#### A Bug's Life

Fixing a MongoDB Replication Protocol Bug with TLA+

William Schultz, Siyuan Zhou



#### **Talk Outline**

- MongoDB Background
- Life of a Replication Bug
- Specifying the Replication Protocol in TLA+
- Model Checking the Specification
- Takeaways
  - Without a formal model, it's nearly impossible to get a complex distributed protocol right. TLA+ and TLC are tools that make this possible for practicing software engineers.



## Background

. and daily



4.11. I.I.I



MongoDB

- MongoDB is a document oriented database
- A MongoDB database consists of a set of collections, where a collection is a set of unique documents e.g.

```
{
    __id: 1,
    name: "Will",
    company: "MongoDB",
    age: 25
}
```





MongoDB Test Infrastructure

- We have an extensive and mature testing infrastructure
- 1000s of hours of testing are run on new commits every day
  - Includes unit/integration tests, randomized fuzzing, concurrency tests, Jepsen, etc.



mongoDB



- MongoDB provides the ability to run a database as a replica set
- This is a set of MongoDB nodes that coordinate to provide high availability using a consensus protocol





- Replica sets use a consensus protocol similar to Raft
- There is a primary node that inserts writes into the oplog
- Secondaries fetch log entries from other nodes and apply them







- Protocol is leader based
- Replica set leaders are totally ordered by term
  - The *term* is a monotonic counter maintained on each node





mongoDB



- Replica set concepts to keep in mind:
  - $_{\circ}$  sync source
  - commit point
  - $_{\circ}$  rollback
  - $_{\circ}$  term
  - "branch of log history"



## Life of a Replication Bug





- Series of safety and liveness bugs in replication protocol
- Stemmed from one bug found in 2016





- 1. 2016 Heartbeat Commit Point Propagation (Safety)
- 2. 2018 No Available Sync Source (*Liveness*)
- 3. 2019 Sync Source Cycles (Liveness)
- 4. 2019 Commit Point Held Back on Stale Nodes (Liveness)
- 5. 2019 Initial Solution Fails in 5 Node Replica Set (Safety)

# Bug 1: Heartbeat Commit Point Propagation

- Correctness bug found in November 2016 by a replication team engineer
- Allowed for nodes to erroneously mark log entries as "committed"
  - Consequence: client could read data it thinks is durable, even if it isn't



## Bug 1: Heartbeat Commit Point Propagation 2016



mongoDB

# Bug 1: Heartbeat Commit Point Propagation

Solution: Only update commit point via sync source spanning tree
 Guarantees commit points are sent between nodes on same branch of log history





## Bug 2: No Available Sync Source

- New liveness bug found in February 2018
  - Discovered in our test infrastructure



## Bug 2: No Available Sync Source



mongoDB

## Bug 2: No Available Sync Source

- Solution: Relax rule sync source selection rules
  - Allow nodes to sync from someone with a higher commit point if logs are the same



## Bug 3: Sync Source Cycle

- Liveness bug of a new variety
- Nodes may get into sync source cycles
  - Prevents them from ever syncing new log entries
- Consequence of the previous alteration to sync source selection rules

## Bug 3: Sync Source Cycle



mongoDB

# Bug 3: Sync Source Cycle

- Solution: Rethink commit point propagation
- Key idea: learn commit point if it is on your branch of history
  - $_{\circ}$  Via heartbeats or sync source



# Bug 4: Commit Point Held Back on Stale Nodes

- New liveness bug noticed in February 2019
- Stale nodes may not be able to advance their commit point



# Bug 4: Commit Point Held Back on Stale Nodes





# Bug 4: Commit Point Held Back on Stale Nodes

- Solution
  - Relax condition for commit point updates from your sync source
  - Sync source guaranteed to be on the same history branch



# Bug 5: Initial Solution Fails in 5 Node Replica Set

- The original solution to Bug 1 was believed safe, even though it had liveness issues
  - Discovered that the solution was not safe in replica sets with > 4 nodes
  - Could lead to nodes erroneously committing log entries in certain cases



### Specifying the Protocol in TLA+



#### **Variables**

 $\*$  The server's term number.

