Computing with Time

Distributed Systems

Sistemi Distribuiti

Andrea Omicini
andrea.omicini@unibo.it

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
Alma Mater Studiorum – Università di Bologna a Cesena

Academic Year 2016/2017

1. Time & Computation
2. Time in Distributed Systems
3. Toward Coordination

About these Slides

- the following slides borrow a lot from the slides coming with the book adopted as the basic one for this course [Tanenbaum and van Steen, 2007]
- material from those slides (including pictures) has been re-used in the following, and integrated with new material according to the personal view of the professor
- every problem or mistake contained in these slides, however, should be attributed to the sole responsibility of this course's professor
- the goal, here, is to be exposed to the classical view on the *evolution* of distributed systems that the scientific community has developed over the last decades

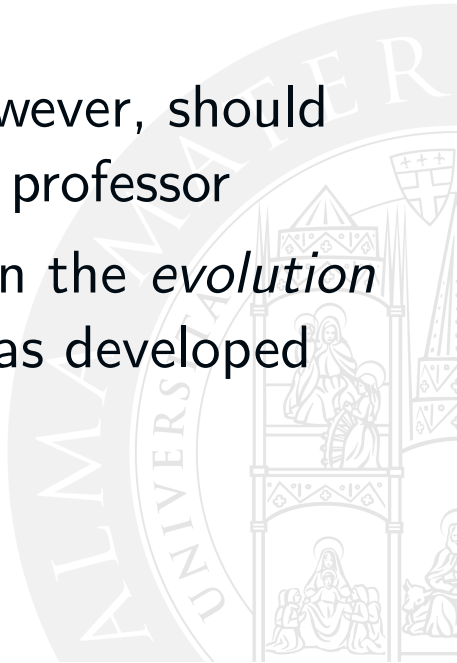

Next in Line...

1. Time & Computation
2. Time in Distributed Systems
3. Toward Coordination

Computing Needs Time

Computation & physical processes [Lee, 2009]

- on the one hand, most computational devices nowadays are *not* just *general-purpose* computers
  - e.g., cars, medical devices, instruments, communication systems, industrial robots, toys, games, ...
  - the explosion of the Internet of things (IoT) asks them to become more and more intelligent and networked [Atzori et al., 2010, Gubbi et al., 2013]
  - more and more *similar* to general-purpose computers
  - possibly harming their dependability
- on the other hand, general-purpose computers are more and more required to *interact* with *physical processes*
  - they integrate media, such as video and audio
  - their migration to personal devices and pervasive systems asks them to sense physical dynamics and control physical devices

Cyber-Physical Systems (CPS) I

Time, computation & CPS [Lee, 2009]

- the foundations of computing are still rooted in Turing [Turing, 1937], Church, and von Neumann [Burks et al., 1982]
  - they are about the *transformation of data*
  - they do *not* deal with the *dynamic* of *physical processes*
- *cyber-physical systems* (CPS) emerge from the integration of physical systems and processes with networked computing
  - CPS deeply embed computations and communication interacting with physical processes to add new capabilities to physical systems

! CPS are about merging computing and networking with physical systems to create new science, techniques, and products

Cyber-Physical Systems (CPS) II

Remarkable example of CPS

- high-confidence medical devices
- assisted living
- traffic control and safety
- advanced automotive systems
- process control
- energy conservation
- environmental control
- avionics, instrumentation
- critical infrastructure control—electric power, water resources, communications systems
- distributed robotics—telepresence, telemedicine
- defence systems
- manufacturing
- smart structures

Cyber-Physical Systems (CPS) III

The passage of time

- the technological basis that engineers and computer scientists have chosen for general-purpose computing and networking does *not support* well the applications dealing with physical processes
  - CPS in particular
- the passage of time is essentially *absent* in computing [Lee, 2009]

Misconceptions about Time in Computing [Lee, 2009] I

computing takes time it is not just the fact that efficiency has its limits — it is the fact that time is aways abstracted away from computing, so that *computing* always *omits time*

time is a resource yes, but a different sort of resource — it is practically unbounded, and is expended anyway: while it is ok having generic ways to deal with resources in programming languages, *time* needs to be a *semantic property*

time is a nonfunctional property with Turing Machine, function can be computed by abstracting from time — however, the function of a computation in a CPS is defined through *actions*, occurring in *physical time*: no less a function than a bit-to-bit transformation

Misconceptions about Time in Computing [Lee, 2009] II

real time is a quality-of-service (QoS) problem time in CPS is not about efficiency, rather about *predictability* and *repeatability* — precision and variability in timing are QoS issues, whereas time itself is much more than that: tie should be part of the semantics of programs

$$\Downarrow$$

core *computing abstractions* should be re-thought to *include time*

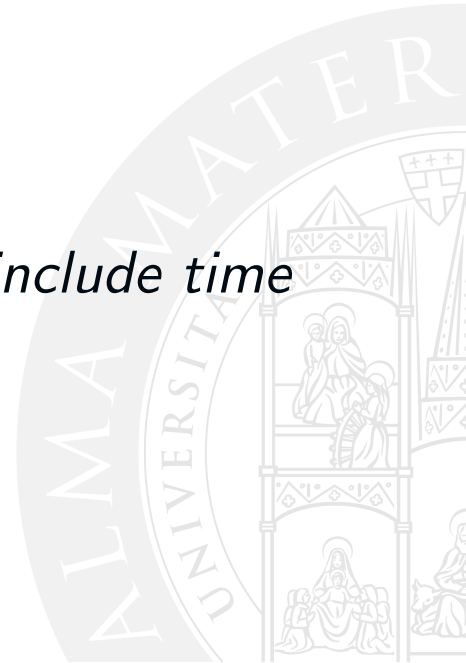

Which Notion of Time?

Time

http://www.britannica.com/science/time

> *Time*, *a measured or measurable period, a continuum that lacks spatial dimensions. Time is of philosophical interest and is also the subject of mathematical and scientific investigation. . . .*

Not a definition that we can easily exploit in our discussion

- too generic to be of some usable consequence in computing
- also, discussion in both philosophy and physics about time are too technical to be of immediate use

→ we resort to a "common sense" notion of time, that is, our notion of time as humans

Next in Line...

