CLOUD COMPUTING CONCEPTS

with Indranil Gupta (Indy)

MULTICAST

Lecture E

VIRTUAL SYNCHRONY

Virtual Synchrony or View Synchrony

- Attempts to preserve multicast ordering and reliability in spite of failures
- Combines a membership protocol with a multicast protocol
- Systems that implemented it (like Isis) have been used in NYSE, French Air Traffic Control System, Swiss Stock Exchange

Views

- Each process maintains a membership list
- The membership list is called a *View*
- An update to the membership list is called a *View Change*
  - Process join, leave, or failure
- Virtual synchrony guarantees that all **view changes are delivered in the same order at all correct processes**
  - If a correct P1 process receives views, say {P1}, {P1, P2, P3}, {P1, P2}, {P1, P2, P4} then
  - Any other correct process receives the *same sequence* of view changes (after it joins the group)
    - P2 receives views {P1, P2, P3}, {P1, P2}, {P1, P2, P4}
- Views may be delivered at different physical times at processes, but they are delivered in the same order

VSync Multicasts

- A multicast M is said to be “delivered in a view V at process P*i*” if
  - P*i* receives view V, and then sometime before P*i* receives the next view it delivers multicast M
- Virtual synchrony ensures that
  1. **The set of multicasts delivered in a given view is the same set at all correct processes that were in that view**
     - *What happens in a View, stays in that View*
  2. The sender of the multicast message also belongs to that view
  3. If a process P*i* does not deliver a multicast M in view V while other processes in the view V delivered M in V, then P*i* will be *forcibly removed* from the next view delivered after V at the other processes

