

eXtreme Modelling in Practice

A. Jesse Jiryu Davis
MongoDB, Inc.
1633 Broadway
New York, NY 10019
jesse@mongodb.com

Max Hirschhorn
MongoDB, Inc.
1633 Broadway
New York, NY 10019
max.hirschhorn
@mongodb.com

Judah Schvimer
MongoDB, Inc.
1633 Broadway
New York, NY 10019
judah@mongodb.com

ABSTRACT

Formal modelling is a powerful tool for developing complex systems. At MongoDB, we use TLA⁺ to model and verify multiple aspects of several systems. Ensuring conformance between a specification and its implementation can add value to any specification; it can avoid transcription errors, prevent bugs as a large organization rapidly develops the specified code, and even keep multiple implementations of the same specification in sync. In this paper, we explore model-based testing as a tool for ensuring specification-implementation conformance. We attempted two case studies: model-based trace-checking (MBTC) in the MongoDB Server’s replication protocol and model-based test-case generation (MBTCG) in MongoDB Realm Sync’s operational transformation algorithm. We found MBTC to be impractical for testing that the Server conformed to a highly abstract specification. MBTCG was highly successful for Realm Sync, however. We analyze why one technique succeeded and the other failed, and advise future implementers making similar attempts at model-based testing.

PVLDB Reference Format:

A. Jesse Jiryu Davis, Max Hirschhorn, Judah Schvimer. eXtreme Modelling in Practice. *PVLDB*, 13(9): 1346 - 1358, 2020.
DOI: <https://doi.org/10.14778/3397230.3397233>

1. INTRODUCTION

It is well established that formal modelling catches bugs, aids software design, and improves code quality [29]. Formal methods are used increasingly to gain confidence in mission-critical, highly complex systems in multiple domains across many companies. Amazon [4, 6, 29], Intel [3, 19], Microsoft [38], and Springer Verlag [28] have written about their uses of formal methods and the value they’ve gained from modelling and verifying industrial software. But software engineers in industry are often skeptical of a model’s value unless it is shown to match the implementation [29, 45], due to the risk of *transcription bugs*; that is, logic errors in the implementation of the specification.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 13, No. 9

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3397230.3397233>

Furthermore, it can be difficult to ensure in a large engineering organization that changes in code are accompanied by changes in the specification. *eXtreme Modelling*, in which “model and implementation are developed in parallel” [10], has been proposed to prevent implementations and specifications from diverging. Of the several guidelines included in this method, we attempted the following:

1. Multiple specifications model aspects of the system
2. Specifications are written just prior to the implementation
3. Specifications evolve with the implementation
4. Tests are generated from the model and/or trace-checking verifies that test traces are legal in the specification

Although eXtreme Modelling was not intended for existing code bases and specifications, we attempted to apply it to two of our existing systems. We intended to lay the groundwork to use this technique on future projects in which we would write the specification before the implementation, then keep the two in conformance as both evolved.

1.1 Contributions

We undertook two case studies simultaneously, using two separate code bases. We attempted model-based trace-checking (MBTC) in the MongoDB Server. One of our requirements was that the marginal cost of implementing MBTC for each future specification should be cheap, following an initial investment in trace-checking infrastructure. After a great deal of effort, we determined that future specifications would not be easily trace-checked, and decided MBTC was impractical for the MongoDB Server. We were impeded by the system’s enormous code size, by its complex internal concurrency control, and by the oversimplifications we had made when writing our TLA⁺ specifications, before we knew they would be used for MBTC. Concurrently, in MongoDB Realm Sync, a separate project with its own code base, we succeeded in using model-based test-case generation (MBTCG) to ensure conformance between an existing implementation, a newly written specification, and a subsequent implementation of the same specification in a different programming language.

This paper shows what factors made MBTCG successful and MBTC unsuccessful, identifies lessons learned in both techniques, and suggests avenues of future research to make MBTC more practical.

We provide background to this research in Section 2. Section 3 reviews related work. Section 4 describes our attempt

at MBTC, and the challenges met. Section 5 describes our success with MBTCG. Finally, Section 6 describes our conclusions and avenues for future work.

2. BACKGROUND

2.1 MongoDB Server

The MongoDB Server is a distributed document database with features that are traditionally found both in NoSQL and Relational DBMSes. It has support for flexible document schemas, as well as ACID compliant transactions and secondary indexes [20]. The MongoDB Server is typically deployed as a redundant group called a *replica set*, which uses a protocol inspired by Raft [32] to elect a leader node and replicate data changes to follower nodes. Each node maintains a copy of the data and a durable log of operations, the *oplog*. Reads and writes offer multiple consistency and durability levels with increasingly strong guarantees [36, 42]. The MongoDB Server has been under development for 12 years, with hundreds of thousands of lines of code, and dozens of active contributors at any time.

2.2 MongoDB Realm Sync

MongoDB Realm is an offline-first database with a full-duplex synchronization protocol; a client is able to upload new changes to a centralized server without needing to download new changes from the server. (The MongoDB Server and MongoDB Realm are two distinct products with no shared code.) MongoDB Realm’s synchronization protocol, called MongoDB Realm Sync, uses *operational transformation* (OT) to provide automatic conflict resolution and ensure all clients and the server converge to the same data [39].

Each of these peers maintains a copy of the data and a durable log of operations, referred to as the *state* and the *history*, respectively. Since the client and server are not in constant communication, it is likely for a client to have already made subsequent changes to its data and history when it receives new changes from the server. The incoming changes from the server must be “rebased” on top of the client’s current state and history. The client identifies the *merge window* for the incoming changes as all client history entries since the last entry integrated by the server. These entries are causally unrelated to the incoming changes from the server. The OT algorithm transforms the incoming changes to produce a modified set of changes that are consistent for all peers. A *last write wins* rule may be used to resolve the lack of a causal ordering for these operations. A similar procedure also occurs on the server when it receives changes from a client.

MongoDB Realm Sync has been under development for four years, with tens of thousands of lines of code, and three active contributors at any time.

2.3 Testing and TLA⁺

To ensure that the MongoDB Server provides its advertised consistency guarantees with good performance, we run thousands of correctness and performance tests in our continuous integration (CI) system [7]. Both the MongoDB Server and MongoDB Realm rely heavily on randomized testing. For the MongoDB Server, tests randomly perturb the topology state and run randomized commands [12].

Engineers on MongoDB’s Server-Replication Team wrote specifications to verify several concurrent and distributed protocols in the Server [37]. We chose the formal modelling language TLA⁺, introduced by Leslie Lamport [22, 23]. A TLA⁺ specification defines a set of variables expressing a system’s state, and a set of actions which modify this state. The specification may include invariants that must hold at each state, as well as *temporal logic* formulas to express properties of the system’s execution over time. TLA⁺ is a precise mathematical language; thus the properties of a TLA⁺ specification can be proven correct, or validated by a finite model checker which exhaustively explores its behaviors. The most commonly used TLA⁺ model checker is TLC [46].

MongoDB chose TLA⁺ due to its success in industry and academic distributed systems research, open source implementation, and rich tooling. Prior to this paper’s research, TLA⁺ specifications for the MongoDB Server’s replication protocol, its initial sync protocol, and some aspects of its hierarchical locking rules had already been written. TLA⁺ has helped us design a complex, novel reconfiguration protocol, reproduce and fix multiple bugs, and verify liveness properties [37]. Our success with TLA⁺ on the MongoDB Server’s replication protocol led to its use for the MongoDB Realm Sync OT specification.

We plan to write specifications for future features and existing distributed protocols, to gain confidence in critical aspects of our system. We use the same GitHub repository and follow the same code review and testing process with our models as we do for our tests and code. We continuously model-check our TLA⁺ code alongside our correctness tests, to ensure that any changes to our specifications are correct, and to document how the model checker should be configured for each specification.

