

CSE 486/586 Distributed Systems Logical Time

Steve Ko
Computer Sciences and Engineering
University at Buffalo

CSE 486/586, Spring 2013

Last Time

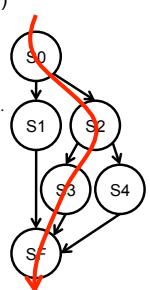
- Clock skews do happen
- External and internal synchronization
 - Cristian's algorithm: external synchronization
 - Berkeley algorithm: internal synchronization
 - Both designed for LAN
- NTP (Network Time Protocol)
 - Hierarchy of time servers
 - Estimates the actual offset between two clocks
 - Designed for the Internet
- Logical time
 - For ordering events, relative time should suffice.
 - Will continue today

CSE 486/586, Spring 2013

2

Basics: State Machine

- State: a **collection of values** of variables
- Event: an occurrence of an action that changes the state, (i.e., **instruction, send, and receive**)
- As a program,
 - We can think of all **possible execution paths**.
- At runtime,
 - There's **only one path** that the program takes.
- Equally applicable to
 - A single process
 - A **distributed set of processes**



CSE 486/586, Spring 2013

3

Ordering Basics

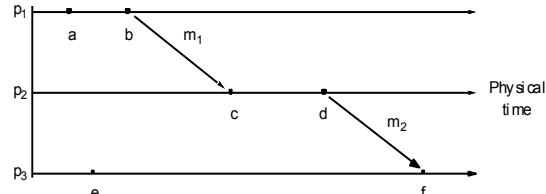
- Why did we want to synchronize physical clocks?
- What we need: Ordering of events.
- Arises in many different contexts...



CSE 486/586, Spring 2013

4

Abstract View

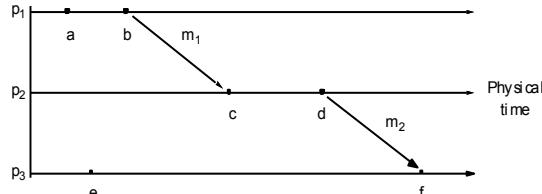


- Above is what we will deal with most of the time.
- Ordering question: what do we ultimately want?
 - Taking two events and determine which one happened before the other one.

CSE 486/586, Spring 2013

5

What Ordering?

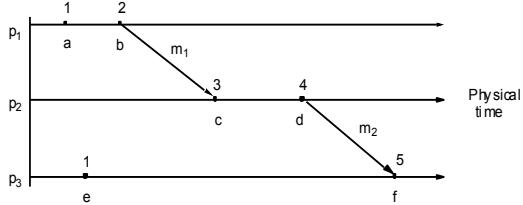


- Ideal?
 - Perfect physical clock synchronization
- Reliably?
 - Events in the same process
 - Send/receive events

CSE 486/586, Spring 2013

6

Lamport Timestamps



CSE 486/586, Spring 2013

7

Logical Clocks

- Lamport algorithm assigns **logical timestamps**:
 - All processes use a counter (clock) with initial value of zero
 - A process increments its counter when a **send** or an **instruction** happens at it. The counter is assigned to the event as its timestamp.
 - A **send (message)** event carries its timestamp
 - For a **receive (message)** event the counter is updated by $\max(\text{local clock}, \text{message timestamp}) + 1$
- Define a logical relation **happened-before** (\rightarrow) among events:
 - On the same process: $a \rightarrow b$, if $\text{time}(a) < \text{time}(b)$
 - If p_1 sends m to p_2 : $\text{send}(m) \rightarrow \text{receive}(m)$
 - (Transitivity) If $a \rightarrow b$ and $b \rightarrow c$ then $a \rightarrow c$
 - Shows **causality** of events

CSE 486/586, Spring 2013

8

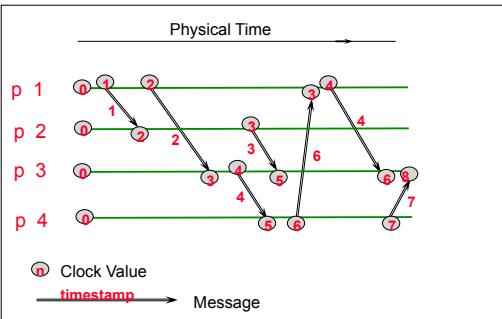
CSE 486/586 Administrivia

- PA2 is out.
 - Due on 3/1
 - Start with the content provider.
- Please understand the flow of PA1.
- Please be careful about your coding style.
- Lecture slides
 - I will try posting them a day before.
 - I will also post a PDF version.
- There is a course website.
 - Schedule, syllabus, readings, etc.

