

# Section 12: Two Phase Commit

November 14, 2015

## Contents

<b>1 Warmup</b>	<b>2</b>
<b>2 Vocabulary</b>	<b>3</b>
<b>3 Problems</b>	<b>4</b>
3.1 Delay of Game . . . . .	4
3.2 Two Pigs with One Bird . . . . .	5

## 1 Warmup

Suppose you had a remote storage system composed of a client (you), a single master server, and a single slave server. All units are separated from each other and communicate using RPC. There is no caching or local memory; all requests are eventually serviced using the backing store (disk) of the slave. The slave guards itself against failure by committing entries to a non-volatile log that never gets deleted in the event of a crash. The system only understands PUT(VALUE) and DEL(VALUE) commands, where VALUE is an arbitrary string. Calls to DEL on values that don't exist cause the slave to VOTE-ABORT.

Suppose you issue the following sequence of commands. Recall that the correct sequence of message passing is CLIENT - MASTER - SLAVE - MASTER - CLIENT. Calls to PUT on values that already exist cause VOTE-ABORT.

- PUT(I LOVE)
- PUT(OPERATING SYSTEMS)
- DEL(I LOVE)
- DEL(I LOVE)
- PUT(GOBEARS)

What is the sequence of messages sent and received by the MASTER server? List communications with the slave only. Your answer should be a list of the form:

SEND: PUT(XXX)

RECEIVE: VOTE-XXX

SEND: DEL(XXX)

...

What is the sequence of messages committed to the log of the slave?

Why don't you simply log the state of the slave instead of the action sequence?

## 2 Vocabulary

- **Transaction** - A transaction is an indivisible set of data operations that must either succeed or fail as a single unit, that takes data from one valid state to another valid state. Commonly, reliable transactions are said to satisfy the ACID properties.
- **ACID** - An acronym standing for the four key properties of a reliable transaction.
  - Atomicity - the transaction must either occur in its entirety, or not at all.
  - Consistency - transactions must take data from one consistent state to another, and cannot compromise data integrity or leave data in an intermediate state.
  - Isolation - concurrent transactions should not interfere with each other; it should appear as if all transactions are serialized.
  - Durability - the effect of a committed transaction should persist despite crashes

- **Idempotent** - Idempotency is the property on operations that defines the ability of an operation to be repeated without effect beyond the first occurrence. If an operation can be repeated multiple times, but only the first such operation changes the outcome, the operation is said to be idempotent.
- **Logging file system** - A logging file system is a file system in which all modifications are done via transactions. Instead of modifying the disk directly, intended changes are written to an append-only log. Once the transaction is committed, it becomes safe to copy those changes onto disk, since in the event of a crash, transactions committed to the log can safely be re-copied onto disk.
- **TPC/2PC** - Two Phase Commit is an algorithm that coordinates transactions between one coordinator (Master) and many slaves. Transactions that change the state of the slave are considered TPC transactions and must be logged and tracked according to the TPC algorithm. TPC ensures atomicity and durability by ensuring that a write happens across ALL replicas or NONE of them. The replication factor indicates how many different slaves a particular entry is copied among. The sequence of message passing is as follows:

```

for every slave replica and an ACTION from the master,
origin [MESSAGE] -> dest :
---
MASTER [VOTE-REQUEST(ACTION)] -> SLAVE
SLAVE [VOTE-ABORT/COMMIT] -> MASTER
MASTER [GLOBAL-COMMIT/ABORT] -> SLAVE
SLAVE [ACK] -> MASTER

```

If at least one slave votes to abort, the master sends a GLOBAL-ABORT. If all slaves vote to commit, the master sends GLOBAL-COMMIT. Whenever a master receives a response from a slave, it may assume that the previous request has been recognized and committed to log and is therefore fault tolerant. (If the master receives a VOTE, the master can assume that the slave has logged the action it is voting on. If the master receives an ACK for a GLOBAL-COMMIT, it can assume that action has been executed, saved, and logged such that it will remain consistent even if the slave dies and rebuilds.)


## 3 Problems

### 3.1 Delay of Game

Consider a system with three worker nodes and a master node. A client connects and performs an action which triggers a two-phase commit. The one-way latency between the client and the master is 100ms. The latency between the master and each of the slaves is 10ms. Let's assume that it takes each slave 20ms to update its log and/or commit, and all other processing/transmission time is negligible.

