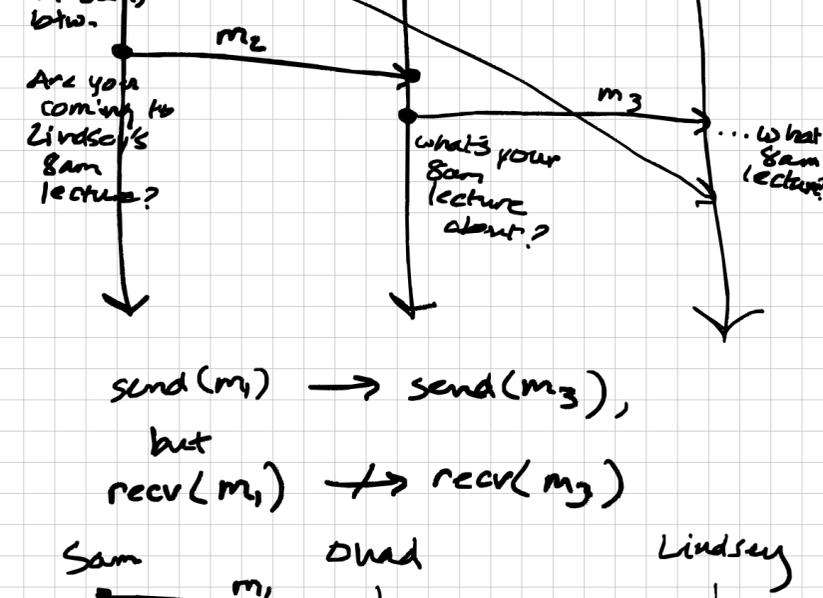


A few ideas from distributed systems for PL folk

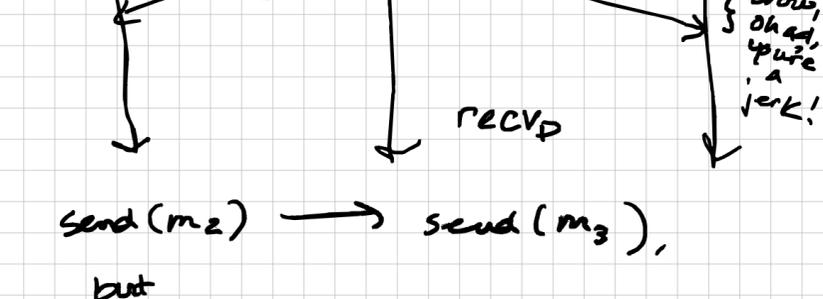
Lindsey Kuper (UC Santa Cruz)

lecture 2 agenda (time permitting):

- Causal message delivery
- Vector clocks
- a causal broadcast protocol
- safety & liveness } these'll probably have to wait until Friday
- fault models



send(m_1) → send(m_3),
but
recv(m_1) ↯ recv(m_3)



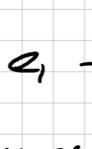
send(m_2) → send(m_3),
but
recv_{Lindsey}(m_2) ↯ recv_{Lindsey}(m_3)

Causal message delivery:

Given messages m, m' in an execution, if $\text{Send}(m) \rightarrow \text{Send}(m')$, then $\exists P$ such that P delivers both m and m' , $\text{deliver}_P(m) \rightarrow \text{deliver}_P(m')$.

What's "delivery"?

- sending is something you do
- receiving is something that happens to you.
- delivering is something you can choose to do with a message you receive.



vector clocks!

$$e_1 \rightarrow e_2 \iff LC(e_1) < LC(e_2)$$

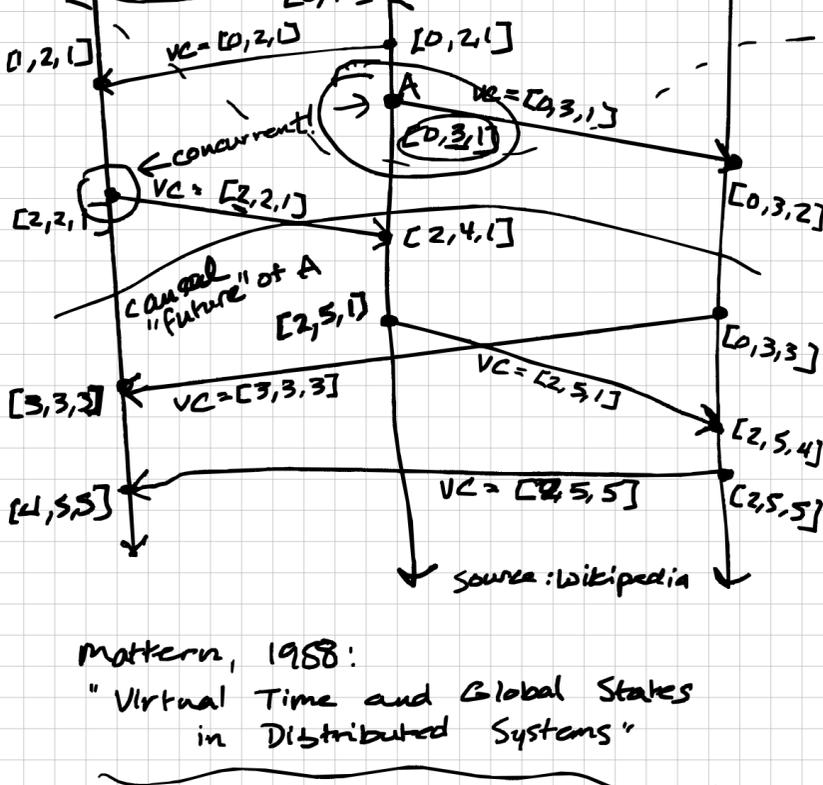
The vector clock algorithm:

