

# JCST

Vol.37 No.1 Jan. 2022

ISSN 1000-9000(Print)  
/1860-4749(Online)  
CODEN JCTEEM

# Journal of Computer Science & Technology



SPONSORED BY INSTITUTE OF COMPUTING TECHNOLOGY  
THE CHINESE ACADEMY OF SCIENCES &



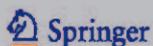
CHINA COMPUTER FEDERATION



SUPPORTED BY NSFC



CO-PUBLISHED BY SCIENCE PRESS &



SPRINGER

COMPUTER

# Checking Causal Consistency of MongoDB

Hong-Rong Ouyang<sup>1</sup> (欧阳鸿荣), Heng-Feng Wei<sup>1,2,\*</sup> (魏恒峰), *Member, CCF*  
Hai-Xiang Li<sup>3,\*</sup> (李海翔), *Member, CCF*, An-Qun Pan<sup>3</sup> (潘安群), *Member, CCF*, and  
Yu Huang<sup>1</sup> (黄宇), *Member, CCF*

<sup>1</sup>State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China

<sup>2</sup>Software Institute, Nanjing University, Nanjing 210093, China

<sup>3</sup>Tencent Distributed SQL Team of Technology and Engineering Group of Tencent, Tencent Inc., Shenzhen 518054, China

E-mail: mf20330056@smail.nju.edu.cn; hfwei@nju.edu.cn; {blueseali, aaronpan}@tencent.com  
yuhuang@nju.edu.cn

Received June 1, 2021; accepted December 20, 2021.

**Abstract** MongoDB is one of the first commercial distributed databases that support causal consistency. Its implementation of causal consistency combines several research ideas for achieving scalability, fault tolerance, and security. Given its inherent complexity, a natural question arises: “Has MongoDB correctly implemented causal consistency as it claimed?” To address this concern, the Jepsen team has conducted black-box testing of MongoDB. However, this Jepsen testing has several drawbacks in terms of specification, test case generation, implementation of causal consistency checking algorithms, and testing scenarios, which undermine the credibility of its reports. In this work, we propose a more thorough design of Jepsen testing of causal consistency of MongoDB. Specifically, we fully implement the causal consistency checking algorithms proposed by Bouajjani *et al.* and test MongoDB against three well-known variants of causal consistency, namely CC, CCv, and CM, under various scenarios including node failures, data movement, and network partitions. In addition, we develop formal specifications of causal consistency and their checking algorithms in TLA<sup>+</sup>, and verify them using the TLC model checker. We also explain how TLA<sup>+</sup> specification can be related to Jepsen testing.

**Keywords** MongoDB, casual consistency, Jepsen, consistency checking, TLA<sup>+</sup>

## 1 Introduction

MongoDB is a general-purpose, document-oriented distributed NoSQL database<sup>①</sup>. A MongoDB database consists of a set of collections, a collection is a set of documents, and a document is an ordered set of keys with associated values<sup>[1]</sup>.

MongoDB achieves scalability by partitioning the data into shards and fault-tolerance by replicating each shard across a set of nodes<sup>[2]</sup>. The most general MongoDB deployment is a sharded cluster, where each shard is a replica set consisting of a primary node and

several secondary nodes (see Fig.1). Client operations are routed to corresponding shards via routers, which have access to config servers that are deployed as a replica set to store metadata for deployment. In a replica set, only the primary can accept writes from clients (via drivers), and it will record the writes in its oplog. Secondaries can accept reads, and they will replicate the primary’s oplog by periodically pulling it from the primary.

According to the PACELC theorem<sup>[3]</sup>, an extension to the CAP theorem<sup>[4,5]</sup>, if there is a network partition (P), a distributed system must trade off availability (A)

---

Regular Paper

Special Section on Software Systems 2021—Theme: Internetwork and Beyond

A preliminary version of the paper was published in the Proceedings of Internetwork 2020.

This work was supported by the CCF-Tencent Open Fund under Grant No. RAGR20200124 and the National Natural Science Foundation of China under Grant Nos. 61702253 and 61772258.

\*Corresponding Author (Heng-Feng Wei and Hai-Xiang Li have contributed significantly to the theoretical and experimental parts of the work, respectively.)

①MongoDB. <https://www.mongodb.com/>, Oct. 2021.

©Institute of Computing Technology, Chinese Academy of Sciences 2022

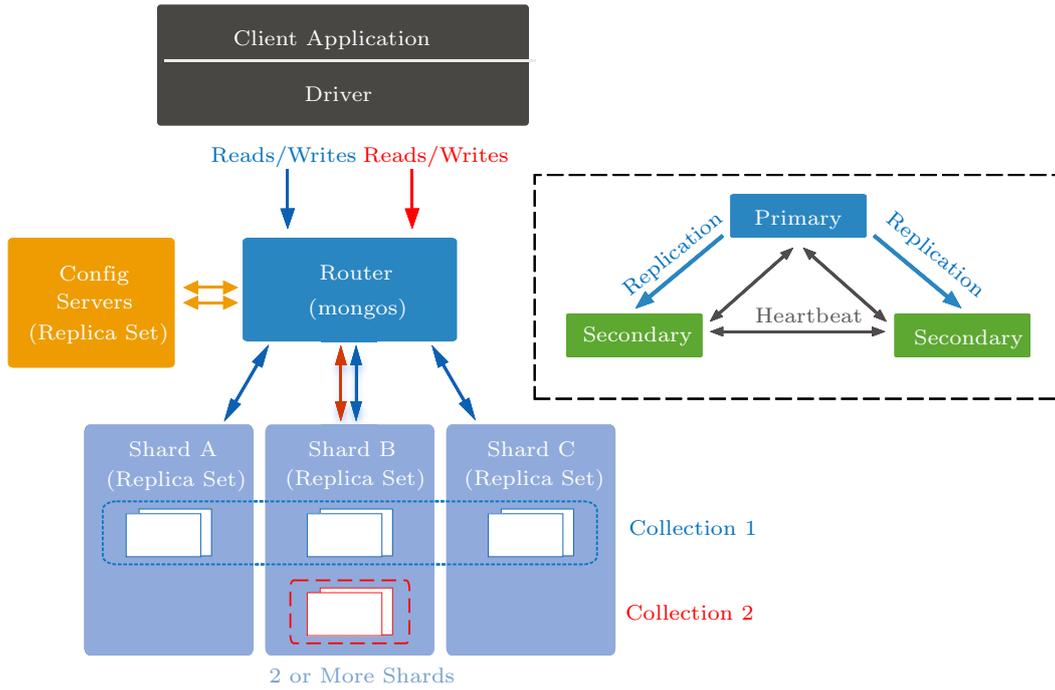


Fig.1. MongoDB deployment as a sharded cluster.

and consistency (C); else (E), it must trade off latency (L) and consistency (C). For high availability and low latency, MongoDB offers relaxed consistency models. Particularly, in version 3.6 released in November 2017, MongoDB introduced causal consistency<sup>②</sup>. It provides clients with session guarantees including read-your-writes, monotonic-reads, monotonic-writes, and writes-follow-reads<sup>③</sup>. As the Jepsen team<sup>④</sup> denoted, MongoDB is one of the first commercial databases that implement causal consistency<sup>[2]</sup>.

Being a production database, MongoDB’s implementation of causal consistency requires multi-dimensional evaluation criteria on performance, scalability, and security<sup>[2]</sup>. It combines several research ideas, including hybrid logical clocks<sup>[7]</sup>, explicit dependency tracking<sup>[8,9]</sup>, Raft-based replication consensus protocol<sup>[10]</sup>, and signature-verification mechanism. Given its inherent complexity, a natural question arises: “Has MongoDB correctly implemented causal consistency as it claimed in docs?” To address this concern, the Jepsen team has conducted black-box testing against MongoDB 3.6.4 and 4.0.0-rc1. The team designed test cases that characterize client operations, ran test cases in various scenarios, collected histories

of executions generated by MongoDB, and utilized an adapted version of the causal consistency checking algorithm proposed by Bouajjani *et al.*<sup>[11]</sup> to check whether these histories satisfy causal consistency.

However, the official Jepsen testing has several drawbacks in terms of specification, test case generation, implementation of causal consistency checking algorithms, and testing scenarios, which undermine the credibility of its reports. Specifically, the drawbacks are as follows.

- There are several variants of causal consistency, including causal consistency (CC)<sup>[12,13]</sup>, causal memory (CM)<sup>[14]</sup>, and causal convergence (CCv)<sup>[13]</sup>. Not all of them are comparable<sup>[13]</sup>. However, the official Jepsen testing does not clearly specify which causal consistency variant it tests against the MongoDB database.
- In terms of test cases, the official Jepsen testing uses independent keys. That is, each session accesses only a single key and different sessions access different keys. Concretely, each session performs a sequence of five operations on its key: an initial read, a write of 1, a read, a write of 2, and a final read. However, causal consistency is not compositional<sup>[15]</sup>, i.e., the composition of a set of keys satisfying causal consistency may

② MongoDB 3.6.0-rc0. <https://www.mongodb.com/blog/post/mongodb-360-rc0-is-released>, Oct. 2021.

