

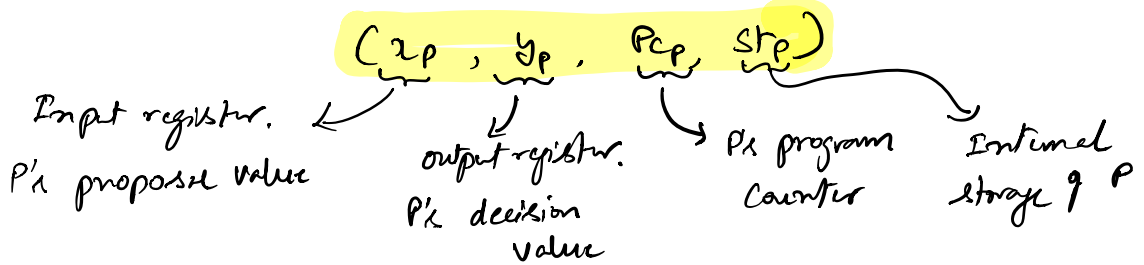
# Welcome to CSCI 7000-001 Lec 5 (Jan 28)!

Recap:

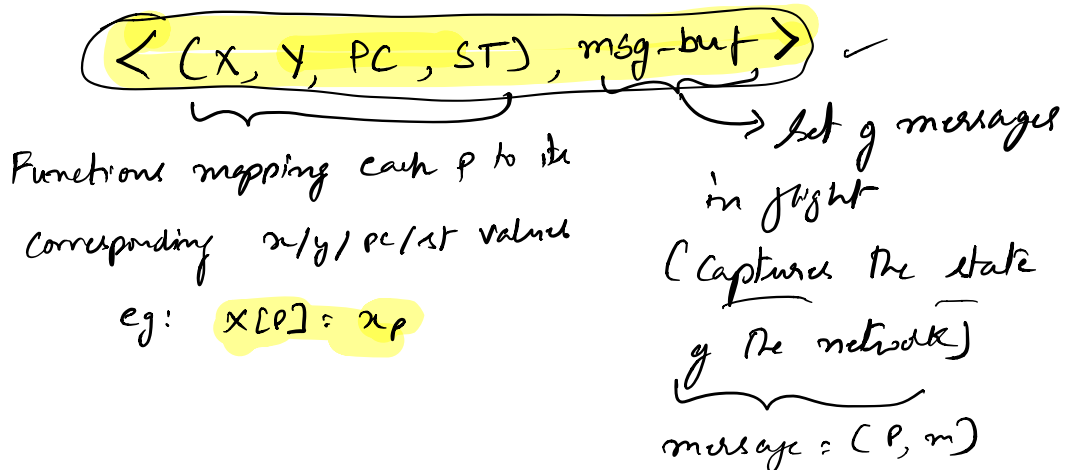
1. Mathematical model of an asynchronous distributed system

\* Set of processes/nodes:  $P = \{P_1, P_2, \dots, P_n\}$

\* State/Configuration of a process  $P$ :

