CLOUD COMPUTING CONCEPTS

with Indranil Gupta (Indy)

MULTICAST

Lecture C

IMPLEMENTING MULTICAST ORDERING 2

Causal Ordering

- Multicasts whose send events are causally related, must be received in the same causality-obeying order at all receivers
- Formally
  - *If multicast(g,m) → multicast(g,m′) then any correct process that delivers m′ would already have delivered m.*
  - *(→ is Lamport's happens-before)*

Causal Multicast: Datastructures

- Each receiver maintains a vector of per-sender sequence numbers (integers)
  - Similar to FIFO Multicast, but updating rules are different
  - Processes P*1* through P*N*
  - P*i* maintains a vector P*i*[1…N] (initially all zeroes)
  - P*i*[*j*] is the latest sequence number P*i* has received from P*j*

Causal Multicast: Updating Rules

- Send multicast at process P*j*:
  - Set P*j*[*j*] = P*j*[*j*] + 1
  - Include new entire vector P*j*[1…N] in multicast message as its sequence number
- Receive multicast: If P*i* receives a multicast from P*j* with vector M[1…N] (= P*j*[1…N]) in message, buffer it until both:
  1. This message is the next one P*i* is expecting from P*j*, i.e.,
     - M[*j*] = P*i*[*j*] + 1
  2. All multicasts, anywhere in the group, which happened-before M have been received at P*i*, i.e.,
     - For all $k \neq j$: M[*k*] $\leq$ P*i*[*k*]
     - i.e., ***Receiver satisfies causality***
  3. When above two conditions satisfied, deliver M to application and set P*i*[*j*] = M[*j*]

