Distributed Systems

ID2201

time

Johan Montelius

the clock is not enough

- In an asynchronous system clocks can note be completely trusted.
- Nodes will not be completely synchronized.
- We still need to
  - talk about before and after
  - order events
  - agree on order

Logical time

- All events in one process are *ordered*.
- The sending of a message occurs *before* the receiving of the message.
- Events in a distributed system are *partially ordered*.
- The order is called “*happened before*”.
- Logical time gives us a tool to talk about ordering without having to synchronize clocks.

Partial order

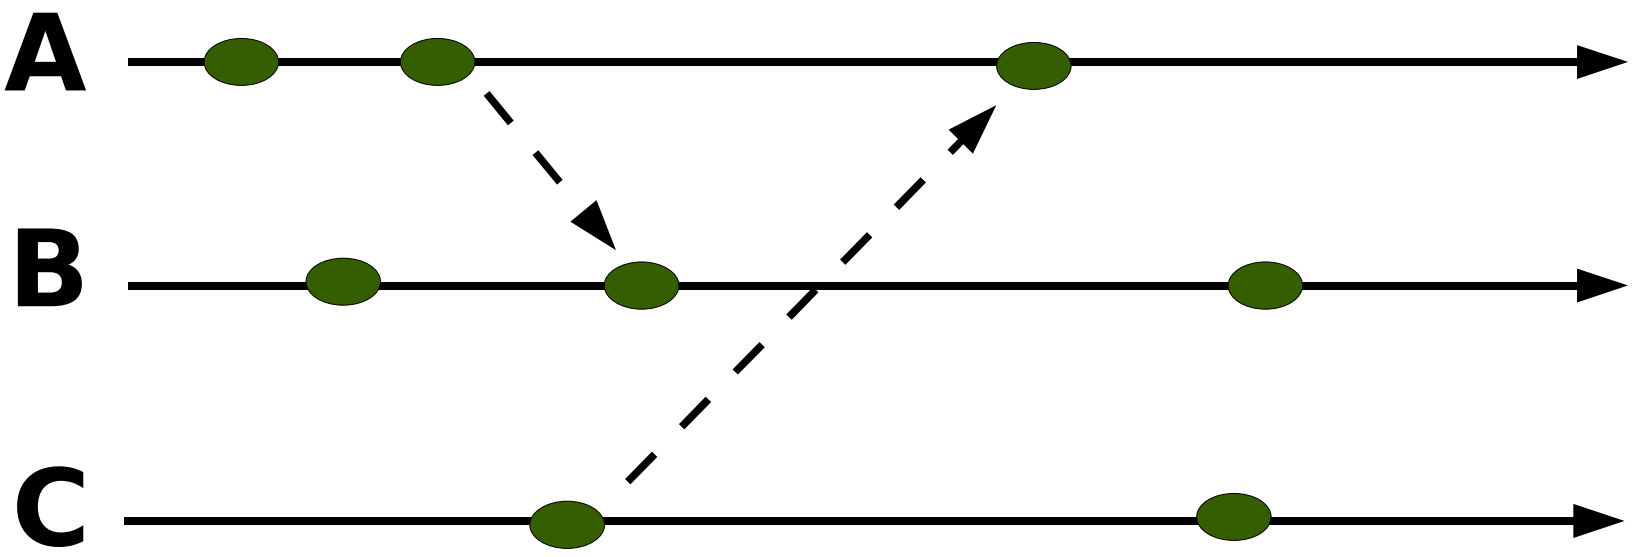

Ordering does not imply causality!

Logical clock - Lamport

- Each process will increment its clock before time stamping an event.
- All messages are tagged with the senders time stamp.
- A receiver will first set its clock to the greatest of its internal clock and the time stamp of the message before setting the time stamp of the receive event.

Logical clock

- If e_1 *happened before* e_2 then the time stamp of e_1 is less than the time stamp of e_2.