③ Causal Consistency. <https://docs.mongodb.com/manual/core/causal-consistency-read-write-concerns/>, Oct. 2021.

④ Jepsen. <https://jepsen.io/>, Oct. 2021.

not be causally consistent. Thus, the test cases are too restrictive for causal consistency checking.

- Given the specific test cases above, the official Jepsen testing presets the expected return value of each read operation in its causal consistency checking algorithm. In other words, it has not fully implemented the causal consistency checking algorithms in [11].

- Although the official Jepsen testing has tested the causal consistency of MongoDB under network partitions, it does not cover the scenarios such as node failures and data movement among shards.

In this work, we propose a more thorough design of Jepsen testing of the causal consistency protocol of MongoDB<sup>⑤</sup>. Specifically, our contributions are as follows.

- We consider three well-known variants of causal consistency, following the formal specification given in [11].

- We generate the most general operation sequences for clients, without any restrictions on keys.

- We fully implement the “bad patterns” based causal consistency checking algorithm proposed by Bouajjani *et al.* in [11].

- We design more testing scenarios, covering network partitions, node failures, and data movement among shards.

Our preliminary experimental results confirm the claim in MongoDB’s documentation that in the presence of node failures or network partitions, causal consistency is guaranteed only for reads with `majority readConcern` (explained shortly in Subsection 2.2) and writes with `majority writeConcern`.

This is an extended version of our conference

paper<sup>[16]</sup> of the same title. In this version, we develop the formal specifications of three causal consistency variants, namely CC, CCv, and CM, and the “bad patterns” based checking algorithms in TLA<sup>+</sup>. We also verify them using the TLC model checker. The model checking results confirm, though on test cases of relatively small scales, the correctness of the checking algorithms. We also explain how TLA<sup>+</sup> specification can be further related to Jepsen testing in Subsection 5.5.

The rest of the paper is organized as follows. Section 2 provides preliminaries on causal consistency, the Jepsen testing framework, and TLA<sup>+</sup>. Section 3 describes the official Jepsen testing of causal consistency of MongoDB and introduces our more thorough design. Section 4 demonstrates our experiments and results. Section 5 shows the formal specifications of causal consistency and checking algorithms in TLA<sup>+</sup> and the model checking results. Section 6 discusses related work and Section 7 concludes the paper.

## 2 Preliminaries

### 2.1 Causal Consistency: Informal Introduction

Causal consistency guarantees that all clients agree on the relative ordering of causally related operations<sup>[14,17]</sup>. However, operations that are not causally related may be observed in different orders by different clients. We informally explain causal consistency in the classic “Lost-Ring” example<sup>[18]</sup> (see Fig.2). Alice first posts that she has lost her ring. After a while, she posts that she has found it. Bob sees Alice’s posts, and comments “Glad to hear it”. We say that there

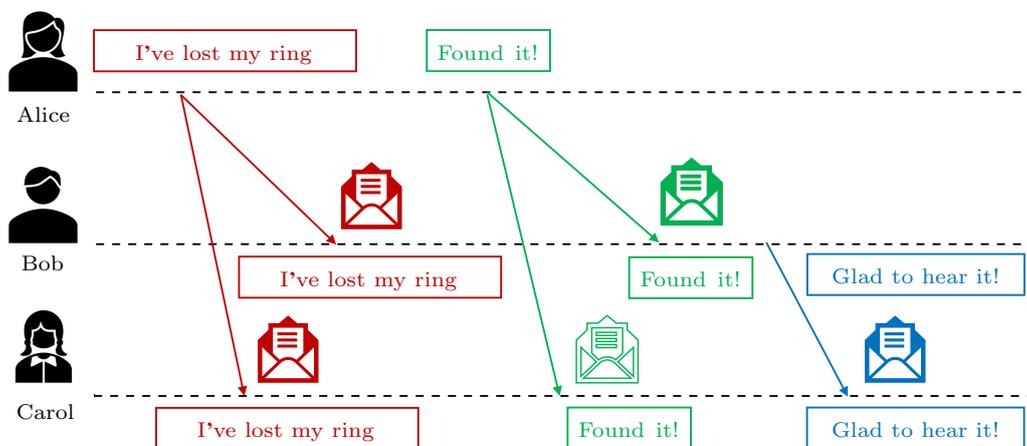


Fig.2. “Lost-Ring” example<sup>[18]</sup> for causal consistency.

<sup>⑤</sup>The project can be found at <https://github.com/Tsunaou/Checking-Causal-Consistency-of-MongoDB>, Oct. 2021.

is a read-from dependency from Alice’s second post to Bob’s get operation, and a session dependency from Bob’s get operation to his own comment. By transitivity, Bob’s comment causally depends on Alice’s second post. Thus, when Carol sees Bob’s comment, she should also see Alice’s second post. Otherwise, she would be quite confused, mistakenly thinking that Bob is glad to hear that Alice has lost her ring.

## 2.2 Causal Consistency in MongoDB

MongoDB enables causal consistency in client sessions. Moreover, MongoDB’s causal consistency can be combined with tunable consistency, which allows clients to select the trade-off between consistency and latency, at a per operation level<sup>[1]</sup>. `writeConcern` specifies the number of replica set members that must acknowledge the write before returning to a client. In particular, `w : majority` requires a write operation to be acknowledged by a majority of the replica set members before being returned to the client. `readConcern` determines what consistency guarantees data returned to a client must satisfy. The default value of `readConcern` is `level : local`, which allows to return the local data in a single replica set member. In contrast, `level : majority` guarantees that the returned data has been written to a majority of the replica set members. As claimed in MongoDB’s documentation, in the presence of node failure or network partitions, causally consistent sessions can only guarantee causal consistency for reads with `majority readConcern` and writes with `majority writeConcern`. In a good condition, however, write operations with `w1 writeConcern` can also provide causal consistency.

## 2.3 Causal Consistency: Formal Specification

We review the formal specification of causal consistency with respect to read-write registers, following [11].

### 2.3.1 Replicated Objects

We focus on read/write registers from  $\mathbb{X}$ , ranged over by  $x, y$ , etc. They support a set of methods  $\mathbb{M} = \{\text{wr}, \text{rd}\}$  for writing to or reading from a register (i.e., key), with input or output values from  $\mathbb{V}$ .

### 2.3.2 Histories

We model the interactions between clients and a distributed database maintaining replicated read/write

registers by histories.

**Definition 1** (History). *A history  $h = (O, \text{PO}, \ell)$  is the poset (partial-ordered set)  $(O, \text{PO})$  labeled by  $\mathbb{M} \times \mathbb{V} \times \mathbb{V}$ , where*

- $O$  is a set of operation identifiers, or simply operations; we use  $R$  and  $W$  to denote the set of read and write operations, respectively;
- $\text{PO}$  is a union of total orders among operations called program order; for  $o_1, o_2 \in O$ ,  $o_1 <_{\text{PO}} o_2$  means that  $o_1$  and  $o_2$  were issued by the same client and  $o_1$  occurred before  $o_2$ ;
- for an operation  $o \in O$ , its label  $\ell(o) = (m, \text{arg}, \text{rv}) \in \mathbb{M} \times \mathbb{V} \times \mathbb{V}$  indicates that  $o$  is an invocation of method  $m$  with input argument  $\text{arg}$ , returning value  $\text{rv}$ . We sometimes denote  $\ell(o)$  by  $m(\text{arg}) \triangleright \text{rv}$ .

We use  $\text{wr}(x, v) \triangleright \perp$  (or simply  $\text{wr}(x, v)$ ) to denote a write of value  $v \in \mathbb{V}$  to register  $x \in \mathbb{X}$  returning  $\perp \notin \mathbb{V}$ , and  $\text{rd}(x) \triangleright v$  to denote a read of  $x$  returning  $v$ . In addition, for an operation  $o$  with  $\ell(o) = \text{wr}(x, v)$  or  $\ell(o) = \text{rd}(x) \triangleright v$ , we define  $\text{var}(o) = x$  and  $\text{val}(o) = v$ .

Let  $\rho = (O, <, \ell)$  be an  $\mathbb{M} \times \mathbb{V} \times \mathbb{V}$  labeled poset and  $O' \subseteq O$  be a set.  $\rho\{O'\}$  is the labeled poset in which only the return values of the operations in  $O'$  are kept. Formally,  $\rho\{O'\}$  is the  $(\mathbb{M} \times \mathbb{V}) \cup (\mathbb{M} \times \mathbb{V} \times \mathbb{V})$  labeled poset  $(O, <, \ell')$  where for all  $o \in O'$ ,  $\ell'(o) = \ell(o)$ , and for all  $o' \in O \setminus O'$ ,  $\ell'(o') = (m, \text{arg})$  if  $\ell(o) = (m, \text{arg}, \text{rv})$ . We denote  $\rho\{O'\}$  by  $\rho\{o\}$  if  $O' = \{o\}$ .