1 Time & Computation

2 Time in Distributed Systems

3 Toward Coordination

Parallel, Concurrent, Distributed Systems: Recap I

Parallel computing & systems

- given a computational system, we talk of parallel computation whenever the *temporal* context is the same for all computational process
- a parallel system is a computational system performing parallel computations

Parallel, Concurrent, Distributed Systems: Recap II

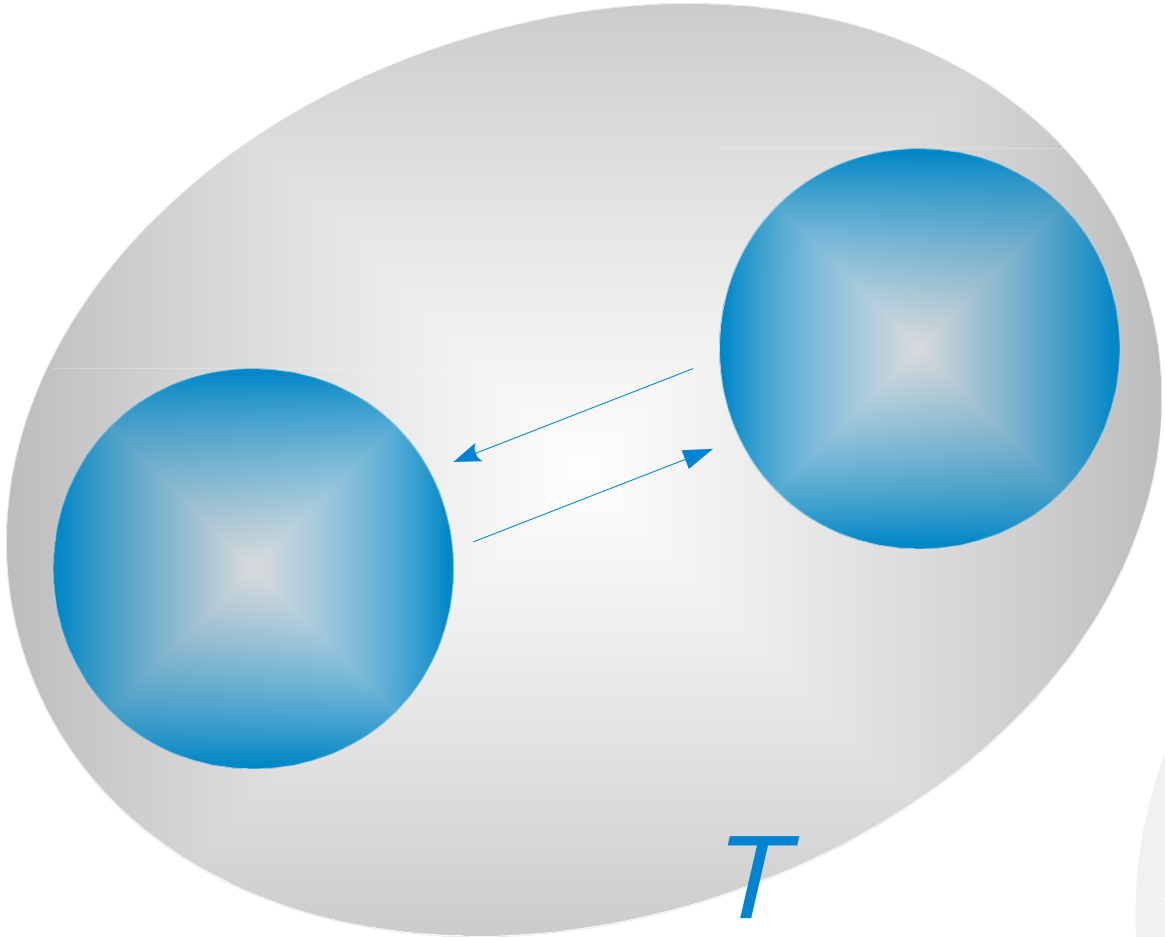

Parallel computing: the same temporal context T for all processes

Parallel, Concurrent, Distributed Systems: Recap III

Concurrent computing & systems

- given a computational system, we talk of concurrent computation whenever at least two computational processes have a different *temporal* context
- a concurrent system is a computational system performing concurrent computations

Parallel, Concurrent, Distributed Systems: Recap IV

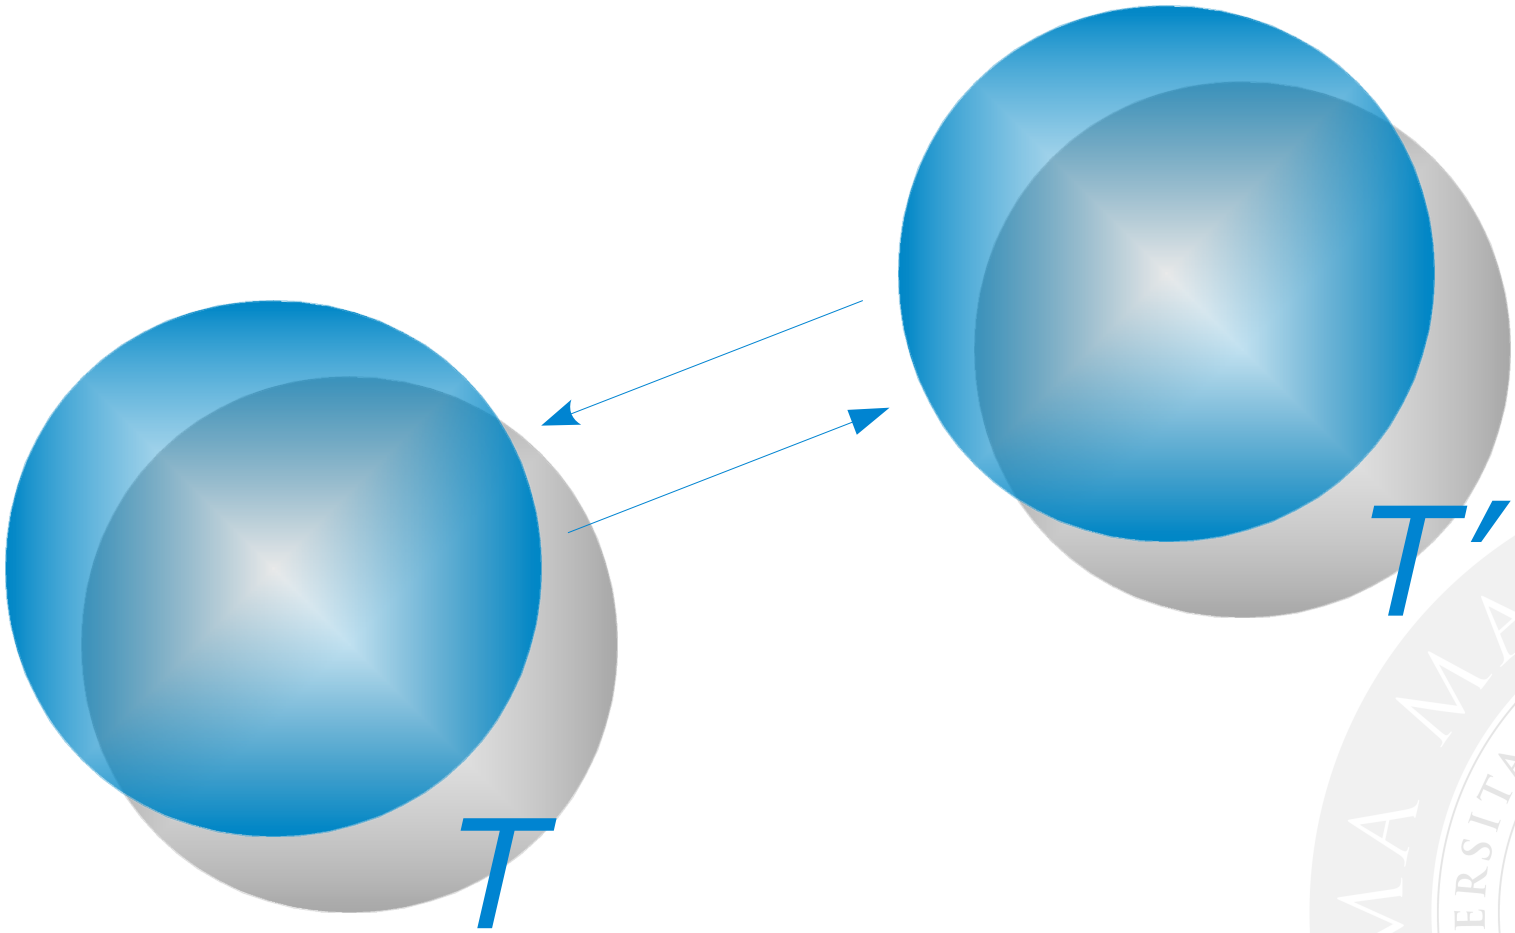

Concurrent computing: different temporal contexts $T \neq T'$ for different processes