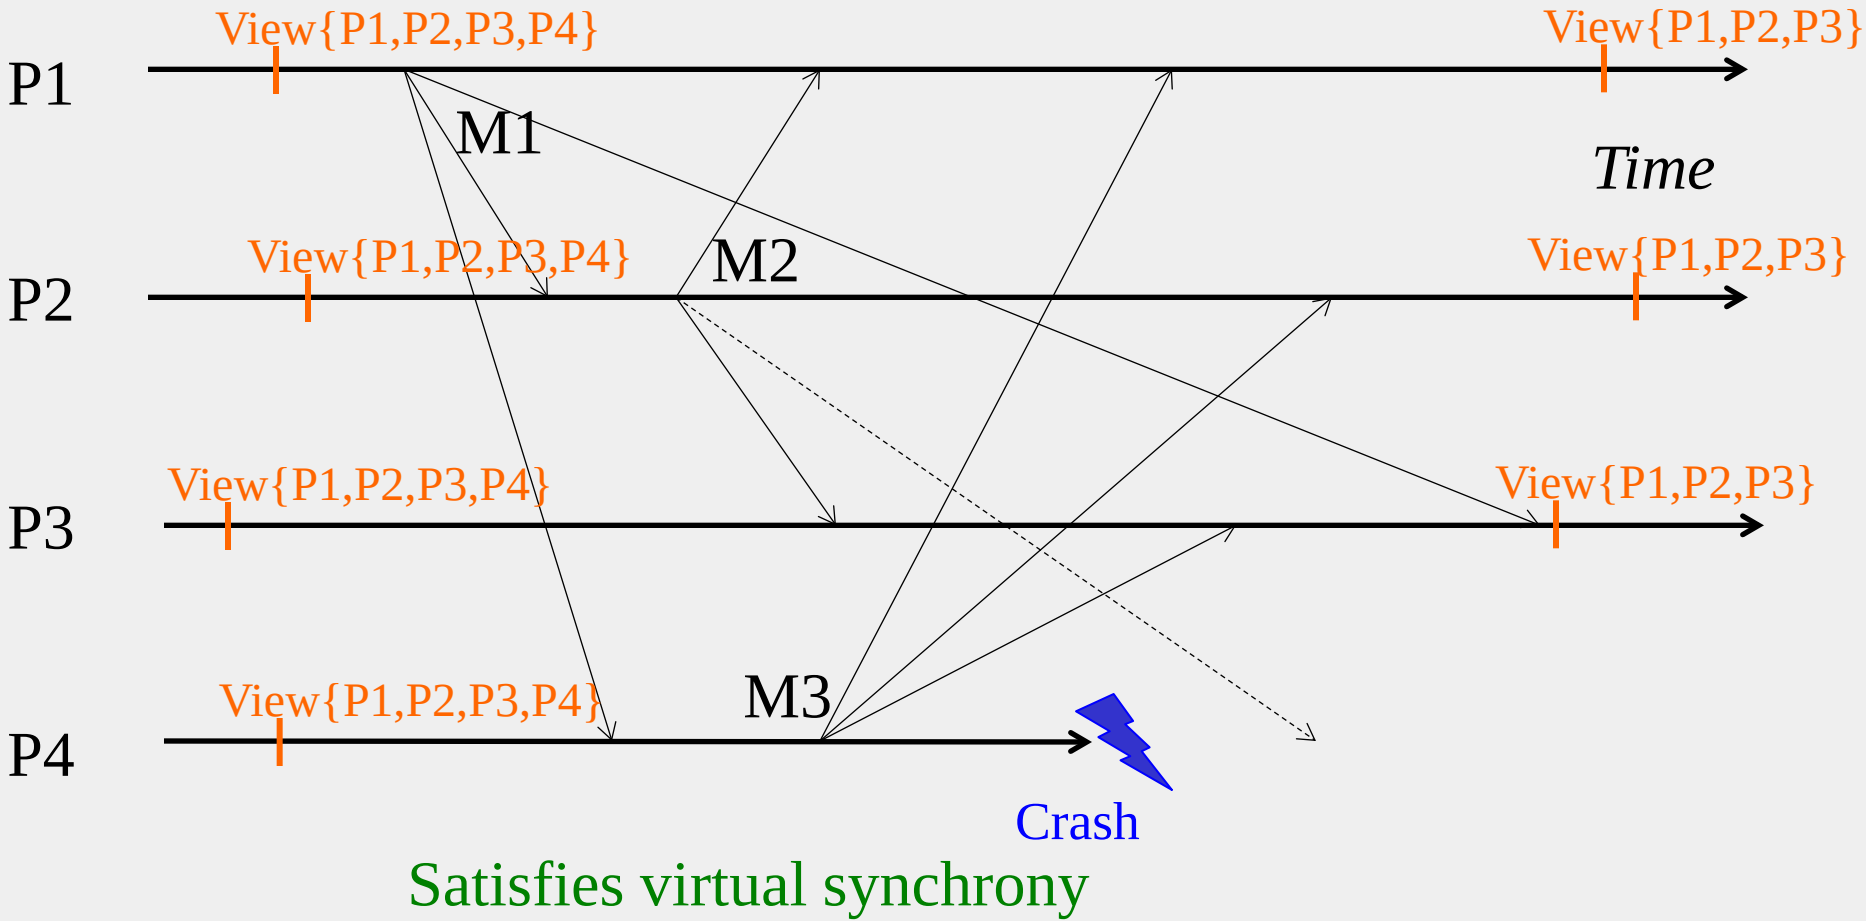

Satisfies virtual synchrony

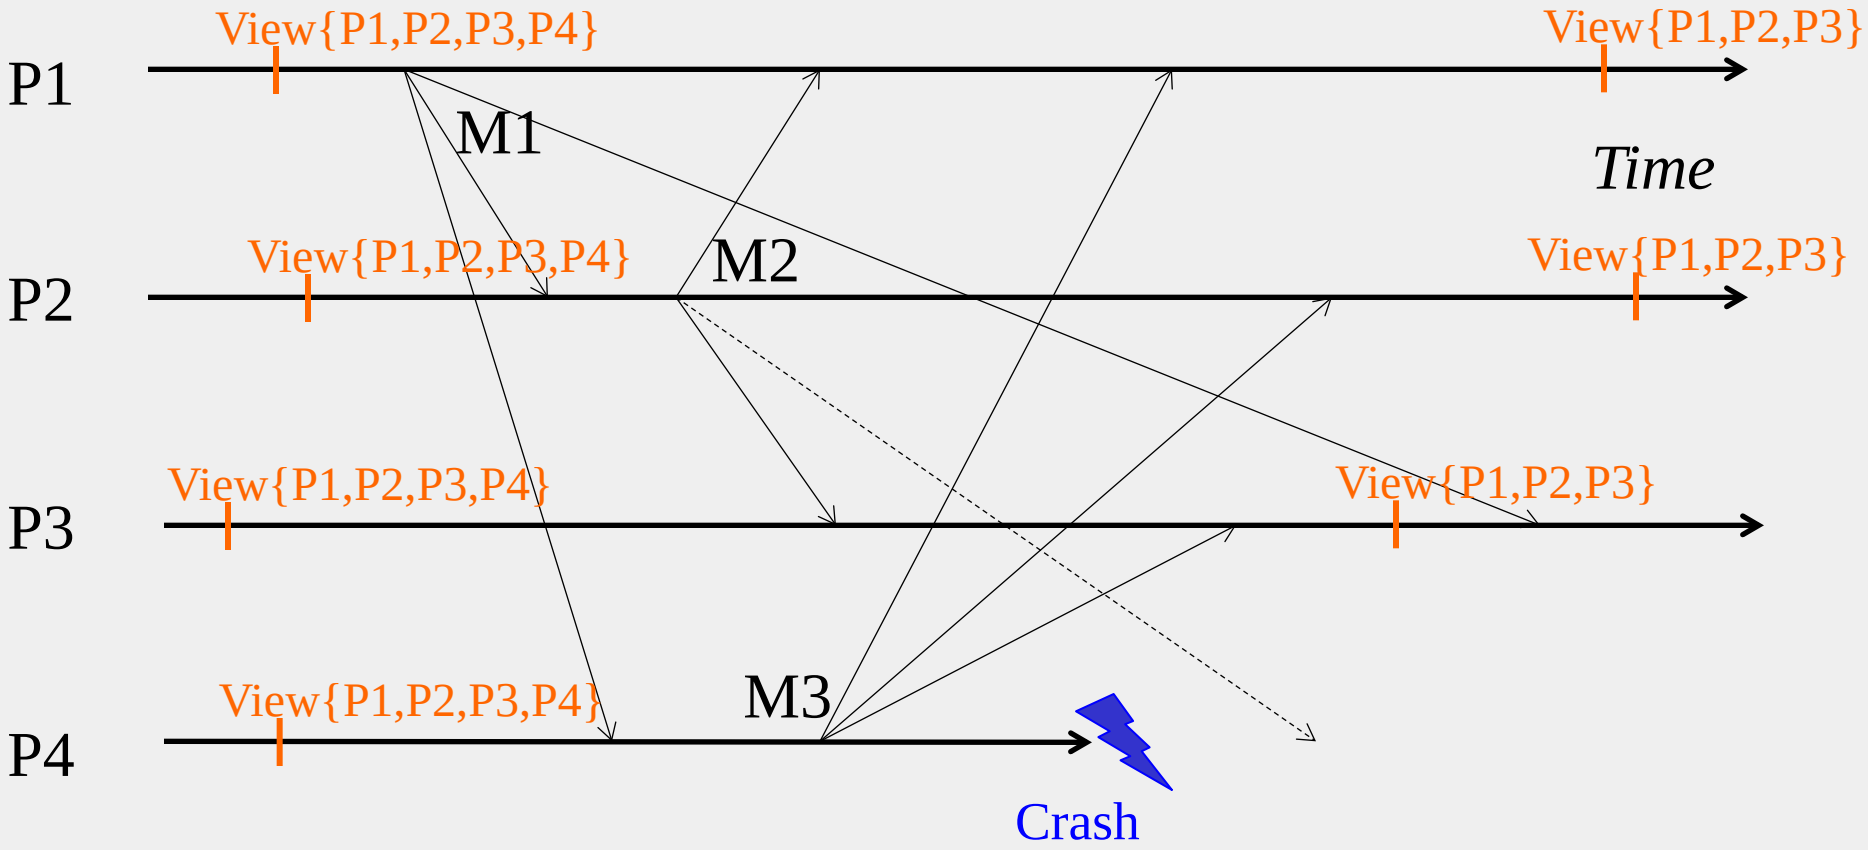

