


Advanced Course


Distributed Systems

Logical Clocks



Paris Carbone

COURSE TOPICS





- ▶ Intro to Distributed Systems
- ▶ Fundamental Abstractions and Failure Detectors
- ▶ Reliable and Causal Order Broadcast
- ▶ Distributed Shared Memory-CRDTs
- ▶ Consensus (Paxos)
- ▶ Replicated State Machines (OmniPaxos, Raft, Zab etc.)
- ▶ Time Abstractions and Interval Clocks (Spanner etc.)
- ▶ Consistent Snapshotting (Stream Data Management)
- ▶ Distributed ACID Transactions (Cloud DBs)

RECAP: CAUSAL ORDER

- Given an execution trace β ,
two events $a, b \in \beta$ are **causally ordered** ($a \rightarrow_{\beta} b$) iff either:
 - **a** occurs before **b** on the same process
 - **a** is a send(m) and **b** deliver(m) event
 - there exists a sequence of causally ordered events from a to b (**transitive**) e.g. If $a \rightarrow_{\beta} c$ and $c \rightarrow_{\beta} b$
- Two events, a and b, are **concurrent** if **not** $a \rightarrow_{\beta} b$ and **not** $b \rightarrow_{\beta} a$
 - Concurrent events are denoted as such: $a || b$

RECAP: CAUSAL ORDER



Key Observations

- We cannot order any two events in a D.S. using physical time.
 - Multiple Physical Clocks cannot keep precise physical time.
- We can order any two events in a D.S. using logical time, i.e., causal order.
 - Logical Clocks can capture causality.

Operating Systems

R. Stockton Gaines
Editor

Time, Clocks, and the Ordering of Events in a Distributed System

Leslie Lamport
Massachusetts Computer Associates, Inc.

The concept of one event happening before another in a distributed system is examined, and is shown to define a partial ordering of the events. A distributed algorithm is given for synchronizing a system of logical clocks which can be used to totally order the events. The use of the total ordering is illustrated with a method for solving synchronization problems. The algorithm is then specialized for synchronizing physical clocks, and a bound is derived on how far out of synchrony the clocks can become.

Key Words and Phrases: distributed systems, computer networks, clock synchronization, multiprocess systems

CR Categories: 4.32, 5.29

A distributed system consists of a collection of distinct processes which are spatially separated, and which communicate with one another by exchanging messages. A network of interconnected computers, such as the ARPA net, is a distributed system. A single computer can also be viewed as a distributed system in which the central control unit, the memory units, and the input-output channels are separate processes. A system is distributed if the message transmission delay is not negligible compared to the time between events in a single process.

We will concern ourselves primarily with systems of spatially separated computers. However, many of our remarks will apply more generally. In particular, a multiprocessing system on a single computer involves problems similar to those of a distributed system because of the unpredictable order in which certain events can occur.

In a distributed system, it is sometimes impossible to say that one of two events occurred first. The relation "happened before" is therefore only a partial ordering of the events in the system. We have found that problems often arise because people are not fully aware of this fact and its implications.

In this paper, we discuss the partial ordering defined by the "happened before" relation, and give a distributed algorithm for establishing a total ordering of all the events. The algorithm is a mechanism for synchronizing physical clocks. It illustrates its use in a distributed system. Prior can occur if the events are not synchronized. It differs from the algorithm given in [1] in that it avoids by introducing a new relation, the "happened before" relation, and gives a distributed algorithm for establishing a total ordering of all the events.



LOGICAL CLOCK INTUITION


- A logical clock is :
 - an algorithm that assigns a timestamp to each event occurring in a distributed system. $t(a), t(b)$, etc.
- Timestamps can be used to derive a relation between events.
- We are interested in preserving the **happen-before** relation.
If $a \rightarrow_{\beta} b$ then $t(a) < t(b)$
- **Two types of clocks: Lamport and Vector clocks**

Lamport Clocks

LAMPORT CLOCKS

- Each process has a local **logical Lamport clock**, kept in variable t_p ,
- initially $t_p = 0$
 - A process p piggybacks t_p on every message sent
- On **internal event** a :
 - $t_p := t_p + 1$; perform internal event a
- **Sending event** of message m :
 - $t_p := t_p + 1$; send(m, t_p))
- Receiving/Delivering **event** of message m with timestamp t_q from q :
 - $t_p := \max(t_p, t_q) + 1$; perform delivery event a

LAMPORT LOGICAL CLOCKS



Lamport logical clocks guarantee that:

If $a \rightarrow_{\beta} b$, then $\mathbf{t}(a) < \mathbf{t}(b)$,

if $\mathbf{t}(a) \geq \mathbf{t}(b)$, then not $(a \rightarrow_{\beta} b)$

Can we do better?