Parallel, Concurrent, Distributed Systems: Recap V

Distributed computing & systems

- given a computational system, we talk of distributed computation whenever at least two computational processes have a different *spatial* context
- a distributed system is a computational system performing distributed computations

Parallel, Concurrent, Distributed Systems: Recap VI

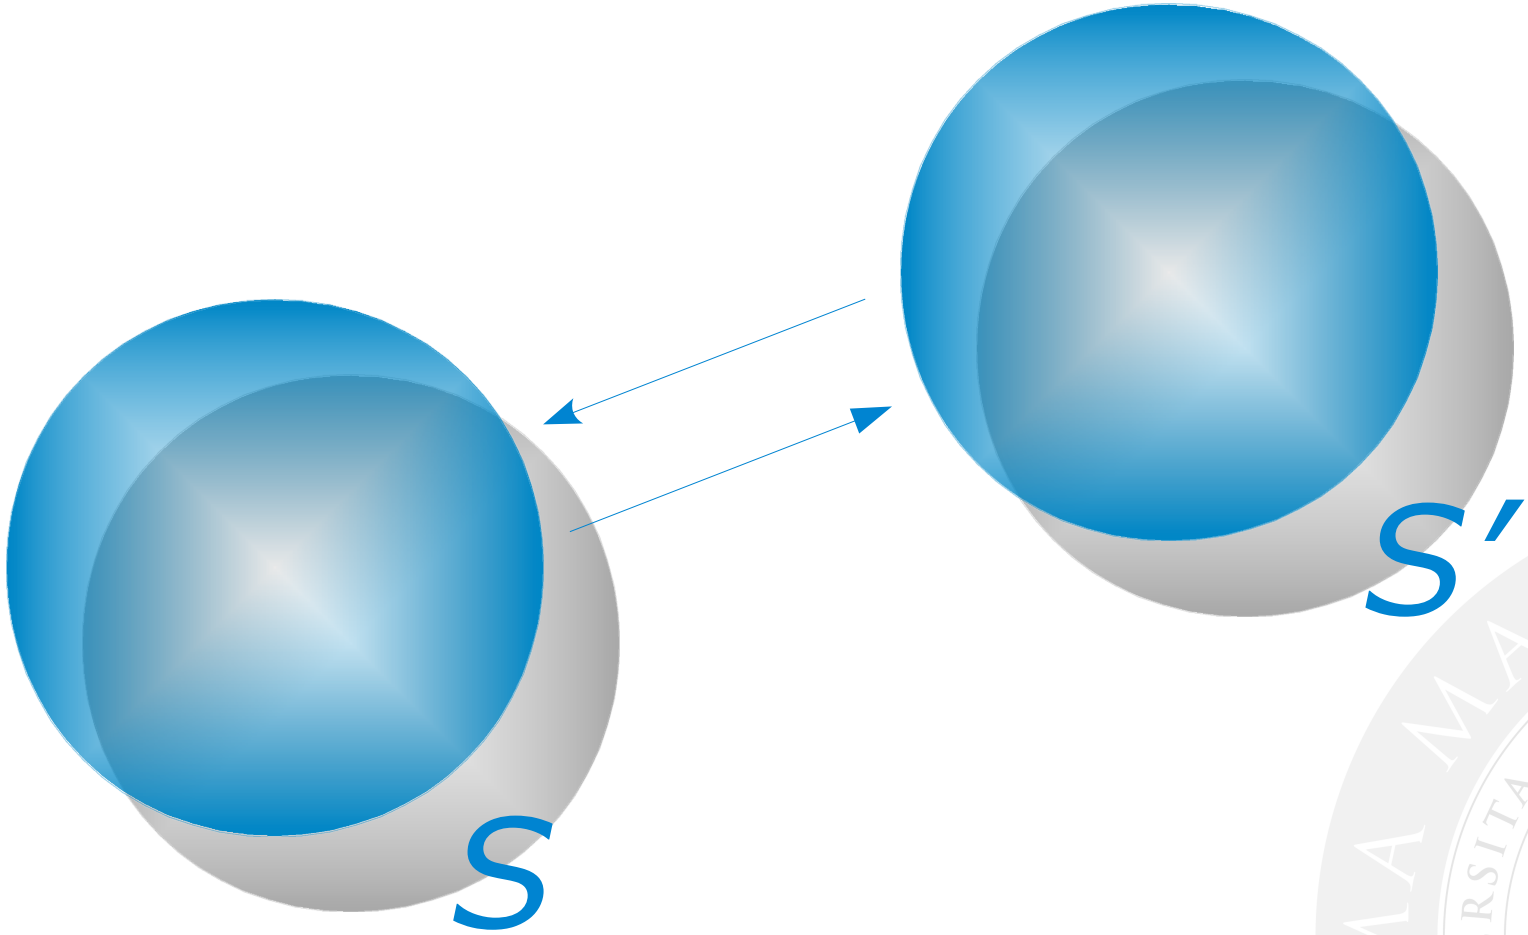

Distributed computing: different spatial contexts $S \neq S'$ for different processes

