




Distributed Transactions

I.O.E.

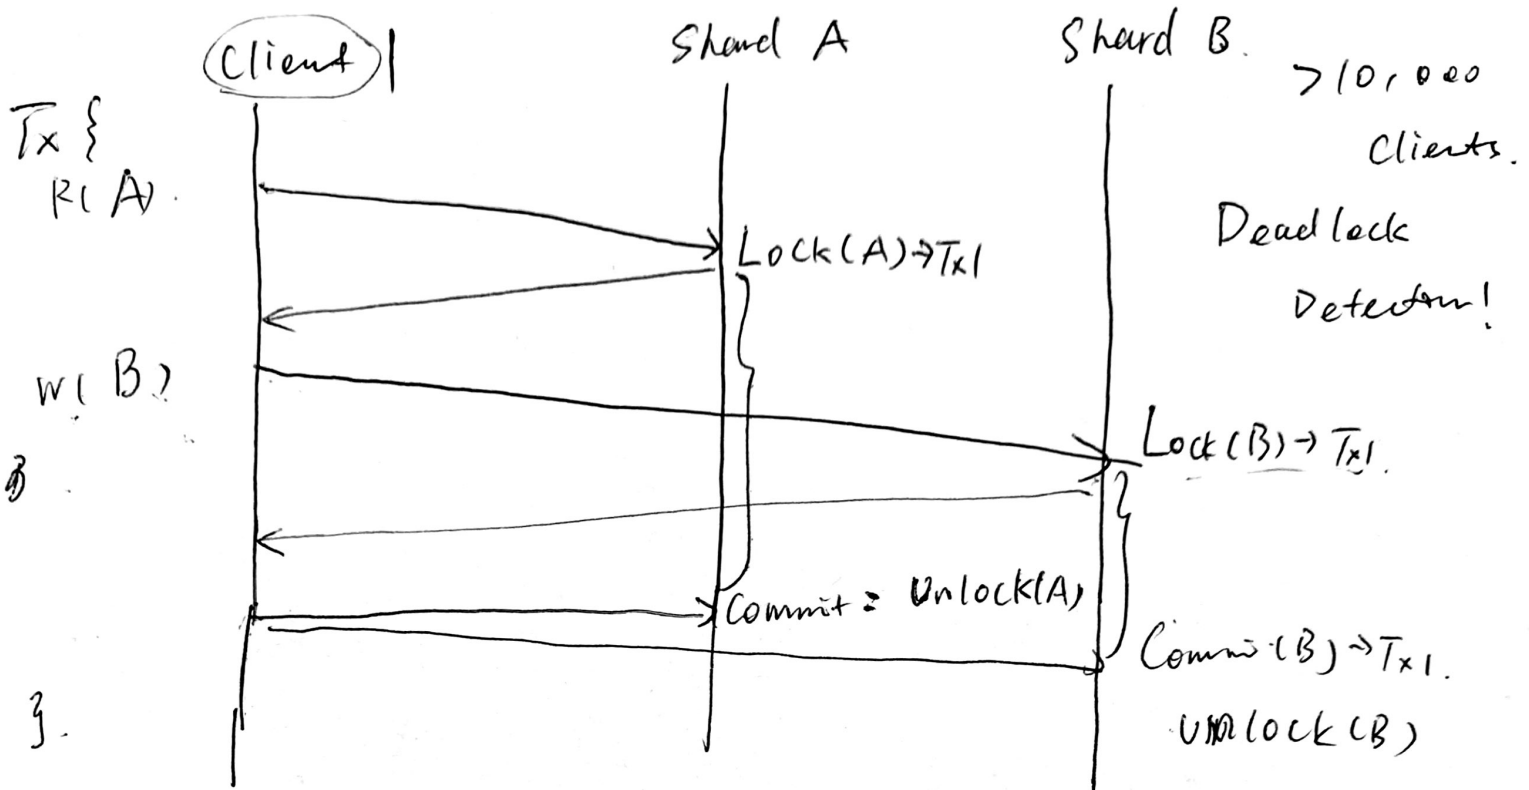


- Mongo DB.
 - RethinkDB.
 - Cassandra
 - Cockroach DB.
 - Aurora.
- Expensive



- ① Scalability / Shard
 - ② Reliability / Replicas
- Consistency / Isolation level through Failures


