

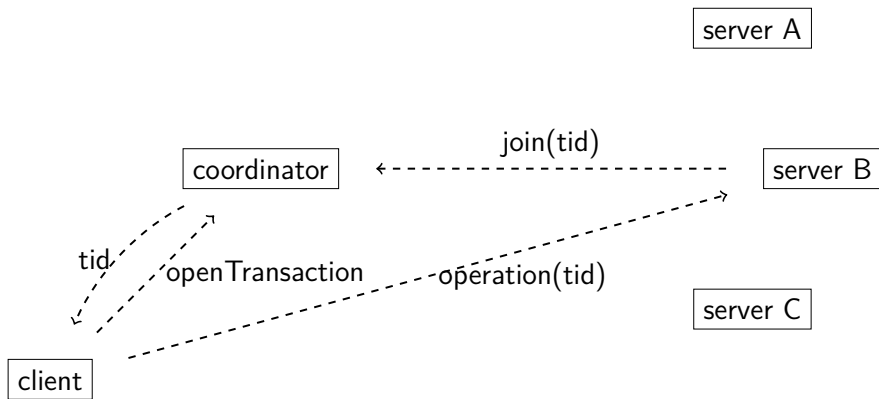
Distributed transactions

Johan Montelius
KTH
HT15

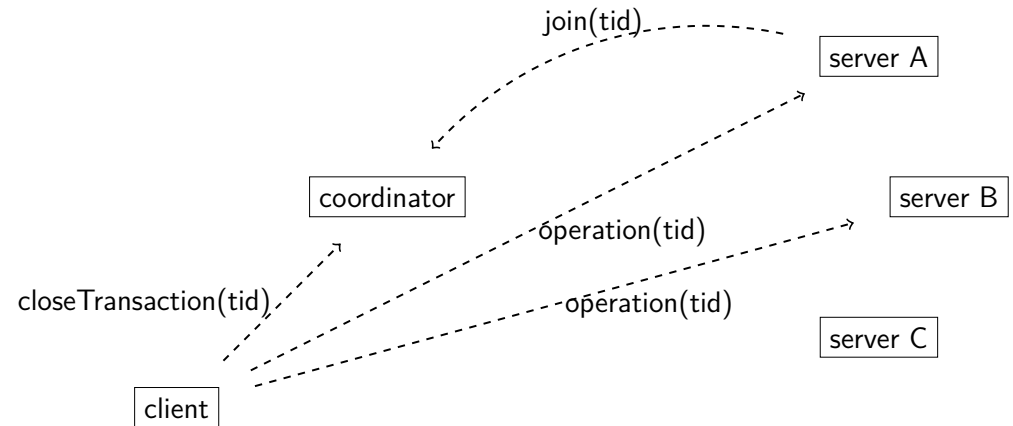
- Several independent transaction servers should be coordinated in one transaction.
- How do we coordinate operations to guarantee serial equivalence?

the architecture

transaction servers



transaction servers



one-phase commit

- Client sends closeTransaction to coordinator.
- Coordinator tells participants to commit the transaction.
- Problem:
 - ?

two-phase commit

- phase one: ask participants to vote for commit or abort
 - if voting for commit one has to be able to commit even after a node crash
 - if anyone aborts all must abort
- phase two: inform all participants of the result

5 / 22

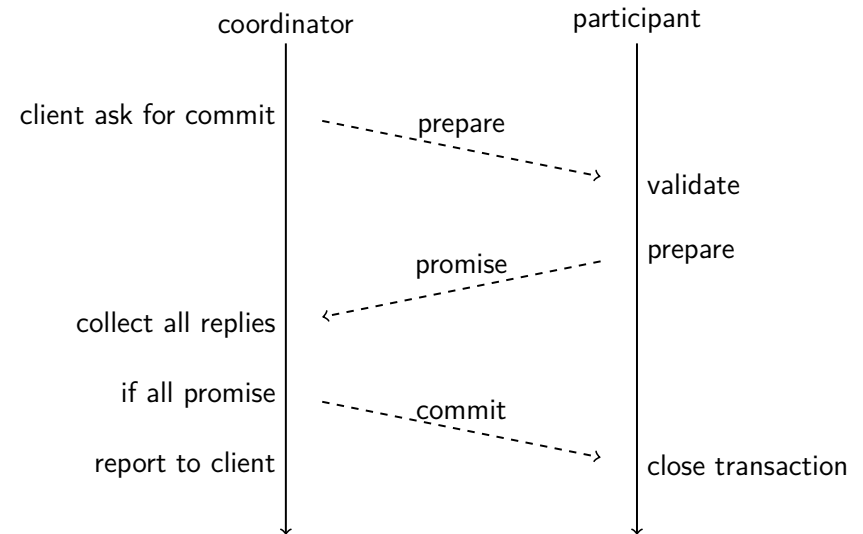
6 / 22

Consensus

Two-phase commit

Two-phase commit is a consensus protocol but:

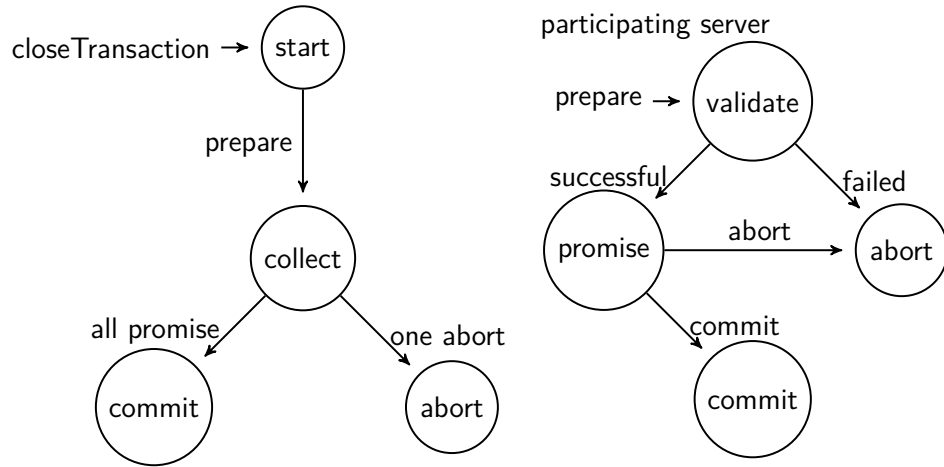
- all servers must vote
- if any server wants to abort then we abort



7 / 22

8 / 22

Two-phase commit



9 / 22

what if ..

- a participating server crashes before making a promise
- a participating server crashes after having promised
- the coordinator crashes before asking for a promise
- the coordinator crashes but you have made a promise

two-phase commit can be suspended waiting for a crashed coordinator

10 / 22

if we know our peers

Assume that the participants know each other.

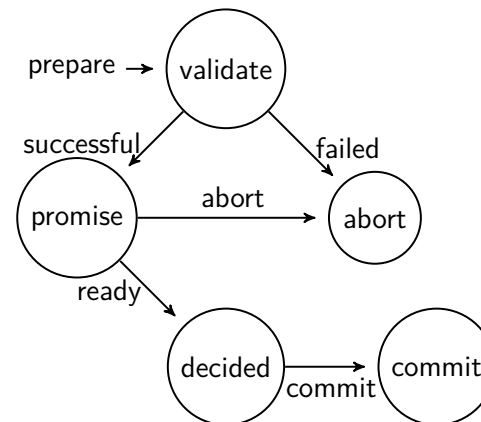
If the coordinator crashes:

- and no participant was told to commit, then it is safe to abort
- if one participant was told to commit, then we should all commit

What if the coordinator and one participant has crashed and none of the surviving participants have received a commit message?

11 / 22

Three-phase commit



- If in the *promised state* and coordinator crashes, and no non-crashed participant is in the *decided state* then abort, otherwise commit.
- If in the *decided state* and coordinator crashes then commit

Relies on perfect failure detectors - and that we know who is in the group.

12 / 22

- locking
- optimistic
- timestamp


Assume we implement *strict two-phase locking* and need to take the locks for foo, bar and zot.

What does it mean and what should we do?

You can either avoid dead-locks or detect them.

We are in a dead-lock if T is waiting for S that is waiting for... that is waiting for T.

Examine the state and look for circular dependencies.



deadlock detection

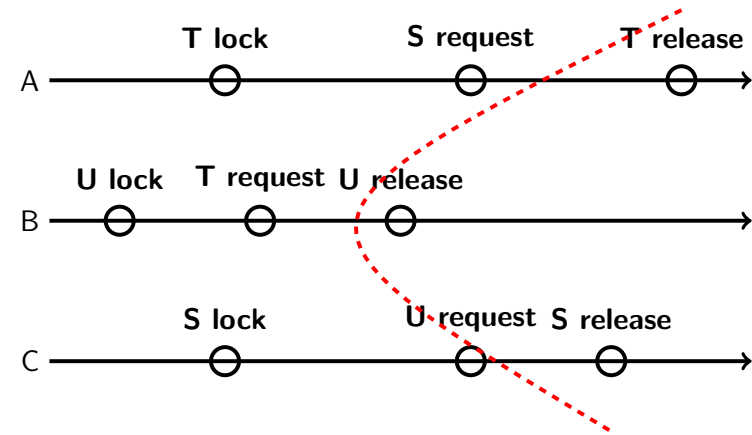
What if:

- server A reports: S is waiting for T
- server B reports: T is waiting for U
- server C reports: U is waiting for S

Deadlock detected, let's do something

17 / 22

phantom deadlock



18 / 22

detection

How do we detect deadlocks?

19 / 22

optimistic concurrency control

Transactions should be validated in a total order.

What if transaction T is validated at A and transaction S at B?

20 / 22

A global timestamp that all transaction servers agree to.

Distributed transactions

- a global total order of transactions
- if one server needs to abort, then all should abort

Two-phase commit

- coordinator asks participants to prepare
- participants promise to commit (or aborts)
- coordinator directs participants to commit

Distributed deadlock

- hard to prevent
- simpler to detect

Concurrency control

- locks
- optimistic
- timestamp