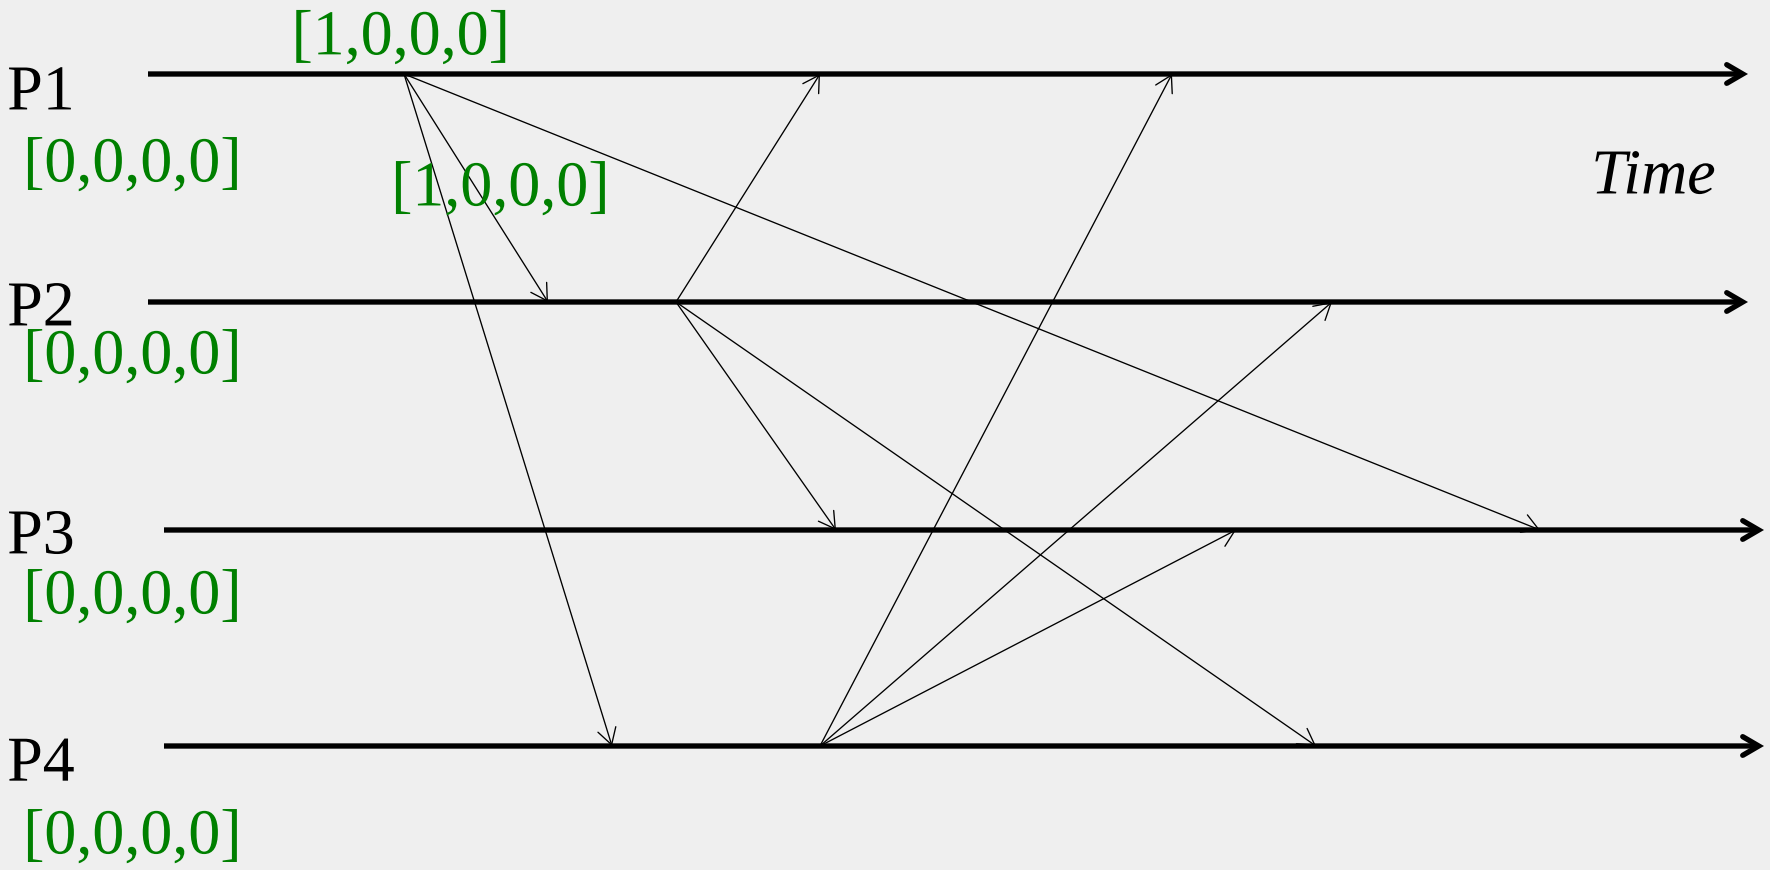

Causal Ordering: Example

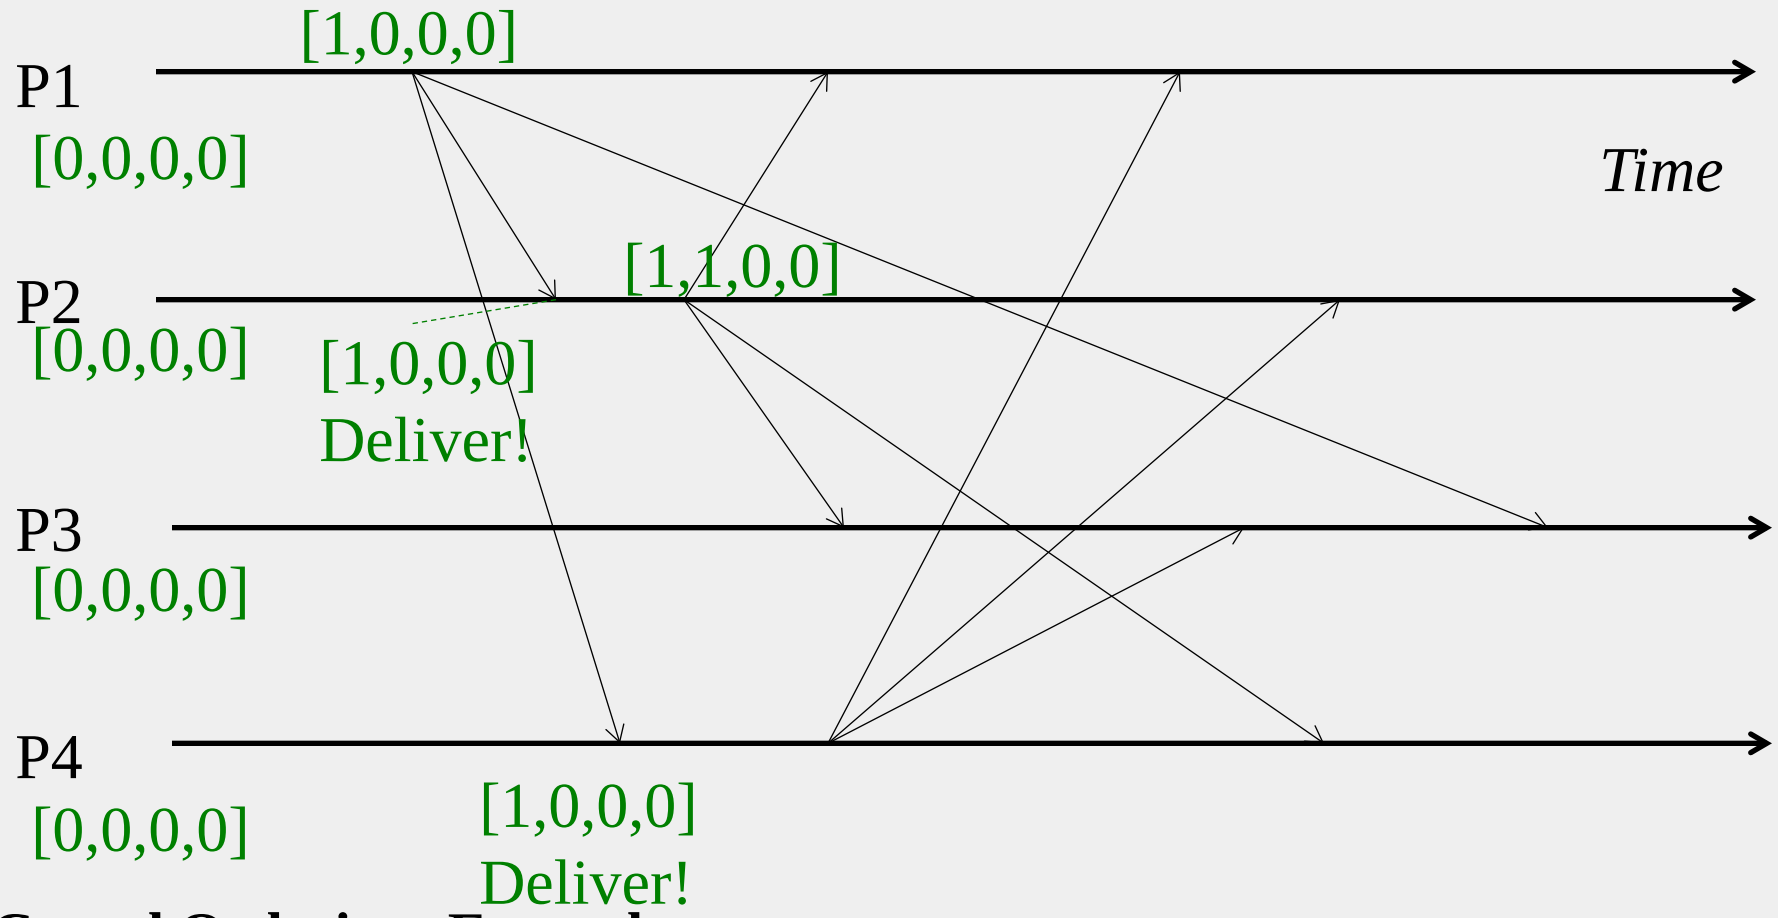