Transaction w/ Sharding (Partition)



Global lock service (MySQL Cluster)

Pro {


- Deadlock Detection easy
- easy develop / debug



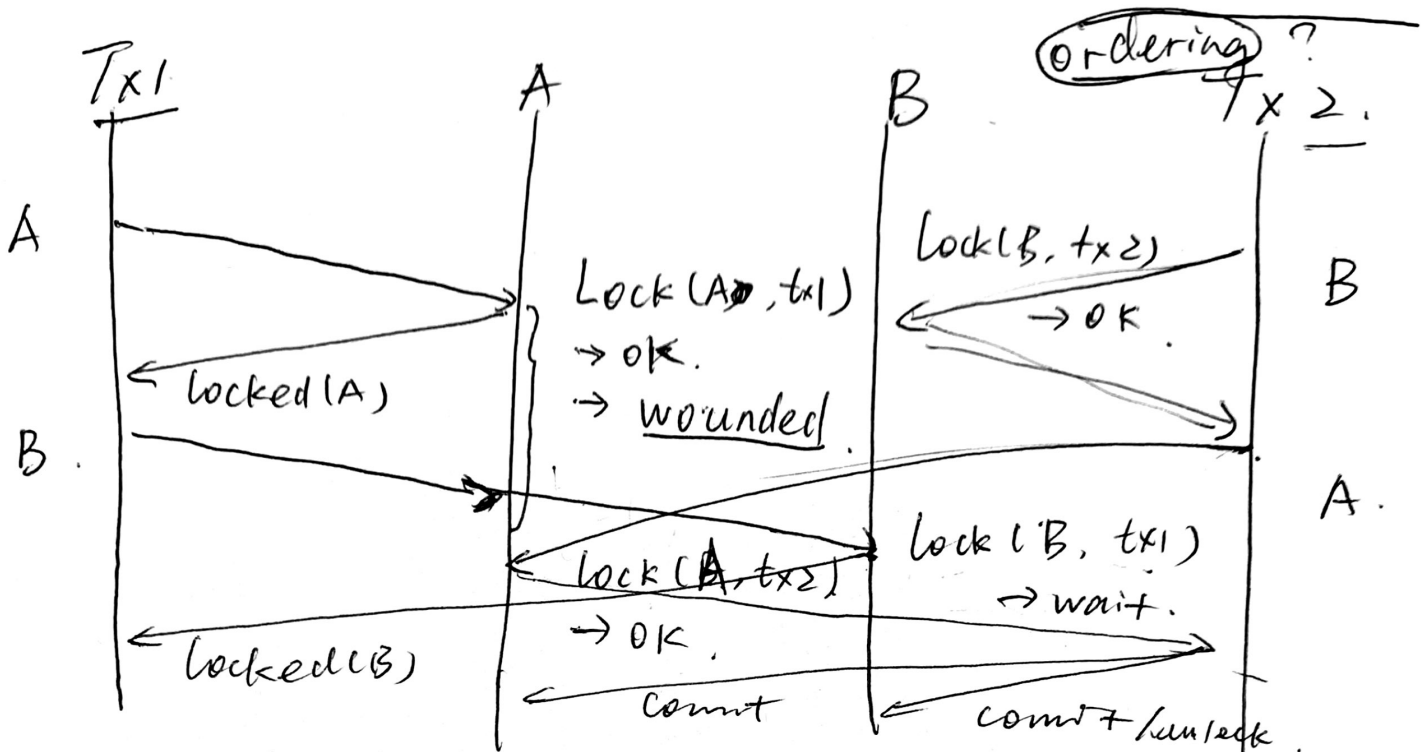
Cons. {

- extra. comm. → lower thp.
- single-point failure


hard to scale



Dead Lock Detection → Prevention



greater id can wound tx w/ smaller id.
 smaller id just wait.





id = read ("next_id")
 write ("id", "buy a pen")

read/write ?
 set prior to execute

N > M. id \approx time.


{ N would M, M waits. \therefore old txs cannot finish.
 M would N, N waits. \therefore old tx slow.



Prepare. → Commit / Abort.

→ 2-Phase-Commit / 2PC (No failures)

(Consensus)



2PC → 3PC (1980s)

(Paxos.)