Does not satisfy virtual synchrony

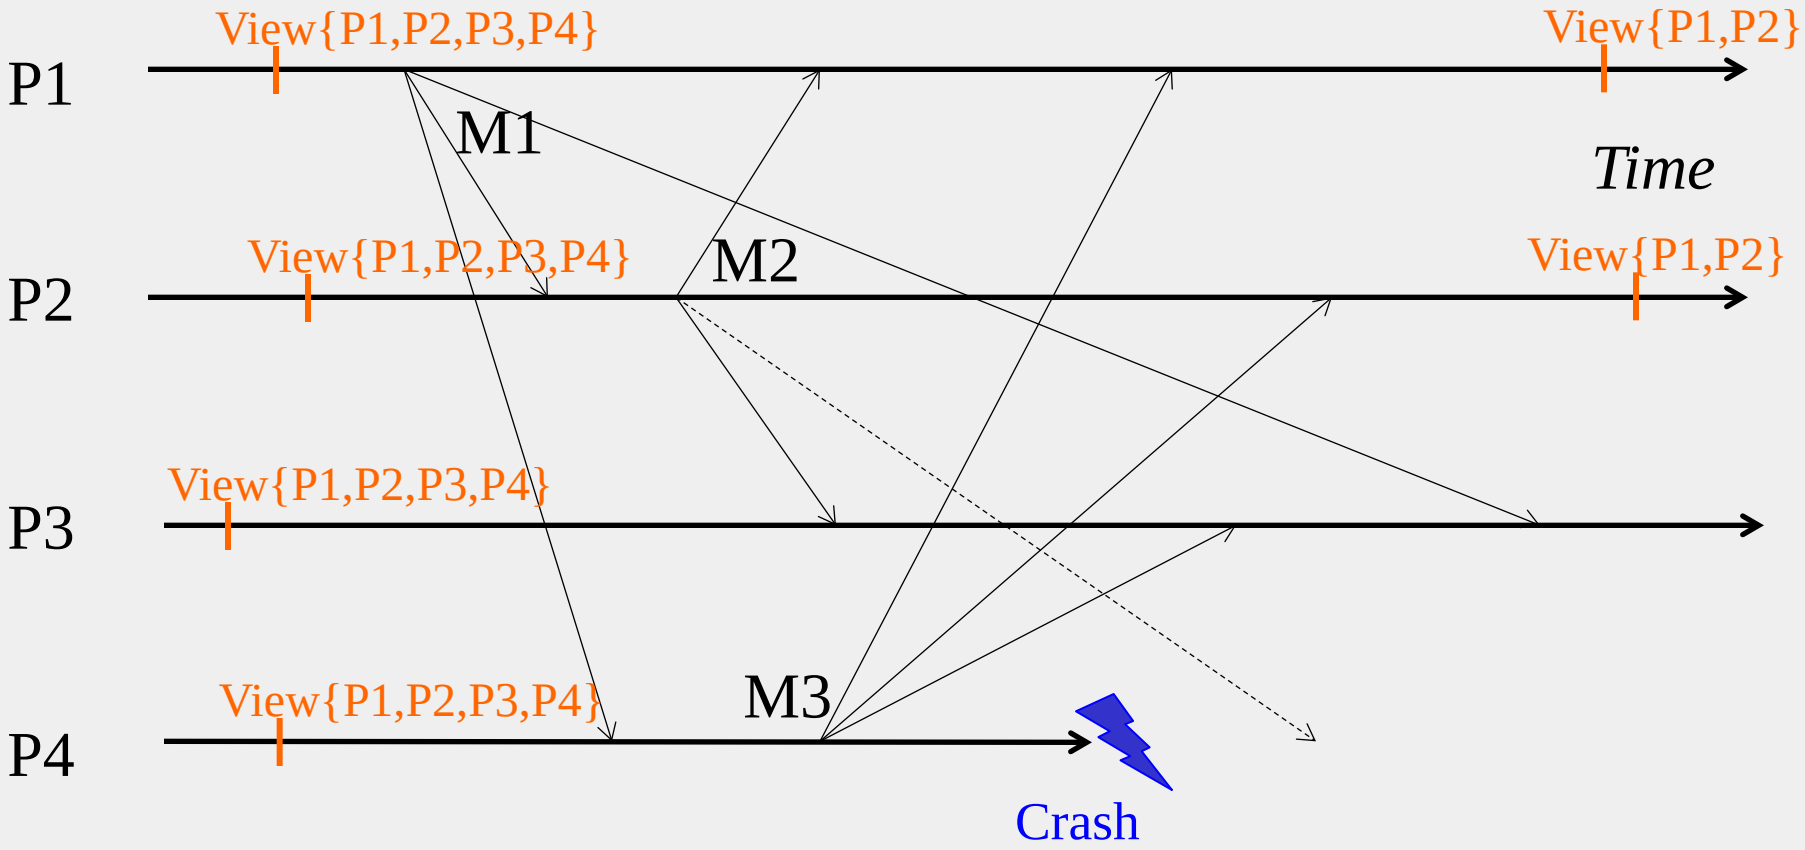

Satisfies virtual synchrony