2.4 MBTC and MBTCG

The eXtreme Modelling approach includes two testing methods that seemed promising for MongoDB: model-based trace-checking (MBTC) [18, 47] and model-based test-case generation (MBTCG) [44].

In MBTC, the behavior of the implementation is observed during testing or in production. The implementation produces an execution trace that captures its sequence of state changes. This trace is checked against the specification, either while the implementation is running or post-hoc, to verify that the trace is a permitted behavior of the specification. The greater the variety of observed behaviors, the more confidence one has in the implementation’s conformance.

In MBTCG, the specification is used to generate test inputs and expected outputs, either exhaustively or from a carefully selected subset [8]. A test framework provides the generated inputs to the implementation and checks whether it responds as specified.

3. RELATED WORK

Our work builds primarily on the eXtreme Modelling approach proposed by Gravel, et. al. [10], and closely related work by Augusto, et. al. [2]. In these papers, the authors discuss the co-evolution of two e-business applications implemented in Java with several small models written in Promela and B. They describe specification engineering methods; for example, holding regular meetings between modellers and implementers, and maintaining a dictionary

to translate between names in the specification and the implementation. Such techniques seem promising, but the authors analyze their effectiveness when practiced on an example application developed by a team of four academics, not on commercial software developed by dozens of engineers. We have attempted to put the ideas of eXtreme Modelling into practice with complex commercial software, and using TLA⁺ and C++ instead of Java, B, and Promela. We especially emphasize the use of MBTC and MBTCG to test that our implementations conform to our specifications. To our knowledge, this is the most thorough attempt to implement eXtreme Modelling in industry.

Besides MBTC and MBTCG, a variety of methods have been proposed to ensure that an implementation conforms to its formal models. Some projects use the same language for the specification and implementation [21], some mechanically generate implementation code from the specification [14], others apply model-checking or formal verification directly to the implementation [5, 13], and yet others perform verified step-wise refinement from the high-level specification all the way down to the implementation [9]. Any of these methods provides high confidence that the implementation conforms to the specification. However, they do not appear to be applicable to the MongoDB Server or to MongoDB Realm. In the research we have reviewed, they are limited to small programs, either written in C or a specialized language such as Prolog; they have not been demonstrated with our implementation language, C++, nor on systems with code sizes as large as ours. They cannot be used when writing new models for a system that has already been implemented.

Compared to the methods above, MBTC and MBTCG are more appropriate for eXtreme Modelling. Considering MBTC first, however, we found that many of the specific uses of MBTC described in the literature are incompatible with eXtreme Modelling.

Ural et. al. in [43] propose to begin with a high-level specification and refine it in stages, using MBTC to gain confidence that each stage’s specification is equivalent to the previous stage. This technique conflicts with two of the eXtreme Modelling guidelines: multiple high-level specifications cannot describe a single system, and co-evolution of specification and implementation is not practical with this technique, because one would have to repeat the effort of multi-stage refinement from top to bottom with each change.

Tasiran et. al. [40] implement MBTC to check that a simulated microprocessor’s behavior conforms to its TLA⁺ specification. After each simulation step they invoke the TLC model checker to determine if the step is permitted by the specification. But before invoking TLC there is a post-processing step implemented in 20,000 lines of C++. This agrees with our experience: we also found that post-processing an execution trace is complex and requires significant engineering investment (Section 4.1). The authors describe a sophisticated method to measure test coverage: they first map the specification’s full state space to a smaller one by eliminating states that are “qualitatively the same” using two TLC features, *symmetries* and *views*. Once the specification state space is reduced they measure the portion of it covered by tests.

The eXtreme Modelling method advocates writing a specification and implementation simultaneously, but in our two case studies we began with a complete imple-

mentation. There is some discussion of modelling existing implementations in [29]. Neubauer et. al. in [28] use machine learning to infer a formal model of an existing system from its observed behaviors. They use MBTC to check if later execution traces of the system still conform to the inferred model; if not, either the system has a bug or the system behaved correctly and the model must be updated. This work approaches the eXtreme Modelling development style, but an inferred model is quite different from one written by humans; the purpose of the former is solely to detect behavior changes in the implementation, the latter expresses the designers’ intent.

The inverse of MBTC is MBTCG, in which a formal specification is used to generate test cases that assert the expected behavior of the implementation. In [8, 25], programs were written in Prolog to perform *partition analysis* and automatically divide the input domain into equivalence classes, such that testing inputs within each class or near the boundary of multiple classes thoroughly explores the implementation’s behavior. Dick and Faivre in [8] automatically ordered sequences of operations to find the shortest test which covers all state transitions. In our MBTCG case study, instead of using partition analysis to reduce the number of generated test cases, it was feasible to exhaustively test *all* combinations of operations subject to a constraint on the length of the initial array.

4. MODEL-BASED TRACE-CHECKING

Because of our growing investment in TLA⁺ specifications of MongoDB Server protocols, we wanted to test whether our C++ implementation conforms to our specifications. MBTCG did not seem appropriate for these protocols due to their non-determinism. A new leader can be elected at any time, for example, and followers can fall behind or catch up unpredictably. We could have injected pauses or other mechanisms into the implementation to force it to follow a sequence of steps that matches a generated test case. This would require a large investment in test code, however, and it would so dramatically change the system’s behavior under test it would cast doubt on the results.

Instead, we attempted to use MBTC. Our proposed method would check execution traces obtained from our existing tests, and we would deploy our trace-checking system to a cluster of continuous integration servers. This would permit us to rapidly develop both the specifications and the implementation, while receiving quick feedback about divergences between them.

4.1 Implementing MBTC

Prior to our research, there were 423 integration tests handwritten in JavaScript that target our replication protocol, and 23 randomized test suites that pause, disconnect, and terminate servers while the replica set is performing regular operations. Our final goal was to capture traces from all our tests and check them against all our specifications.

For our prototype we chose one specification, a 345-line TLA⁺ file called `RaftMongo.tla` in the MongoDB Server repository [27]. The specification’s primary concern is to describe the gossip protocol by which nodes learn the *commit point*: the newest oplog entry that has been replicated by a majority. This specification is based on Diego Ongaro’s TLA⁺ specification for Raft [31]. The same as in Raft, nodes in our specification can assume roles Leader or Follower.

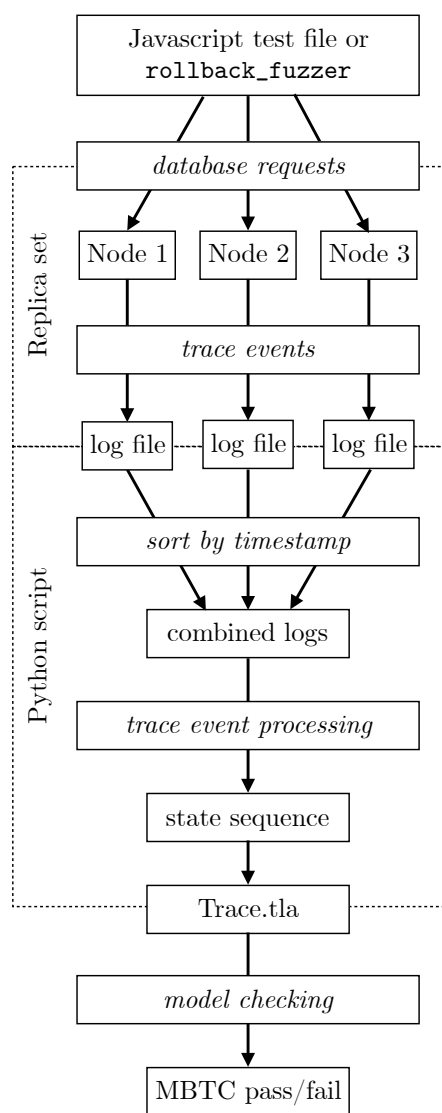


Figure 1: MBTC data pipeline

Time is divided into consecutively numbered election terms. A new term begins whenever a node calls for an election to choose a new leader. One difference from Raft is that the MongoDB Server uses a pull protocol for replication, whereas Raft uses a push protocol; the MongoDB Server’s follower nodes request oplog entries from the leader or other nodes, rather than the leader sending entries to followers. `RaftMongo.tla` is very high-level with few state variables and invariants in order to focus on how the commit point is gossiped. Unlike Ongaro’s work, our specification does not explicitly model messages passed between nodes.