Causal Ordering: Example

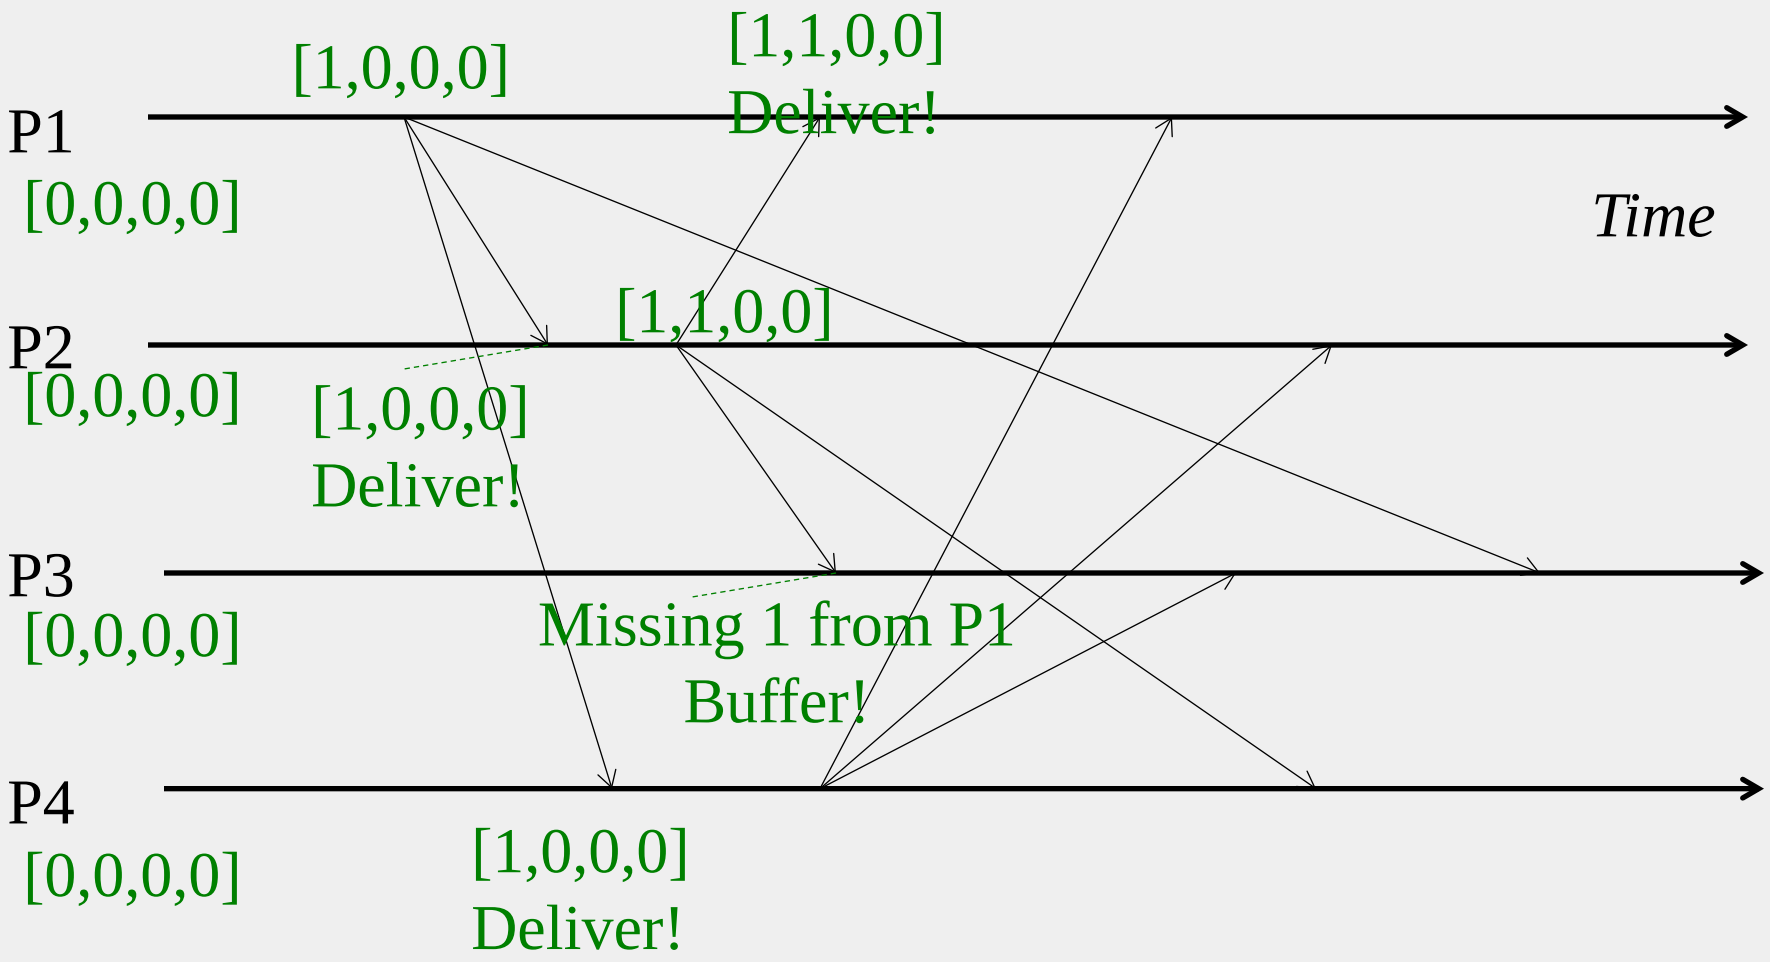

Causal Ordering: Example

