

CSE 452

Distributed Systems

Virtual + Vector clocks

(Lamport)

(logical) (vs physical clocks)

Simultaneity

e_1 happens before e_2 ($e_1 \text{ HB } e_2$)

- they happen on one machine, and e_1 is first
- e_2 receives a msg that e_1 sent
- or any transitive combination


e_1 is concurrent w/ e_2 (simultaneous)


if e_1 did not HB e_2
and e_2 — e_1

EVENT

- node executed a local instruction
 - sending a msg
 - receiving a msg
-

some events may appear to occur.
in different orders
to different observers






time

process P

process Q

process R



Can we compute HB from
inside the distributed system?

- logical clocks $C: \text{events} \rightarrow \text{timestamps}$
nonneg. ints.

$C(e) = \text{logical timestamp of } e$

clock condition:

if $e_1 \text{ HB } e_2$ then $C(e_1) < C(e_2)$

Lamport Clock Algorithm

- every node will timestamp its events
- add to every message, the logical time of the send event
- ON RECV:
 - timestamp of prev event on this node
 - timestamp on the message (of send)
 - take the max, add 1