Each node’s state is modelled with four variables:

- **role**: “Leader” or “Follower”
- **term**: Newest election term the node knows
- **commitPoint**: Newest majority-committed oplog entry it knows
- **oplog**: Contents of its oplog

The specification has seven named state transitions:

- **AppendOplog**: A node receives entries from any node
- **RollbackOplog**: A node removes divergent entries
- **BecomePrimaryByMagic**: A node is elected leader instantaneously—the election protocol is abstracted away
- **Stepdown**: A leader becomes a follower
- **ClientWrite**: A leader executes write operations
- **AdvanceCommitPoint**: The leader advances the commit point
- **UpdateTermThroughHeartbeat**: A node learns the election term from any node
- **LearnCommitPointWithTermCheck**: A node learns the commit point from any node
- **LearnCommitPointFromSyncSourceNeverBeyondLastApplied**: A node learns the commit point from a more up-to-date node

When we configure this specification with 3 nodes and constrain the model checker to at most 3 terms and oplogs up to 3 entries long, TLC successfully model-checks it and discovers 371,368 distinct states. TLC validates an invariant that committed writes are not rolled back and a temporal property that the commit point is eventually propagated.

We wrote a C++ procedure `logTlaPlusTraceEvent`, enabled only in testing, which used our logging framework to emit as JSON the values of the four state variables above. We found code locations in the MongoDB Server corresponding to the seven state transitions and added calls to this procedure.

MBTC with a distributed system requires a partial order of trace events; we achieved a strict order by running all processes on one machine and sleeping before logging the trace event. Since all log messages include the current timestamp with millisecond precision, it was sufficient to sleep until the system clock’s millisecond digit changed (Figure 2). (Jard et. al. describe a more general solution with vector clocks [17], which was not necessary in our case.)

We enabled tracing for our 423 handwritten JavaScript tests. Of these, 120 failed due to incompatibilities with tracing. (Section 4.2.2 describes these incompatibilities.) In total, these tests produced 42,262 trace events. We also selected one of our randomized tests, called `rollback_fuzzer`,

```

PROCEDURE logTlaPlusTraceEvent(event) {
  /* Timestamp has millisecond precision */
  Timestamp beforeTime = getCurrentTimestamp()
  Timestamp afterTime = getCurrentTimestamp()
  while (afterTime == beforeTime) {
    sleep 1 millisecond
    afterTime = getCurrentTimestamp()
  }

  assert(afterTime > beforeTime,
         "Clock went backwards")

  log(event, afterTime)
}
  
```

Figure 2: Pseudocode for `logTlaPlusTraceEvent`

for MBTC. This test orchestrates network partitions which cause nodes to temporarily diverge, then to roll back writes and re-synchronize when the partitions are healed. Random CRUD and DDL operations are run against leader nodes in the set to test that, with high probability, all combinations of operations and their behavior on rollback work consistently. Nodes are also randomly restarted to test that clean and unclean restarts during rollback procedures do not cause data corruption. A representative run of `rollback_fuzzer` produced 2,683 trace events.

We wrote a Python script [35] to post-process trace logs. The script merges the nodes’ logs and sorts them by timestamp to obtain a sequence of trace events. Each event describes the state of only *one* node at the moment after it executes a state transition. In order to construct a sequence of states describing the *entire* replica set, the script begins with a known initial state, combines it with the first trace event to determine the next state, and so on. The logic to combine a current state S with a trace event E from node N to produce next state S' is as follows:

- **role:** The script assumes there are never two leaders at once (although this is unrealistic, see Section 4.2.2). Thus if the role of N in E is “Leader” then its role in S' is “Leader” and all others’ roles in S' are set to “Follower”. If N had role “Leader” in S and now has role “Follower” in E , its role in S' is “Follower” and the other nodes’ role values are unchanged.
- **term:** The term of N in S' is set to its term in E . The other nodes’ terms are unchanged.
- **commitPoint** and **oplog:** Similar to term.

For example, in Figure 3, the replica set’s current state has Node 1 as the leader in term 1. The script processes a trace event from Node 2 announcing it has become leader in term 2, arriving at the next state.

Once the script constructs this sequence of states, it implements MBTC following a method proposed by Pressler [33]: it generates a TLA⁺ module called `Trace.tla` which includes the sequence of states (Figure 4), and uses TLC to check that the sequence is permitted by the `RaftMongo.tla` specification. The entire data pipeline can be seen in Figure 1. It was implemented by two engineers in 10 weeks.

Task	Effort	Lines of Code
Event tracing	4 weeks	570 C++
Update RaftMongo.tla	3 weeks	252 TLA ⁺
Python post-processor	3 weeks	484 Python
Total	10 weeks	

4.2 Analysis

We had intended to trace-check all of our specifications against traces from all of our handwritten and randomized tests, deploy the trace-checker to our continuous integration system, and measure accumulated state space coverage over all tests. We had hoped that the marginal cost of checking each specification after the first would decrease. Had we achieved this, we would have built much of the test infrastructure required for eXtreme Modelling. However, we applied trace-checking to only 5 handwritten tests and one randomized test. Only one handwritten test generated traces that passed the trace-checker; the other 4 produced traces that violated the specification due to two implementation

discrepancies (see the discussion of initial sync and term in Section 4.2.2). We did not deploy to continuous integration nor measure coverage. The effort to implement MBTC proved so costly, we abandoned the project after 10 weeks of engineering effort. We faced complexities with concurrency and locking, discrepancies between our models and implementation, and incomplete support in TLC. With further effort we might have achieved MBTC for one specification, but we estimated that the marginal cost of checking each additional specification would approach the cost of the first.

In the following sections we describe our issues trace-checking `RaftMongo.tla`, argue that additional specifications would be just as costly, and offer advice for future implementers of MBTC.

4.2.1 Visibility, hierarchical locking

Visibility: When we began to add tracing to the MongoDB Server, we realized each trace event must be logged after it has occurred, but *before* the change is visible to other nodes. For example, when a leader receives a write from a client application and creates an oplog entry, it must log the `ClientWrite` event after the entry appears in its own oplog, but before any followers can replicate the entry and log an `AppendOplog` event. If a follower logged an `AppendOplog` event with an earlier timestamp than the leader’s `ClientWrite` event, the trace would violate the causal relationship described in the specification. To ensure each state change was logged before it became visible to other nodes, our logging code had to hold several locks.

Hierarchical locking: Formal specifications of distributed systems algorithms, such as Raft, model a