Spatial vs. Temporal Contexts: Recap I

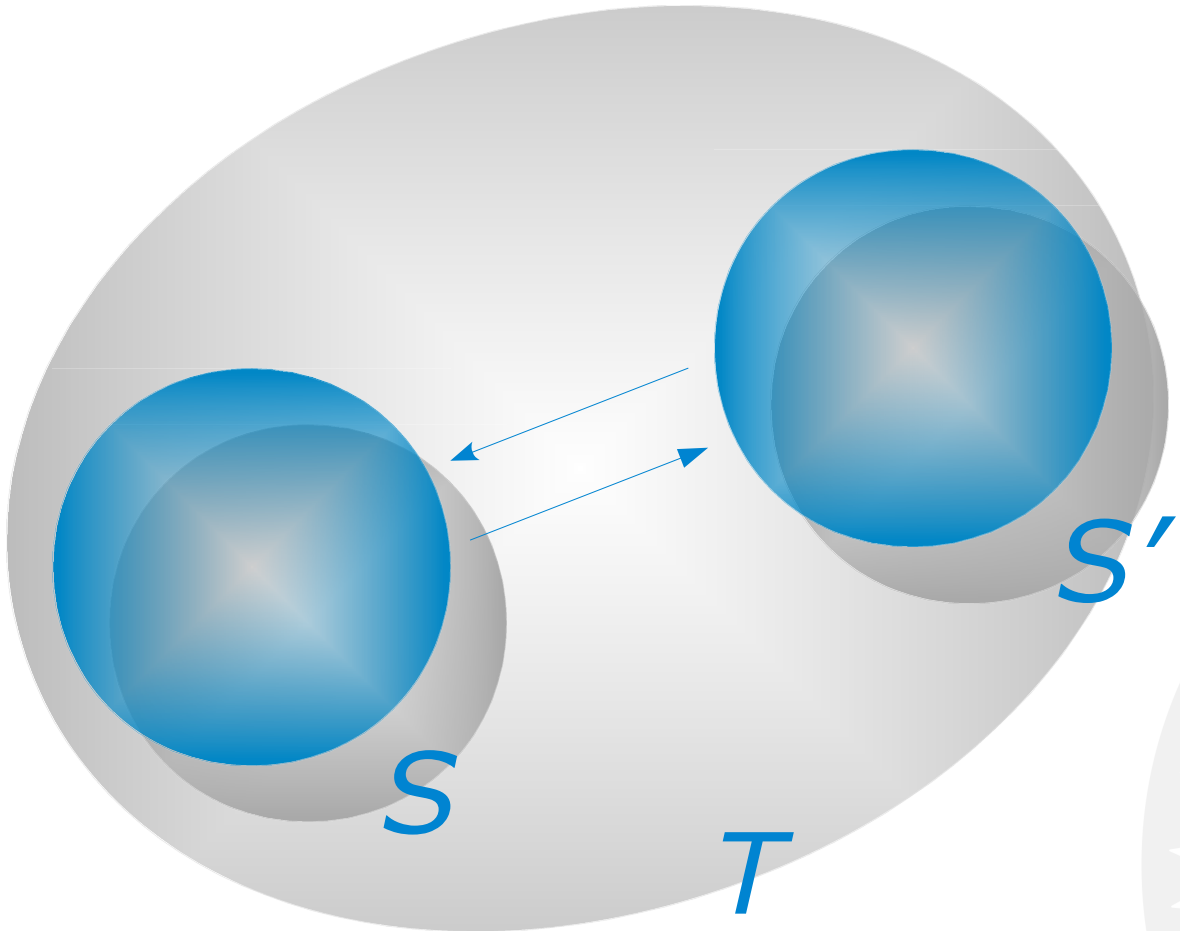

Distributed parallel computing: $S \neq S'$, same T

Spatial vs. Temporal Contexts: Recap II

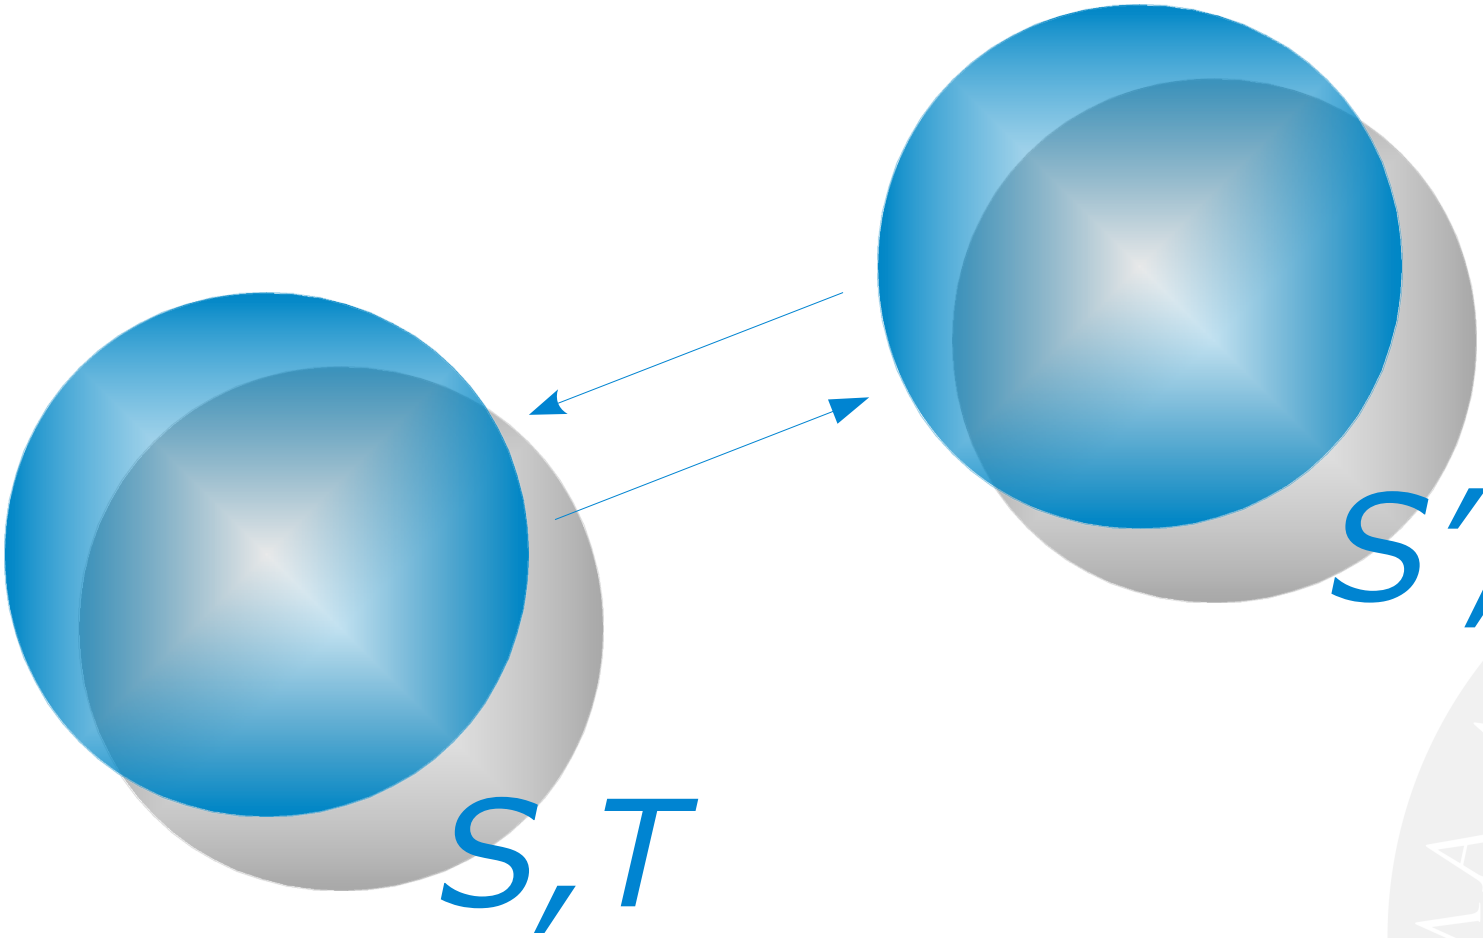

Distributed concurrent computing: $S \neq S'$, $T \neq T'$

Basic Questions

- is there any useful and well-founded notion of *time* in a distributed system?
- is there any coherent notion of *global* time in a distributed system?
- if not, what can we do about this?
- if not, how can we *synchronise* activities within a distributed system?

