

Welcome to CSCI 7000-001 Lec 8 (Feb 9)!

Announcements:

* Paper presentations start from Feb 23rd

* Submit your preferences on Google form.

$$\langle 10, 0, 1, 3, \dots, 7 \rangle$$
$$\begin{array}{cccccccc} C_1 & C_{11} & C_{12} & \dots & \dots & \dots & \dots & C_{16} \\ \vdots & C_{21} & \dots & \dots & \dots & \dots & \dots & C_{216} \\ & C_{81} & \dots & \dots & \dots & \dots & \dots & C_{8,16} \end{array}$$

$$\forall i. C_{i1} + C_{i2} + \dots + C_{i6} = 2 \quad C_{ij} \in \{0, 1\}$$

$$\forall j. C_{1j} + C_{2j} + \dots + C_{8j} = 1$$

$$C_{ij} \quad S_j \quad \left(\sum_{i,j} S_j * C_{ij} \right)$$


Maximize

PAXOS

→ An "algorithm" for Consensus in an async. dist. Sys.

↓
All the participants have to agree to choose exactly one of the given finite set of values.

→ FLP: choose two of three:



→ Paxos fixes & generalizes 2PC

1. Transaction Manager (TM)

is SPOF in 2PC

2. 2PC doesn't guarantee

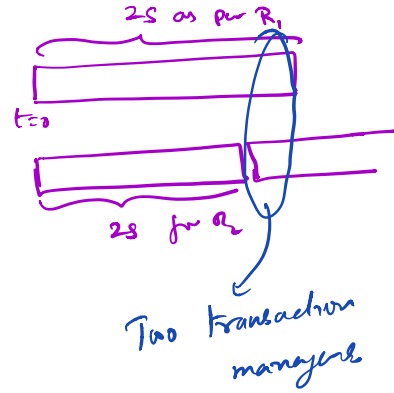
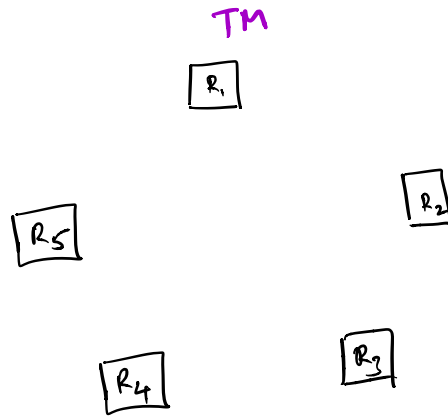
Validity

1. 2PC decides between Commit / Abort

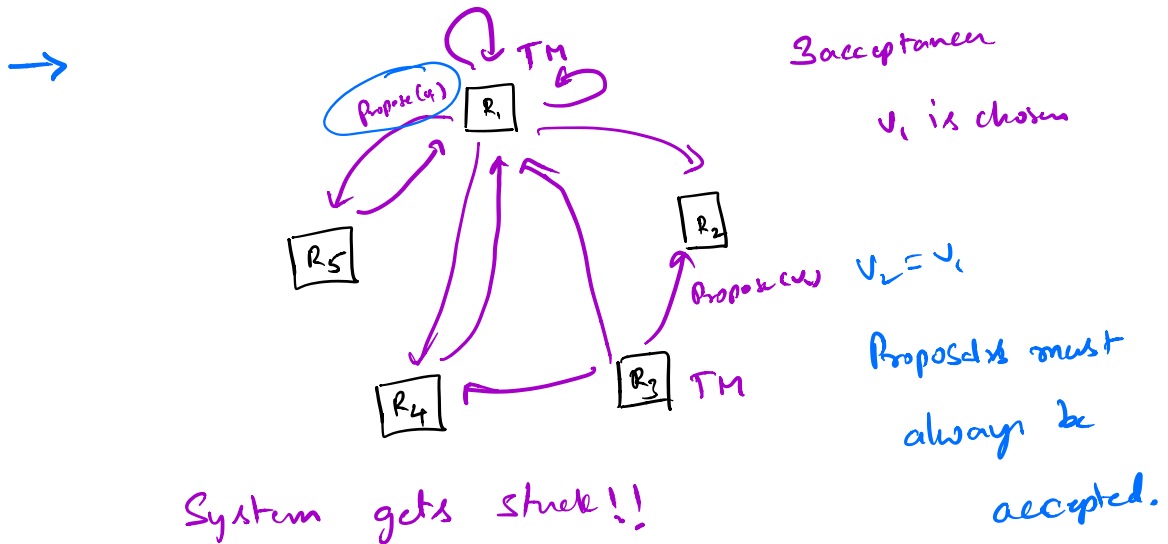
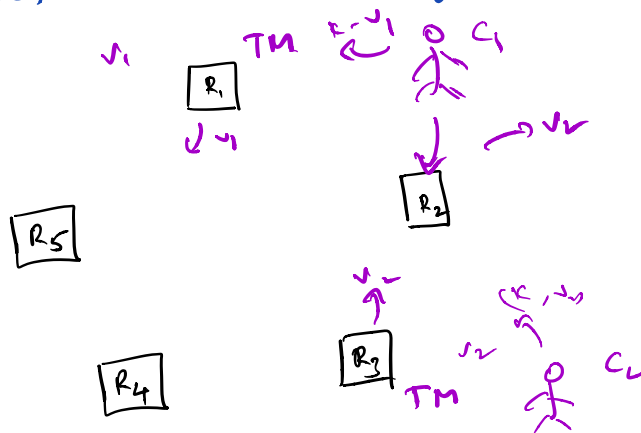
What if we want to decide on

one among many unrelated values?

v_1, \dots, v_3




→ We cannot avoid multiple proposers. Our Consensus algorithm has to deal with many proposers.



Quorum: majority of nodes

→ Propose Accept Decide



○ Learner
v2 at time

→ 2 phases → Prepare/Promise (N) Numbered.
Propose / Accept

"I will not accept any proposals with"

number $< n$ "

Paxos made Simple .

RAFT