#### VARIABLE globalCurrentTerm

\\* The server's state (Follower, Candidate, or Leader).

#### VARIABLE state

 $\$  The commit point learned by each server.

VARIABLE commitPoint

 $\*$  A sequence of log entries.

#### VARIABLE log

 $\$  The current sync source of each server, if it has one.

enanduk kenili di kerdiki

VARIABLE syncSource



**Initial State Predicate** 

a man tot the difference of the second second



a marchitel and the little

#### **Next State Relation**

 $\$  Defines how the variables may transition.

#### Next ==

- $\*$  -- Replication protocol
- \/ AppendOplogAction
- \/ RollbackOplogAction
- \/ BecomePrimaryByMagicAction
- \/ ClientWriteAction
- \/ ChooseNewSyncSourceAction
- $\*$  -- Commit point learning protocol
- \/ AdvanceCommitPoint
- \/ LearnCommitPointAction

**Statistics** 

- Original <u>spec</u>
  - o **295 lines** of TLA+ including comments & model checking helpers
- Extended <u>spec</u> that models sync source selection
  - o **378 lines** of TLA+ including comments & model checking helpers



a na chiata



11

신지

1

**Bug Timeline** 

- 1. 2016 Heartbeat Commit Point Propagation (Safety)
- 2. 2018 No Available Sync Source (*Liveness*)
- 3. 2019 Sync Source Cycles (Liveness)
- 4. 2019 Commit Point Held Back on Stale Nodes (Liveness)
- 5. 2019 Initial Solution Fails in 5 Node Replica Set (Safety)

**Bug Timeline** 

#### 1. 2016 - Heartbeat Commit Point Propagation (Safety)

- 2. 2018 No Available Sync Source (*Liveness*)
- 3. 2019 Sync Source Cycles (Liveness)
- 4. 2019 Commit Point Held Back on Stale Nodes (Liveness)
- 5. 2019 Initial Solution Fails in 5 Node Replica Set (Safety)

**Bug 1: Heartbeat Commit Point Propagation, 3 nodes** 

• The action used for propagating commit points via heartbeat

 $\$  Node i learns the commit point from j via heartbeat.

LearnCommitPoint(i, j) ==

- /\ CommitPointLessThan(i, j)
- /\ commitPoint' = [commitPoint EXCEPT ![i] = commitPoint[j]]
- /\ UNCHANGED <<electionVars, logVars, syncSource>>

**Bug 1: Heartbeat Commit Point Propagation, 3 nodes** 

• The invariant we want to check:

```
RollbackCommitted(i) ==
```

- \E j \in Server:
  - /\ CanRollbackOplog(i, j)
  - /\ IsCommitted(i, Len(log[i]))

```
\ The invariant to check.
```

NeverRollbackCommitted == \A i \in Server: ~RollbackCommitted(i)



Bug 1: Heartbeat Commit Point Propagation, 3 nodes, LearnCommitPoint action

- Model checking stats:
  - 3 nodes, a symmetry set
  - Propagate commit point via heartbeats (LearnCommitPoint action)
  - State constraints: globalCurrentTerm ≤ 3, ∀ i ∈ Server : Len(log[i]) ≤ 3
  - o Invariant: NeverRollbackBeforeCommitPoint
  - 2015 Macbook Pro, 3 CPU Cores (3.1 GHz Intel Core i7)
  - Invariant violation found in ~2 seconds after generating ~500 distinct states
  - 9177 distinct states in full state space
  - 。 <u>TLC config</u>

**Bug 1: Heartbeat Commit Point Propagation, 3 nodes** 

• We can try fix the protocol with a new action definition:

LearnCommitPointFromSyncSource(i, j) ==

/\ ENABLED AppendOplog(i, j) \\* only learn commit point from sync source.

