

# 14 Replication

\* redundancy, replication

MTTF = mean time to fail

MTTR = mean time to repair

avail =  $MTTF / (MTTF + MTTR)$

RAID 1: 2 mirrored disks  $O(2N)$

RAID 4: one parity disk (XOR'd), but all writes write to parity disk  $O(N+1)$  T.T

RAIDS: spread that out!  $O(N+1)$

## GFS

### 15 Transactions

We can't replicate EVERYTHING...

- atomicity - actions that happen completely or not at all  $\rightarrow$  single sector writes
- shadow copies - make a copy and then edit the copy  $\rightarrow$  bad performance
- logs???

- Transaction BEGIN+END, wb if they run concurrently

$\rightarrow$  - Isolation - locks???. really bad perf

$\rightarrow$  "when multiple transactions run concurrently, but they seem to run serially"

### 16 Logging

Better atomicity for txns

\* log change: whether tx commit/aborts - BEGIN, WRITES, READS, UPDATES, COMMITS, ABORTS  
reads are kella slow

LOG + INSTALL write ahead logging

recovery of storage is bad bc whole log. writes are 2X - if storage is always fully updated, then you have to go back to every thing to see whether tx was committed after updated (# updates)

- cache writes then flush - writes aren't as bad bc cache

- log truncation [CHECKPOINT] in log

$\rightarrow$  recovery is better

### 17 ISOLATION

We've enforced atomicity. Now how do we make sure concurrent transactions don't fck up?

Answer: 2PL

\* final state serializability - for some seq schedule,  $O \circ O$  leads to same final answer

\* conflict serializability - order of all conflicts = order in some seq sched

- conflict graph: CS iff  $\exists$  acyclic conflict graph

\* 2PL: each shared var has a lock

Before any operation, tx must acquire lock

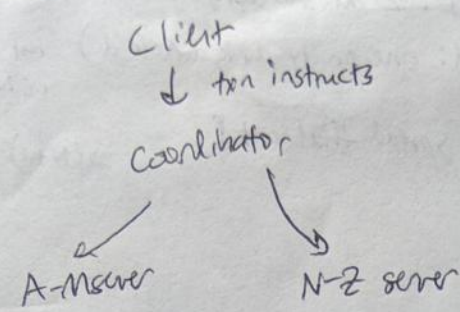
After releasing a lock, it can't acquire any new ones

Dead lock? Abort a tx

\* reader/writer locks

Lec 18: Multi Site Atomicity

Problem: Maintaining atomicity on multiple machines. Consistency



Problem: What if A-M server commits, but N-Z does not?

\*2PC - Two phase commit - PREPARE + COMMIT

Remaining Question - When worker is down, some data is unavailable? Handle via single copy consistency? Consistency Guarantees?

Lec 19:

Problem: We have multiple machines, but how do we maintain consistency across copies?  
 Sol: REPLICATED STATE MACHINES.

Players: clients, coordinators, primary/backup servers

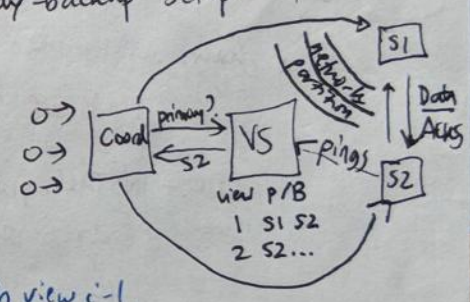
VIEW SERVER - determines which replica server is primary, alerts servers about roles

COORDINATORS - ask view server who is primary/backup

servers - ping VS so VS can detect failures. primary-backup setup


Network partitions create problems!

- 1) Primary must wait for backup to accept each request
- 2) Backup must reject direct coordinator requests.
- 3) Primary must reject forwarded requests.
- 4) Primary in view i must have been primary or backup in view i-1.



Lec 20: Intro to Security

- 1) Policy - who can r/w data - Confidentiality, Integrity, Availability
- 2) Threat Model - who are we protecting against



"Principle of least Privilege" - limit # of trusted components in system

Lec 21: Auth

How to verify someone is who they say they are?

- 1) Auth w/ pws -> salt hashes
- 2) Session Cookies - includes serverkey ST ppl can't pretend to be someone else and use a forged cookie
- 3) Phishing - Challenge response