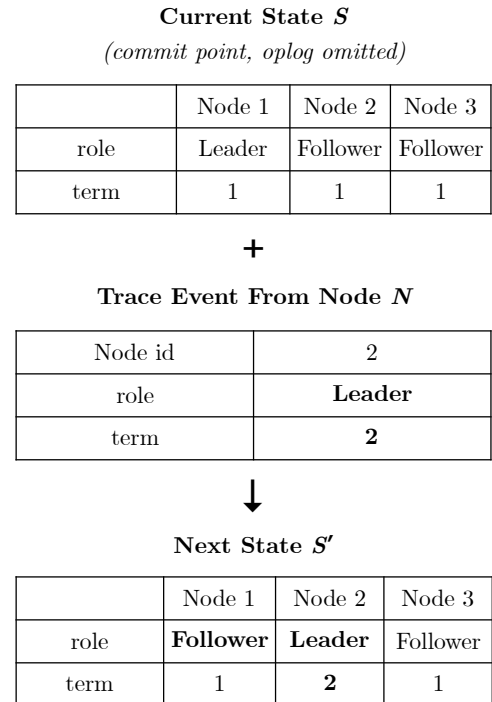


Figure 3: Trace Event Processing

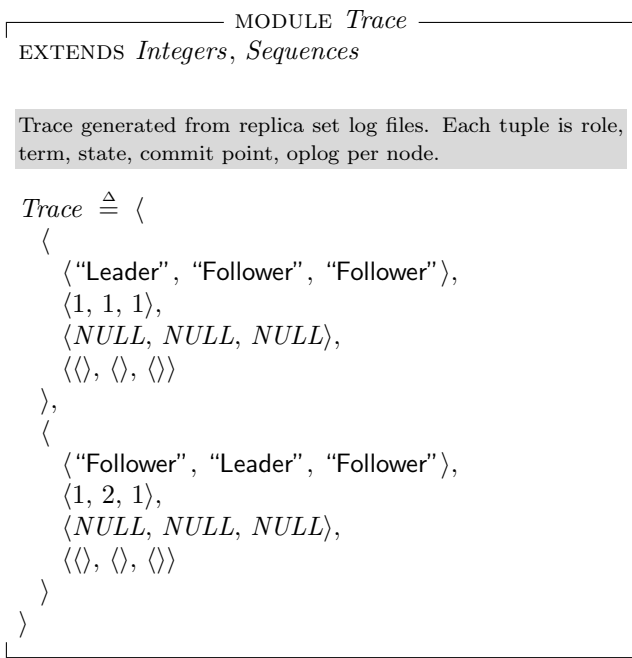


Figure 4: State sequence as TLA⁺ tuple (simplified)

concurrent system of interacting processes, but they typically model each process as single-threaded. Production database systems such as the MongoDB Server, however, almost always have high intra-process concurrency and employ some degree of hierarchical locking [11]. The MongoDB Server specifically combines hierarchical locking with storage-engine Multi-Version Concurrency Control (MVCC), C++ latches, and higher level concurrency control primitives like futures. It may not be feasible to log a consistent snapshot of such a process’s state at the moment of a trace event.

One approach we considered, was to only log the state that actually changed with each event, and use post-processing to fill in the missing state. We already used this approach to only log one process’s state at a time rather than attempting to snapshot multiple processes. We chose not to do this for fear of masking bugs if we incorrectly post-processed trace events. In hindsight it may have made trace logging more cost effective with little loss in value.

To implement MBTC for `RaftMongo.tla`, a node must include the contents of its oplog with each trace event, but acquiring the locks to obtain a snapshot of the oplog proved difficult. Suppose our `logTlaPlusTraceEvent` procedure (Figure 5) must acquire locks A, B, and C, in that order, to read the oplog. Consider a procedure `becomeLeader` that acquires locks A and C, then changes the node’s role to Leader. If we add a call from `becomeLeader` to `logTlaPlusTraceEvent`, the latter must acquire lock B, but this is the wrong order and risks deadlocking with other threads. If `becomeLeader` first *dropped* lock C, then acquired locks B and C, that would be the correct order. But dropping lock C could allow a thread servicing an external request to communicate the node’s new role to another process, thus making the role change *visible* and violating the visibility rule described above. General solutions are unpalatable: callers of `logTlaPlusTraceEvent` could be

```

PROCEDURE becomeLeader() {
  acquire Lock A
  acquire Lock C
  role := Leader
  logTraceEvent()
}

PROCEDURE logTlaPlusTraceEvent() {
  /* Wrong acquisition order if called by
   * becomeLeader, risks deadlock */
  acquire Lock A if not yet acquired
  acquire Lock B if not yet acquired
  acquire Lock C if not yet acquired
  read oplog
  /* Remainder of procedure as in Fig. 2 */
}

```

Figure 5: Pseudocode for a node becoming a Leader

responsible for acquiring all the locks it will need, but this encodes intimate knowledge about `logTlaPlusTraceEvent` into its callers, and it significantly alters the system’s behavior under test.

Solving both the visibility and the locking challenges, for each of the seven named state transitions in `RaftMongo.tla`, was the single most difficult aspect of our MBTC implementation, costing roughly a month of engineering effort. Eventually we discovered code locations that obeyed each transition’s visibility requirements. We managed to avoid the need for acquiring locks out of order by exploiting the MongoDB Server’s MVCC features: our storage engine can serve reads from a stale snapshot of the oplog, instead of locking the oplog to read its current contents. Fortunately, in each location where `logTlaPlusTraceEvent` could not acquire the locks needed to read the current oplog, a stale snapshot was permitted by the specification. In each code location where the specification required the most recent oplog, on the other hand, it was possible to lock it.

In contrast to our difficulties adding tracing to the MongoDB Server, Gravel et. al. [10] emphasize how easily they added tracing. We believe their task was easier because 1) their application was a small prototype, 2) they had deliberately written their specification and implementation to closely correspond (see Section 4.2.3), and 3) their application did not support the complex intra-process concurrency of a typical database. We were surprised at the difficulty of adding tracing to the MongoDB Server, both because it had been easy for Gravel et. al. and because we had not anticipated how visibility and locking would complicate our code. We expect any MBTC implementation for a concurrent database would encounter similar complexities with hierarchical locking and visibility of state changes.

4.2.2 Implementation discrepancies

As soon as we checked our tests’ execution traces against `RaftMongo.tla`, we encountered violations.

Arbiters: MongoDB Server nodes can be configured as *arbiters* which vote in elections but have no data. `RaftMongo.tla` does not model arbiters and we did not implement tracing for them; arbiters crash when tracing is enabled. It was trivial to skip tests that use arbiters, but this gave us a hint that our specification was an

oversimplification of the real system, and would not be suitable for MBTC.

Initial sync: The second discrepancy was in the behavior of *initial sync*, the process a newly-added replica set member uses to obtain a copy of the data and oplog. Trace-checking the output of `rollback_fuzzer` immediately reproduced a known violation of the specification: our implementation considers an oplog entry to be majority-committed once a quorum of nodes, including initial-syncing nodes, has replicated it. However, initial-syncing nodes should not be considered quorum members because their oplog entries are not durable until initial sync completes. When we first wrote `RaftMongo.tla` we modelled the correct behavior. We did not have MBTC in mind, so we deliberately wrote an idealized specification with the intention to eventually bring our implementation into conformance. This violation increased our confidence in trace-checking. It illustrated how well MBTC works in practice, and had we not already known about this behavior, it would have been a significant discovery. However, it made our MBTC project nearly worthless: the violation came only 4 steps from the trace’s start and left the remaining 2,683 steps unchecked.

To permit deeper trace-checking, we could 1) fix the implementation of initial sync sooner than planned, 2) avoid triggering the non-conforming behavior in testing, 3) update the specification to match the flawed implementation, or 4) post-process the traces to simulate a conformant implementation. Fixing the implementation would require substantial work and is already scheduled for a future release. Instead, we chose solution 2 and modified `rollback_fuzzer` so all followers were fully synced before the test began any writes.

Two leaders: As in Raft, MongoDB’s election protocol permits two leaders to co-exist briefly. However, we chose in `RaftMongo.tla` to assume there is at most one leader, since the election protocol is not the specification’s main concern. To avoid trace-checking failures, we again chose solution 2, and avoided tests exhibiting two leaders.

Term: Other discrepancies required solution 3, updating the specification. Since the election protocol was not its focus, `RaftMongo.tla` originally modelled the election term as a single global number known by all nodes. In reality, election terms are gossiped among nodes and each learns the new term at a different time. MBTC required us to make the specification more complex to match reality. We chose to update the specification to resolve many other discrepancies, as well: over the course of our research we added or changed 252 of the 345 lines of TLA⁺ in `RaftMongo.tla`, costing three weeks of effort.