- Every process maintains a vector of natural numbers, all initially 0, one entry for each process. $[0, 0, 0]$
- On every event, a process increments its own entry in its vector.
- when sending a message, a process includes its current vector (after incrementing!)
- when receiving a message, a process updates its vector to the pointwise maximum of the received vector and its own vector. (and increment its own entry, too!)

$$\max([2, 0, 3], [1, 3, 4]) = [2, 3, 4]$$

$$e_1 \rightarrow e_2 \iff VC(e_1) < VC(e_2)$$

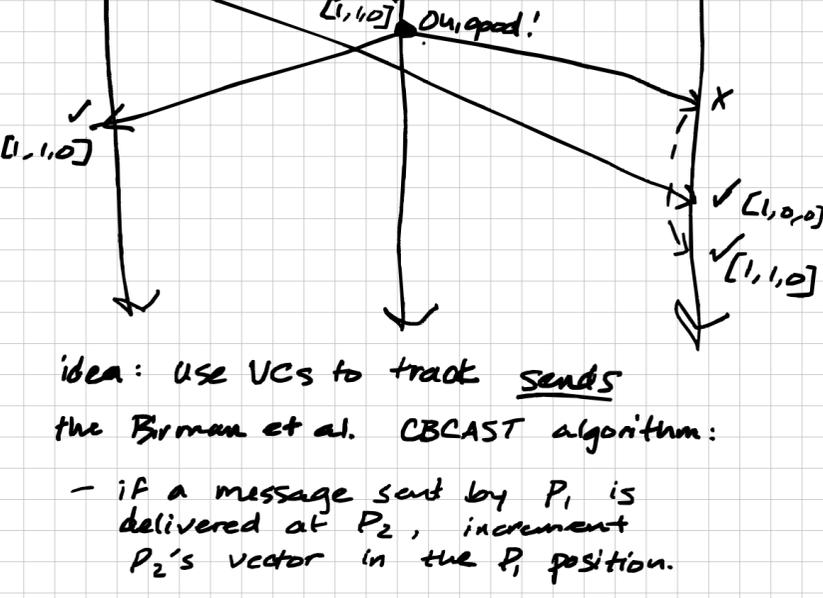
$$VC(e_1) < VC(e_2) \stackrel{\text{def}}{=} \text{for all indices } i, VC(e_1)_i \leq VC(e_2)_i \text{ and } VC(e_1) \neq VC(e_2).$$



Mattern, 1988: "Virtual Time and Global States in Distributed Systems"

A causal broadcast protocol

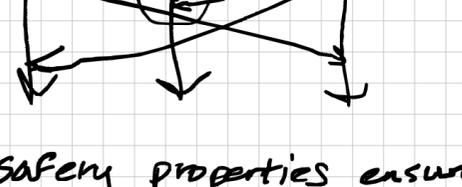
Birman et al. (1991), "Lightweight causal and atomic group multicast"



idea: use VCs to track sends
the Birman et al. CBCAST algorithm:

- if a message sent by P_1 is delivered at P_2 , increment P_2 's vector in the P_1 position.
- when a process sends a message, increment its own entry and include its vector with the message.
- (Deliverability policy): A message sent by P_1 to P_2 is only deliverable at P_2 if, let VC_m be the message's VC.
 - $VC_m[P_1] = VC_{P_2}[P_1] + 1$
 - and
 - $VC_m[P_k] \leq VC_{P_2}[P_k]$ for all $k \neq 1$

deliverability condition!



Safety properties ensure that a "~~bad~~" thing doesn't happen.

Liveness properties ensure that a "~~good~~" thing (eventually) happens.