

REPLICATED

STATE

MACHINES

Q: How do we build fault tolerant services?

- What types of fault
- How severe?
- What types of service?
- What correctness guarantees?
- Where will the service run? What assumptions can we make?


Failure Models (What type of failure)

Fail-Stop



Byzantine

of faulty nodes




Quorum

Asynchronous? [where will the service run]

FLP

Partial Synchrony?



Clock Synchronization?



↓
F.D.

اس کا

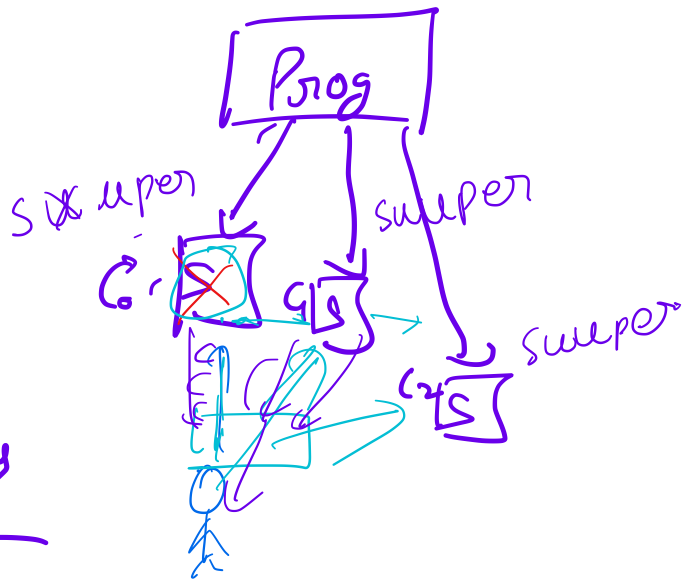
□R

□R

State Machines [What types of services]


Deterministic

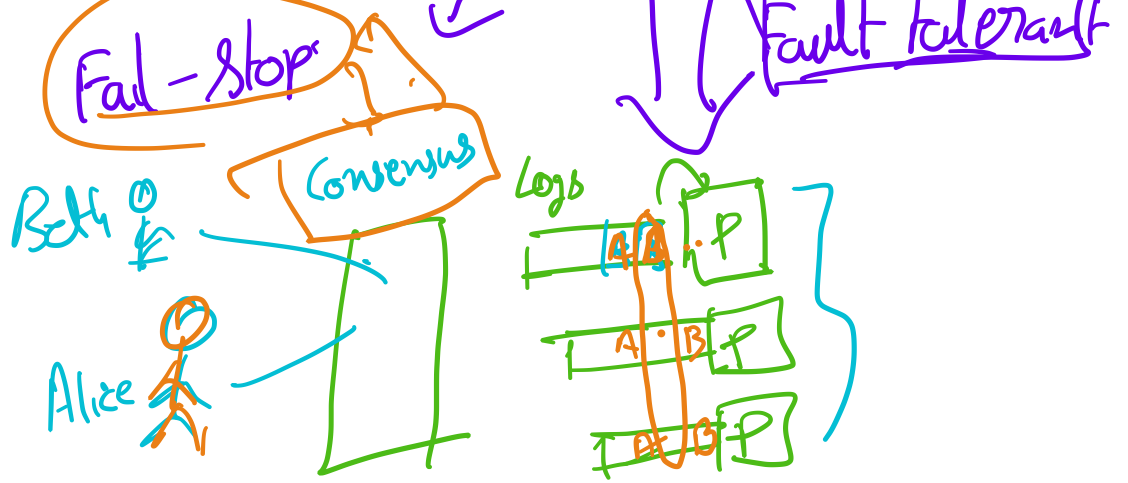
→ A program whose ~~state~~ depends on seq. of commands
Behavior



~~Deterministic behavior~~

State Machine Replication





Dealing with I/O


State Machine Replication

Agreement + Ordering → Agreement


Validity:

Termination:


Logical Construction



Replica



PHYSICAL CHALLENGES



(C)

R_2

(D)

R_3

PHYSICAL REQUIREMENTS

→ SEQUENCER FAULT TOLERANCE

→ QUORUMS

R2



Quorum Intersectors

n machines

k -

~~at~~