Copying the oplog: In several cases we chose solution 4, post-processing the logs to simulate conformance. For example, in `RaftMongo.tla`, when a node performs an initial sync, it copies the leader’s entire oplog. In the implementation, a new node copies only recent entries. The `RaftMongo.tla` behavior more closely matches the Raft protocol that inspired it, and makes it simpler to express invariants such as “an entry is never rolled back after it is replicated to a majority of nodes’ oplogs.” We resolved the discrepancy by adding logic to our Python script that filled in the missing entries while it generated the state sequence. Such an intrusive modification made us concerned that a mistake in the Python script might mask a harmful transcription bug in the MongoDB Server code. On the other hand, choosing solution 3 and adding detail to the specification might ex-

plode the state space and make model-checking intractable. At the least, it would make the specification harder to read.

Each discrepancy we discovered through MBTC required us to judge which of the four solutions to employ. Ideally, when an organization practices eXtreme Modelling, they resolve each MBTC failure either by fixing the implementation or, if the failure is not a bug, by updating the specification. In practice, we often concluded it was expedient to work around violations. We chose to avoid the non-conforming behavior (solution 2) or to simulate a conformant behavior (solution 4) depending on which took less effort and least undermined our confidence in the test. Unfortunately, our choices were based on estimates and speculation.

We might have resolved some discrepancies with a refinement mapping [1] from `RaftMongo.tla` to a lower-level specification which would be used for MBTC. In the “copying the oplog” discrepancy above, it seems possible to use a history variable [24] to bridge the gap between our specification and a more realistic one. In other cases (two leaders, term gossiping) either no refinement mapping exists or it would require great effort to find.

4.2.3 Modelling for trace-checking

There appears to be a conflict between writing a specification for documenting and model-checking a design, and writing one that can be trace-checked.

We tried to implement MBTC with models that had been written for documentation and model-checking, but we found that MBTC is only practical if the specification is written with MBTC in mind. If we began again, we would rewrite our specifications to closely correspond to the implementation. The specification’s major components would be structured similarly to the implementation’s, and critical sections of the implementation would correspond with single actions in the specification. In the case of multi-step events such as elections, we would model them with multiple actions, instead of expressing them as instantaneous single actions as we did with `BecomePrimaryByMagic` in `RaftMongo.tla`. We would faithfully model the flaws in the implementation if we did not plan to fix them immediately.

We would rewrite our specification to model events that are easily observed in the implementation: for example, we might model protocol messages, akin to the original Raft specification. We would try to avoid modelling state that is difficult to snapshot, especially if it is protected with complex locking. If some variables in the specification cannot be logged by the implementation, Pressler proposes a refinement mapping technique [33] in which TLC checks whether there is any sequence of values for the missing variables that would permit the trace to match the specification.

Most of these changes in our modelling style would be unobjectionable. However, there is the danger that a less abstract specification has a state space too large to model-check in reasonable time. The limited changes we made to `RaftMongo.tla` increased the state space from 42,034 states to 371,368, and model-checking time from 2 seconds to 14 minutes. To complete our MBTC project would have required a more detailed specification, or applying Pressler’s refinement mapping technique, both of which would threaten to explode the state space even more.

4.2.4 Tooling and TLC

Tool support for MBTC with TLA⁺ and TLC is a work in progress. Pressler’s method worked well to check traces of hundreds of events, but for thousands of events it was impractically slow. Pressler proposed, and Markus Alexander Kuppe has begun to implement, features to check long traces by bypassing the TLA⁺ parser in favor of a special-purpose Java extension to TLC [41]. MBTC with TLA⁺ will be more convenient once these features are released, with publicly available example code. Another missing feature is the ability to combine state-space coverage reports over multiple TLC executions on different traces, which would permit engineers to calculate the total coverage achieved by deploying MBTC to continuous integration.

4.2.5 Marginal cost of trace-checking a specification

Our experience with `RaftMongo.tla` convinced us that trace-checking additional specifications would be nearly as costly as the first. For example, we imagined a scenario in which we had succeeded at trace-checking `RaftMongo.tla` and we moved on to `Locking.tla` [26], which models aspects of the MongoDB Server’s lock hierarchy. Once again we would begin with a highly abstract specification written for documentation and model-checking, and we would repeat the effort of adding detail and resolving discrepancies to make the specification suitable for MBTC. The state variables of `Locking.tla` are disjoint from those modelled by `RaftMongo.tla`, so there would be little code reuse in the tracing implementation or the post-processing script. The locking specification applies only to one process, not to an entire replica set, so the log post-processing script would be simpler for `Locking.tla` than for `RaftMongo.tla`; however, this difference would further preclude code reuse.

We canceled the MBTC project because we had achieved disappointingly little utility from our efforts, and our revised estimate of the marginal cost made it clear that the project was not worthwhile.

5. MODEL-BASED TEST-CASE GENERATION

Our second case study involved a distinct product, MongoDB Realm Sync. Architectural decisions have led us to re-implement the MongoDB Realm Sync server in Golang while keeping the clients in C++. Originally, a shared C++ code base ensured the client’s and server’s behavior were identical. Now that there are two implementations, it is paramount they resolve conflicts identically. Designing operational transformation (OT) functions that ensure peers converge to the same data is already challenging and error-prone [15, 16, 34]; implementing it the same way twice is even more difficult.

Model-based test-case generation (MBTCG) [10] provided means to confidently assess whether parity had been achieved while both the C++ and Golang code bases were under active development. We wrote a new TLA⁺ specification for this purpose and used the state space explored by the TLC model checker to generate C++ and Golang test cases for aspects of the OT algorithm. These test cases can be automatically generated and run after any change to the model, and ensure that both the C++ and Golang implementations make appropriate changes to maintain equivalence.

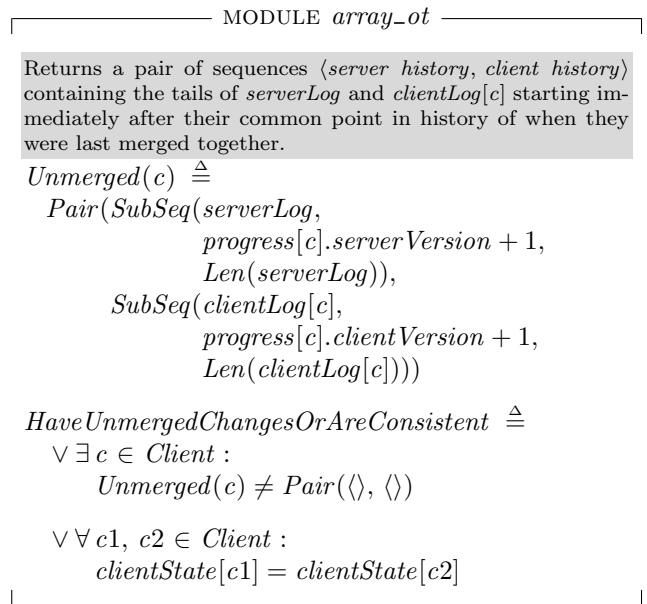


Figure 6: TLA⁺ invariant for MongoDB Realm Sync

MongoDB Realm Sync has 19 distinct operations which can be performed on a group of tables, an individual table, an object, or a list of values. These include operations such as deleting a table, creating a new object in a table, setting the property of an object, and inserting a new element into a list. Each type of operation must define how it merges with all other types of operations. This yields $19(19 + 1)/2 = 190$ merge rules that must be defined, with the remaining $19^2 - 190 = 171$ merge rules inferred by symmetry. Approximately three-quarters of the merge rules have trivial implementations where the incoming operation is applied unchanged by both peers.

The most complex merge rules are for the six array-based operations:

- **ArraySet**: replacing the value of an existing element
- **ArrayInsert**: inserting a new element at a position within the list, or growing the size of the list by one
- **ArrayMove**: moving an element from one position to another
- **ArraySwap**: swapping the position of two elements
- **ArrayErase**: removing an element from the list
- **ArrayClear**: removing all of the elements from the list

The $6(6 + 1)/2 = 21$ merge rules for these six operations are implemented in approximately 1,000 lines of C++.