Synchronous

`http://www.oxforddictionaries.com/definition/english/synchronous`

synchronous,

1. *Existing or occurring at the same time*

   ...

The Issue of Time

Time in distributed systems

- in non-distributed systems, time is unambiguous
- in a distributed system, there is no *natural* notion of time
  - distributed concurrent computing is the most natural and general computational model for distributed systems

? is it *possible* to build up a global notion of time in any distributed system?

? is it *useful* to build up a global notion of time in any distributed system?

Focus on...

1. Time & Computation
2. Time in Distributed Systems
   - Physical Time
   - Logical Time
3. Toward Coordination

Physical Clocks I

Timers

- a clock in a computer is actually a *timer*—typically, an oscillating quartz with a *counter* and a *holding register*
- when the counter gets to zero, an interrupt is generated, and the counter is reloaded from the holding register
- each interrupt is a clock tick

Physical Clocks II

Multiple CPUs

- no way to ensure two different crystals oscillate exactly at the same frequency
- different clocks gradually get out of synch—*clock skew* is the difference in time, *clock drift* the phenomenon observed
- need for synchronising algorithms
- two approaches
  - *global absolute time*
  - *global relative time*

Global Absolute Time I

Absolute time

- absolute time is handled by BIH (Bureau International de l'Heure) in Paris
- expressed in terms of Universal Coordinated Time (UTC)
- broadcasted as a short radio pulse (WWV) by NIST (National Institute of Standard Time) every UTC second, and by satellites providing UTC service
- if one machine in the system has access to an UTC service, an algorithm can be used that synchronises all machines based on this

Global Absolute Time II

Example: NTP

- Network Time Protocol (NTP)
  - RFC 5905 for NTP v. 4
- a time server has the global absolute time, and other machines have to synchronise

! clocks can only run *forward* – corrections cannot bring clocks backward

Global Relative Time

Relative time

- sometimes, the only thing needed is that there is a *common time*, regardless of absolute time
- so, algorithms based on active servers polling other servers to find out the average time, and the required estimated corrections as well
- no machine is required to have UTC time

Examples

- the Berkeley Algorithm: time daemons in all machines poll and respond to each other, and agree on a common time

  [Gusella and Zatti, 1989]
- Reference Broadcast Synchronisation (RBS): global relative time in wireless networks [Elson et al., 2002]

Focus on...

1. Time & Computation
2. Time in Distributed Systems
   - Physical Time
   - Logical Time
3. Toward Coordination

Physical vs. Logical Time

Physical time not always needed

- till now, we have implicitly assumed that synchronisation is related to physical time
- however, we have also seen the case where the only actual requirement is *a shared notion of time* among the processes of a distributed system, with no need for it to be exactly the "real" time
- as a step further, we may observe that often the only need for a distributed system is a shared clock, even *unrelated* to real time
- a notion of logical time is both possible and useful

Logical Clocks [Lamport, 1978]

Synchronisation is possible with no need to be absolute

- if two processes do not interact, there is no need of synchronisation—lack of synchronisation would not be observable
- often, what really matters is not the exact time when events occur, but the *order* in which events occur
- example: UNIX `make`

Logical clocks

- synchronisation of non-physical, logical clocks is then both *admissible* and *useful*

Notation

Relation *happens-before*

- $a \rightarrow b$ represents the *relation* happens before, "a happens before b"
- it means that all processes agree that a occurs first, then b occurs
- $a \rightarrow b$ can be directly observed in two situations
  1. if a and b are events of the same process, and a comes before b, then $a \rightarrow b$ — local events are ordered by local time
  2. if a message is sent by process with an event a, and received by another process with an event b, then $a \rightarrow b$ — a message takes a finite, positive, non-zero amount of time to propagate from sender to receiver
- $a \rightarrow b$ is a transitive relation: $a \rightarrow b$, $b \rightarrow c$ imply $a \rightarrow c$
- *happens-before* defines a partial ordering over the events in a distributed system: when neither $a \rightarrow b$ nor $b \rightarrow a$ can be observed, then nothing can be said on their ordering — a and b are said to be concurrent

Logical Time

Measuring time with logical clocks: *time values*

- a shared notion of time for an event a: *time value* $C(a)$ is such that *every process agrees* upon it
- time value should be thought as the value of a logical clock upon which processes agree
- time values are such that $a \rightarrow b$ implies $C(a) < C(b)$ — that is, time values should be assigned so that $C(a) < C(b)$
  1. if a and b are events of the same process, and a comes before b, then $C(a) < C(b)$
  2. if a message is sent by process with an event a, and received by another process with an event b, then $C(a) < C(b)$
- since neither physical nor logical clocks can run backward, any correction to clock time should go forward (increasing), never backward (decreasing)

Lamport's Algorithm I

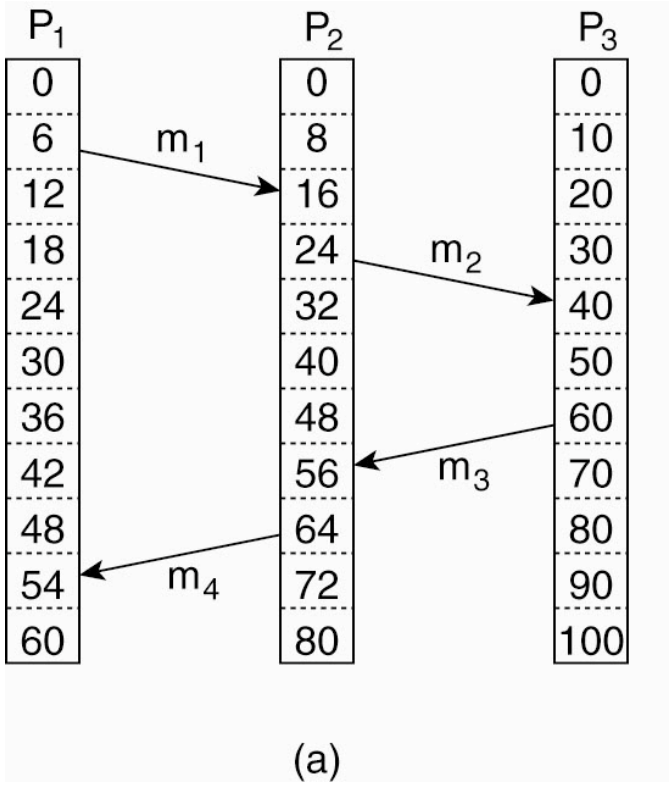

Concurrent message transmission using logical clocks

[Tanenbaum and van Steen, 2007]