Let  $\rho = (O, <, \ell)$  and  $\rho' = (O, <', \ell')$  be two  $(\mathbb{M} \times \mathbb{V}) \cup (\mathbb{M} \times \mathbb{V} \times \mathbb{V})$  labeled posets.  $\rho' \preceq \rho$  means that  $\rho'$  has less order and label constraints on the set  $O$ . Formally,  $\rho' \preceq \rho$  if  $<' \subseteq <$  and for all  $o \in O$ ,  $\ell'(o) = \ell(o)$  or  $\ell'(o) = (m, \text{arg})$  if  $\ell(o) = (m, \text{arg}, \text{rv})$ .

### 2.3.3 Sequential Semantics

The consistency of replicated read-write registers is defined with respect to the sequential semantics of read-write registers. Intuitively, in any operation sequence on read-write registers, an rd operation returns the value of the latest preceding wr on the same register, or the initial value 0 if there are no such prior writes. Formally, the sequential semantics  $S_{\text{RW}}$  of read-write registers is the smallest set of sequences labeled by  $\mathbb{M} \times \mathbb{V} \times \mathbb{V}$  satisfying

- $\epsilon \in S_{\text{RW}}$ , where  $\epsilon$  is the empty sequence;
- if  $\rho \in S_{\text{RW}}$ , then  $\rho \cdot \text{wr}(x, v) \in S_{\text{RW}}^{\textcircled{c}}$ ;
- if  $\rho \in S_{\text{RW}}$  contains no writes on  $x$ , then  $\rho \cdot \text{rd}(x) \triangleright 0 \in S_{\text{RW}}$ ;
- if  $\rho \in S_{\text{RW}}$  and the last write in  $\rho$  on register  $x$  is  $\text{wr}(x, v)$ , then  $\rho \cdot \text{rd}(x) \triangleright v \in S_{\text{RW}}$ .

<sup>①</sup>The symbol  $\cdot$  means the connection between operations.

### 2.3.4 Causal Consistency

Following [11], we consider three well-known variants of causal consistency, namely CC (Causal Consistency), CCv (Causal Consistency Convergence), and CM (Causal Memory).

A history is CC if there exists a causal order that explains the return value of each operation.

**Definition 2** (Causal Consistency). *A history  $h = (O, PO, \ell)$  is CC with respect to specification  $S_{RW}$  if there exists a strict partial order  $co \subseteq O \times O$  called the causal order such that for each operation  $o \in O$ , there exists a sequence  $\rho_o \in S_{RW}$  satisfying*

$$\begin{aligned} \text{AxCausal} &\triangleq PO \subseteq co, \\ \text{AxCausalValue} &\triangleq (co^{-1}(o), co, \ell)\{o\} \preceq \rho_o. \end{aligned}$$

Here  $co^{-1}(o)$  is the set of operations that precede  $o$  in causal order. Formally,  $co^{-1}(o) \triangleq \{o' \mid o' \leq_{co} o\}$ .

CCv ensures eventual convergence via a total arbitration order.

**Definition 3** (Causal Convergence). *A history  $h = (O, PO, \ell)$  is CCv with respect to specification  $S_{RW}$  if there exists a strict partial order  $co \subseteq O \times O$  called the causal order and a strict total order  $arb \subseteq O \times O$  called the arbitration order such that for each operation  $o \in O$ , there exists a sequence  $\rho_o \in S_{RW}$  satisfying*

$$\begin{aligned} \text{AxCausal} &\triangleq PO \subseteq co, \\ \text{AxArb} &\triangleq co \subseteq arb, \\ \text{AxCausalArb} &\triangleq (co^{-1}(o), arb, \ell)\{o\} \preceq \rho_o. \end{aligned}$$

CM requires each client to be consistent with respect to the returned values it has observed before.

**Definition 4** (Causal Memory). *A history  $h = (O, PO, \ell)$  is CM with respect to specification  $S_{RW}$  if there exists a strict partial order  $co \subseteq O \times O$  called the causal order such that for each operation  $o \in O$ , there exists a sequence  $\rho_o \in S_{RW}$  satisfying*

$$\begin{aligned} \text{AxCausal} &\triangleq PO \subseteq co, \\ \text{AxCausalSeq} &\triangleq (co^{-1}(o), co, \ell)\{PO^{-1}(o)\} \preceq \rho_o. \end{aligned}$$

Here  $PO^{-1}(o) \triangleq \{o' \mid o' \leq_{PO} o\}$ .

## 2.4 Causal Consistency Checking

The general decision problem of checking whether a history over read-write registers is causally consistent is NP-complete<sup>[11]</sup>. However, for differentiated histories in which the values written to the same register are

distinct, it is polynomial time<sup>[11]</sup>. Differentiated histories can be achieved by attaching unique timestamps to writes in implementation. We consider only differentiated histories below.

The polynomial-time checking algorithms proposed by Bouajjani *et al.* are based on the notion of “bad patterns”<sup>[11]</sup>. Each causal consistency variant can be precisely characterized by lacking a set of certain bad patterns. The bad patterns are expressed in terms of program order PO, read-from relation RF, causal order CO, conflict relation CF, and happened-before relation HB on operations.

**Definition 5** (Read-From Relation). *The read-from relation  $RF \subseteq W \times R$  associates a read with the write from which it obtains the value. Formally,*

$$\begin{aligned} \forall w \in W, r \in R. (w, r) \in RF &\iff \\ \text{var}(w) = \text{var}(r) \wedge \text{val}(w) = \text{val}(r). \end{aligned}$$

**Definition 6** (Causal Order). *The causal order  $CO \subseteq O \times O$  is defined as the transitive closure of program order and read-from relation. Formally,*

$$CO = (PO \cup RF)^+.$$

**Definition 7** (Conflict Relation). *The conflict relation  $CF \subseteq W \times W$  orders two writes on the same register according to a third read operation. Formally,*

$$\begin{aligned} \forall w, w' \in W. (w, w') \in CF &\iff \\ \exists r' \in R. (w', r') \in RF \wedge \text{var}(w) = \text{var}(r') \wedge (w, r') \in CO. \end{aligned}$$

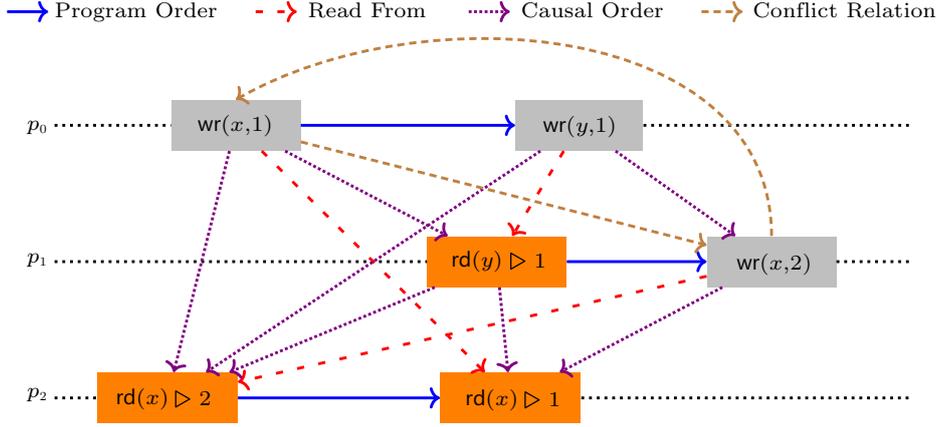
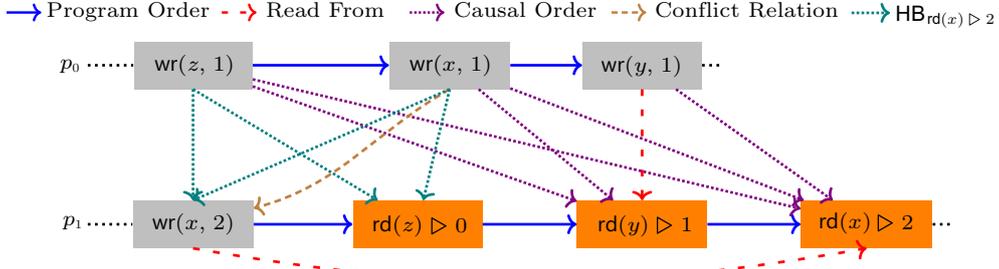
*Example 1.* Let us consider the history  $h$  of Fig.3; since  $wr(x, 2) <_{RF} rd(x) \triangleright 2$  and  $rd(x) \triangleright 2 <_{PO} rd(x) \triangleright 1$ , we have  $wr(x, 2) <_{CO} rd(x) \triangleright 1$ . In addition,  $wr(x, 1) <_{RF} rd(x) \triangleright 1$ . Thus,  $wr(x, 2) <_{CF} wr(x, 1)$ .