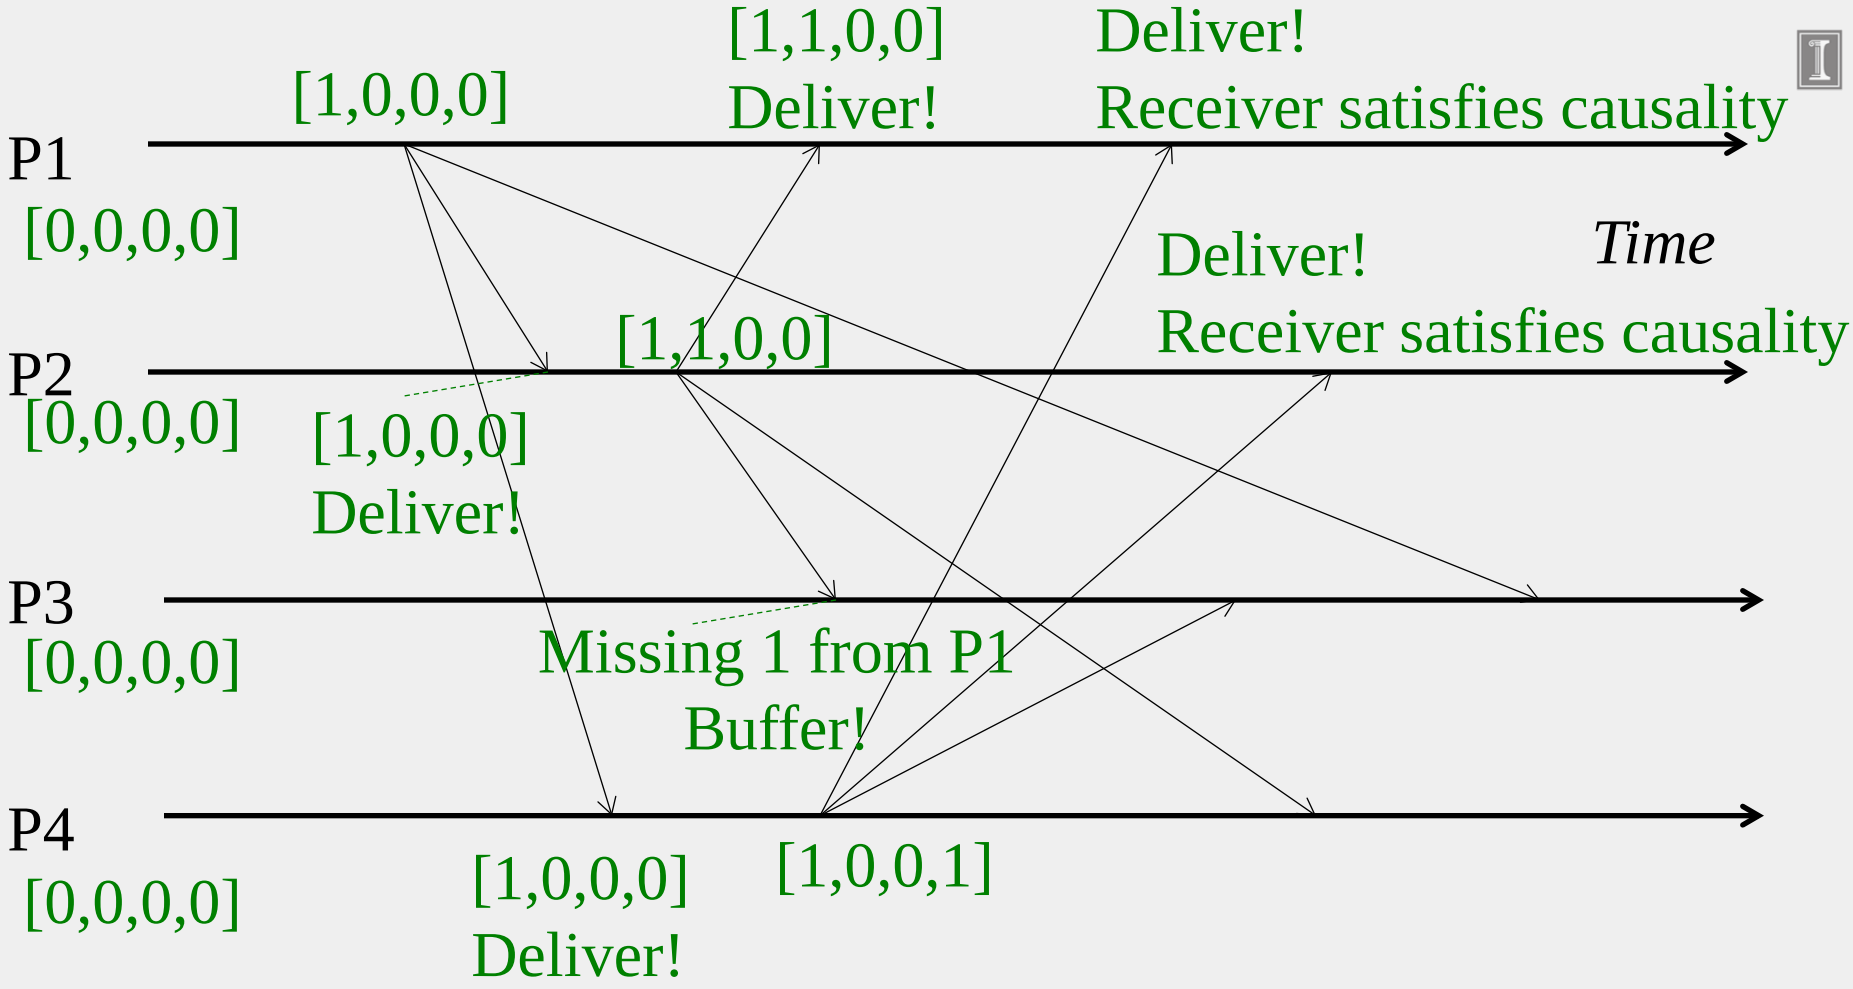

Causal Ordering: Example

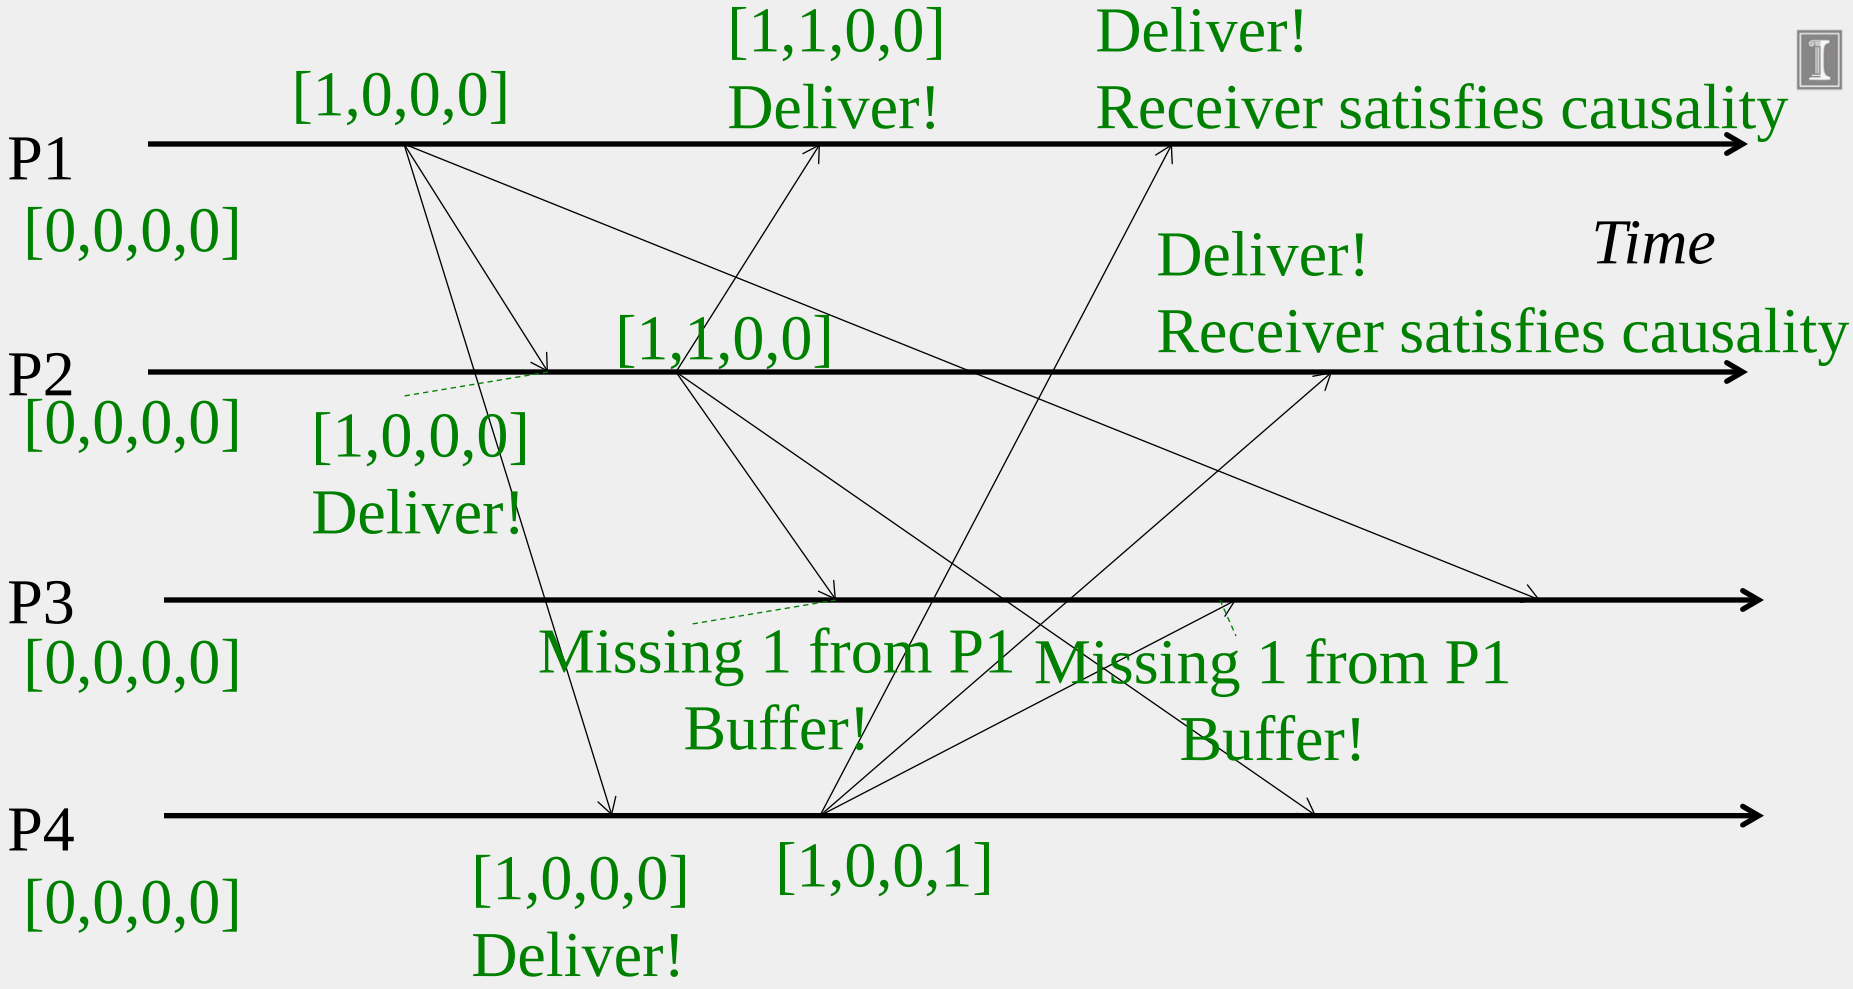