For each scenario, please give the amount of time that passes from when the client first sends its request to the time the slaves receive the final GLOBAL-COMMIT or GLOBAL-ABORT message AND draw the diagram corresponding to the communications sent between the master and the workers.

- a.) Every transmission occurs correctly (and no nodes fail).



b.) One worker node disconnects and misses the VOTE-REQ, and does not reconnect fast enough to stop the master from timing out. The master times out in 5s. The worker comes back online just as the master times out, for simplicity.



### 3.2 Two Pigs with One Bird

You are playing a game of Irascible Avians, a turn based game of loose physics simulations. The game state is hosted entirely on remote machines located somewhere underground in Switzerland. Your machine communicates with the host server via RPC. Because of the high popularity of the game, the master server does not perform any computations or host any storage and offloads those tasks to worker nodes. The master and worker nodes are synchronized via Two Phase Commit.

The only legal message you may send to the remote master is SHOOT(ANGLE, POWER). Because the developers are obsessive about data integrity, (for competitive leaderboards and such) your particular game session is replicated across two different worker servers. When you (the client) make a move in the game, the master forwards the SHOOT message to BOTH workers. The workers then compute the resulting game state at the end of your turn and forward the resulting STATE to the master. The master then communicates the new state of the game back to you. (Note that this is an example of a recursive query.)

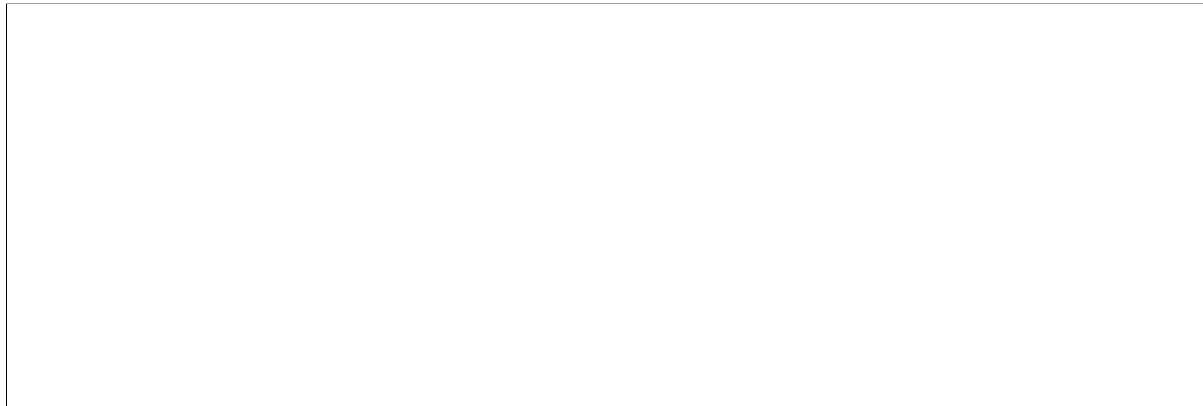
If the action fails to commit on at least one of the workers, the worker propagates an ERROR back to the user and the game behaves as if the current turn did not occur. Each worker knows the initial state of its game session and logs the sequence of committed actions specified by the user. The log is non-volatile and invincible to all known disasters. A crashed worker can recreate the current state of the game by executing the series of logged actions on its initial state (the game is deterministic).

(a) First, draw and label a diagram of the setup. Use boxes to denote machines/nodes and arrows to denote remote communications for a successful turn. Label each arrow with numbers to denote the order of message passing in a complete sequence. (Remember that the master coordinates the workers with 2PC.) The master contacts each slave in serial. The game session is unique to you (your messages will not be confused with other players) and only one message from the client may be processed at a time (the game is discrete/turn based). On a successful commit, the worker sends the game state data with its final ACK. Also, CIRCLE the commands that the slave needs to save to its log.



(b) Suppose the second worker dies during PHASE 1 of 2PC and does not rebuild until after the master times out. What happens if worker death occurs before the phase 1 action is written to the log? How about after? List the sequence of actions that occurs for both cases and indicate which actions are logged. Your answer should be of the form:

1. client -> master: SHOOT(POWER, ANGLE)
2. master -> slave1: <MESSAGE> + LOG?
- ...



(c) Now repeat the last question assuming that the same worker dies in PHASE 2. (Both before and after logging the phase 2 command) What edge case do you need to guard against to ensure proper worker behavior in this case?