**Definition 8** (Happened-Before Relation). *For each operation  $o \in O$ , the happened-before relation  $HB_o \subseteq O \times O$  of  $o$  is the smallest transitive relation satisfying that  $CO|_{CO^{-1}(o)} \subseteq HB_o$  and*

$$\begin{aligned} \forall w, w' \in W. (w, w') \in HB_o &\iff \exists r' \in R. \\ (r' \leq_{PO} o \wedge w' <_{RF} r' \wedge \text{var}(w) = \text{var}(r') \wedge (w, r') \in HB_o). \end{aligned}$$

Here  $CO|_{CO^{-1}(o)}$  is the relation CO restricted on the set  $co^{-1}(o)$ .

*Example 2.* Let us consider the history  $h$  of Fig.4. Since  $wr(x, 1) <_{PO} wr(y, 1) <_{RF} rd(y) \triangleright 1$  and  $rd(y) \triangleright 1 <_{PO} rd(x) \triangleright 2$ , we have  $wr(x, 1) <_{CO} rd(x) \triangleright 2$ . In addition, for operation  $rd(x) \triangleright 2$ ,  $CO^{-1}(rd(x) \triangleright 2) = CO$ , therefore we have  $CO \subseteq HB_{rd(x) \triangleright 2}$ . And since  $wr(x, 2) <_{RF} rd(x) \triangleright 2$  and  $wr(x, 1) <_{HB_{rd(x) \triangleright 2}} rd(x) \triangleright 2$ , we have  $wr(x, 1) <_{HB_{rd(x) \triangleright 2}} wr(x, 2)$ . For the transitivity, we can also get  $wr(z, 1) <_{HB_{rd(x) \triangleright 2}} rd(z) \triangleright 0$ .


 Fig.3. A history  $h$  that is not CCv. (The arrows for CO that are implied by transitivity are not shown.)

 Fig.4. A history  $h$  that is not CM. (The arrows for CO that are implied by transitivity are not shown.)

The following theorem characterizes CC, CCv, and CM in terms of bad patterns defined in Table 1.

**Theorem 1** [11]. *A history  $h$  is CC if and only if  $h$  does not exhibit any bad patterns of CyclicCO, WriteCOInitRead, ThinAirRead or WriteCOWrite. A history  $h$  is CCv if and only if it is CC and does not exhibit any bad patterns of CyclicHF. A history  $h$  is CM if and only if it is CC and does not exhibit any bad patterns of WriteHBInitRead or CyclicHB.*

*Example 3.* Let us consider the history  $h$  of Fig.3. It is not CCv. First, since  $wr(x,1) <_{CO} wr(x,2) <_{CO} rd(x) > 1$  and  $wr(x,1) <_{RF} rd(x) > 1$ , it exhibits the bad

pattern WriteCOWrite. In addition, there is a cycle in CF:  $wr(x,1) <_{CF} wr(x,2) <_{CF} wr(x,1)$ . Thus, it also exhibits the bad pattern CyclicCF.

*Example 4.* Let us consider the history  $h$  of Fig.4. It is not CM. Since we have  $wr(z,1) <_{HB_{rd(x) > 2}} rd(z) > 0$  and  $rd(z) > 0 <_{PO} rd(x) > 2$ , it exhibits the bad pattern WriteHBInitRead.

## 2.5 Jepsen

Jepsen<sup>Ⓢ</sup> is a library for black-box testing of distributed systems. A typical Jepsen testing of a distributed database consists of a deployment of the

Table 1. Definitions of Bad Patterns [11]

Bad Pattern	Description
CyclicCO	$PO \cup RF$ is cyclic
ThinAirRead	$\exists r \in R. \text{val}(r) \neq 0 \wedge (\nexists w \in W. w <_{RF} r)$
WriteCOInitRead	$\exists r \in R, w \in W. w <_{CO} r \wedge \text{var}(w) = \text{var}(r) \wedge \text{val}(r) = 0$
WriteCOWrite	$\exists w_1, w_2 \in W, r_1 \in R. \text{var}(w_1) = \text{var}(w_2) \wedge w_1 <_{CO} w_2 <_{CO} r_1 \wedge w_1 <_{RF} r_1$
CyclicCF	$CF \cup CO$ is cyclic
WriteHBInitRead	$\exists o \in O, r \in R, w \in W. r \leq_{PO} o \wedge w <_{HB_o} r \wedge \text{var}(w) = \text{var}(r) \wedge \text{val}(r) = 0$
CyclicHB	$\exists o \in O. HB_o$ is cyclic

<sup>Ⓢ</sup>Jepsen Library. <https://github.com/jepsen-io/jepsen>, Oct. 2021.

database and a control node. The control node starts several worker processes called clients. A generator is responsible for continuously generating operations and dispatching them to clients, according to user-defined rules. Clients interact with the database by issuing operations. The invocations and responses produced are recorded in a history. When the test finishes, the history is checked by a checker against a desired consistency model.

To test the fault-tolerant capability of the database, special worker processes called nemeses continuously inject faults or rare events (such as data movement among shards) into the database deployment.

## 2.6 TLA<sup>+</sup>

TLA<sup>+</sup> is a high-level formal specification language developed by Lamport [19]. It was designed for modeling and reasoning about programs and systems, especially concurrent and distributed ones.

TLA<sup>+</sup> is based on TLA, the Temporal Logic of Actions [20]. With TLA, a system can be modeled as a state machine which is described by its initial states and actions. Since we focus on the specification of causal consistency, we omit the temporal operators in TLA<sup>+</sup> here.

TLA<sup>+</sup> combines TLA with the first-order logic and the Zermelo-Fraenkel set theory. Table 2 summarizes the (non-temporal) operators that we use in [21]. Interested readers are referred to the complete version of Summary of TLA<sup>+</sup> [8].

A specification in TLA<sup>+</sup> consists of modules. In

a module, we can declare constants (CONSTANTS) and variables (VARIABLES), and define operators like  $Op(p_1, \dots, p_n) \triangleq exp$ . We can also import the declarations, definitions, and operators from other modules  $M_1, \dots, M_n$ , by writing EXTENDS  $M_1, \dots, M_n$  in  $M$ .

TLC is an explicit-state model checker for TLA<sup>+</sup> [22]. It verifies the TLA<sup>+</sup> specifications by exploring the whole state space of finite-state instances of them. In this paper, we use TLC only to evaluate constant expressions.

## 3 Jepsen Testing of Causal Consistency of MongoDB

In this section we first describe the official Jepsen testing of causal consistency of MongoDB 3.6.4 and 4.0.0-rc1, from the perspectives of specification, test case generation, implementation of causal consistency checking algorithms, and testing scenarios. To overcome its drawbacks identified in Section 1, we then design a more thorough Jepsen testing of causal consistency of MongoDB.

### 3.1 Official Jepsen Testing

The MongoDB deployment under test consists of two shards, each of which is a replica set of five nodes.

#### 3.1.1 Specification

The Jepsen team claimed that they had tested MongoDB against causal consistency [9]. However, they did not clearly specify the variant of causal consistency.

**Table 2.** Summary of TLA<sup>+</sup> Operators Used in This Paper

Category	Operator	Meaning
Set	SUBSET $S$	Powerset of $S$
	UNION $S$	Union of all elements of $S$
	$\{e : x \in S\}$ $\{x \in S : p\}$	Set of elements $e$ such that $x$ is in $S$ Set of elements $x$ in $S$ satisfying $p$
Function	DOMAIN $f$	Domain of function $f$
	$f[e]$	Function application
	$[x \in S \mapsto e]$	Function $f$ such that $f[x] = e$ for $x \in S$
Record	$e.h$	$h$ -field of record $e$
	$[h_1 \mapsto e_1, \dots, h_n \mapsto e_n]$	Record whose $h_i$ field is $e_i$
	$[h_1 : S_1, \dots, h_n : S_n]$	Set of all records with $h_i$ field in $S_i$
Tuple	$e[i]$	The $i$ -th component of tuple $e$
	$\langle e_1, \dots, e_n \rangle$	The $n$ -tuple whose $i$ -th component is $e_i$
Sequence	$SubSeq(s, m, n)$	Sequence $\langle s[m], s[m+1], \dots, s[n] \rangle$
	$Range(s)$	Set of elements of sequence $s$

[8] Leslie Lamport. Summary of TLA<sup>+</sup>. <http://lamport.azurewebsites.net/tla/summary-standalone.pdf>, May 2021.