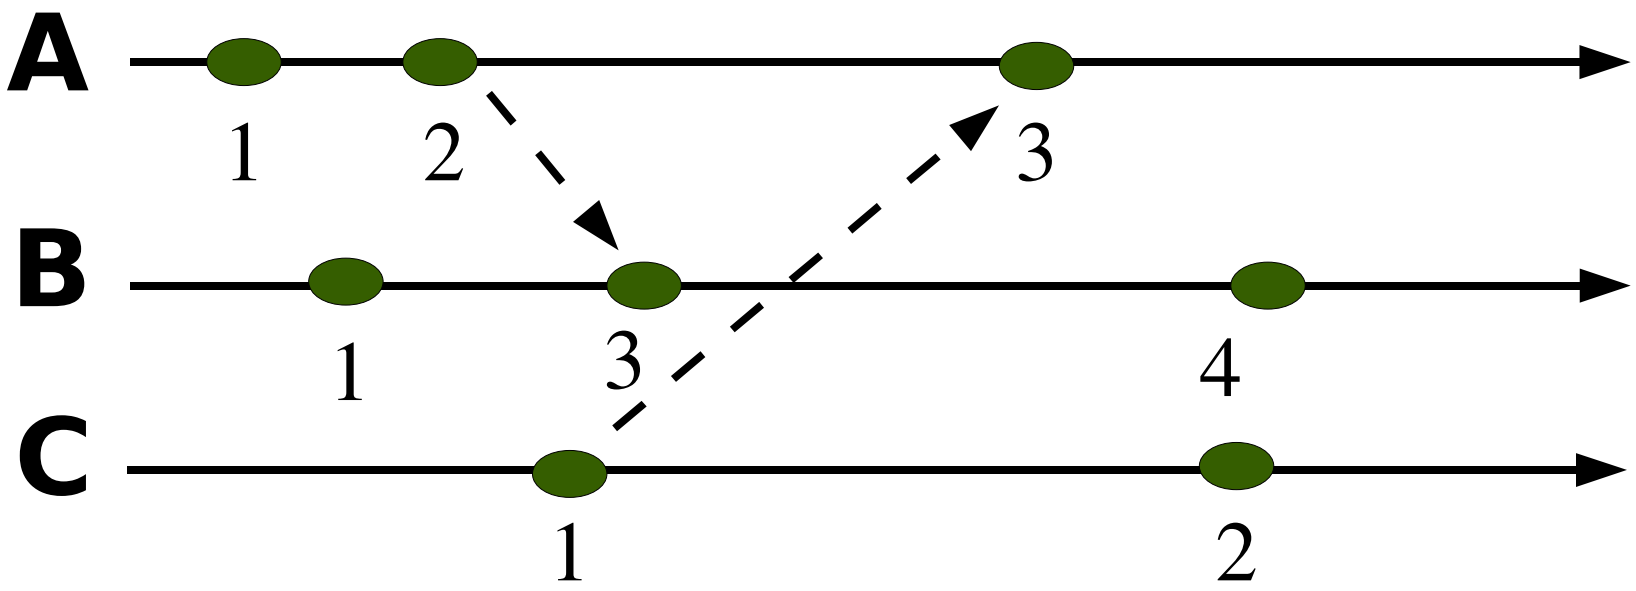

Logical clock

- What do we know if the time stamp of f e_1 is less than the time stamp of e_2?