Lamport's Algorithm II

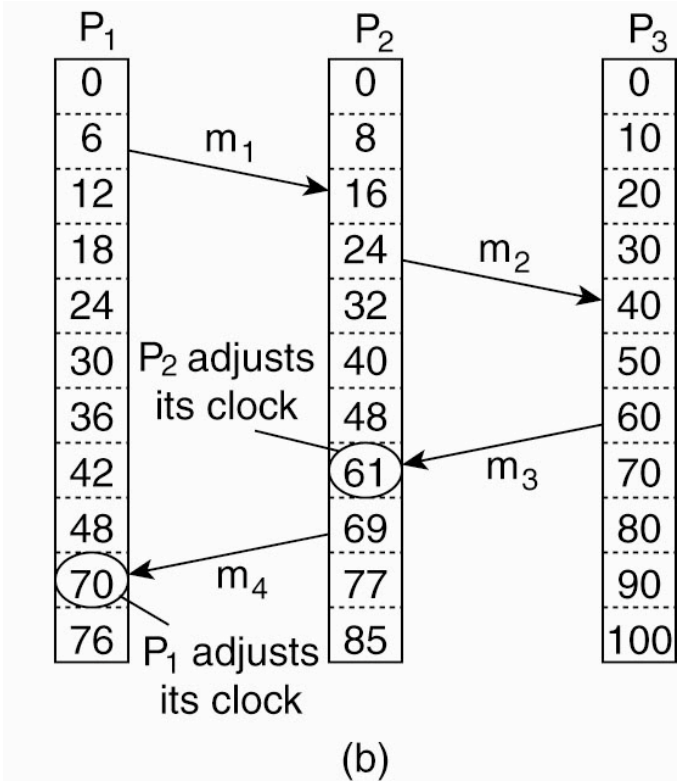

Lamport's algorithm corrects the clocks

[Tanenbaum and van Steen, 2007]

Lamport's Algorithm III

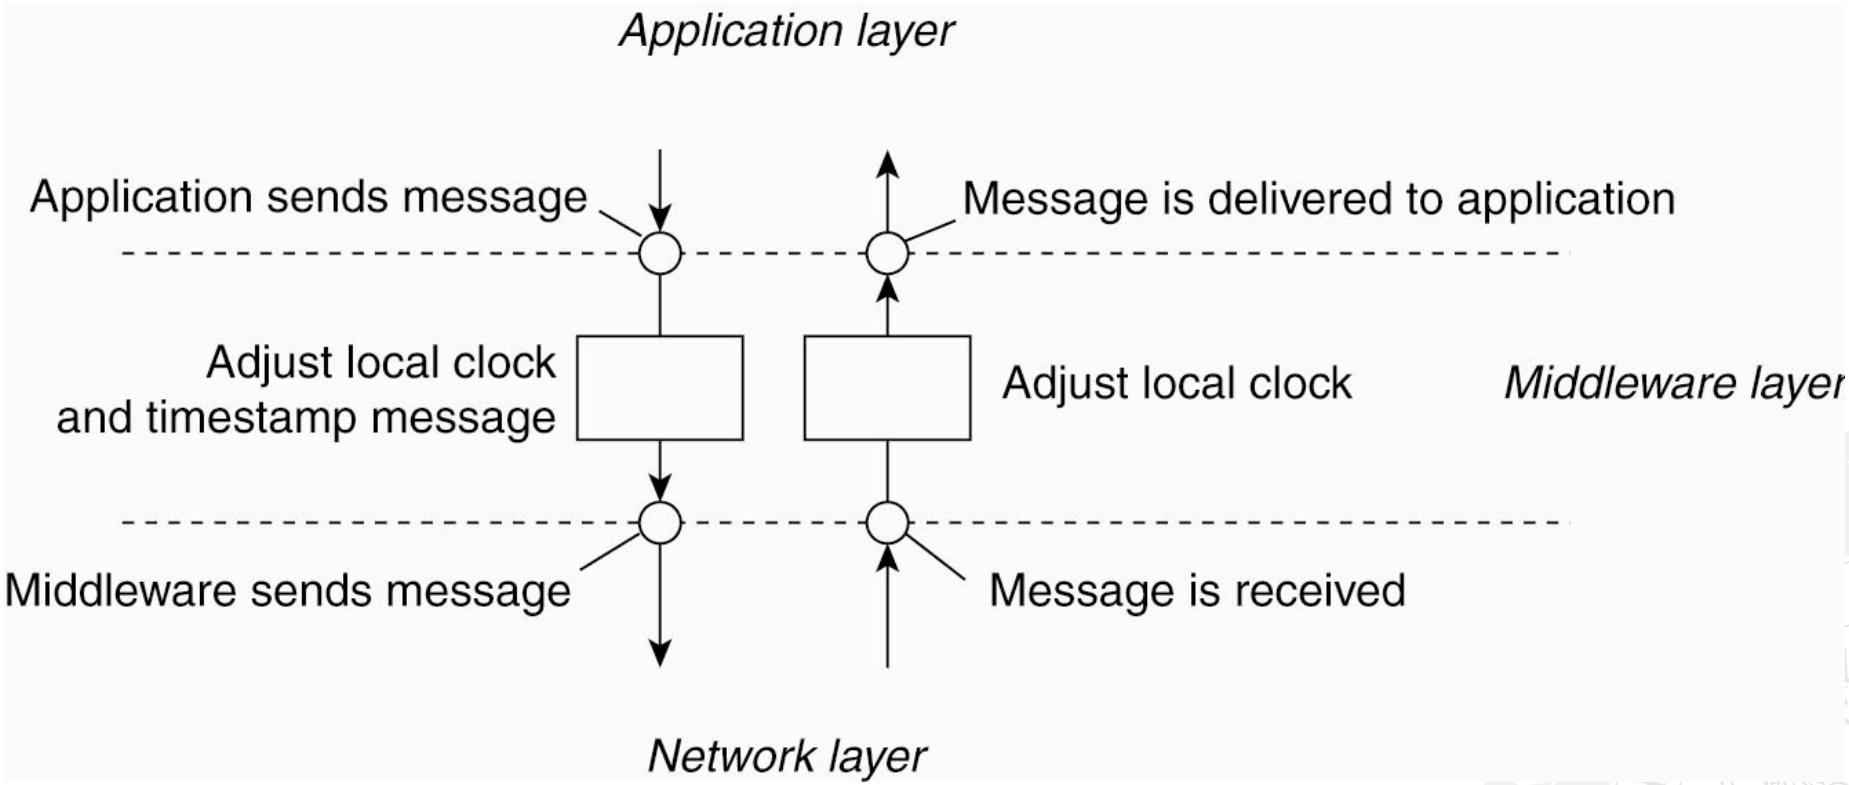

Middleware support for Lamport's logical clocks

[Tanenbaum and van Steen, 2007]

Lamport's Algorithm IV

Implementation of Lamport's logical clocks

- each process P_i maintains a *local* counter C_i
- local counters are updated following three steps
  1. before executing an event, P_i executes $C_i \leftarrow C_i + 1$
  2. when sending a message m to P_j, process P_i sets m's timestamp $ts(m)$ to C_i after updating its counter (see step above)
  3. upon reception of a message m, process P_j adjusts its local counter such that $C_j \leftarrow max(C_j, ts(m))$, then goes back to step (1) and delivers the message to the application