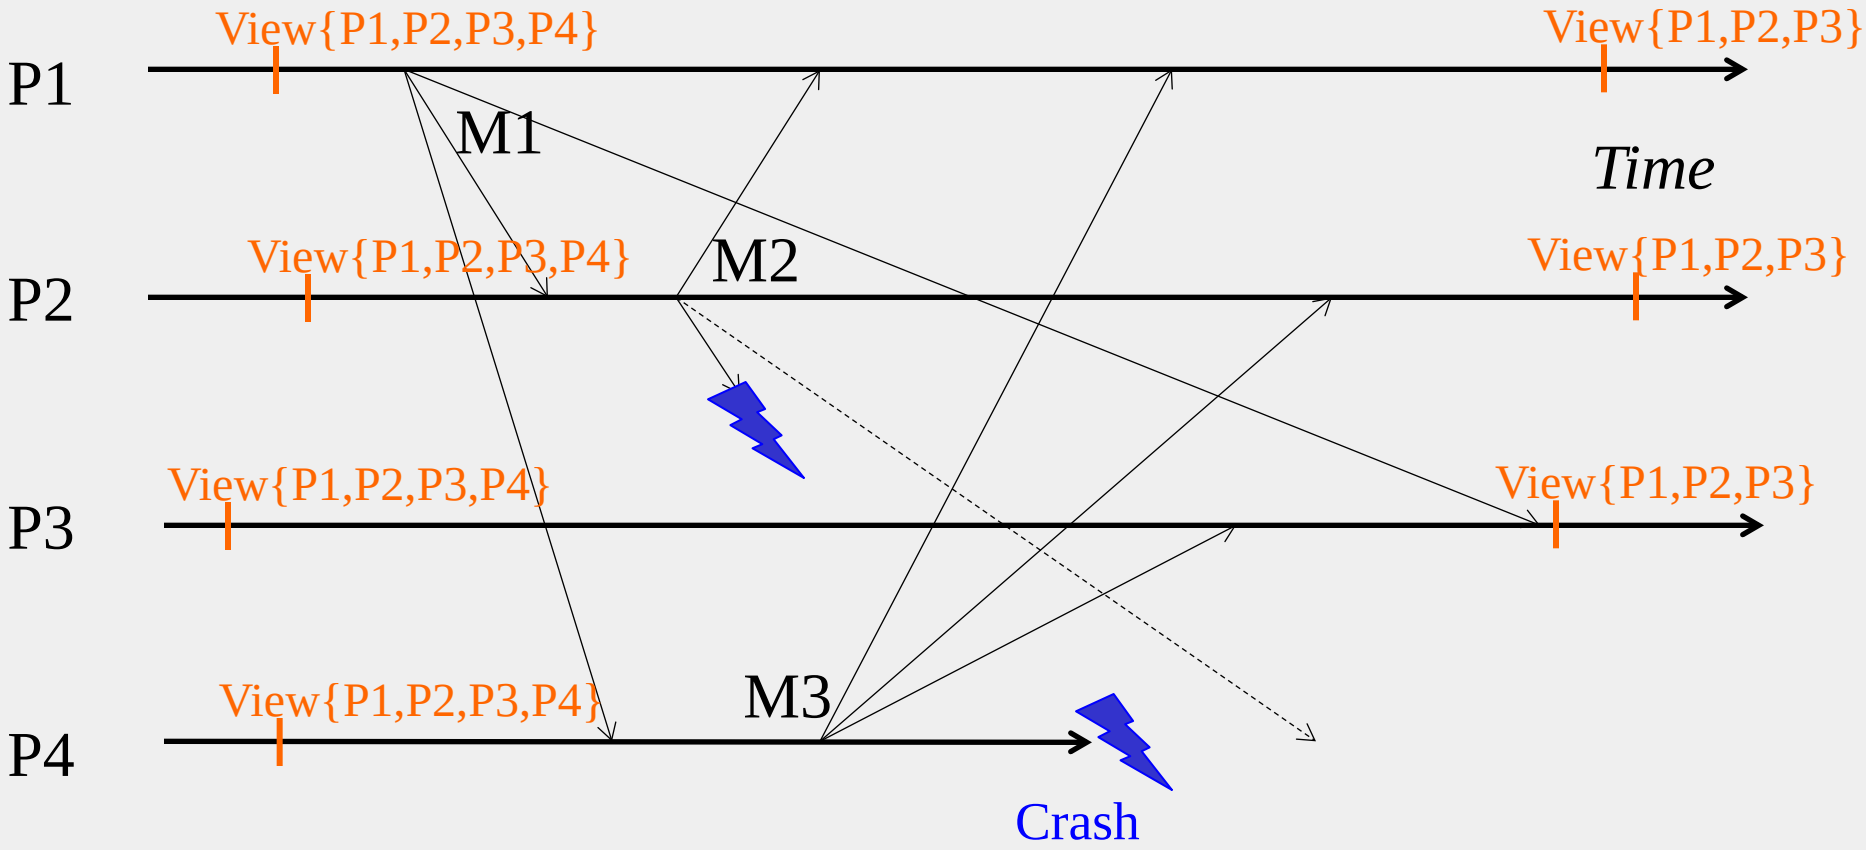

Does not satisfy virtual synchrony