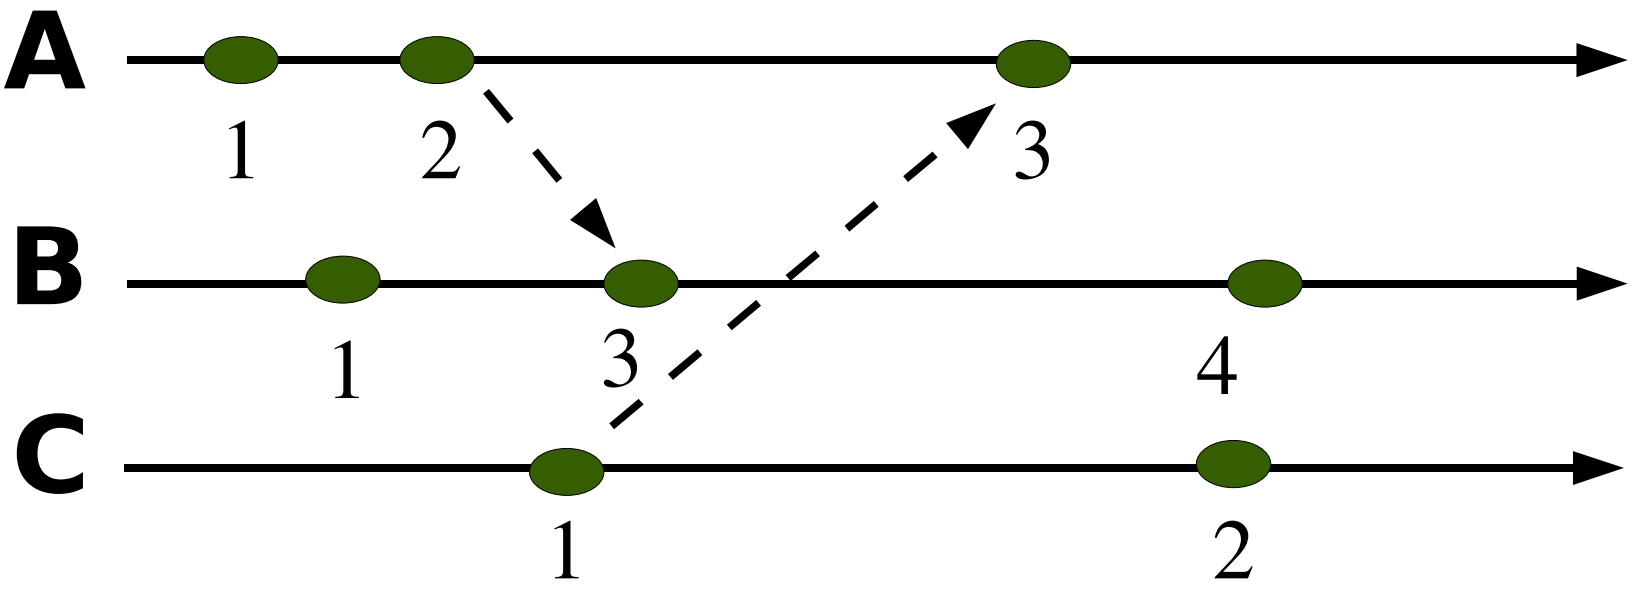

Can we do better?

- We should be able to time stamp events so that we can capture the partial order.
- We want to look at two time stamps and
  - if they are ordered
  - then the events are ordered

Vector clocks

- a vector with one clock per process
- initially set to <0,....>
- each process increment its own index
- sent messages are tagged with vector
- receiver
  - *merges* (max of each index)
    - time stamp of message
    - and own time stamp
  - increment own index