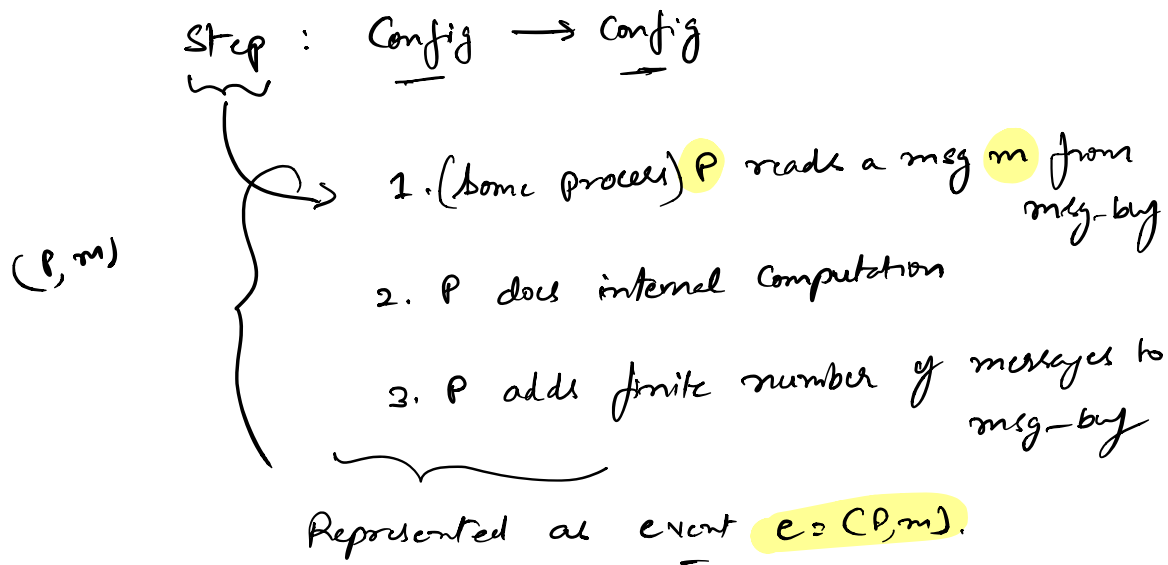


\* State/Configuration of Entire System:



\* Execution of the system  $\equiv$  Sequence of steps, where each step changes the configuration of the entire system

Run (S)



Hence an execution  $\sigma$  of the system =

Series of steps / events  $(e_1, e_2, \dots, e_m)$   
 m could be  $\infty$

\* There can be many different executions of the system starting from an initial configuration depending on the order in which events occur.

Execution<sub>1</sub> =  $\sigma_1 = \langle \text{process } p_1 \text{ runs first, sends msg } m_2 \text{ to } p_2, \dots \rangle$

Execution<sub>2</sub> =  $\sigma_2 = \langle \text{process } p_2 \text{ runs first, sends msg } m_1 \text{ to } p_1, \dots \rangle$

\* A theorem about the system = a property that is true for every execution starting from every initial config

---

### FLP Impossibility

\* A theorem about a system running a consensus protocol where one process is allowed to fail.

\* What would be legal executions of such a system?

Agreement

1. Two different processes  $P_1$  &  $P_2$  cannot make two different decisions. In any config  $C$  that occurs during  $\sigma$ , it cannot be the case that  $C.V[P_1] \neq C.V[P_2]$

Validity

2. Every proposable value must be decidable. There must be some execution  $\sigma_0$  deciding a value 0, and  $\exists \sigma_1$  deciding a value 1.

Robustness

3. No more than one process must "fail".  
 E.g: If  $P_2$  fails, then we won't see events of form  $(P_2, m)$  in  $\sigma$ , but every other process must take steps

FLP Impossibility = Given a legal <sup>A decision is made</sup> deciding run of the system, one can always construct a legal non-deciding run that is infinitely long.


No decision is made



No consensus algorithm is guaranteed to terminate if we allow even one process to fail.

Today!

CAP theorem & its repercussions



webservice ⇒ Internet for communication

Asynchronous System

1. Each machine runs its own clock
2. No timing guarantees
3. Assumptions: All messages are "eventually" delivered.

1. Consistency: Distributed Execution must be similar to a sequential Execution

Each user request must appear to atomically execute  
Linearizability.

2. Availability: Every user request must receive a response.

3. Partition Tolerance: well-defined behaviour even when network partitions.

### CAP Theorem

In an asynchronous distributed system, it is impossible to simultaneously ensure Consistency, Availability, & Partition Tolerance.

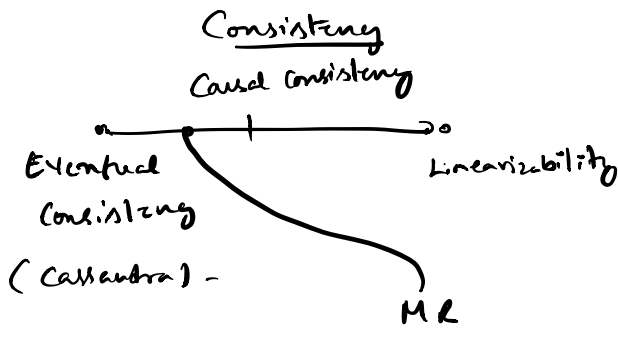
\*  $\checkmark\checkmark$  AP: Cassandra, Riak, Dynamo

\* CP: VoltDB,

\* AC: -----

→ WAN Systems: Given P: chose b/w A & C.

→



Availability

0%  $\longrightarrow$  100% available



Corruption of Execution

latency of requests

PACELL  
abc } Latency vs Consistency