! sometimes, it is required that no two events occur exactly at the same time – process label can be added as a decimal number to the timestamp

Lamport's Algorithm V

Distributed implementation of global time

- C_i is local time at process P_i
- a is an event in a distributed system
- $\forall a \in P_i, C \leftarrow C_i(a)$

$\rightarrow$ C is the *global time* for the distributed system

An Example I

Totally-ordered multicast

- a replicate database exists of the accounts of a bank in LA and NY
- a customer adds $100 to his account, while at the same time a bank employee applies a 1% increment to the account
- given that the original account contained $1000, it may easily happens that, say, the LA replica records $1110, the NY one $1111

→ inconsistency due to concurrent updates over a distributed replicated database

An Example II

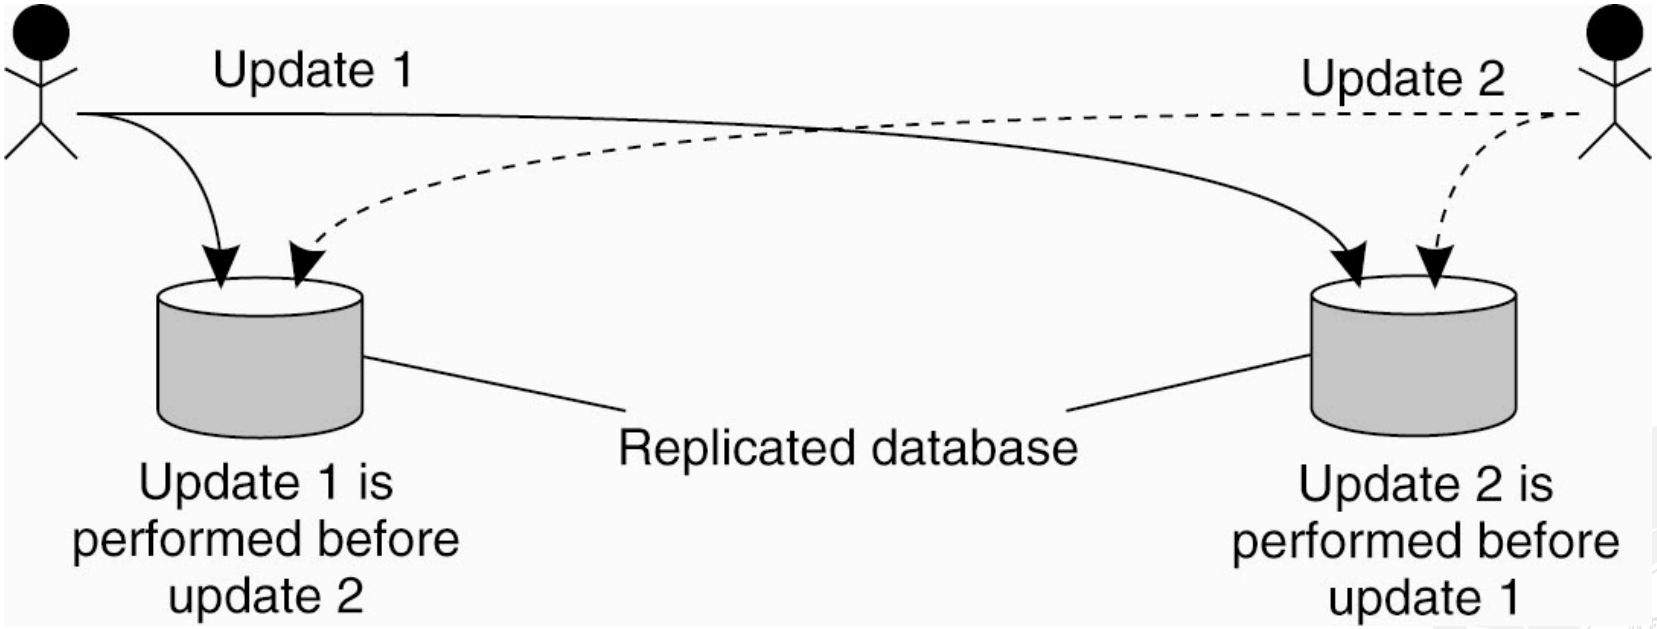

Inconsistency in a replicated database after two concurrent updates

[Tanenbaum and van Steen, 2007]

The Solution: Totally-ordered Multicast I

Assumptions

- a group of processes multicasting each other
- each messaged is timestamped by the sender with its local logical time
- also the sender conceptually receives the multicasted message
- messages from the same sender are received in the same order they are sent, and no message is lost

The Solution: Totally-ordered Multicast II

Algorithm

- each process maintains a local queue of all messages received, ordered according to its timestamp
- every message received is acknowledge with a multicasted message, timestamped according to Lamport's algorithm

→ timestamp of a received message is lower than the timestamp of the acks

→ every process has essentially the same queue

- only when all acknowledgements have been received, the middleware can deliver a queued message to the application
- since all queues are equal, all messages are delivered to the application level at the same time on all the machines in the distributed system

The Solution: Totally-ordered Multicast III

Result

- a totally-ordered multicasting is perceived at the application level — as provided by the middleware layer
- in the example above, either the client or the employee command is issued first on all replicas

→ all replicas will be consistently updated

→ no idea, however, on whether the final record will be \$1110 or \$1111...

Vector Clocks I

The problem

- in essence, $a \rightarrow b$ implies $C(a) < C(b)$, whereas $C(a) < C(b)$ does not imply $a \rightarrow b$
  - so that, for instance, time values could be totally ordered when events are not
  - when events are unrelated, comparison of time values is meaningless
- Lamport's logical clocks say nothing about that
- something more is needed
- to say in particular whether a and b are (un)related

Vector Clocks II

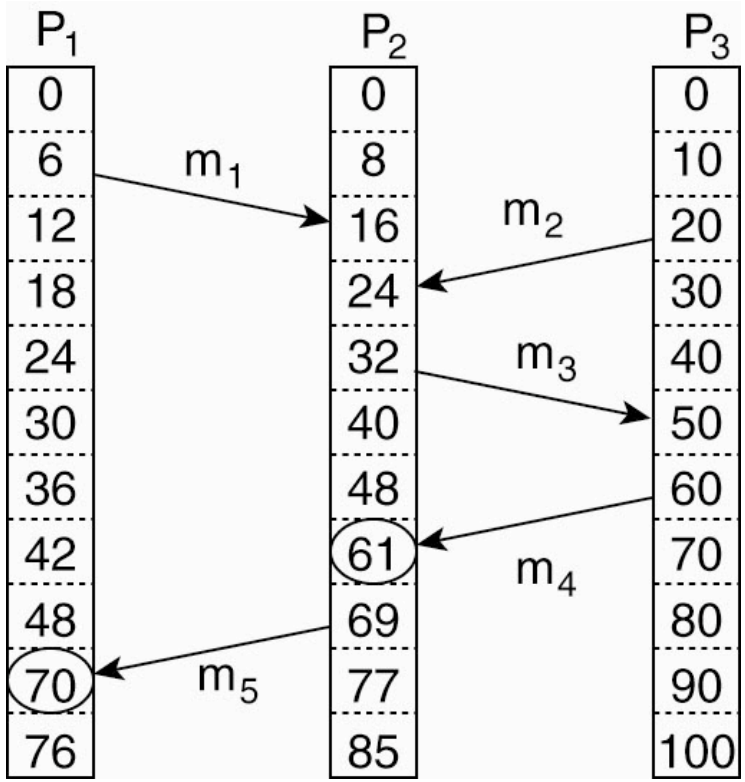

Concurrent message transmission using logical clocks

[Tanenbaum and van Steen, 2007]