/\ LearnCommitPoint(i, j)



Bug 1: Heartbeat Commit Point Propagation, 3 nodes, LearnCommitPointFromSyncSource action

- Model checking stats:
  - 3 nodes, a symmetry set
  - Propagate commit point via sync source (LearnCommitPointFromSyncSource action)
  - State constraints: globalCurrentTerm ≤ 3, ∀ i ∈ Server : Len(log[i]) ≤ 3
  - o Invariant: NeverRollbackBeforeCommitPoint
  - 2015 Macbook Pro, 3 CPU Cores (3.1 GHz Intel Core i7)
  - No errors found, TLC finished in ~3 seconds
  - 7402 distinct states in full state space
    - 1775 fewer states than the previous model
  - <u>TLC config</u>

**Bug 1: Heartbeat Commit Point Propagation, 3 nodes** 

- The protocol appears to be safe when using the sync source propagation rule
- Is it safe with more than 3 nodes, though?



Bug 1: Heartbeat Commit Point Propagation, 5 nodes, LearnCommitPointFromSyncSource action

- Model checking stats:
  - 5 nodes, a symmetry set
  - Propagate commit point via sync source (LearnCommitPointFromSyncSource action)
  - State constraints: globalCurrentTerm ≤ 3,  $\forall$  i ∈ Server : Len(log[i]) ≤ 3
  - o Invariant: NeverRollbackBeforeCommitPoint
  - Ubuntu 16.10 workstation, 10 CPU cores (3.40GHz Intel Core i7)
  - Invariant violation found in ~2 seconds after generating ~3000 distinct states
  - **230,091 distinct states** in full state space, TLC finished in ~1 min.
  - This bug was never found in production or testing
  - TLC config

#### **Bug 1: Heartbeat Commit Point Propagation, 5 nodes**



en andra Sterille Hiller at the fo

mongoDB

**Bug Timeline** 

- 1. 2016 Heartbeat Commit Point Propagation (Safety)
- 2. 2018 No Available Sync Source (Liveness)
- 3. 2019 Sync Source Cycles (Liveness)
- 4. 2019 Commit Point Held Back on Stale Nodes (Liveness)
- 5. 2019 Initial Solution Fails in 5 Node Replica Set (Safety)

**Bug 2: No Available Sync Source** 

 We define a liveness property of commit points that we want to hold true

 $\*$  At any time, if two nodes' commit points are not the same, they

 $\*$  will be the same eventually.

CommitPointEventuallyPropagates ==

/\ \A i, j \in Server:

(commitPoint[i] /= commitPoint[j] ~> commitPoint[i] = commitPoint[j])

**Bug 2: No Available Sync Source** 

#### Slight modification needed to account for perpetual rollbacks

 $\$  At any time, if two nodes' commit points are not the same, they

 $\$  will be the same eventually.

 $\$  This is checked after all possible rollback is done.

CommitPointEventuallyPropagates ==

```
/\ \A i, j \in Server:
```

(commitPoint[i] /= commitPoint[j] ~>

(~ENABLED RollbackOplogAction => commitPoint[i] = commitPoint[j]))

a marchild and the life

**Bug 2: No Available Sync Source** 

- Demonstrated the original liveness bug with TLC
  - $_{\circ}$  3 nodes
  - o Property: CommitPointEventuallyPropagates
  - 2015 Macbook Pro, 3 CPU Cores (3.1 GHz Intel Core i7)
  - **Temporal Property Violation** found in 1 min. 06s
  - 19,694 distinct states generated
  - <u>TLC config</u>



**Bug Timeline** 

- 1. 2016 Heartbeat Commit Point Propagation (Safety)
- 2. 2018 No Available Sync Source (*Liveness*)
- 3. 2019 Sync Source Cycles (Liveness)
- 4. 2019 Commit Point Held Back on Stale Nodes (Liveness)
- 5. 2019 Initial Solution Fails in 5 Node Replica Set (Safety)