Causal Ordering: Example

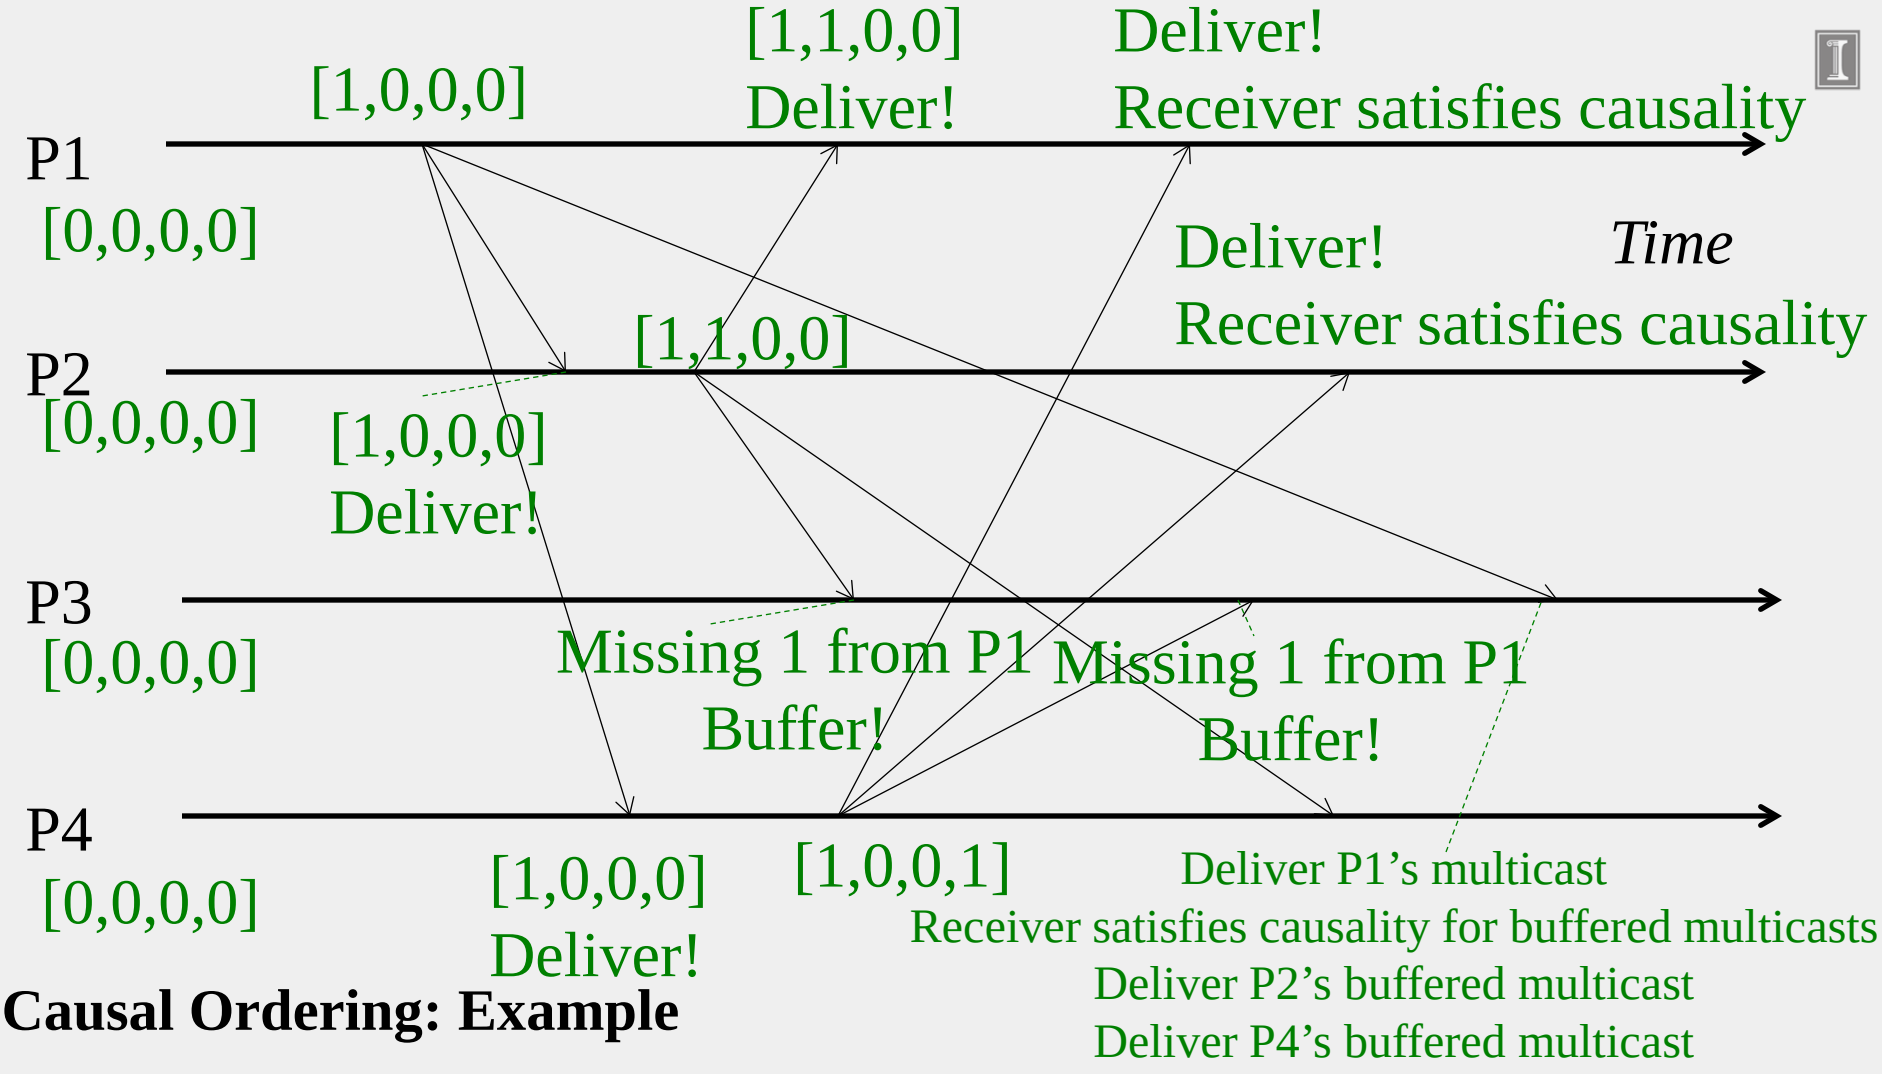

Causal Ordering: Example

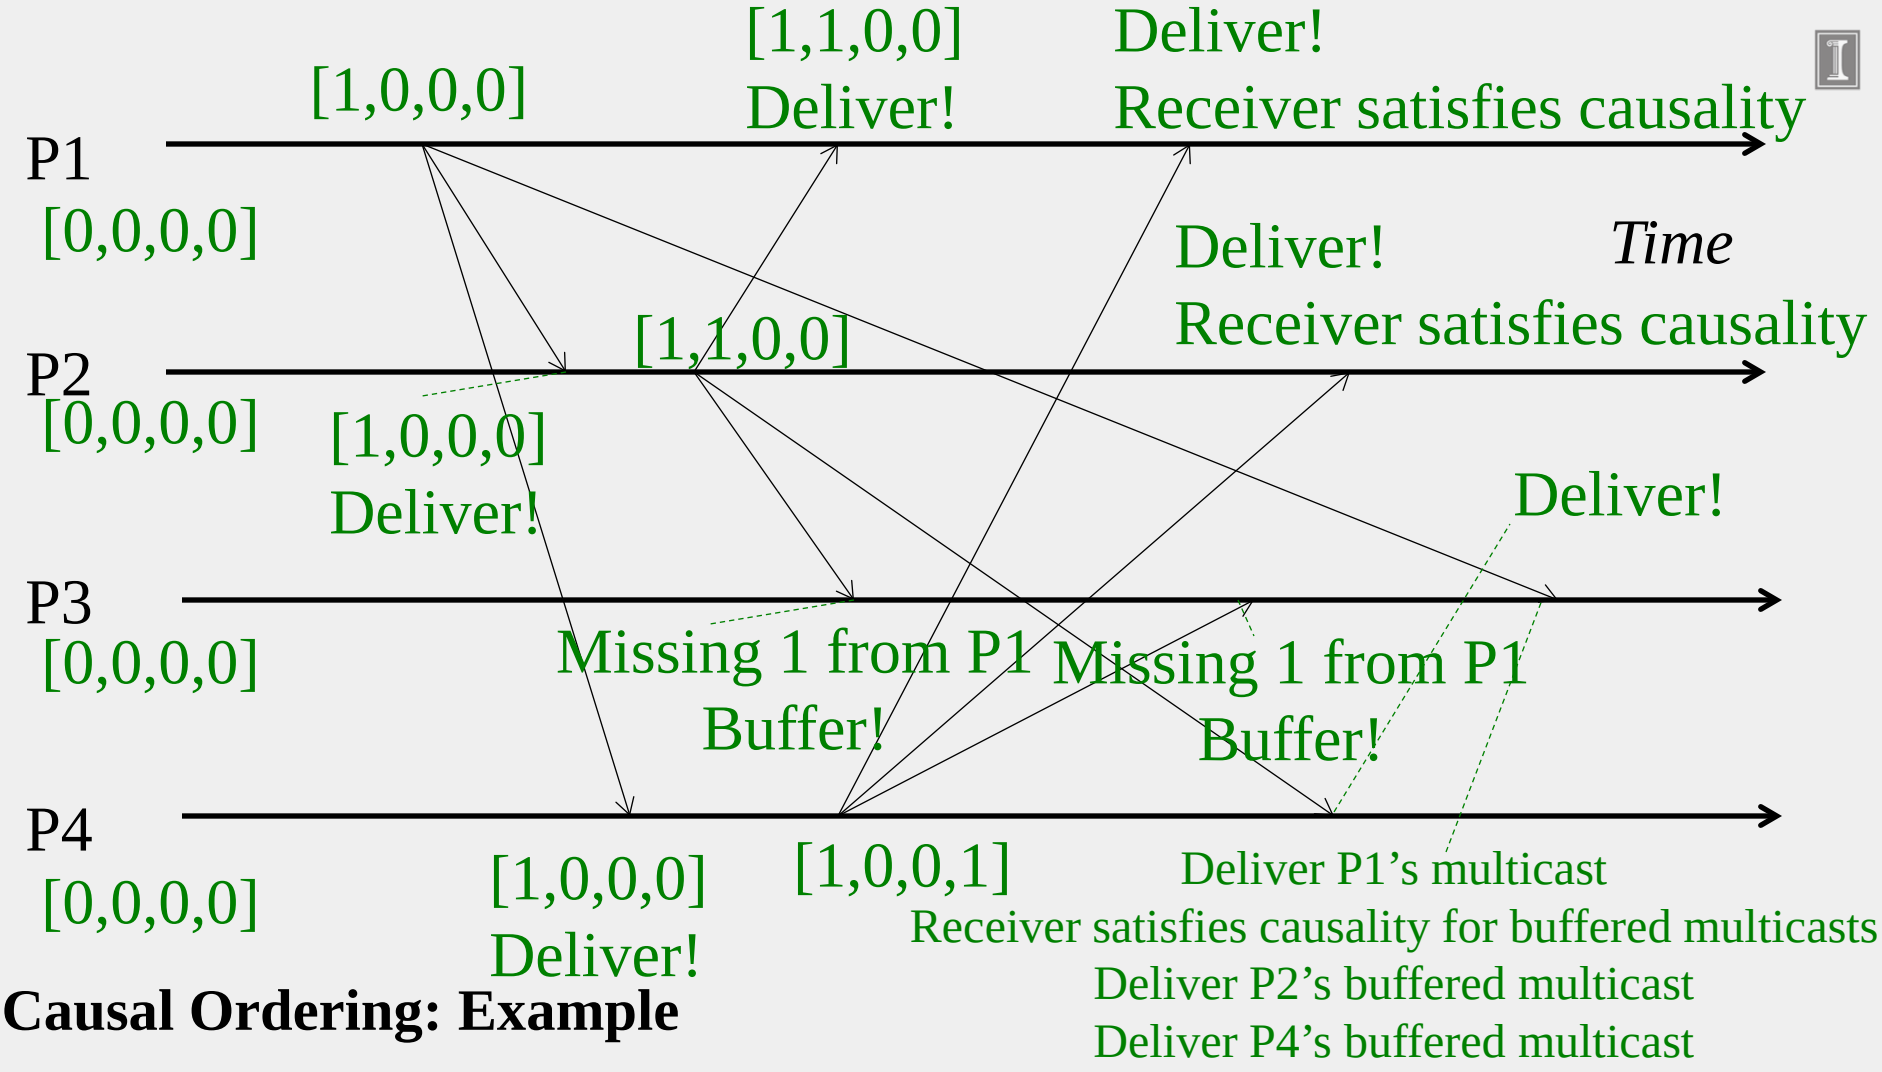

Causal Ordering: Example

Summary: Multicast Ordering

- Ordering of multicasts affects correctness of distributed systems using multicasts
- Three popular ways of implementing ordering
  - FIFO, Causal, Total
- And their implementations
- What about reliability of multicasts?
- What about failures?