[9] Jepsen Testing of MongoDB 3.6.4. <https://jepsen.io/analyses/mongodb-3-6-4>, Oct. 2021.

### 3.1.2 Test Case Generation

Treating a MongoDB collection as a set of read-write registers, the generator generates read and write operations for clients. The dispatch rule ensures that each client accesses only a single register and different clients access different registers. Specifically, the operation sequence of each client consists of five operations as follows:

$$(r, w1, r, w2, r),$$

where  $r$  denotes a read of the register that belongs to the client,  $w1$  a write of value 1 to the register, and  $w2$  a write of value 2 to the register.

### 3.1.3 Checking Algorithms

Since the test cases are quite restrictive, it is sufficient for the checker to verify whether the three reads of each client return 0, 1, and 2 in order.

### 3.1.4 Testing Scenarios

The official Jepsen testing has designed a kind of nemesis called partition-random-halves to trigger network partitions randomly. Specifically, in the five-node deployment of MongoDB, the network will be split into two disconnected parts: one (denoted as  $P_1$ ) consists of two nodes, one of which is the original primary node, and the other (denoted as  $P_2$ ) consists of three nodes. Since three nodes in  $P_2$  constitute a majority (of five nodes), one of them will be elected as a new primary. Consequently, there would temporarily be two nodes that consider themselves as the primary of the cluster.

After the network recovers, the writes performed on the original primary node during network partition will be rolled back. The Jepsen testing revealed that in the presence of network partitions, causally consistent sessions can only guarantee causal consistency for reads with `majority readConcern` and writes with `majority writeConcern`.

## 3.2 Our Design of Jepsen Testing

As shown in Table 3, we improve the official Jepsen testing in the following aspects.

### 3.2.1 Specification

We test MongoDB against three well-known variants of casual consistency, namely, CC, CM, and CCv. Specifically, we adopt the formal specification given in [11].

### 3.2.2 Test Case Generation

In our design, the generator generates an arbitrary differentiated operation sequence for each client using YCSB [23]<sup>Ⓓ</sup>. Particularly, we impose no restrictions on keys as the official Jepsen testing does, only controlling the range and distribution of generated keys, and the ratio of read and write operations.

The generated keys follow a uniform distribution. To ensure that all writes on the same register write unique values, the generator attaches values 1, 2, ... to them in order. We record necessary information about each operation during generation and execution, including its type (i.e., read or write), the value it reads or writes, the client that issues the operation, and the index indicating the order in which the operation is generated.

### 3.2.3 Checking Algorithms

To check an arbitrary differentiated history against several variants of causal consistency, we fully implement the “bad patterns” based causal consistency checking algorithms for CC, CM, and CCv [11].

### 3.2.4 Testing Scenarios

Besides partition-random-halves in the official Jepsen testing, we introduce two additional nemeses called node-failure and data-mover. The node-failure nemesis randomly selects a database node, suspends it for a while, and then recovers it. This may trigger

**Table 3.** Comparison Between the Official Jepsen Testing and Our Design

Perspective	Official Jepsen Testing	Our Design of Jepsen Testing
Specification	Unspecified	Three well-known variants: CC, CM, and CCv
Test case generation	Restricted on keys and operation sequences	General for differentiated histories
Checking algorithms	Ad hoc for restricted test cases	Full implementation of [11]
Testing scenarios	Network partition	Network partition, data movement, node failure

<sup>Ⓓ</sup>YCSB. <https://github.com/brianfrankcooper/YCSB>, Oct. 2021.

leader election. The data-mover nemesis periodically moves data among shards. In an execution, partition-random-halves, node-failure, and data-mover are generated and scheduled by the generator, according to user-defined rules.

## 4 Preliminary Evaluations

We implement the checking algorithms of [11] and check histories produced by MongoDB 4.2.3 against CC, CM, and CCv. We use the Jepsen testing framework of version 0.1.17<sup>①</sup>. Table 4 shows the hardware configurations of the control node, the database nodes, and the checker server.

### 4.1 Experimental Setup

We adopt the same MongoDB deployment as that in the official Jepsen testing: it consists of two shards, each of which is a replica set of five nodes.

In each experiment, we fix 100 registers and 10 clients. The generator generates read or write operations and appends them into a queue. For each register, the ratio between the number of read operations and that of write operations is 3 : 1. Each client creates a causally consistent session, extracts operations from the operation queue, and issues them to MongoDB servers.

For each experiment, we tune the total number of operations and the `readConcern` and `writeConcern` levels for operations. To handle possible exceptions thrown by MongoDB during write operations, we restart a new causally consistent session in the corresponding client. Moreover, we cover both the scenarios with and without nemesis. For each history produced by MongoDB, we check whether it satisfies CC, CM, and CCv.

### 4.2 Experimental Results

Table 5 shows the experimental results of checking causal consistency of MongoDB.

#### 4.2.1 Causal Consistency Checking

The preliminary experimental results confirm the claim in MongoDB’s documentation that in the presence of nemesis (such as partition-random-halves, node-failure, and data-mover), causally consistent sessions guarantee causal consistency only for reads with `majority readConcern` and writes with `majority writeConcern`. In contrast, in the presence of nemesis, the histories with `local readConcern` and `w1 writeConcern` may violate any of three causal consistency variants. On the other hand, without nemesis, MongoDB can provide all three variants of causal consistency even with `local readConcern` and `w1 writeConcern`.

#### 4.2.2 Performance

Fig. 5 demonstrates the performance of checking whether histories satisfy causal consistency. According to [11], it takes  $O(n^3)$  to check a differentiated history with  $n$  operations against CC or CCv. In contrast, it takes  $O(n^5)$  against CM. The experimental results in Fig. 5 exhibit such a substantial performance gap.

### 4.3 Unexpected ThinAirRead Bad Patterns

We observe some unexpected ThinAirRead bad patterns in our preliminary evaluations, marked  $\otimes$  in Table 5. They appear in some histories that are produced without nemesis and consist of reads with `majority readConcern` and writes with `majority writeConcern`. Table 6 shows a snippet of such a history. Note that the write operation `wr(85, 5)` of No. 1128 incurs a runtime exception called `com.mongodb.MongoWriteException`. Since the causal consistency checking algorithms in [11] implicitly assume that all write operations are successful, this write operation is considered failed and discarded from the history. However, a later read operation `rd(85)` of No. 1266 obtains the value 5 from key 85, indicating that the write operation `wr(85, 5)` has actually written its value to the database. This gives rise to a ThinAirRead bad pattern during checking.

Table 4. Hardware Configurations

Component	Configuration
Control node	Intel® Core™ i5-9500 CPU @ 3.00 GHz, 16 GB, Ubuntu 20.04
Database node	Intel® Xeon® Platinum 8269CY CPU @ 2.50 GHz, 4 GB, Ubuntu 16.04
Checker server	Intel® Core™ i9-9900X CPU @ 3.50 GHz, 32 GB, Ubuntu 16.04

<sup>①</sup>Jepsen Library 0.1.17. <https://github.com/jepsen-io/jepsen/tree/0.1.17>, Oct. 2021.

**Table 5.** Experimental Results of Causal Consistency Checking of MongoDB

Number of Operations	With Nemesis						Without Nemesis					
	(majority, majority)			(w1, local)			(majority, majority)			(w1, local)		
	CC	CM	CCv	CC	CM	CCv	CC	CM	CCv	CC	CM	CCv
100	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
200	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
300	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
400	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
500	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
600	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
700	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓	✓	✓
800	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
900	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓	✓	✓
1 000	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
1 100	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓	✓	✓
1 200	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓	✓	✓
1 300	⊗	⊗	⊗	✓	✓	✓	✓	✓	✓	✓	✓	✓
1 400	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
1 500	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
1 600	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓	✓	✓
1 700	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓	✓	✓
1 800	⊗	⊗	⊗	✗	✗	✗	✓	✓	✓	✓	✓	✓
1 900	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓	✓	✓
2 000	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
2 500	⊗	⊗	⊗	✗	✗	✗	✓	✓	✓	✓	✓	✓
3 000	⊗	⊗	⊗	✗	✗	✗	✓	✓	✓	✓	✓	✓
3 500	⊗	⊗	⊗	✗	✗	✗	✓	✓	✓	✓	✓	✓
4 000	⊗	⊗	⊗	✗	✗	✗	✓	✓	✓	✓	✓	✓
4 500	⊗	⊗	⊗	✗	✗	✗	✓	✓	✓	✓	✓	✓
5 000	✓	✓	✓	✗	✗	✗	✓	✓	✓	✓	✓	✓

Note: ✓: satisfaction; ✗: violation; ⊗: unexpected ThinAirRead bad patterns discussed in [Subsection 4.3](#).

