Vector Clocks

COS 418 + 518: (Advanced) Distributed Systems
Lecture 5

Mike Freedman & Wyatt Lloyd

Lamport Clocks Review

Q: a → b => LC(a) < LC(b)

Q: LC(a) < LC(b) => b -/-> a (a → b or a || b)

Q: a || b => nothing

Lamport Clocks and Causality

- Lamport clock timestamps do not capture causality
- Given two timestamps C(a) and C(z), want to know whether there's a chain of events linking them:

a → b → ... → y → z

Vector clock: Introduction

- One integer can't order events in more than one process
- So, a Vector Clock (VC) is a vector of integers, one entry for each process in the entire distributed system
  - Label event e with VC(e) = [c_1, c_2 …, c_n]
    - Each entry c_k is a count of events in process k that causally precede e

Vector clock: Update rules

- Initially, all vectors are [0, 0, …, 0]
- Two update rules:

1. For each local event on process i, increment local entry c_i
2. If process j receives message with vector $[d_1, d_2, …, d_n]$:
   - Set each local entry $c_k = \max\{c_k, d_k\}$
   - Increment local entry c_j

5

Vector clock: Example

- All processes' VCs start at [0, 0, 0]
- Applying local update rule
- Applying message rule
  - Local vector clock piggybacks on inter-process messages

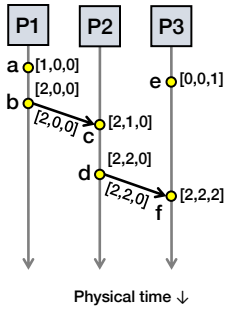

6

Comparing vector timestamps

- Rule for comparing vector timestamps:
  - V(a) = V(b) when $a_k = b_k$ for all k
  - V(a) < V(b) when $a_k \leq b_k$ for all k and V(a) ≠ V(b)
- Concurrency:
  - V(a) || V(b) if $a_i < b_i$ and $a_j > b_j$, some i, j

7

Vector clocks capture causality

- V(w) < V(z) then there is a chain of events linked by Happens-Before (→) between a and z
- V(a) || V(w) then there is no such chain of events between a and w

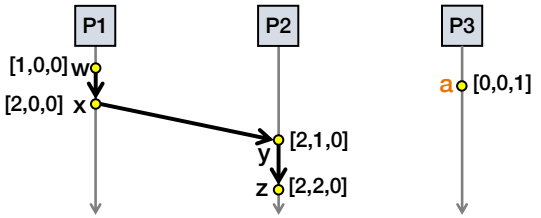

8

Comparing vector timestamps

- Rule for comparing vector timestamps:
  - V(a) = V(b) when $a_k = b_k$ for all k
    - They are the same event
  - V(a) < V(b) when $a_k \leq b_k$ for all k and V(a) ≠ V(b)
    - a → b
- Concurrency:
  - V(a) || V(b) if $a_i < b_i$ and $a_j > b_j$, some i, j
    - a || b

9

Two events a, z

Lamport clocks: C(a) < C(z)

Conclusion: z -/-> a, i.e., either a → z or a || z

Vector clocks: V(a) < V(z)

Conclusion: a → z

Vector clock timestamps precisely capture happens-before relation (potential causality)

10