5.1 Writing the TLA⁺ specification

MongoDB’s Realm Sync Team wrote a TLA⁺ specification for these array-based operations to assess the soundness of the existing C++ implementation and to exhaustively generate test cases for a new Golang implementation.

Writing the TLA⁺ specification took approximately 40 hours over the course of two weeks. It was written by copy-pasting the C++ code and manually updating the syntax to be valid TLA⁺. The merge rule for two **ArrayInsert** operations was written first along with the


```

MODULE array_ot
Transform_ArrayErase_ArraySet(eraseOp, setOp) ≜
CASE setOp.ndx = eraseOp.ndx →
    Set the removed element.
    Pair(⟨eraseOp⟩, ⟨⟩)
□ setOp.ndx > eraseOp.ndx →
    Set an element after the position of
    the removed element.
    Pair(⟨eraseOp⟩,
        ⟨[setOp EXCEPT !.ndx = @ - 1]⟩)
□ OTHER → Pair(⟨eraseOp⟩, ⟨setOp⟩)

```

Figure 7: TLA⁺ code for merging ArrayErase and ArraySet operations

HaveUnmergedChangesOrAreConsistent invariant (see Figure 6). The merge rules for each subsequent operation type were then added and checked with TLC after each addition. TLC found transcription errors multiple times during this process. The counterexample produced was manually analyzed by associating the branches taken in the series of if-statements in the C++ code with the nested IF and CASE expressions taken in the TLA⁺ specification.

This transcription process was challenging in two ways. First, the C++ code relies on mutable variables for expressing how the incoming operations must be modified to be correctly applied by the non-originating peer, whereas TLA⁺ variables cannot be mutated in the same way. Second, the order in which clients perform operations locally and then later merge with the server doesn't affect their final state, but, if left unconstrained, would quickly lead to a state space explosion.

5.1.1 Handling C++ mutability and flow control

TLA⁺ is not a general-purpose programming language. *Operators* in TLA⁺ (which behave most closely to a function in a general-purpose programming language) cannot reassign an argument or have multiple statements as in C++ functions. While PlusCal offers an imperative syntax, it only supports reassigning state variables—and not arguments within an operator—because it transpiles to TLA⁺. As shown in Figure 8, the C++ variable `set_op` is mutated to produce the version of these operations to be applied on the non-originating peer. This mismatch between the two languages makes it more difficult to formalize a C++ source of truth into a TLA⁺ specification, especially in the presence of nested if-statements. Additionally, a series of multiple if-statements must be transcribed to explicitly define all branch combinations that may or may not be taken consecutively. Fortunately, the TLC model checker was readily able to catch human transcription errors as safety violations. Some examples of transcription errors were forgetting to substitute the updated index number in later comparisons, or entirely forgetting to embed a deeply-nested or later occurring if-statement as another nested IF or CASE expression in TLA⁺.

```

DEFINE_MERGE(ArrayErase, ArraySet)
{
  if (same_array(erase_op, set_op)) {
    // LCOV_EXCL_STOP
    if (set_op.ndx == erase_op.ndx) {
      // CONFLICT: Update of a removed element.
      // RESOLUTION: Discard the ArraySet
      // operation.
      set_op.discard();
    }
    else if (set_op.ndx > erase_op.ndx) {
      set_op.ndx -= 1;
    }
    // LCOV_EXCL_START
  }
}

```

Figure 8: C++ code for merging ArrayErase and ArraySet operations

5.1.2 Constraining the state space

The state constraint was progressively refined as the TLA⁺ specification was being implemented in order to reduce the run-time required by the TLC model checker. One simplification was a `MergeAction` operator which simultaneously uploads any changes from the client and downloads any changes from the server. Modelling these as distinct actions would only have been beneficial when explicitly modelling time. Instead, the model's purpose is to generate C++ and Golang test cases where clients perform a sequence of operations starting from the same initial array. A client performing a sequence of operations that are causally dependent on another client's sequence of operations is equivalent (in terms of both its state and history) to the former client performing the concatenation of the two sequences on the same initial array. However, a longer sequence of operations necessarily has more combinations and would lead to a state space explosion. We therefore shifted to focus on varying the initial array and restricted all clients to performing a single operation. Since the TLC model checker would explore all combinations of the clients performing a single operation, starting with a sufficiently large initial array would still enable the TLC model checker to exercise every case in the merge rules.

We artificially constrained the state space to have clients perform and merge operations in ascending order of their ID to avoid exploring redundant states. In particular, it doesn't matter which order clients 1 and 2 perform operations locally until they communicate with the server. We aren't able to define the clients as a symmetry set because the ID is used to order operations when their timestamps are equal. As mentioned previously, the specification doesn't model time, so we need at least three clients to capture a client merging both with an earlier operation and with a later operation. Using the minimum of three clients ensures we don't increase the state space unnecessarily. The TLC model checker was run with constraints of three clients each performing a single operation on an array already containing three elements.

5.1.3 Results from the TLC model checker

In addition to detecting safety violations caused by transcription errors, the TLC model checker also encountered

a `StackOverflowError` due to a case in the merge rule for the `ArraySwap` and `ArrayMove` operations leading the merge function to never terminate. This issue was found to also exist in the C++ code; the bug had faithfully been transcribed from C++ to TLA⁺. It was surprising to discover a bug in the C++ implementation due to the maturity of the product and the amount of randomized testing it had undergone. The discovery of this issue became the deciding factor to not support a dedicated `ArraySwap` operation in the new Golang server implementation and to deprecate it in the existing C++ implementation. The `ArraySwap` operation was therefore excluded from testing for the remainder of this experiment.

Even if the model checker does not find a violation of the `HaveUnmergedChangesOrAreConsistent` invariant, this does not imply that the specification accurately describes the C++ implementation. For example, a trivial way to achieve convergence (at the cost of user-intent preservation) would be to have every merge rule clear the contents of the array. We must test that the behavior of the specification conforms to the C++ code.

5.2 Generating C++ test cases

The state space explored by the TLC model checker provides an exhaustive set of test cases for how a given pair of operations should be transformed when applied by the non-originating peer. The TLC model checker supports writing the graph of all reachable states to a GraphViz DOT file, where each edge is labeled with the action taken to progress from one state to the next. We wrote a Golang program to parse this file and generate C++ test cases composed of (1) the initial array, (2) the operations each client performed, (3) the transformed version of the operations each client applied after merging the original operations from the other clients, and (4) the final state of the array. The Golang program then produces a C++ program which uses MongoDB Realm Sync’s unit test framework. The `fixture.sync_all_clients()` call takes the place of performing the merge action for all the clients (in the same order the TLA⁺ specification performs the merge).

For an initial array containing three elements and with three clients each performing a single operation, the Golang program generated 4,913 C++ test cases. The coverage of the specific sections of the C++ code was measured by adding in `LCOV_EXCL_STOP` and `LCOV_EXCL_START` comments within the block of the `same_array()` condition. This was done because the TLA⁺ specification was designed to only exercise cases where the clients act on the same array. The 36 handwritten C++ test cases covered 18 of the 86 branches (21%) within the merge rules for the array-based operations. This low percentage is excusable because we also rely on fuzz testing. The `fuzz-transform` test executable, which uses the AFL fuzzer [48] to produce randomized inputs that are then mapped to randomized operations, covered 79 of 86 branches (92%) after approximately 8 million total executions. The generated C++ test cases covered all 86 branches (100%).