We remark that the unexpected ThinAirRead bad patterns above do not necessarily imply bugs in the causal consistency protocols of MongoDB. However, to better explain such unexpected results, it needs to design checking algorithms for histories which may contain failed write operations.

## 5 TLA<sup>+</sup> Specification of Causal Consistency and Checking Algorithms

In this section, we formally specify both the specification of causal consistency and the “bad patterns” based checking algorithms in [11] in TLA<sup>+</sup>, and verify them using the TLC model checker. We explain how TLA<sup>+</sup> specification can be further related to Jepsen testing in [Subsection 5.5](#). [Table 7](#) summarizes the auxiliary operators we define in this paper.

### 5.1 TLA<sup>+</sup> Specification of Causal Consistency

We follow the way how the specification of causal consistency is developed in [Subsection 2.3](#).

#### 5.1.1 Replicated Objects

In module *ReplicatedObjects* ([Fig.6\(a\)](#)), we assume each key is a single-character and take values from natural numbers for read/write registers. Following [11], we set the initial value of each key to 0. We assume that each operation is associated with a unique identifier.

#### 5.1.2 History

We define *Session* and *History* in module *History* ([Fig.6\(b\)](#)). A session  $s \in \text{Session}$  is a sequence of operations issued by the same client, and a history  $h \in \text{History}$  consists of a set of sessions. The program order  $PO(h)$  of a history  $h$  is a union of strict total

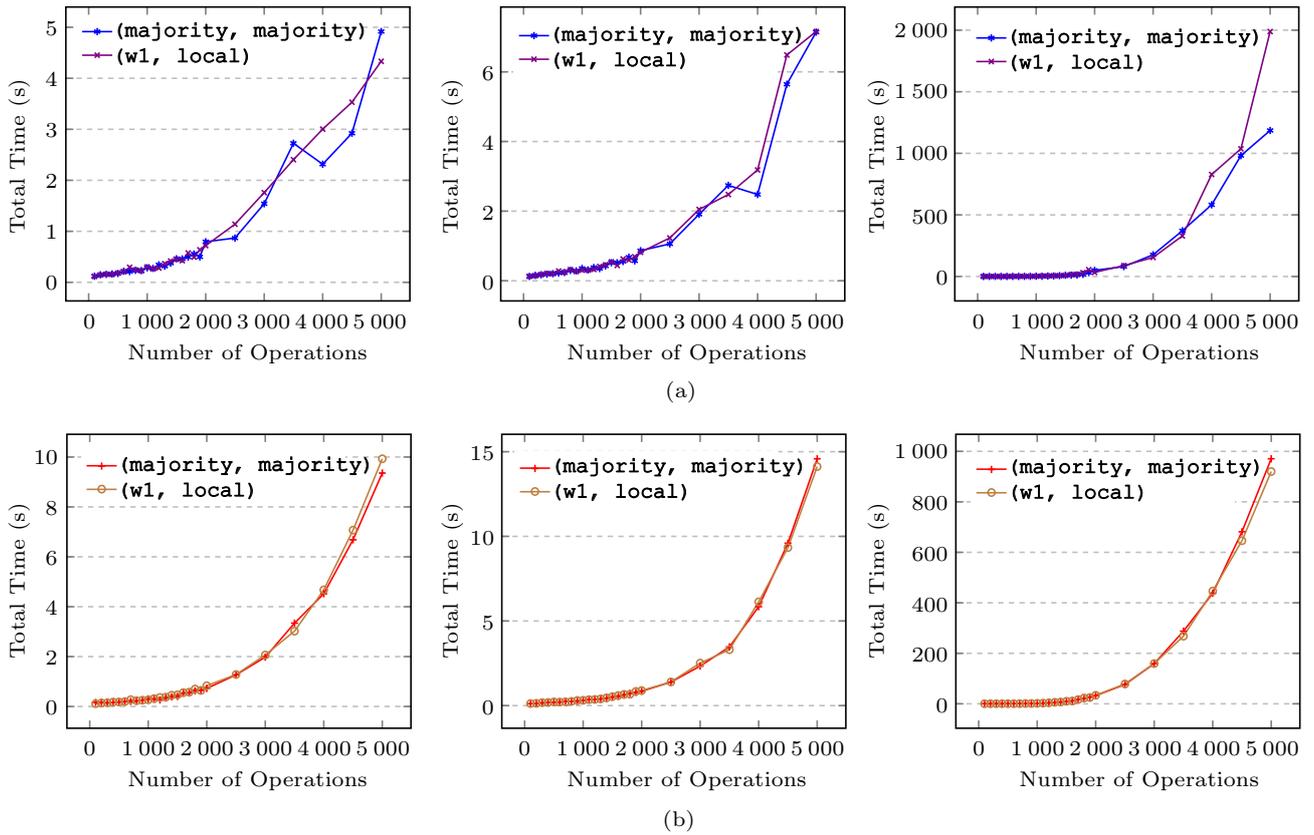


Fig.5. Time of checking whether histories satisfy causal consistency. (a) Time for causal consistency checking with nemesis. (b) Time for causal consistency checking without nemesis.

Table 6. Snippet of a History Exhibiting an Unexpected ThinAirRead Bad Pattern

No.	Operation	Exception	readConcern
1128	wr(85, 5)	MongoWriteException	majority
1129	wr(20, 5)	MongoWriteException	majority
1149	rd(20) $\triangleright$ 5	No Exception	majority
1266	rd(85) $\triangleright$ 5	No Exception	majority
1336	rd(20) $\triangleright$ 5	No Exception	majority
3756	rd(20) $\triangleright$ 5	No Exception	majority

orders among operations in the same session.

### 5.1.3 Sequential Semantics

The operator  $RWRegSemantics(seq, o)$  in module  $RWRegSemantics$  (Fig. 6(c)) checks whether the operation  $o$  is legal with respect to the sequential semantics when it is appended to the operation sequence  $seq$ .

### 5.1.4 Causal Consistency

The module  $Axioms$  (Fig. 7(a)) defines the axioms used in the specification of variants of causal consistency, which are shown in the module  $CausalDefinition$

(Fig. 7(b)).

The axiom  $AxCausalValue$  requires that for an operation  $o$ , there exists a linear extension  $seq$  of the causal order  $co$  when restricted on the set of operations preceding  $o$  such that  $RWRegSemantics(seq, o)$  is satisfied.

The axiom  $AxCausalArb$  requires that for an operation  $o$ , the arbitration order  $arb$  when restricted on the set of operations preceding  $o$  in causal order  $co$  is legal with respect to the sequential semantics.

The axiom  $AxCausalSeq$  requires that for an operation  $o$ , there exists a linear extension  $seq$  of the causal order  $co$  when restricted on the set of operations preced-

**Table 7.** Summary of Auxiliary Operators Defined in This Paper

Operator	Meaning
$PreSeq(s, e)$	The prefix of the sequence $s$ ending with element $e$ (which is unique in $s$ )
$Seq2Rel(s)$	Convert a sequence $s$ into a strict total order relation
$SelectSeq(s, Test(-))$	The subsequence of $s$ consisting of all elements $s[i]$ such that $Test(s[i])$ is true
$R S$	Restriction of relation $R$ on set $S$
$AllLinearExtensions(R, S)$	All possible linear extensions of the partial order $R$ defined on the set $S$
$AnyLinearExtension(R, S)$	An arbitrary linear extension of the partial order $R$ defined on the set $S$
$Respect(R, T)$	Does the relation $R$ respect relation $T$ ?
$TC(R)$	Transitive closure of the relation $R$
$POPast(h, o)$	The set of operations that precede $o \in Operation$ in program order in history $h \in History$ (including $o$ )
$StrictCausalPast(co, o)$	The set of operations that precede $o \in Operation$ in causal order $co$
$CausalPast(co, o)$	The set of operations that precede $o \in Operation$ in causal order $co$ (including $o$ )
$StrictCausalHist(co, o)$	The restriction of causal order $co$ to the operations in $StrictCausalPast(co, o)$
$CausalHist(co, o)$	The restriction of causal order $co$ to the operations in $CausalPast(co, o)$
$StrictCausalArb(co, arb, o)$	The restriction of arbitration $arb$ to the operations in $StrictCausalPast(co, o)$
$CausalArb(co, arb, o)$	The restriction of arbitration $arb$ to the operations in $CausalPast(co, o)$
$Ops(h)$	The set of all operations in history $h \in History$
$ReadOps(h)$	The set of all read operations in history $h \in History$
$ReadOpsOnKey(h)$	The set of all read operations on key $k \in Key$ in history $h \in History$
$WriteOps(h)$	The set of all write operations in history $h \in History$
$WriteOpsOnKey(h, k)$	The set of all write operations on key $k \in Key$ in history $h \in History$
$KeyOf(h)$	The set of keys read or written in $h \in History$

```

MODULE ReplicatedObjects
Key  $\triangleq$  Range("abcdefghijklmnopqrstuvwxy")
Val  $\triangleq$  Nat
InitVal  $\triangleq$  0
Oid  $\triangleq$  Nat

Operation  $\triangleq$  [type : {"read", "write"}, key : Key, val : Val, oid : Oid]
R(k, v, oid)  $\triangleq$  [type  $\mapsto$  "read", key  $\mapsto$  k, val  $\mapsto$  v, oid  $\mapsto$  oid]
W(k, v, oid)  $\triangleq$  [type  $\mapsto$  "write", key  $\mapsto$  k, val  $\mapsto$  v, oid  $\mapsto$  oid]