Vector Clocks III

Causality

- $m1$ is received before $m2$ is sent, according to Lamport's clock: can we conclude anything about $m1$ and $m2$?
- in general, the problem is that Lamport's clocks do not capture *causality*
- *vector clocks* capture causality

Vector Clocks IV

Definition

- a vector clock $VC(a)$ assigned to an event a is such that $\exists b, VC(a) < VC(b) \longrightarrow a$ causally precedes b
- each process P_i maintain a vector VC_i such that
  - $VC_i[i]$ is the number of events occurred so far at P_i — basically, the logical clock of P_i
  - ← Every new event occurring in P_i increments $VC_i[i]$
  - $VC_i[j] = k$ means that P_i knows that k events have occurred at P_j — basically, the logical clock of P_j according to P_i's best knowledge
- ← every message from P_i is timestamped with vector VC_i

Vector Clocks V

Algorithm

- before any event is executed at P_i, $VC_i[i] \leftarrow VC_i[i] + 1$
- a message m from P_i to P_j timestamped with vector VC — $ts(m) = VC$
- a message m received by P_j makes it adjust VC_j such that $\forall k, VC_j[k] \leftarrow max(VC_j[k], ts(m)[k]$ — then m is delivered up to the application level

Vector Clocks VI

Result

- every process knows how many events have preceded the sending of the received message at the sender process—information about the "chain of events" is preserved and shared among processes
- each $ts(m)[i]$ refers to the events causally preceding m within each process P_i
- $ts(m)$ tells how many events may causally precede the sending of m, on which m may causally depend

$\rightarrow$ as a result, for instance, the delivery of a message to the application level could be suspended until all preceding messages from the same source are received

Enforcing Causal Communication

Causally-ordeded multicasting

- by using vector clocks, a message could be delivered only when all messages causally preceding it have been received
- . . . assuming that all messages are multicasted in a group

! weaker than totally-ordered multicasting: if two messages are not causally related, they could be delivered to applications in any order

Next in Line...

1 Time & Computation

2 Time in Distributed Systems

3 Toward Coordination

Communication & Interaction in Distributed System

Communication is just half of the story

- interaction is a more general issue
- governing (inter)action is a fundamental issue in (distributed) systems
- doing the right thing at the right time is essential
- "at the right *time*" is the critical problem

Beyond Synchronisation I

Ordering events is not enough

- sometimes, more articulated policies are required
- for instance, to ensure that concurrent accesses to a shared resource could harm its consistency, or corrupt it

Mutual exclusion

- a number of algorithms — centralised, decentralised, distributed — for instance, Token Ring
- we do not review them here
- the main point: some of them are based on a coordinator, all of them are *coordination* algorithms

Beyond Synchronisation II

Election algorithms

- many distributed algorithms requires a *coordinator* to be elected
- again, we do not review them: election algorithms are (used by) coordination algorithms

It is not merely a matter of time

- synchronisation is about *when* things happen
- actions are more than sending messages
- interaction does not merely translate into suitably-ordered distributed actions — undifferentiated actions
- actions have a nature, and meaningful interaction within a distributed system typically depends on such a nature

Beyond Synchronisation III

The problem of coordination

- governing interaction based both on time, and on the nature of actions, and aimed at the achievement of some global objective for the distributed system
- this is the problem of *coordination*

Summing Up

Time in distributed systems

- the issue of time
- physical time / clock
- logical time / clock
- causality and vector clocks

Toward coordination

- what do we do when we have some coherent notion of time?
- coordinators and distributes algorithms

References I

Atzori, L., Iera, A., and Morabito, G. (2010).
The Internet of Things: A survey.
Computer Networks, 54(15):2787–2805.

Burks, A. W., Goldstine, H. H., and von Neumann, J. (1982).
Preliminary discussion of the logical design of an electronic computing instrument.
In Randell, B., editor, *The Origins of Digital Computers*, Texts and Monographs in Computer Science, pages 399–413. Springer Berlin Heidelberg, Berlin, Heidelberg.

Elson, J., Girod, L., and Estrin, D. (2002).
Fine-grained network time synchronization using reference broadcasts.
ACM SIGOPS Operating Systems Review, 36(SI):147.

Gubbi, J., Buyya, R., Marusic, S., and Palaniswami, M. (2013).
Internet of Things (IoT): A vision, architectural elements, and future directions.
Future Generation Computer Systems, 29(7):1645–1660.

Gusella, R. and Zatti, S. (1989).
The accuracy of the clock synchronization achieved by TEMPO in Berkeley UNIX 4.3BSD.
IEEE Transactions on Software Engineering, 15(7):847–853.

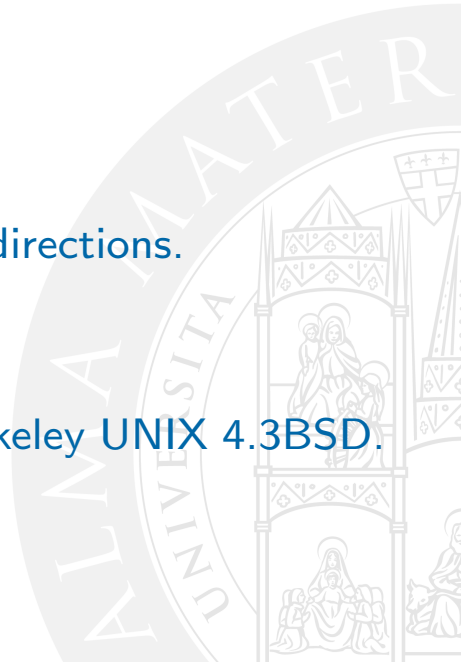

References II

Lamport, L. (1978).
Time, clocks, and the ordering of events in a distributed system.
Communications of the ACM, 21(7):558–565.

Lee, E. A. (2009).
Computing needs time.
Communications of the ACM, 52(5):70–79.

Tanenbaum, A. S. and van Steen, M. (2007).
Distributed Systems. Principles and Paradigms.
Pearson Prentice Hall, Upper Saddle River, NJ, USA, 2nd edition.

Turing, A. M. (1937).
On computable numbers, with an application to the entscheidungsproblem.
Proceedings of the London Mathematical Society, s2-42(1):230–265.

Computing with Time

Distributed Systems

Sistemi Distribuiti

Andrea Omicini

andrea.omicini@unibo.it

Dipartimento di Informatica – Scienza e Ingegneria (DISI)

Alma Mater Studiorum – Università di Bologna a Cesena

Academic Year 2016/2017