Achieving 100% branch coverage, with all the generated C++ test cases passing, gives us confidence that the TLA⁺ specification conforms to the C++ code. There may still be programming bugs in the C++ implementation stemming from undefined behavior. However, for the well-formed inputs from the TLC model checker, we can be certain the OT

```
TEST(Transform_Node__6971023528664242108)
{
    size_t num_clients = 2;
    TransformArrayFixture fixture{
        test_context, num_clients, {1, 2, 3}};

    fixture.transaction(0, [](TableRef array) {
        array->set_int(0, 2, 4);
    });
    fixture.transaction(1, [](TableRef array) {
        array->remove(1);
    });

    fixture.sync_all_clients();
    fixture.check_array({1, 4});

    fixture.check_ops(0, {ArrayErase{1}});
    fixture.check_ops(1, {ArraySet{1, 4}});
}
```

Figure 9: C++ test case for merging `ArrayErase` and `ArraySet` operations

algorithm converges and the C++ and Golang implementations of the merge rules for the array-based operations always agree. It was worthwhile for a single engineer to spend one month writing the TLA⁺ specification and implementing the C++ test case generator.

Task	Effort	Lines of Code
array_ot.tla	2 weeks	795 TLA ⁺
C++ test case generator	2 weeks	755 Golang
Total	4 weeks	

6. CONCLUSIONS

MBTC was impractical for our system because we lacked a close correspondence between the MongoDB Server implementation and `RaftMongo.tla`. This was partly by necessity: the MongoDB Server is hundreds of thousands of lines of code; the replication protocol alone is tens of thousands of lines. Any close correspondence with the MongoDB Server’s implementation would have made the specification overly complicated for its purpose, and likely it would have inflated the state space beyond what is feasible to model-check. In the absence of a close correspondence, MBTC requires significant post-processing of traces. We learned that in a system of the MongoDB Server’s complexity, post-processing should be used early and often. Our trace-logging code was hampered by the Server’s intra-process concurrency control. Future engineers in our situation should consider logging only those state variables the implementation can access at the moment it executes a state transition. A post-processing script can fill in the missing variables, so long as there is some reliable way to infer their values post-hoc. For example, it might be guaranteed that the missing variables’ values have not changed since the last time they were logged. Bugs in post-processing could mask implementation bugs, so it is critical to think about what kinds of bugs the MBTC effort aims to catch and what kinds are targeted by other tests.

An avenue we did not explore is to generate traces from implementation modules running in a unit test framework,

rather than an integration test of the entire multi-process system, as we did. Modules are simpler than the whole, and non-determinism can be more easily eliminated, for example by simulating the system clock. The difficulties we experienced with visibility and hierarchical locking would likely be more manageable. By testing modules in isolation, one could sacrifice realism in exchange for implementing MBTC cost-effectively. Especially if there are multiple specifications modelling aspects of the system, as in our case, it would be appropriate to pair a specification with each module.

Regardless, MBTC in a system of the Server’s scale will require significant effort for every specification. No initial infrastructure investment can make the marginal cost of trace-checking additional specifications cheap.

MBTCG succeeded because the specification was deliberately written with a close correspondence to the implementation, via direct transcription. The implementation under test was a self-contained function with defined inputs and output, of approximately one thousand lines of code. By design of the OT algorithm, the order of clients exchanging operations they performed locally does not change the transformed version of these operations or the resulting array. As a result of this symmetry, model-checking the specification is feasible, and the implementation can be completely tested by a constrained set of tests that can be generated, compiled, and run fairly quickly.

MBTC and MBTCG are both promising approaches to ensuring conformance between a specification and its implementation. In a small system, MBTCG was highly effective and accomplished its goal. When MBTCG is possible, and when the state space is not too large, MBTCG is the gold standard of ensuring specification-implementation conformance. If the state space becomes too large though, or if any individual test case would take too long to run, we must start sampling the space, and MBTC can be effective. MBTC is also a good approach when the implementation’s behavior is non-deterministic and cannot be forced to execute specific test sequences. For a very large system, however, either technique may require great effort per specification.

Future work should make both model-based testing approaches more practical. TLC performance was an obstacle, and any improvements would help. Better research and tooling to post-process partial states and merge them together under certain assumptions could make the post-processing safer and easier. More research should determine best practices for writing specifications intended to be used in model-based testing. It should also determine best practices for logging trace events and generating test-cases from a model. Developing tooling for whole-process snapshotting could have greatly simplified MBTC trace logging, since we could have used the snapshots to create trace events. Research into record-and-replay debuggers [30] might be applicable to this area. Making MBTC and MBTCG first-class features in TLC could have saved us a great deal of time. The state space representation in TLC could be improved to be more easily parsed, which would have eased MBTCG. Trace-checking could be built in where users only need to provide a trace and a specification and TLC efficiently checks the traces for them.

Going forward, we plan to continue using TLA⁺ to give us confidence and increase our own understanding in our

complex protocols. MongoDB Realm Sync is actively using and improving its use of MBTCG. More teams and projects throughout MongoDB are starting to use TLA⁺ to gain similar benefits in their own development. These projects may find interest in ensuring specification-implementation conformance, and this research and any further research will help us determine when the investment is worth it, and what model-based testing method we should use.

7. ACKNOWLEDGMENTS

Many experts have advised us, proposed workarounds for obstacles, reviewed our code, and reviewed drafts of this paper, including David Bradford, Mark Callaghan, Mike O’Brien, Siyuan Zhou, Will Schultz, Michael Cahill, Simon Ulsnes, Tess Avitabile, Henrik Edin, Andy Schwerin, and Ron Pressler. We are embarrassingly indebted to Markus Alexander Kuppe for answering our questions about TLC within minutes, and implementing features we requested within hours.

8. REFERENCES

- [1] M. Abadi and L. Lamport. The Existence of Refinement Mappings. Technical Report 29, Digital Equipment Corporation, Systems Research Centre, 14 Aug. 1988.
- [2] J. C. Augusto, Y. Howard, A. M. Gravell, C. Ferreira, S. Gruner, and M. Leuschel. Model-Based Approaches for Validating Business Critical Systems. In *STEP*, pages 225–233. IEEE Computer Society, 2003.
- [3] R. Beers. Pre-RTL Formal Verification: An Intel Experience. In *Proceedings of the 45th Annual Design Automation Conference*, DAC ’08, page 806–811, New York, NY, USA, 2008. Association for Computing Machinery.
- [4] A. Chudnov, N. Collins, B. Cook, J. Dodds, B. Huffman, C. MacCárthaigh, S. Magill, E. Mertens, E. Mullen, S. Tasiran, A. Tomb, and E. Westbrook. Continuous Formal Verification of Amazon s2n. In H. Chockler and G. Weissenbacher, editors, *Computer Aided Verification*, pages 430–446, Cham, 2018. Springer International Publishing.
- [5] A. Chudnov, N. Collins, B. Cook, J. Dodds, B. Huffman, C. MacCárthaigh, S. Magill, E. Mertens, E. Mullen, S. Tasiran, A. Tomb, and E. Westbrook. Continuous Formal Verification of Amazon s2n. In H. Chockler and G. Weissenbacher, editors, *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14–17, 2018, Proceedings, Part II*, volume 10982 of *Lecture Notes in Computer Science*, pages 430–446. Springer, 2018.
- [6] B. Cook. Formal Reasoning About the Security of Amazon Web Services. In H. Chockler and G. Weissenbacher, editors, *Computer Aided Verification*, pages 38–47, Cham, 2018. Springer International Publishing.
- [7] D. Daly, W. Brown, H. Ingo, J. O’Leary, and D. Bradford. Industry Paper: The Use of Change Point Detection to Identify Software Performance Regressions in a Continuous Integration System. 2020.
- [8] J. Dick and A. Faivre. Automating the Generation and Sequencing of Test Cases from Model-Based

