

Time and Logical Clocks 2



جامعة الملك عبد الله
للعلوم والتقنية
King Abdullah University of
Science and Technology

CS 240: Computing Systems and Concurrency
Lecture 4

Marco Canini

Lamport Clocks Review

- Happens-Before relationship
 - Event **a** *happens before* event **b** ($\mathbf{a} \rightarrow \mathbf{b}$)
 - **c**, **d** not related by \rightarrow so *concurrent*, written as $\mathbf{c} \parallel \mathbf{d}$
- Lamport clocks is a logical clock construction to capture the order of events in a distributed systems (disregarding the precise clock time)
 - Tag every event **a** by $C(\mathbf{a})$
 - If $\mathbf{a} \rightarrow \mathbf{b}$, then ?
 - If $C(\mathbf{a}) < C(\mathbf{b})$, then ?
 - If $\mathbf{a} \parallel \mathbf{b}$, then ?

Lamport Clocks Review

- Happens-Before relationship
 - Event **a** *happens before* event **b** ($\mathbf{a} \rightarrow \mathbf{b}$)
 - **c**, **d** not related by \rightarrow so *concurrent*, written as $\mathbf{c} \parallel \mathbf{d}$
- Lamport clocks is a logical clock construction to capture the order of events in a distributed systems (disregarding the precise clock time)
 - Tag every event **a** by $C(\mathbf{a})$
 - If $\mathbf{a} \rightarrow \mathbf{b}$, then $C(\mathbf{a}) < C(\mathbf{b})$
 - If $C(\mathbf{a}) < C(\mathbf{b})$, then **NOT** $\mathbf{b} \rightarrow \mathbf{a}$ ($\mathbf{a} \rightarrow \mathbf{b}$ or $\mathbf{a} \parallel \mathbf{b}$)
 - If $\mathbf{a} \parallel \mathbf{b}$, then nothing

Lamport Clocks and causality

- Lamport clock timestamps **don't capture causality**
- Given two timestamps $C(\mathbf{a})$ and $C(\mathbf{z})$, want to know whether there's a chain of events linking them:

$\mathbf{a} \rightarrow \mathbf{b} \rightarrow \dots \rightarrow \mathbf{y} \rightarrow \mathbf{z}$

Take-away points: Lamport clocks

- Can **totally-order** events in a distributed system: that's useful!
 - We saw an application of Lamport clocks for totally-ordered multicast
- **But:** while by construction, $a \rightarrow b$ implies $C(a) < C(b)$,
 - The converse is not necessarily true:
 - $C(a) < C(b)$ does not imply $a \rightarrow b$ (possibly, $a \parallel b$)

Can't use Lamport clock timestamps to infer **causal relationships** between events

Today

1. Logical Time: Vector clocks

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 $VC(\mathbf{e}) = [c_1, c_2 \dots, c_n]$
 - Each entry c_k is a **count of events** in process **k** that **causally precede e**

Vector clock: Update rules

- Initially, all vectors are $[0, 0, \dots, 0]$
- **Two update rules:**
 1. For each **local event** on process i , increment local entry c_i
 2. If process j **receives** message with vector $[d_1, d_2, \dots, d_n]$:
 - Set each local entry $c_k = \max\{c_k, d_k\}$, for $k = 1 \dots n$
 - Increment local entry c_j

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




Comparing vector timestamps

- Rule for comparing vector timestamps:
 - $V(\mathbf{a}) = V(\mathbf{b})$ when $\mathbf{a}_k = \mathbf{b}_k$ for all k
 - $V(\mathbf{a}) < V(\mathbf{b})$ when $\mathbf{a}_k \leq \mathbf{b}_k$ for all k and $V(\mathbf{a}) \neq V(\mathbf{b})$
- Concurrency:
 - $\mathbf{a} \parallel \mathbf{b}$ if $\mathbf{a}_i < \mathbf{b}_i$ and $\mathbf{a}_j > \mathbf{b}_j$, some i, j

Vector clocks capture causality

- $V(\mathbf{w}) < V(\mathbf{z})$ **then** there is a chain of events linked by Happens-Before (\rightarrow) between \mathbf{w} and \mathbf{z}
- If $V(\mathbf{a}) \parallel V(\mathbf{w})$ then there is **no such chain of events** between \mathbf{a} and \mathbf{w}



Two events a, z

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

Conclusion: NOT $z \rightarrow a$ (either $a \rightarrow z$ or $a \parallel z$)

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

Conclusion: $a \rightarrow z$

Vector clock timestamps precisely capture Happens-Before relationship (potential causality)

Disadvantage of vector timestamps


- Compared to Lamport timestamps, vector timestamps $O(n)$ overhead for storage and communication, $n = \text{no. of processes}$

Take-away points

- Vector Clocks
 - Precisely capture happens-before relationship

VC Quiz

- Suppose these processes maintain vector clocks. Write the vector clock of each event starting from clock time 0.



Safety and liveness properties

Reasoning about fault tolerance

- This is hard!
 - How do we design fault-tolerant systems?
 - How do we know if we're successful?
- Often use “properties” that hold true for every possible execution
- We focus on **safety** and **liveness** properties

Properties

- **Property**: a predicate that is evaluated over a run of the system
 - “every message that is received was previously sent”
- Not everything you may want to say about a system is a property:
 - “the program sends an average of 50 messages in a run”

Safety properties

- “Bad things” don’t happen, ever
 - No more than k processes are simultaneously in the critical section
 - Messages that are delivered are delivered in causal order
- A safety property is “prefix closed”:
 - if it holds in a run, it holds in every prefix

Liveness properties

- “Good things” eventually happen
 - A process that wishes to enter the critical section eventually does so
 - Some message is eventually delivered
 - Eventual consistency: if a value doesn't change, two servers will eventually agree on its value
- Every run can be extended to satisfy a liveness property
 - If it does not hold in a prefix of a run, it does not mean it may not hold eventually

Often a trade-off

- “Good” and “bad” are application-specific
- Safety is very important in banking transactions
 - May take some time to confirm a transaction
- Liveness is very important in social networking sites
 - See updates right away