CSE 486/586, Spring 2013

9

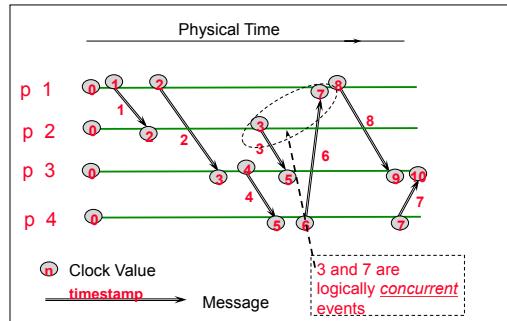
Find the Mistake: Lamport Logical Time



CSE 486/586, Spring 2013

10

Corrected Example: Lamport Logical Time

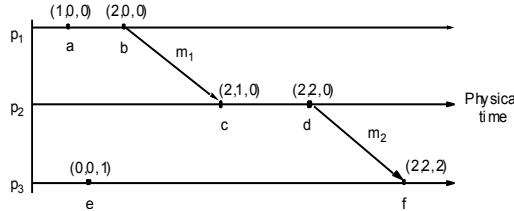


CSE 486/586, Spring 2013

11

Vector Timestamps

- With Lamport clock
 - e "happened-before" $f \Rightarrow \text{timestamp}(e) < \text{timestamp}(f)$, but
 - $\text{timestamp}(e) < \text{timestamp}(f) \nrightarrow e$ "happened-before" f
- Idea?



CSE 486/586, Spring 2013

12

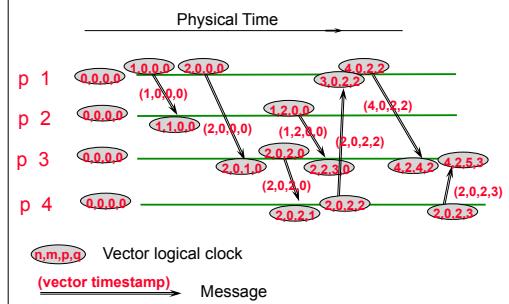
Vector Logical Clocks

- Vector Logical time addresses the issue:
 - All processes use a vector of counters (logical clocks), i^{th} element is the clock value for process i , initially all zero.
 - Each process i increments the j^{th} element of its vector upon an **instruction** or **send** event. Vector value is timestamp of the event.
 - A **send(message)** event carries its vector timestamp (counter vector)
 - For a **receive(message)** event, $V_{\text{receiver}[i]} =$
 - $\text{Max}(V_{\text{receiver}[i]}, V_{\text{message}[i]})$, if j is not self,
 - $V_{\text{receiver}[i]} + 1$, otherwise

CSE 486/586, Spring 2013

13

Find a Mistake: Vector Logical Time



CSE 486/586, Spring 2013

14

Comparing Vector Timestamps

- $VT_1 = VT_2$,
 - iff $VT_1[i] = VT_2[i]$, for all $i = 1, \dots, n$
- $VT_1 \leq VT_2$,
 - iff $VT_1[i] \leq VT_2[i]$, for all $i = 1, \dots, n$
- $VT_1 < VT_2$,
 - iff $VT_1 \leq VT_2 \text{ & } \exists j (1 \leq j \leq n \text{ & } VT_1[j] < VT_2[j])$
- VT_1 is concurrent with VT_2
 - iff $(\text{not } VT_1 \leq VT_2 \text{ AND not } VT_2 \leq VT_1)$

CSE 486/586, Spring 2013

15

The Use of Logical Clocks

- Is a design decision
- NTP error bound
 - Local: a few ms
 - Wide-area: 10's of ms
- If your system **doesn't care about this inaccuracy**, then NTP should be fine.
- Logical clocks impose an arbitrary order over concurrent events anyway
 - Breaking ties: process IDs, etc.

CSE 486/586, Spring 2013

16

Summary

- Relative order of events enough for practical purposes
 - Lamport's logical clocks
 - Vector clocks
- Next: How to take a global snapshot

CSE 486/586, Spring 2013

17

Acknowledgements

- These slides contain material developed and copyrighted by Indranil Gupta at UIUC.

CSE 486/586, Spring 2013

18