```

(a)

```

MODULE History
Session  $\triangleq$  Seq(Operation)
History  $\triangleq$  SUBSET Session
PO(h)  $\triangleq$  UNION {Seq2Rel(s) : s  $\in$  h}

```

(b)

```

MODULE RWRegSemantics
RWRegSemantics(seq, o)  $\triangleq$ 
  IF o.type = "write" THEN TRUE ELSE
  LET wseq  $\triangleq$  SelectSeq(seq, LAMBDA op : op.type = "write"  $\wedge$  op.key = o.key)
  IN IF wseq =  $\langle \rangle$  THEN o.val = InitVal
  ELSE o.val = wseq[Len(wseq)].val

```

(c)

Fig.6. TLA<sup>+</sup> modules for replicated read-write registers. (a) TLA<sup>+</sup> module *ReplicatedObjects*. (b) TLA<sup>+</sup> module *History*. (c) TLA<sup>+</sup> module *RWRegSemantics*.

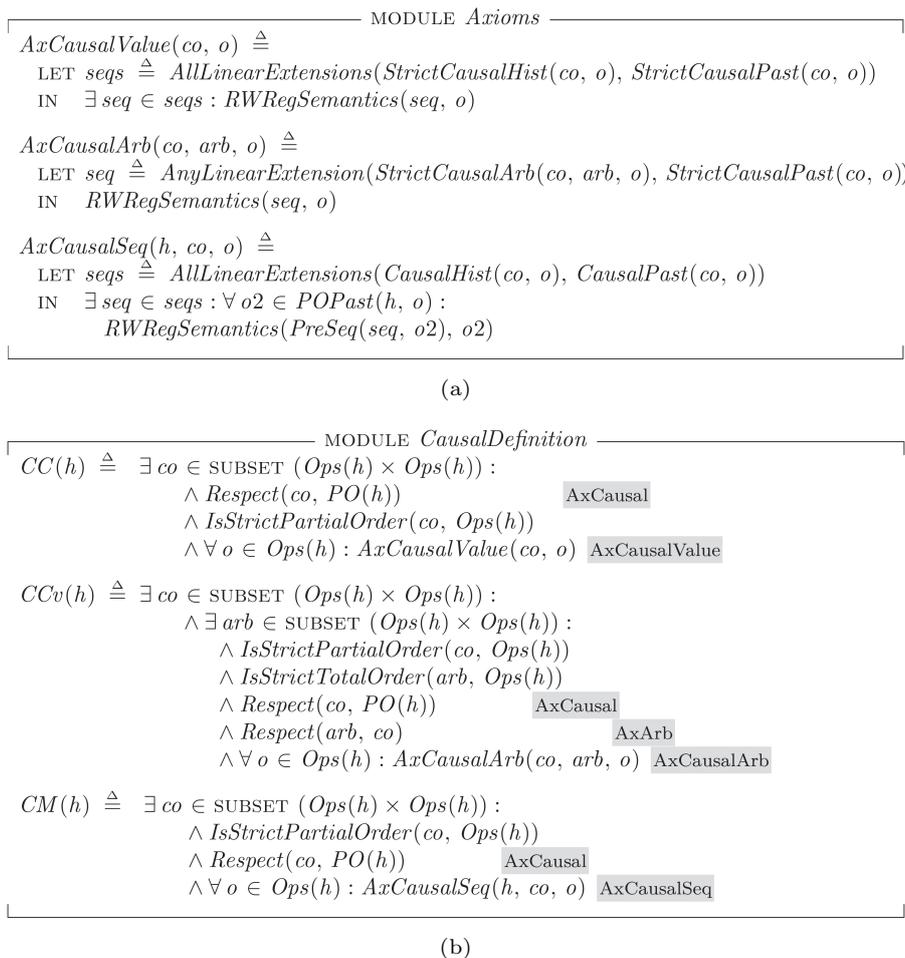


Fig.7. TLA<sup>+</sup> modules for the definition of variants of causal consistency. (a) TLA<sup>+</sup> module *Axioms*. (b) TLA<sup>+</sup> module *CausalDefinition*.

ing  $o$  in  $co$  such that for each operation  $o2$  preceding  $o$  in program order,  $\text{RWRegSemantics}(\text{PreSeq}(seq, o2), o2)$  is satisfied.

## 5.2 TLA<sup>+</sup> Specification of Causal Consistency Checking Algorithms

The module *Relations* (Fig.8) defines the relations including RF, CO, CF, and HB on the set of operations in histories. The module *BadPatterns* (Fig.9(a)) then defines all the bad patterns mentioned in Subsection 2.4. Finally, the module *Algorithm* (Fig.9(b)) specifies the “bad patterns” based checking algorithms for CC, CCv, and CM.

## 5.3 Optimizations

We observe that model checking histories against  $CC$ ,  $CCv$ , or  $CM$  as defined in Fig.7(b) are prohibitively inefficient. In this subsection, we propose

several optimizations, taking  $CCv$  as an example (see Fig.10).

### 5.3.1 CCv1: Rearranging Clauses

In  $CCv$ , we first enumerate all possible relations on  $ops$  as candidates for  $co$  and  $arb$ . In this way, for a history with  $n$  operations, the number of all possible combinations of  $co$  and  $arb$  is  $2^{2n^2}$ . To eliminate undesired  $co$  candidates as early as possible, we move the two constraints  $\text{IsStrictPartialOrder}(co, ops)$  and  $\text{Respect}(co, \text{PO}(h))$  on  $co$  to the front, before enumerating  $arb$  (see  $CCv1$  in Fig.10).

### 5.3.2 CCv2: Computing Linear Extensions of $co$ As Candidates for $arb$

The axiom  $\text{AxArb}$  requires  $co \subseteq arb$ . Therefore, we can directly compute the linear extensions of  $co$  as candidates for  $arb$ , instead of enumerating all possible relations on  $ops$  (see  $CCv2$  in Fig.10).

$$\begin{array}{l}
\text{MODULE } \textit{Relations} \\
\text{RF}(h) \triangleq \{ \langle w, r \rangle \in \textit{WriteOps}(h) \times \textit{ReadOps}(h) : w.\textit{key} = r.\textit{key} \wedge w.\textit{val} = r.\textit{val} \} \\
\text{CO}(h) \triangleq \textit{TC}(\textit{PO}(h) \cup \textit{RF}(h)) \\
\text{CF}(h) \triangleq \text{LET } \textit{co} \triangleq \textit{CO}(h) \textit{rf} \triangleq \textit{RF}(h) \\
\quad \text{IN } \{ \langle w1, w2 \rangle \in \textit{WriteOps}(h) \times \textit{WriteOps}(h) : \\
\quad \quad \wedge w1.\textit{key} = w2.\textit{key} \\
\quad \quad \wedge w1.\textit{val} \neq w2.\textit{val} \\
\quad \quad \wedge \exists r \in \textit{ReadOps}(h) : \langle w1, r \rangle \in \textit{co} \wedge \langle w2, r \rangle \in \textit{rf} \} \\
\text{HBo}(h, o) \triangleq \text{LET } \textit{base} \triangleq \textit{CO}(h) \mid \textit{CausalPast}(\textit{co}, o) \\
\quad \text{RECURSIVE } \textit{HBoRE}(-) \\
\quad \quad \textit{HBoRE}(\textit{hbo}) \triangleq \\
\quad \quad \quad \text{LET } \textit{update} \triangleq \{ \\
\quad \quad \quad \quad \langle w1, w2 \rangle \in \textit{WriteOps}(h) \times \textit{WriteOps}(h) : \\
\quad \quad \quad \quad \quad \wedge w1.\textit{key} = w2.\textit{key} \\
\quad \quad \quad \quad \quad \wedge w1.\textit{val} \neq w2.\textit{val} \\
\quad \quad \quad \quad \quad \wedge \exists r2 \in \textit{ReadOpsOnKey}(h, w2.\textit{key}) : \\
\quad \quad \quad \quad \quad \quad \wedge r2.\textit{val} = w2.\textit{val} \\
\quad \quad \quad \quad \quad \quad \wedge \langle w1, r2 \rangle \in \textit{hbo} \\
\quad \quad \quad \quad \quad \quad \wedge \vee r2 = o \vee \langle r2, o \rangle \in \textit{PO}(h) \} \\
\quad \quad \quad \quad \textit{hbo2} \triangleq \textit{update} \cup \textit{hbo} \\
\quad \quad \quad \quad \text{IN } \text{IF } \textit{hbo2} = \textit{hbo} \text{ THEN } \textit{hbo} \text{ ELSE } \textit{HBoRE}(\textit{TC}(\textit{hbo2})) \\
\quad \text{IN } \textit{TC}(\textit{HBoRE}(\textit{base}))
\end{array}$$