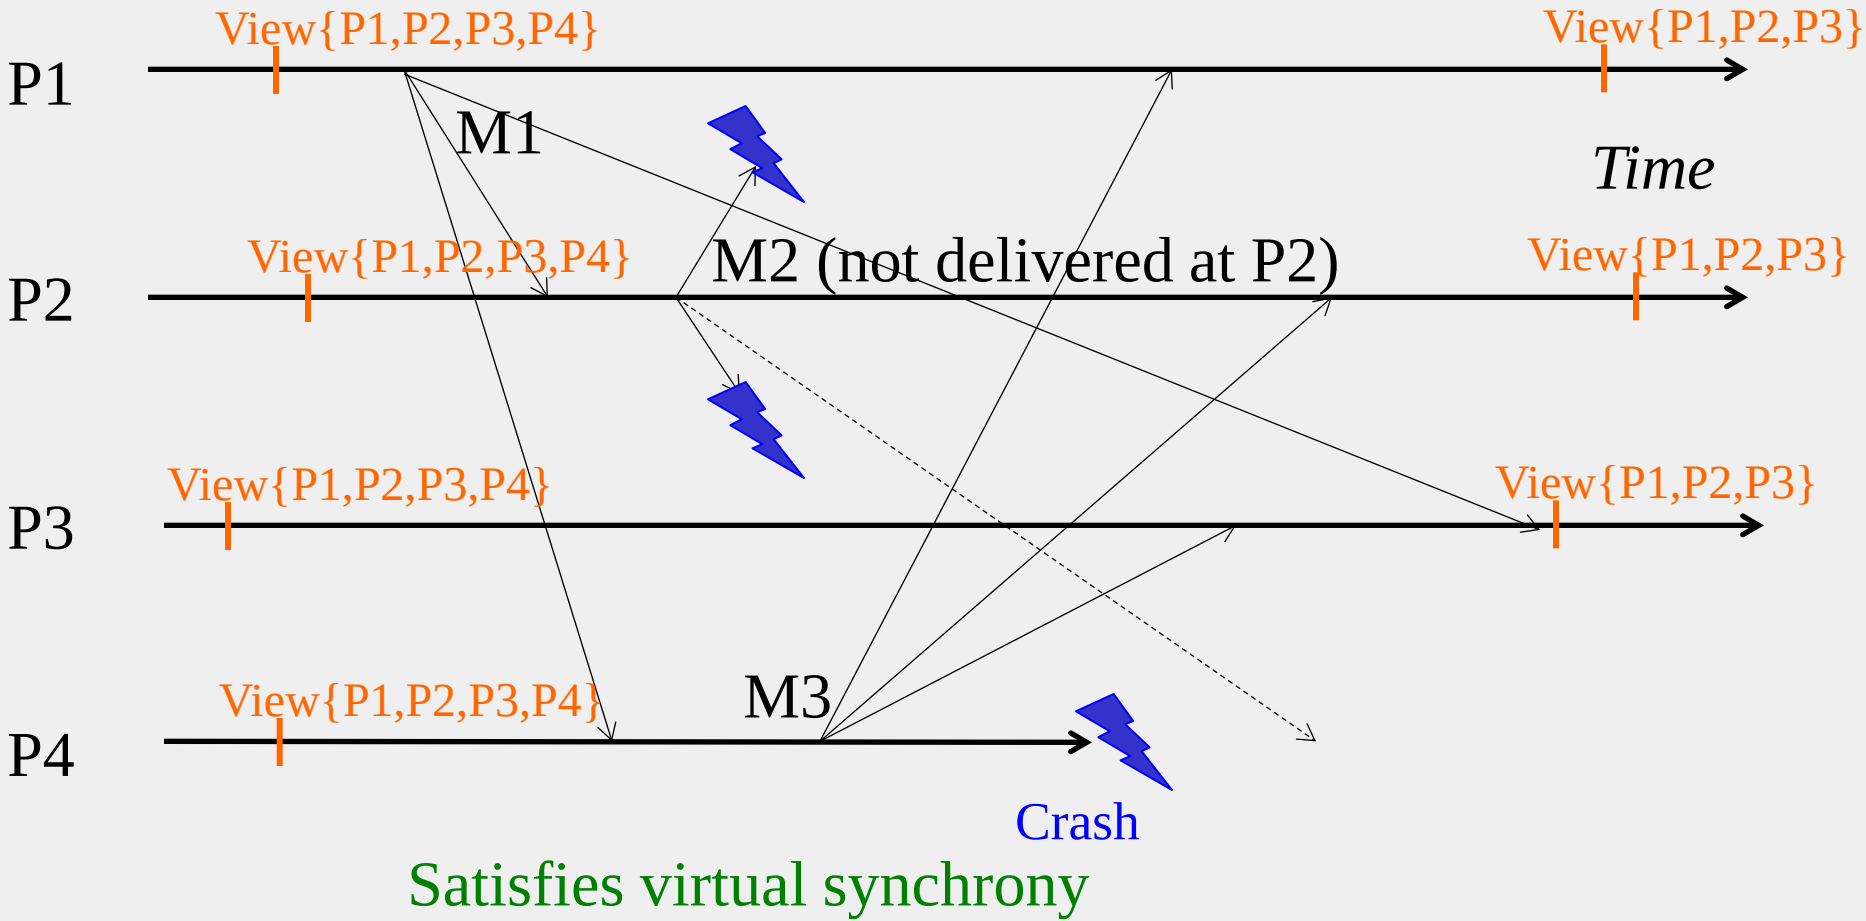

Satisfies virtual synchrony

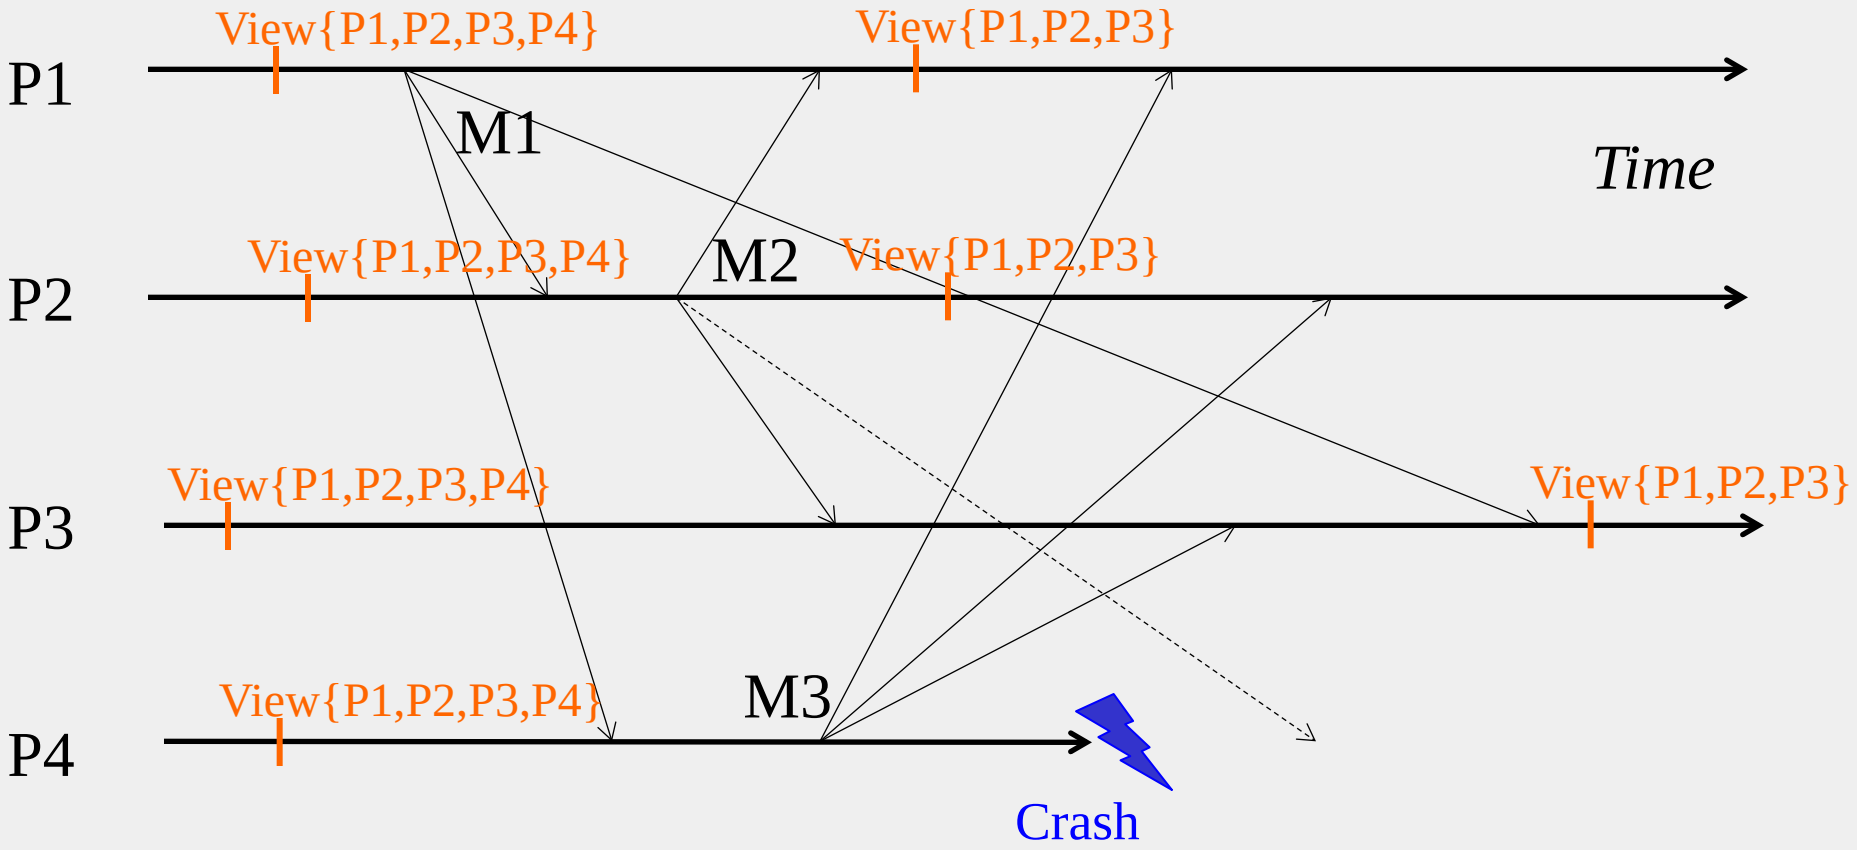

