Bayou + weaker Consistency

Distribution
Replication

reliability

(but pert. too)

"Strong"

"Data" perspective

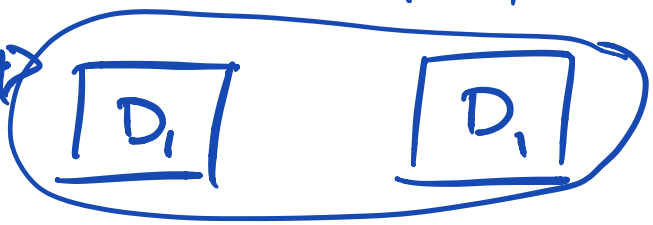

no longer true

"one copy" reads unites

prime reason: weaker models?

(performance, availability)

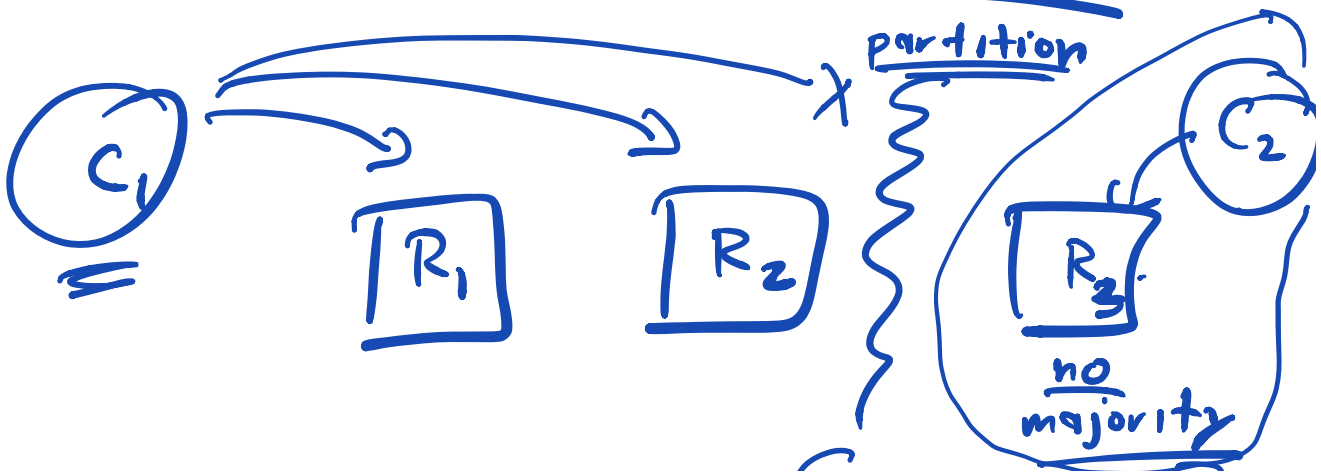

one approach: (quorum-based)

=> CAP theorem
truism

"pick two"

Consistency

Availability — Partitions

Bayou: early, important

=> What is Ev./W. C.?

=> When can it be painful? how?

=> could weaken durability

=> [resolution of conflicts]

=> often

directory, mailbox -> automatic

-> left to application

Weaker consistency systems:

=> Google Drive

=> Dropbox → avail. matters

general: hard

conflict:

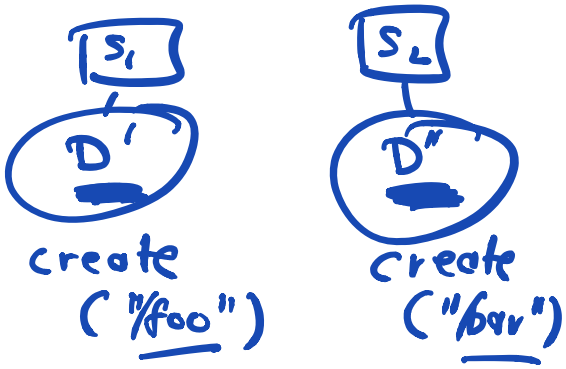

Conflict resolution:

=> version vectors

=> ~ vector clocks

=> Bayou approach

file/object versions

dep. checks, merge proc.

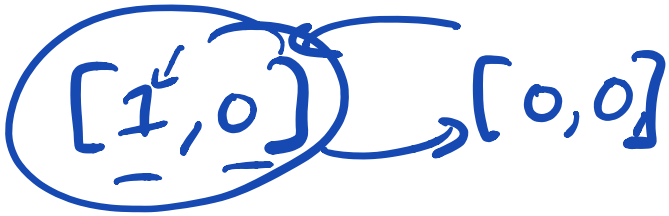

[2, 3] <-/-> [3, 2]

granularity affects:

=> what conflicts are detected

=> amount of resolution

Dep Check: app-specific constraint

nice vs. "traditional" optimistic ?

Specific ⟸ (generic)

DB => exact state

Merge Procedure:

how hard is it to write this code?

(replication: make this harder)

=> order how?

determinism:

↳ how could this

go wrong!

=> mem alloc: must succeed/fail in some way

=> versions

=> convergence: ordered

could use time (from all nodes) => commit

why not? must see/comm. w/ all machines

=> primary: per dataset

"anti-entropy": randomized updates prop.