**Bug 3: Sync Source Cycles** 

- We add an action to model sync source selection
- And then specify our correctness property
  - We will specify this particular liveness property as an invariant



**Bug 3: Sync Source Cycles** 

```
\ Server i chooses server j as its new sync source.
\uparrow i can choose j as a sync source if log[i] is a prefix of log[j] and log[j] is longer than log[i].
ChooseNewSyncSource(i, j) ==
   /\ \/ IsLogPrefix(i, j)
      \ If logs are equal, allow choosing sync source if it has a newer commit point.
      // / \log[i] = \log[j]
         /\ commitPoint[j].index > commitPoint[i].index
   /\ state[i] = Follower \* leaders don't need to sync oplog entries.
   /\ syncSource' = [syncSource EXCEPT ![i] = j]
   /\ UNCHANGED <<electionVars, logVars, commitPoint>>
```

**Bug 3: Sync Source Cycles** 

• Specifying the case of a 2 node cycle is much easier:

```
\* Does a 2 node sync source cycle exist?
SyncSourceCycleTwoNode ==
  \E s, t \in Server :
    /\ s /= t
    /\ syncSource[s] = t
    /\ syncSource[t] = s
```



**Bug 3: Sync Source Cycles** 

- We can also specify the general case i.e. a multi-node cycle
- Core idea: model the sync source spanning tree/graph in TLA+





**Bug 3: Sync Source Cycles** 

 $\$  The set of all paths (with bounded length) in the node graph that consists  $\$  of edges <<s,t>> where s has t as a sync source.

The shipt of the state

SyncSourcePaths ==

- {p \in BoundedSeq(Server, Cardinality(Server)) :
  - \A i \in 1..(Len(p)-1) : syncSource[p[i]] = p[i+1]}

#### **Bug 3: Sync Source Cycles**

\\* Is there a non-trivial path in the sync source graph from node i to node j?
\\* This rules out trivial paths i.e. those of length 1.
SyncSourcePath(i, j) ==
 \E p \in SyncSourcePaths :
 /\ Len(p) > 1
 /\ Len(p) > 1
 /\ p[1] = i \\* the source node.
 /\ p[Len(p)] = j \\* the target node.

CONTRACTOR OF THE

```
\* Does a general (multi-node) sync source cycle exist?
SyncSourceCycle == \E s \in Server : SyncSourcePath(s, s)
```



**Bug 3: Sync Source Cycles** 

• Finally, we can ask specifically for cycles of size > 2

 $\*$  The sync source cycle predicate.

NonTrivialSyncCycle == SyncSourceCycle /\ ~SyncSourceCycleTwoNode

 $\$  The invariant.

NoNonTrivialSyncCycle == ~NonTrivialSyncCycle



**Bug 3: Sync Source Cycles** 

- Model checking stats:
  - 4 nodes, a symmetry set
  - State constraints: globalCurrentTerm ≤ 3, ∀ i ∈ Server : Len(log[i]) ≤ 3
  - o Invariant: NoNonTrivialSyncCycle
  - <sup>o</sup> 2015 Macbook Pro, 3 CPU Cores (3.1 GHz Intel Core i7)
  - Invariant violation found in ~6 seconds
  - 226,262 distinct states in full state space
  - Multi-node sync source cycle was never seen in production or testing
  - <u>TLC config</u>

### Takeaways

. nation!



4.11.11



- Hard to know if a protocol is really correct without a formal model
  - It's very difficult for humans to reason about edge cases of concurrent/distributed algorithms
- Even very simple and abstract models can help catch non-trivial bugs
  - No models allowed more than 3 log entries per node
  - Asynchronous message passing was not modeled explicitly in our spec





- We expect that formally modeling our system upfront could have saved 100s of hours of engineering time
  - Multi-year effort to root out all these bugs
  - Only took a few weeks to model and check the protocol using TLA+
- Future goal is to integrate TLA+ into design process at MongoDB





 The specs and models used can be found here: <u>https://github.com/will62794/mongo-repl-tla-models</u>