- Specifications. In *First International Symposium of Formal Methods Europe on Industrial-Strength Formal Methods (FME '93)*, volume 670 of *LNCS*, pages 268–284. Springer-Verlag, Apr. 1993.
- [9] A. T. Eiriksson and K. L. McMillan. Using Formal Verification/Analysis Methods on the Critical Path in System Design. *Lecture Notes in Computer Science*, 939:367–380, 1995.
- [10] A. M. Gravell, Y. Howard, J. C. Augusto, C. Ferreira, and S. Gruner. Concurrent Development of Model and Implementation. *CoRR*, abs/1111.2826, 2011.
- [11] J. Gray, G. Putzolu, and I. Traiger. Granularity of locks and degrees of consistency in a shared data base. In *Modeling in Data Base Management Systems*. North Holland, Amsterdam, 1976.
- [12] R. Guo. MongoDB’s JavaScript fuzzer. *Communications of the ACM*, 60:43–47, 04 2017.
- [13] G. J. Holzmann and R. Joshi. Model-Driven Software Verification. In S. Graf and L. Mounier, editors, *Model Checking Software, 11th International SPIN Workshop, Barcelona, Spain, April 1-3, 2004, Proceedings*, volume 2989 of *Lecture Notes in Computer Science*, pages 76–91. Springer, 2004.
- [14] S. Houhou, L. Kahloul, S. Benharzallah, and R. Bettira. Framework for Wireless Sensor Networks Code Generation from Formal Specification. pages 35–52, 09 2017.
- [15] A. Imine, P. Molli, G. Oster, and M. Rusinowitch. Proving Correctness of Transformation Functions in Real-Time Groupware. In *Proceedings of the Eighth Conference on European Conference on Computer Supported Cooperative Work, ECSCW’03*, page 277–293, USA, 2003. Kluwer Academic Publishers.
- [16] A. Imine, M. Rusinowitch, G. Oster, and P. Molli. Formal Design and Verification of Operational Transformation Algorithms for Copies Convergence. *Theor. Comput. Sci.*, 351(2):167–183, Feb. 2006.
- [17] C. Jard, T. Jérón, G.-V. Jourdan, and J.-X. Rampon. A General Approach to Trace-Checking in Distributed Computing Systems. pages 396–403, Los Alamitos, CA, USA, June 1994. IEEE Computer Society Press.
- [18] C. Jard and G. von Bochmann. An approach to testing specifications. *Journal of Systems and Software*, 3(4):315–323, 1983.
- [19] R. Kaivola, R. Ghughal, N. Narasimhan, A. Telfer, J. Whittemore, S. Pandav, A. Slobodová, C. Taylor, V. Frolov, E. Reeber, and A. Naik. Replacing Testing with Formal Verification in Intel Core™ i7 Processor Execution Engine Validation. In A. Bouajjani and O. Maler, editors, *Computer Aided Verification*, pages 414–429, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [20] A. Kamsky. Adapting TPC-C Benchmark to Measure Performance of Multi-Document Transactions in MongoDB. *PVLDB*, 12(12):2254–2262, 2019.
- [21] Kern and Greenstreet. Formal Verification In Hardware Design: A Survey. *ACMTDAES: ACM Transactions on Design Automation of Electronic Systems*, 4, 1999.
- [22] L. Lamport. Specifying Concurrent Systems with TLA⁺. In *Calculational System Design*, Amsterdam, 1999. IOS Press.
- [23] L. Lamport. *Specifying Systems: The TLA⁺ Language and Tools for Hardware and Software Engineers*. Addison-Wesley Longman Publishing Co., Inc., USA, 2002.
- [24] L. Lamport and S. Merz. Auxiliary Variables in TLA⁺. *CoRR*, abs/1703.05121, 2017.
- [25] B. Legeard, F. Peureux, and M. Utting. Automated Boundary Testing from Z and B. In L.-H. Eriksson and P. A. Lindsay, editors, *Formal Methods—Getting IT Right*, volume 2391 of *LNCS*, pages 21–40. Springer-Verlag, July 2002.
- [26] TLA⁺ specification for MongoDB Locking Behaviors. <https://github.com/visualzhou/mongo-locking-tla>.
- [27] MongoDB. <https://github.com/mongodb/mongo>.
- [28] J. Neubauer, B. Steffen, O. Bauer, S. Windmüller, M. Merten, T. Margaria, and F. Howar. Automated continuous quality assurance. In S. Gnesi, S. Gruner, N. Plat, and B. Rumpe, editors, *Proceedings of the First International Workshop on Formal Methods in Software Engineering - Rigorous and Agile Approaches, FormSERA 2012, Zurich, Switzerland, June 2, 2012*, pages 37–43. IEEE, 2012.
- [29] C. Newcombe, T. Rath, F. Zhang, B. Munteanu, M. Brooker, and M. Deardeuff. Use of formal methods at Amazon Web Services. 2014.
- [30] R. O’Callahan, C. Jones, N. Froyd, K. Huey, A. Noll, and N. Partush. Engineering Record And Replay For Deployability: Extended Technical Report. *CoRR*, abs/1705.05937, 2017.
- [31] D. Ongaro. TLA⁺ specification for the Raft consensus algorithm. <https://github.com/ongardie/raft.tla>, 2014.
- [32] D. Ongaro and J. K. Ousterhout. In Search of an Understandable Consensus Algorithm. In G. Gibson and N. Zeldovich, editors, *2014 USENIX Annual Technical Conference, USENIX ATC ’14, Philadelphia, PA, USA, June 19-20, 2014*, pages 305–319. USENIX Association, 2014.
- [33] R. Pressler. Verifying Software Traces Against a Formal Specification with TLA⁺ and TLC. <https://pron.github.io/files/Trace.pdf>, 2018.
- [34] A. Randolph, H. Boucheneb, A. Imine, and A. Quintero. On Consistency of Operational Transformation Approach. *Electronic Proceedings in Theoretical Computer Science*, 107:45–59, Feb 2013.
- [35] repl-trace-checker. <https://github.com/mongodb-labs/repl-trace-checker>.
- [36] W. Schultz, T. Avitabile, and A. Cabral. Tunable Consistency in MongoDB. *PVLDB*, 12(12):2071–2081, 2019.
- [37] W. Schultz and S. Zhou. A Bug’s Life: Fixing a MongoDB Replication Protocol Bug with TLA⁺. <https://conf.tlapl.us/>, 2019.
- [38] D. Shukla. Azure Cosmos DB: Pushing the frontier of globally distributed databases. <https://azure.microsoft.com/en-us/blog/azure-cosmos-db-pushing-the-frontier-of-globally-distributed-databases/>, 2018.
- [39] A. Stigsen. Conflict Resolution in Distributed Computing, Nov 2019. Patent No. US20190354540A1,

- Filed May 14th., 2019, Issued Nov. 21st., 2019.
- [40] S. Tasiran, Y. Yu, and B. Batson. Using a formal specification and a model checker to monitor and direct simulation. pages 356–361.
 - [41] TLA⁺ issue 413.
<https://github.com/tlaplus/tlaplus/issues/413>.
 - [42] M. Tyulenev, A. Schwerin, A. Kamsky, R. Tan, A. Cabral, and J. Mulrow. Implementation of Cluster-Wide Logical Clock and Causal Consistency in MongoDB. In *Proceedings of the 2019 International Conference on Management of Data, SIGMOD '19*, page 636–650, New York, NY, USA, 2019. Association for Computing Machinery.
 - [43] H. Ural and R. L. Probert. Automated testing of protocol specifications and their implementations. *Computer Communication Review*, 14(2):149–155, 1984.
 - [44] M. Utting and B. Legeard. *Practical Model-Based Testing: A Tools Approach*. Morgan-Kaufmann, 2006.
 - [45] H. Wayne. Augmenting Agile with Formal Methods. <https://www.hillelwayne.com/post/augmenting-agile/>, 2018.
 - [46] Y. Yu, P. Manolios, and L. Lamport. Model Checking TLA⁺ Specifications. In *Proceedings of the 10th IFIP WG 10.5 Advanced Research Working Conference on Correct Hardware Design and Verification Methods, CHARME '99*, page 54–66, Berlin, Heidelberg, 1999. Springer-Verlag.
 - [47] A. M. G. C. F. Yvonne Howard, Stefan Gruner and J. C. Augusto. Model-Based Trace-Checking. *CoRR*, abs/1111.2825, 2011.
 - [48] M. Zalewski. American fuzzy lop.
<https://lcamtuf.coredump.cx/af1/>.