Vector clocks

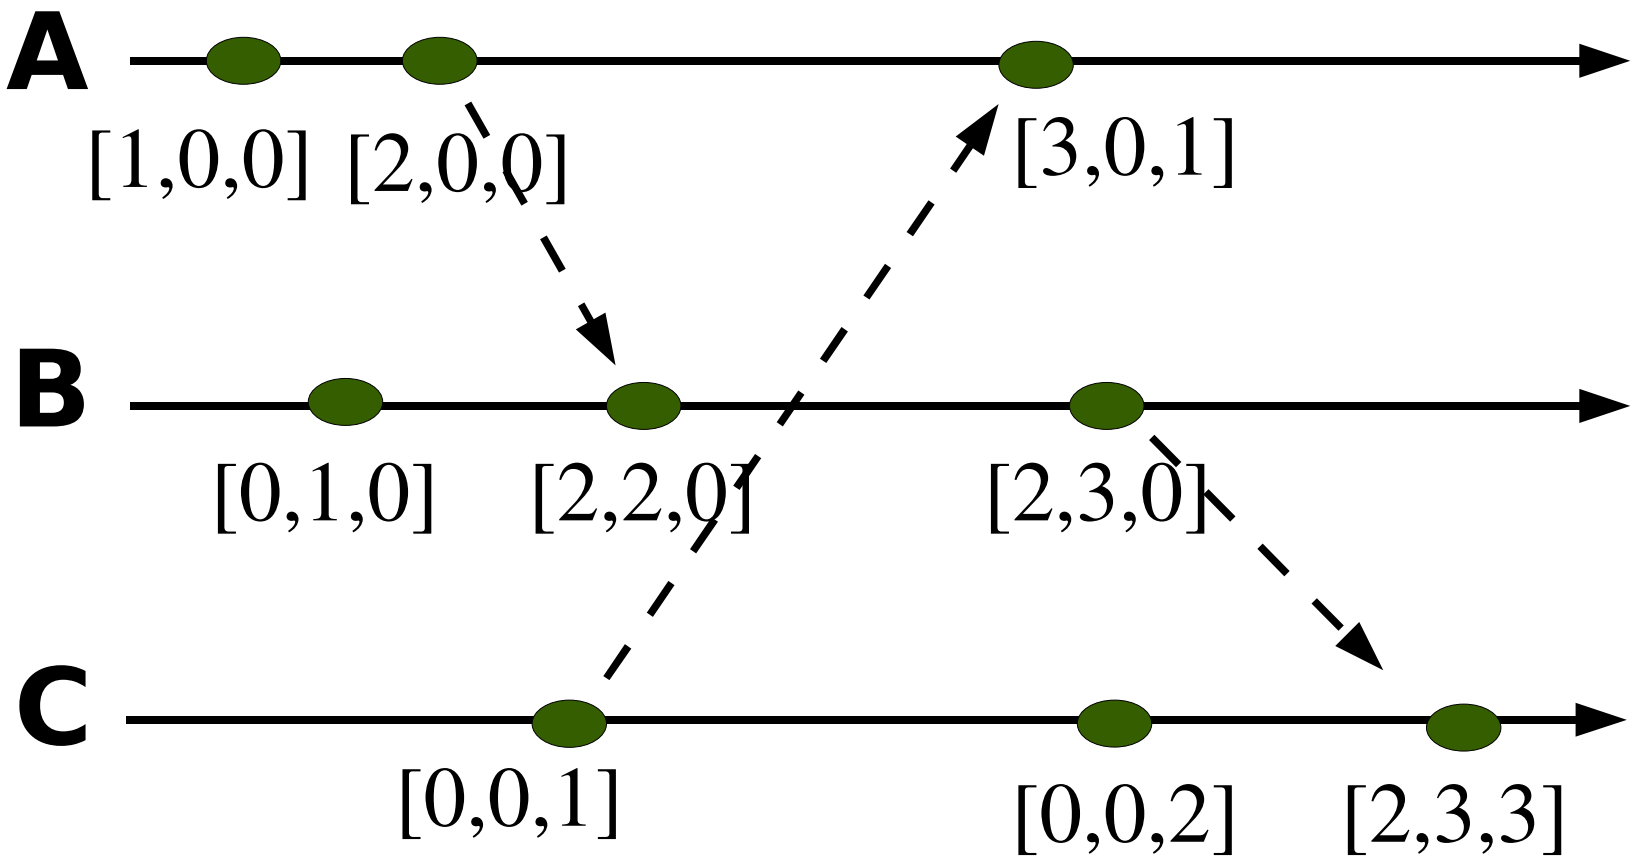

Vector clocks

- If, and only if, e_1 *happened before* e_2 then the time stamp of e_1 is less than the time stamp of e_2.

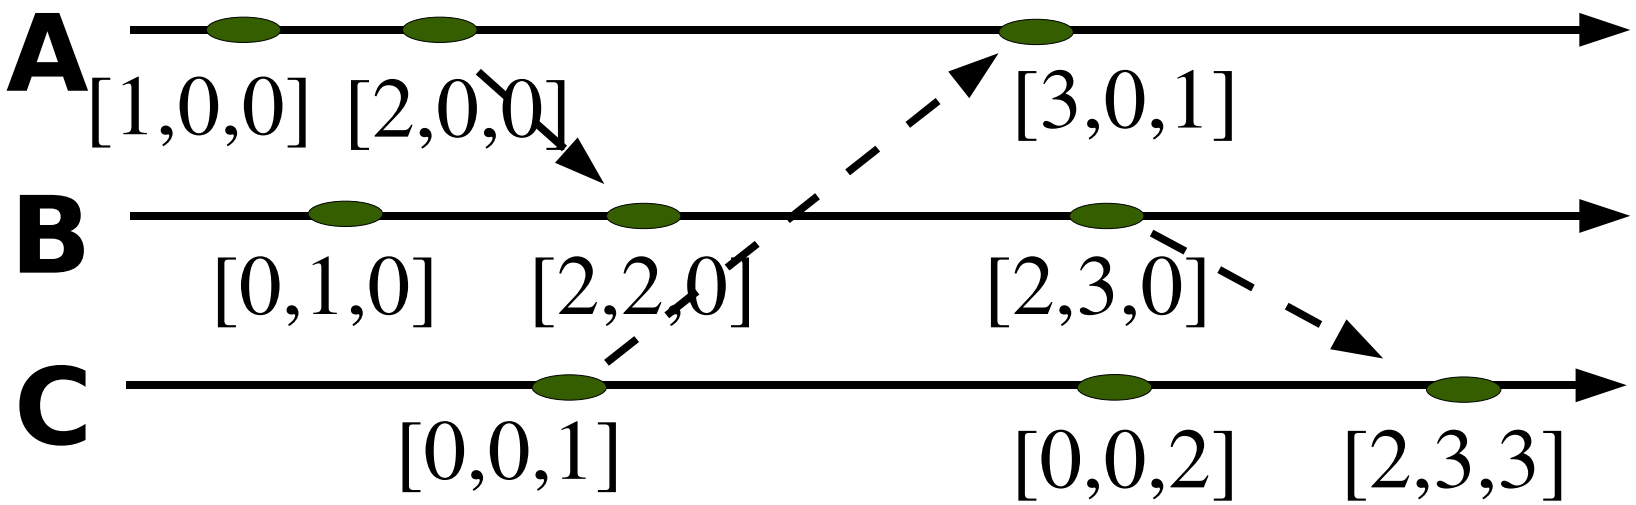

Vector clocks pros and cons

- The partial order is complete, that is we can look at the time stamp and determine if two events are ordered.
- The vectors will take up a some space and could become a problem.
- What should we do if more processes come and leave, there is no mechanism to add new clocks to the system.