


# BRINGING ORDER TO THE GALAXY



# SYNCHRONY VS ASYNCHRONY

## Synchronous systems

- Known bound on message delivery
- Known bound on processing speed
- Considered a strong assumption

## Asynchronous systems

- **No bound** on message delivery
- **No bound** on processing speed
- Weak assumption = less vulnerable
- Asynchronous  $\neq$  slow

This lecture: asynchronous + no process failures

# ORDERING EVENTS IN A DISTRIBUTED SYSTEM

What does it mean for an event to  
“happen before” another event?

# What is a distributed system?

A collection of distinct processes that:

- are spatially separated
- communicate with one another by exchanging messages
- have non-negligible communication delay
- do not share fate
- have separate physical clocks

*(imperfect, unsynchronized)*



## Non-distributed system

- A single clock
- Each event has a timestamp
- **Compare timestamps to order events**




## Distributed system

- Each process has its own clock
- Each clock runs at a different speed
- **Cannot directly compare clocks**

# ORDERING EVENTS WITHOUT PHYSICAL CLOCKS

Modeling a process:

- A set of instantaneous events with an a priori total ordering
- Events can be local, sends, or receives.




# ORDERING EVENTS WITHOUT PHYSICAL CLOCKS

“Happened-before” relation, denoted:  $\rightarrow$

## Part I

- If  $a$  and  $b$  are events on the same process and  $a$  comes before  $b$ , then  $a \rightarrow b$




# ORDERING EVENTS WITHOUT PHYSICAL CLOCKS

“Happened-before” relation, denoted:  $\rightarrow$

## Part 2

- If  $a$  is the sending of a message by one process and  $b$  is the receipt of the same message by another process, then  $a \rightarrow b$






# ORDERING EVENTS WITHOUT PHYSICAL CLOCKS

“Happened-before” relation, denoted:  $\rightarrow$


## Part 3

- If  $a \rightarrow b$  and  $b \rightarrow c$ , then  $a \rightarrow c$




# ORDERING EVENTS WITHOUT PHYSICAL CLOCKS

Putting it all together




# ORDERING EVENTS WITHOUT PHYSICAL CLOCKS

Can arrows go backwards?



# ORDERING EVENTS WITHOUT PHYSICAL CLOCKS


Can cycles be formed?



No, because an event would happen before itself


# ORDERING EVENTS WITHOUT PHYSICAL CLOCKS

Are all events related by  $\rightarrow$ ?




# A PARTIAL ORDER

The set of events  $q$  such that  $q \rightarrow p$  are the events that could have influenced  $p$  in some way



# A ~~PARTIAL~~ CAUSAL ORDER

If two events could not have influenced each other, it doesn't matter when they happened relatively to each other



$h$  and  $d$  are concurrent:  $h \not\rightarrow d, d \not\rightarrow h$

Goal: generate a **total** order that is consistent with the happened-before partial order



# LAMPORT CLOCKS

Define a function **LC** such that:

$$p \rightarrow q \Rightarrow LC(p) < LC(q)$$

(the Clock condition)

# LAMPORT CLOCKS

Define a function **LC** such that:


$$p \rightarrow q \Rightarrow LC(p) < LC(q)$$

(the Clock condition)

Implement **LC** by keeping a local **LC<sub>i</sub>** at each process *i*


# LAMPORT CLOCKS

Single process



# LAMPORT CLOCKS


Across processes




$$b \rightarrow h \Rightarrow LC(b) < LC(h)$$

$$g \rightarrow h \Rightarrow LC(g) < LC(h)$$


# PUTTING IN ALL TOGETHER




IS THIS CORRECT?



# GENERATING A TOTAL ORDER



- Order messages by LC
- Ties are broken by unique process ID



Lamport clocks implement the Clock condition

$$p \rightarrow q \Rightarrow LC(p) < LC(q)$$

But is that all we need?




# ADMINISTRIVIA

- Make sure to subscribe to our Piazza forum
  - Announcements, discussion, etc.
- At capacity, will issue more overrides as more people drop
- Remember to send me a selfie of you!

# FIFO DELIVERY

FIFO delivery


$$\text{send}_i(m) \rightarrow \text{send}_i(m') \Rightarrow \text{deliver}_j(m) \rightarrow \text{deliver}_j(m')$$



# FIFO DELIVERY

FIFO delivery


$$\text{send}_i(m) \rightarrow \text{send}_i(m') \Rightarrow \text{deliver}_j(m) \rightarrow \text{deliver}_j(m')$$



# CAUSAL DELIVERY


When more processes are involved, causal delivery is needed:

$$send_i(m) \rightarrow send_j(m') \Rightarrow deliver_k(m) \rightarrow deliver_k(m')$$



# GAP DETECTION


Should  $r$  deliver  $m'$ ?



**Gap detection:** Given two events  $e$  and  $e'$ , where  $LC(e) < LC(e')$ , determine whether some other event  $e''$  exists such that

$$LC(e) < LC(e'') < LC(e')$$

# GAP DETECTION



**Gap detection:** Given two events  $e$  and  $e'$ , where  $LC(e) < LC(e')$ , determine whether some other event  $e''$  exists such that

$$LC(e) < LC(e'') < LC(e')$$

Lamport clocks don't provide gap detection!

# HOW TO IMPLEMENT CAUSAL DELIVERY?

(in other words, when is it safe to deliver  $m'$ ?)

a) Wait to receive a message with higher LC from each channel

b) Implement better clocks!

# FROM CLOCKS TO STRONG CLOCKS

$$p \rightarrow q \Rightarrow LC(p) < LC(q)$$

Clock condition


$$p \rightarrow q \Leftrightarrow LC(p) < LC(q)$$

Strong clock condition



# CAUSAL HISTORIES

The set of events  $q$  such that  $q \rightarrow p$  are the events that could have influenced  $p$  in some way



$$\theta(g) = \{a, b, e, f, h, g\}$$


# IMPLEMENTING STRONG CLOCKS

(the hard way)

- Initialize  $\theta := \emptyset$
- For send and local events  $e$ ,  $\theta(e) := \theta \cup \{e\}$
- For receive events  $e = \text{recv}(m)$ ,  $\theta(e) := \theta \cup \{e\} \cup \theta(m)$

# IMPLEMENTING STRONG CLOCKS

(the hard way)



Strong clock condition:  $p \rightarrow q \Leftrightarrow \theta(p) \subset \theta(q)$