Does not satisfy virtual synchrony

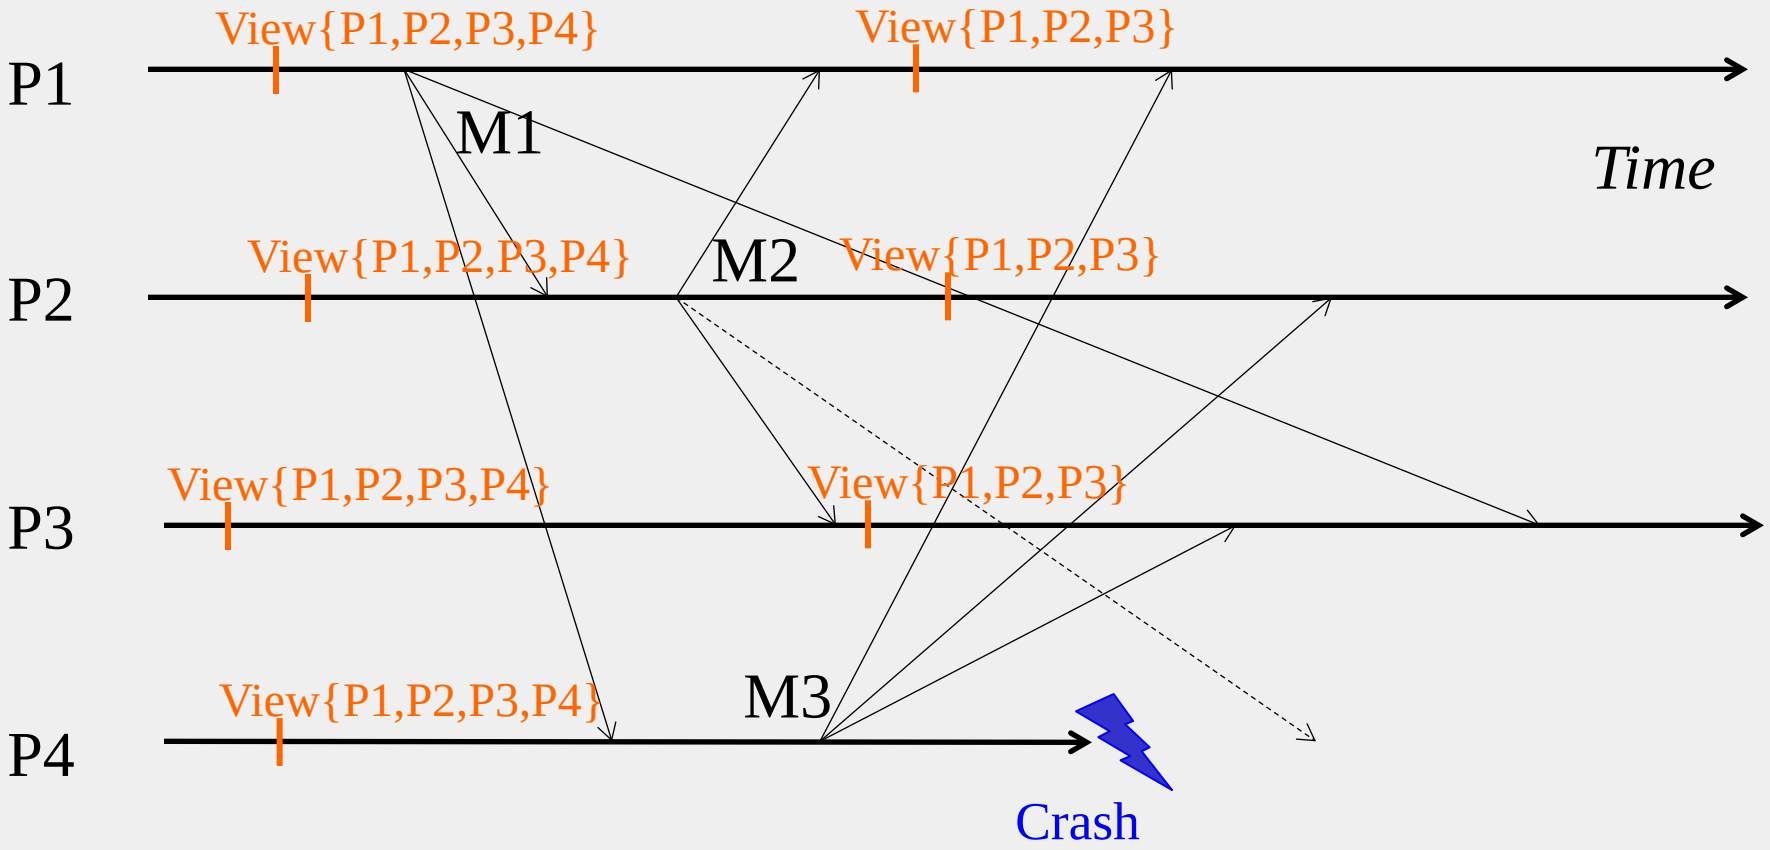

Does not satisfy virtual synchrony