TOTAL ORDER WITH LAMPORT CLOCKS

- We can timestamp with process identifiers.
- The pair (t, p) is unique

Total Order Relation (\prec) : $(t_p, p) \prec (t_q, q)$ iff either

- $t_p < t_q$
- $t_p = t_q \wedge p < q$

i.e. break ties using process identifiers

e.g. $(5, p_5) < (7, p_2)$, $(4, p_2) < (4, p_3)$

LAMPORT CLOCKS

- Each process has a local **logical lamport clock**, kept in variable t_p ,
- initially $t_p = 0$
 - A process p piggybacks (t_p, p) on every message sent
- On **internal event** a :
 - $t_p := t_p + 1$; perform internal event a
- **Sending event** of message m :
 - $t_p := t_p + 1$; send($m, (t_p, p)$)
- Receiving/Delivering **event** of message m with timestamp (t_q, q) from q :
 - $t_p := \max(t_p, t_q) + 1$; perform delivery event a

REMARKS


- The **total order** ($<$) defined on (t, p_i) pairs is:
 - A method to deterministically derive a total order of events using local and relative process information.
 - Always respecting **causal order**.
 - A **convention** rather than the actual order of events (whatever that means).

Vector Clocks

NON-CAUSALITY AND CONCURRENT EVENTS

- Two events a and b are **concurrent** ($a \parallel_{\beta} b$) in an execution E ($\text{trace}(E) = \beta$) if
 - **not** $a \rightarrow_{\beta} b$ and **not** $b \rightarrow_{\beta} a$
- Computation theorem implies that if $(a \parallel_{\beta} b)$ in β then there are **two executions** (with traces β_1 and β_2) that are **similar** where a occurs before b in β_1 , b occurs before a in β_2

NON-CAUSALITY AND CONCURRENT EVENTS




VECTOR CLOCKS

- **Lamport Clock Limitation:** We cannot tell by looking to the timestamps of event a and b whether there is a causal relation between the events, or they are **concurrent**.
- **Vector Clocks** guarantee that:
 - 1. if $\mathbf{v}(a) < \mathbf{v}(b)$ then $a \rightarrow_{\beta} b$
 - 2. if $a \rightarrow_{\beta} b$ then $\mathbf{v}(a) < \mathbf{v}(b)$ (same as Lamport Clocks)
 - where $\mathbf{v}(a)$ is a vector clock of event a

VECTOR CLOCK DEFINITION

- Vector clock for an event a
 - $\mathbf{v}(a) = [x_1, \dots, x_n]$
 - x_i is the number of events at p_i that happened-before a



3 events at p_1

1 event at p_2

0 events at p_3

$[3, 1, 0]$


VECTOR TIME IMPLEMENTATION

- Processes $p_1, \dots, p_n \in P$
- Each process p_i has local vector \mathbf{v} of size \mathbf{n} (number of processes)
 - $\mathbf{v}[i] = 0$ for all i in $1 \dots n$
 - Piggyback \mathbf{v} on every sent message
- For each transition (on each event) update local \mathbf{v} at p_i :
 - $\mathbf{v}[i] := \mathbf{v}[i] + 1$ (internal, send or deliver)
 - $\mathbf{v}[j] := \max(\mathbf{v}[j], \mathbf{v}_q[j])$, for all $j \neq i \in |P|$ (deliver)
 - where \mathbf{v}_q is clock in message received from process q


COMPARING VECTOR CLOCKS

- $\mathbf{v_p} \leq \mathbf{v_q}$ iff
 - $v_p[i] \leq v_q[i]$ for $\forall i \in |P|$
$$[3,0,0] \leq [3,1,0]$$
- $\mathbf{v_p} < \mathbf{v_q}$ iff
 - $\mathbf{v_p} \leq \mathbf{v_q}$ and for some i , $v_p[i] < v_q[i]$
$$[3,0,0] < [3,1,0]$$
- $\mathbf{v_p}$ and $\mathbf{v_q}$ are concurrent ($\mathbf{v_p} \parallel \mathbf{v_q}$) iff
 - not $\mathbf{v_p} < \mathbf{v_q}$, and not $\mathbf{v_q} < \mathbf{v_p}$
$$[3,1,0] \nlessgtr [4,0,0]$$
- Vector clocks guarantee
 - If $v(a) < v(b)$ then $a \rightarrow b$, and
 - If $a \rightarrow b$, then $v(a) < v(b)$
 - where $v(a)$ is the vector clock of event a

EXAMPLE OF VECTOR CLOCKS




$v(a) < v(b)$ implies $a \rightarrow b$



$v(a) < > v(b)$ implies $a \parallel b$

LIMITATIONS OF VECTOR CLOCKS?



- Vectors need to be statically defined of size n
- Insufficient for problems that require total event ordering

ORDERING - SUMMARY

- the relation \rightarrow_{β} on events in executions
 - **Partial:** \rightarrow_{β} doesn't order concurrent events
- the relation $<$ on Lamport logical clocks
 - **Total:** any two distinct clock values are ordered (adding pid)
- the relation $<$ on vector timestamps
 - **Partial:** timestamp of concurrent events not ordered

LOGICAL CLOCK INVARIANTS

Lamport clock

If $a \rightarrow_{\beta} b$ then $t(a) < t(b)$ (1)

If $t(a) < t(b)$ then **not** $b \rightarrow_{\beta} a$ (2)

Vector clock

If $a \rightarrow_{\beta} b$ then $v(a) < v(b)$ (1)

If $v(a) < v(b)$ then $a \rightarrow_{\beta} b$ (2)