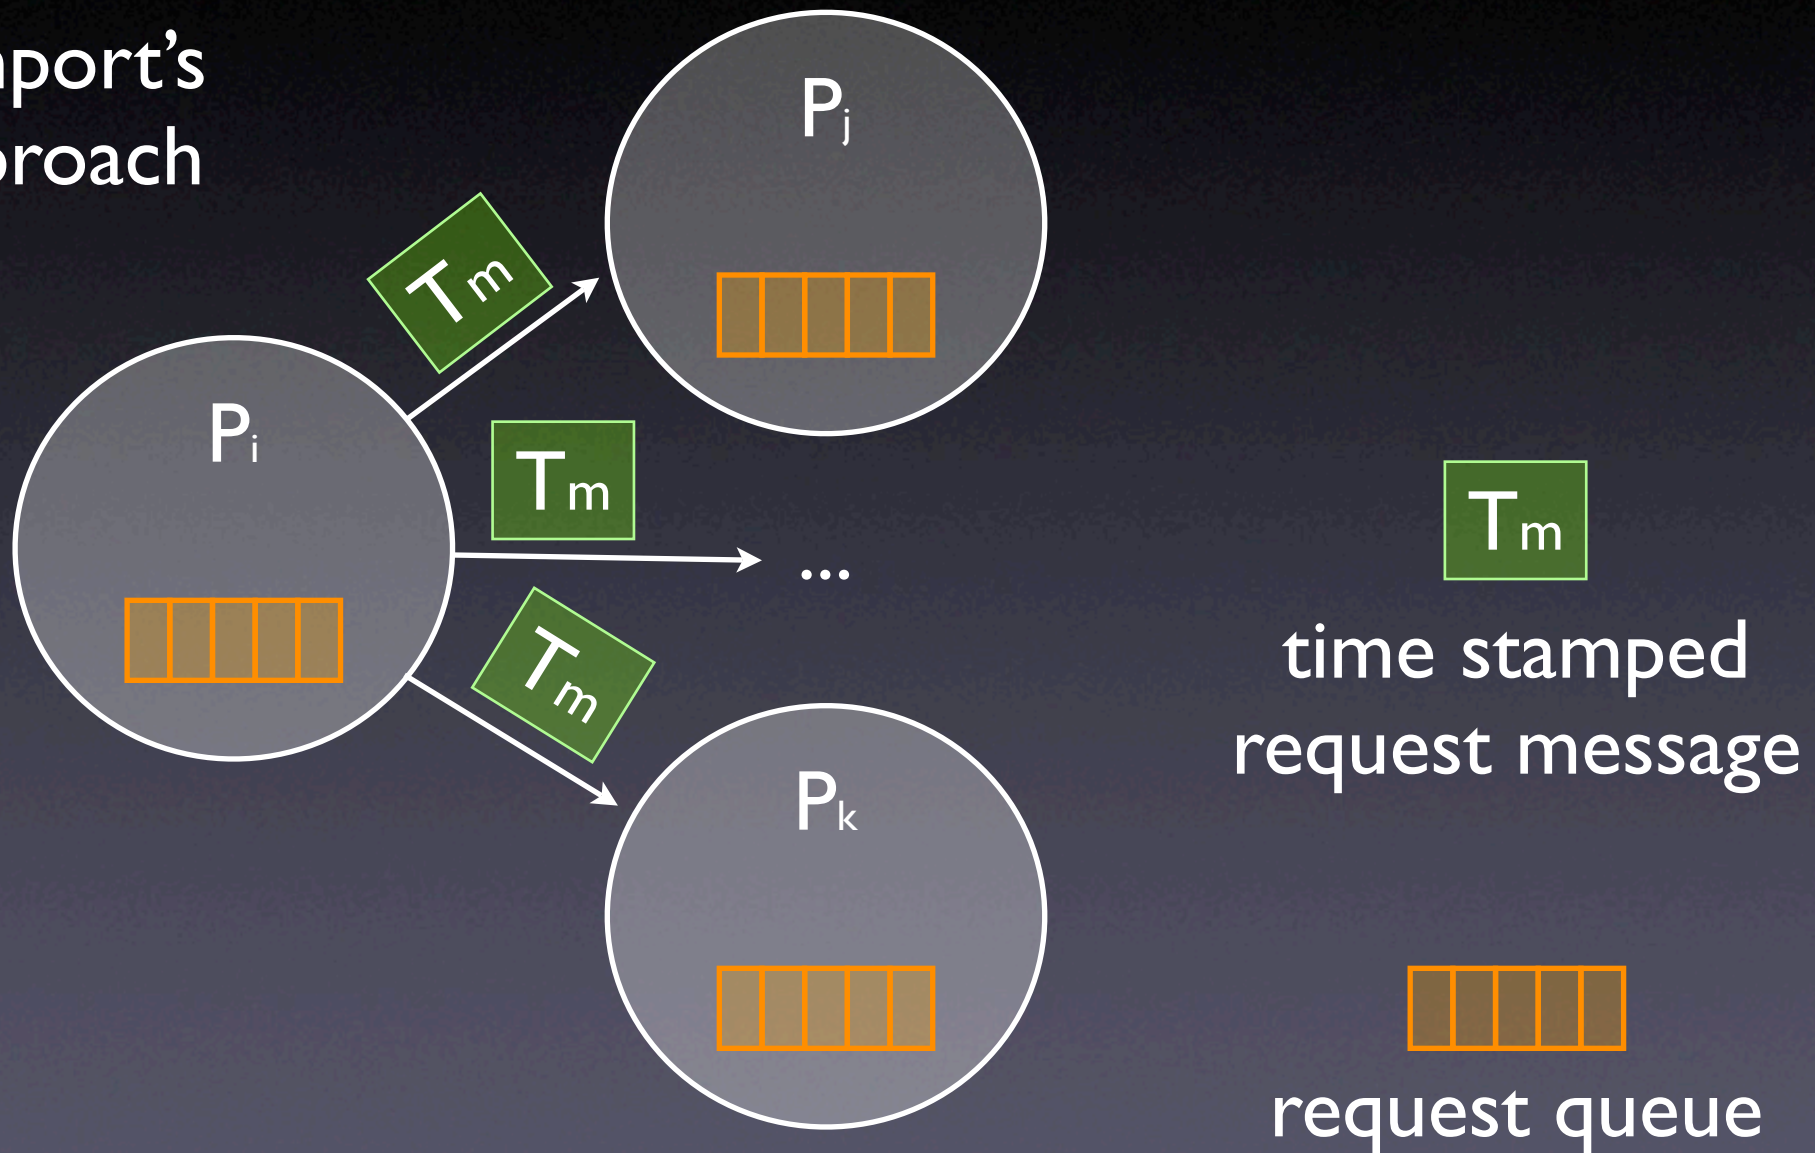

request queue

Mutual Exclusion

- Centralized Approach
 - does not guarantee order
 - suppose P_j sends a request to P_0 then sends a message to P_k
 - upon receiving the message, P_k sends a request that arrives before P_j 's request
 - order is broken


Mutual Exclusion

Lamport's
Approach



Mutual Exclusion

Lamport's
Approach



Mutual Exclusion

- Lamport's Approach (Obtaining)
 - P_i broadcast request message to all processes with time stamp
 - each process adds the request to its request queue and sends a time stamped acknowledgment
 - obtains access if request is first on P_i 's request queue and P_i receives a message from all processes with a time stamp no later than original broadcasted time stamp

Mutual Exclusion

- Lamport's Approach (Release)
 - broadcast a message to all other processes releasing the resource
 - process removes resource request from queue

Physical Clocks

- Similar notion, except that difference between time must be minimal
- Update the clock
 - $t = \max(t', T_m + T_s)$
 - t' - current time of the clock
 - T_m - time stamp of the message
 - T_s - minimum delay for the message to travel

References

- Lamport, L., Time, Clocks, and Ordering of Events in a Distributed System, Communications of the ACM, July 1978
- Bic, L., Shaw, A., Operating Systems Principles Prentice Hall, 2003