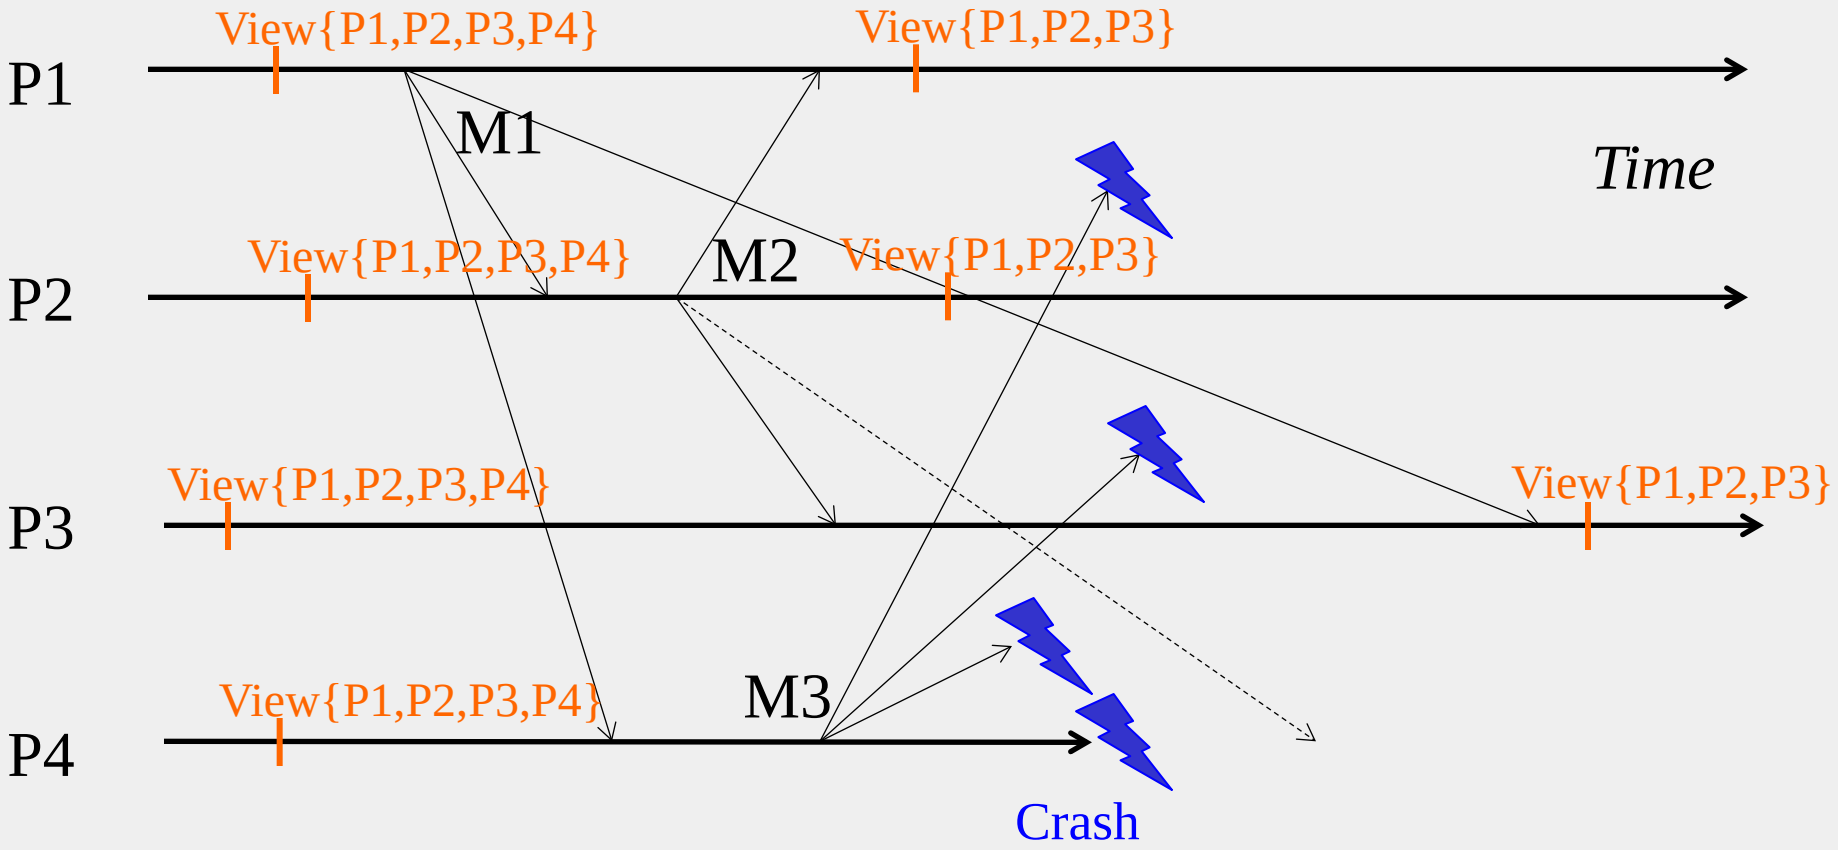

Satisfies virtual synchrony

What about Multicast Ordering?

- Again, orthogonal to virtual synchrony
- The set of multicasts delivered in a view can be ordered either
  - FIFO
  - Or Causally
  - Or Totally
  - Or using a hybrid scheme

About that name

- Called “virtual synchrony” since in spite of running on an asynchronous network, it gives the appearance of a synchronous network underneath that obeys the same ordering at all processes
- So can this virtually synchronous system be used to implement consensus?
- No! VSync groups susceptible to partitioning
  - E.g., due to inaccurate failure detections

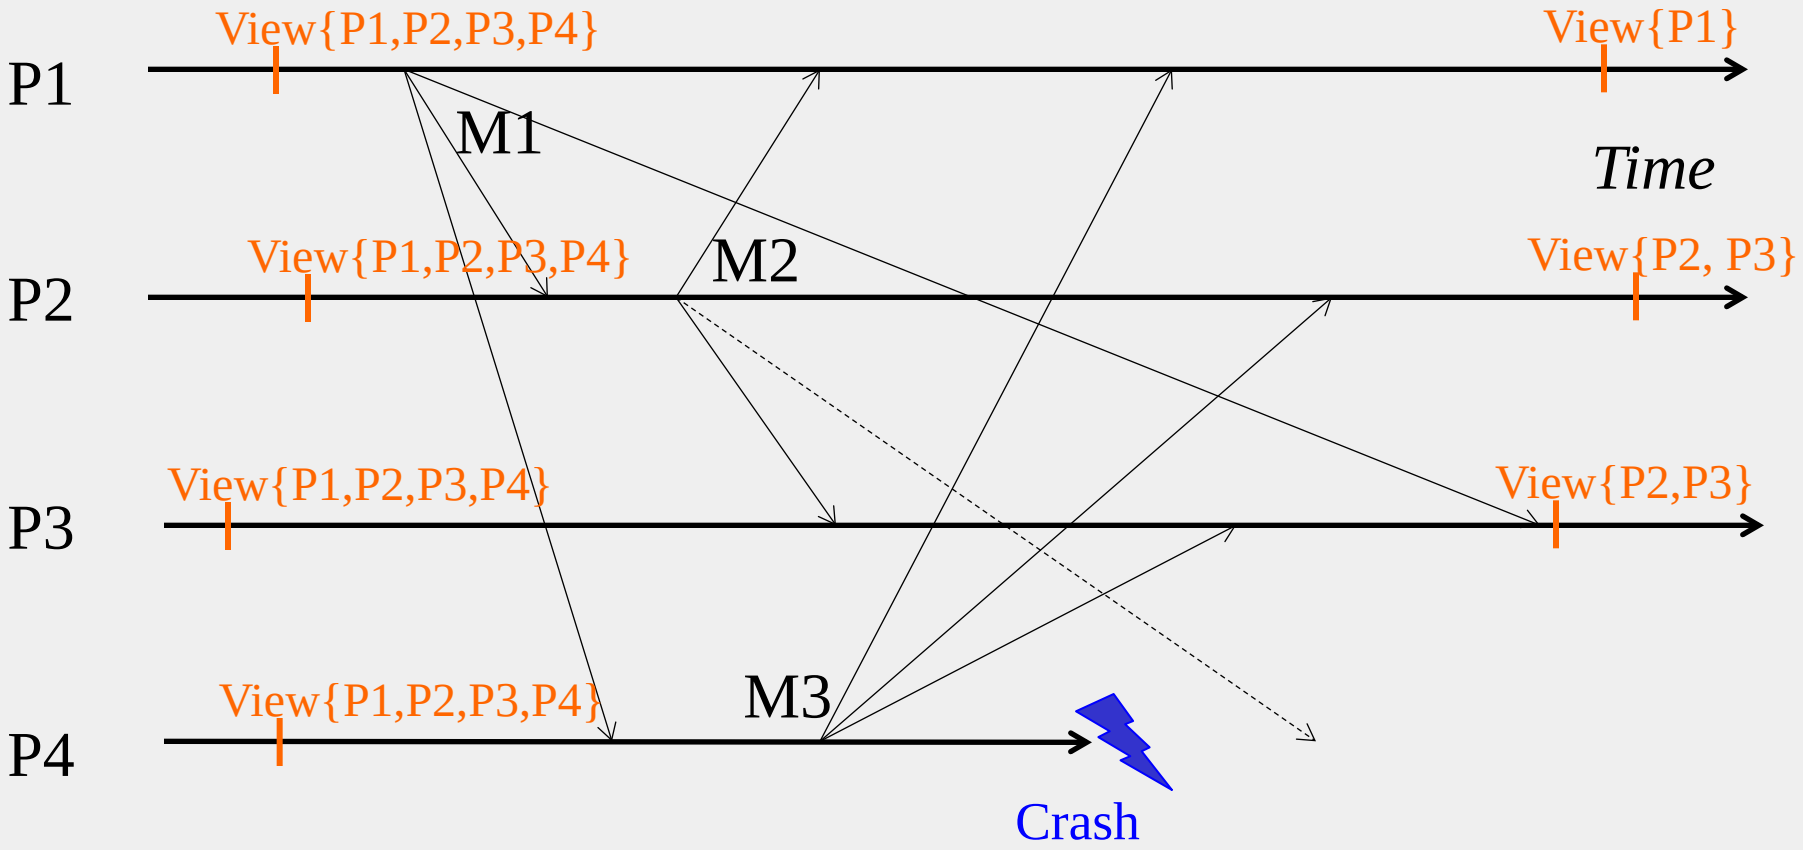

Partitioning in View synchronous systems

Summary

- Multicast an important building block for cloud computing systems
- Depending on application need, can implement
  - Ordering
  - Reliability
  - Virtual synchrony