Fig.8. TLA<sup>+</sup> module *Relations*.

$$\begin{array}{l}
\text{MODULE } \textit{BadPatterns} \\
\text{CyclicCO}(h) \triangleq \textit{Cyclic}(\textit{PO}(h) \cup \textit{RF}(h)) \\
\text{WriteCOInitRead}(h) \triangleq \exists k \in \textit{KeyOf}(h) : \exists r \in \textit{ReadOpsOnKey}(h, k), \\
\quad \quad \quad w \in \textit{WriteOpsOnKey}(h, k) : \\
\quad \quad \quad \wedge \langle w, r \rangle \in \textit{CO}(h) \wedge r.\textit{val} = \textit{InitVal} \\
\text{ThinAirRead}(h) \triangleq \exists k \in \textit{KeyOf}(h) : \\
\quad \quad \exists r \in \textit{ReadOpsOnKey}(h, k) : \\
\quad \quad \quad \wedge r.\textit{val} \neq \textit{InitVal} \\
\quad \quad \quad \wedge \forall w \in \textit{WriteOpsOnKey}(h, k) : \langle w, r \rangle \notin \textit{RF}(h) \\
\text{WriteCORead}(h) \triangleq \exists k \in \textit{KeyOf}(h) : \\
\quad \quad \exists w1, w2 \in \textit{WriteOpsOnKey}(h, k), \\
\quad \quad \quad r1 \in \textit{ReadOpsOnKey}(h, k) : \\
\quad \quad \quad \wedge \langle w1, w2 \rangle \in \textit{CO}(h) \wedge \langle w2, r1 \rangle \in \textit{CO}(h) \\
\quad \quad \quad \wedge \langle w1, r1 \rangle \in \textit{RF}(h) \\
\text{CyclicCF}(h) \triangleq \textit{Cyclic}(\textit{CF}(h) \cup \textit{CO}(h)) \\
\text{WriteHBInitRead}(h) \triangleq \exists o \in \textit{Ops}(h) : \\
\quad \quad \exists r \in \textit{POPast}(h, o) : \\
\quad \quad \quad \wedge r.\textit{val} = \textit{InitVal} \\
\quad \quad \quad \wedge \text{LET } \textit{writes} \triangleq \textit{WriteOpsOnKey}(h, r.\textit{key}) \\
\quad \quad \quad \text{IN } \exists w \in \textit{writes} : \langle w, r \rangle \in \textit{HBo}(h, o) \\
\text{CyclicHB}(h) \triangleq \exists o \in \textit{Ops}(h) : \textit{Cyclic}(\textit{HBo}(h, o))
\end{array}$$

(a)

$$\begin{array}{l}
\text{MODULE } \textit{Algorithm} \\
\text{CCAlg}(h) \triangleq \wedge \neg \textit{CyclicCO}(h) \wedge \neg \textit{WriteCOInitRead}(h) \\
\quad \quad \wedge \neg \textit{ThinAirRead}(h) \wedge \neg \textit{WriteCORead}(h) \\
\text{CCvAlg}(h) \triangleq \wedge \textit{CCAlg}(h) \wedge \neg \textit{CyclicCF}(h) \\
\text{CMAlg}(h) \triangleq \wedge \textit{CCAlg}(h) \wedge \neg \textit{WriteHBInitRead}(h) \wedge \neg \textit{CyclicHB}(h)
\end{array}$$

(b)

Fig.9. TLA<sup>+</sup> modules for the “bad patterns” based checking algorithm of variants of causal consistency. (a) TLA<sup>+</sup> module *BadPatterns*. (b) TLA<sup>+</sup> module *Algorithm*.

MODULE <i>Optimization</i>
$  \begin{aligned}  CCv1(h) &\triangleq \text{LET } ops \triangleq Ops(h) \\  &\text{IN } \exists co \in \text{SUBSET } (ops \times ops) : \\  &\quad \wedge \text{Respect}(co, PO(h)) \\  &\quad \wedge \text{IsStrictPartialOrder}(co, ops) \\  &\quad \wedge \exists arb \in \text{SUBSET } (ops \times ops) : \\  &\quad \quad \wedge \text{Respect}(arb, co) \\  &\quad \quad \wedge \text{IsStrictTotalOrder}(arb, ops) \\  &\quad \quad \wedge \forall o \in ops : AxCausalArb(co, arb, o)  \end{aligned}  $
$  \begin{aligned}  CCv2(h) &\triangleq \\  &\text{LET } ops \triangleq Ops(h) \\  &\text{IN } \exists co \in \text{SUBSET } (ops \times ops) : \\  &\quad \wedge \text{Respect}(co, PO(h)) \\  &\quad \wedge \text{IsStrictPartialOrder}(co, ops) \\  &\quad \wedge \exists arb \in \{Seq2Rel(le) : le \in AllLinearExtensions(co, ops)\} : \\  &\quad \quad \wedge \forall o \in ops : AxCausalArb(co, arb, o)  \end{aligned}  $
$  \begin{aligned}  CCv3(h) &\triangleq \\  &\text{LET } ops \triangleq Ops(h) \\  &\text{IN } \exists co \in \text{StrictPartialOrderSubset}(ops) : \\  &\quad \wedge \text{Respect}(co, PO(h)) \\  &\quad \wedge \exists arb \in \{Seq2Rel(le) : le \in AllLinearExtensions(co, ops)\} : \\  &\quad \quad \wedge \forall o \in ops : AxCausalArb(co, arb, o)  \end{aligned}  $

Fig.10. TLA<sup>+</sup> module *Optimization*.

### 5.3.3 CCv3: Enumerating Strict Partial Order As Candidates for *co*

In *CCv2*, we still need to enumerate all possible relations on *ops* as candidates for *co*, and then eliminate the ones that are not strict partial orders. In *CCv3* (Fig.10), we directly compute all possible strict partial orders on *ops*. To this end, we implement the efficient partial order enumeration algorithm of [24] in Python,

and let TLC call it when necessary<sup>⑫</sup>.

## 5.4 Model Checking Results

We verify the TLA<sup>+</sup> specification of causal consistency and their “bad patterns” based checking algorithms against five sample histories from [11] using the TLC model checker. The sample histories are described in TLA<sup>+</sup> in module *Samples* (Fig.11). It is quite easy

MODULE <i>Samples</i>
$  \begin{aligned}  hasa &\triangleq \langle W("x", 1, 1), R("x", 2, 2) \rangle \\  hasb &\triangleq \langle W("x", 2, 3), R("x", 1, 4) \rangle \\  ha &\triangleq \{hasa, hasb\} \text{ CM but not CCv}  \end{aligned}  $
$  \begin{aligned}  hbsa &\triangleq \langle W("z", 1, 1), W("x", 1, 2), W("y", 1, 3) \rangle \\  hbsb &\triangleq \langle W("x", 2, 4), R("z", 0, 5), R("y", 1, 6), R("x", 2, 7) \rangle \\  hb &\triangleq \{hbsa, hbsb\} \text{ CCv but not CM}  \end{aligned}  $
$  \begin{aligned}  hcsa &\triangleq \langle W("x", 1, 1) \rangle \\  hcsb &\triangleq \langle W("x", 2, 2), R("x", 1, 3), R("x", 2, 4) \rangle \\  hc &\triangleq \{hcsa, hcsb\} \text{ CC but not CM nor CCv}  \end{aligned}  $
$  \begin{aligned}  hdsa &\triangleq \langle W("x", 1, 1), W("y", 2, 2), R("y", 2, 3) \rangle \\  hdsb &\triangleq \langle W("y", 1, 4), R("x", 1, 5), R("y", 1, 6) \rangle \\  hd &\triangleq \{hdsa, hdsb\} \text{ CC, CM, and CCv}  \end{aligned}  $
$  \begin{aligned}  hesa &\triangleq \langle W("x", 1, 1), W("y", 1, 2) \rangle \\  hesb &\triangleq \langle R("y", 1, 3), W("x", 2, 4) \rangle \\  hesc &\triangleq \langle R("x", 2, 5), R("x", 1, 6) \rangle \\  he &\triangleq \{hesa, hesb, hesc\} \text{ not CC (nor CM, nor CCv)}  \end{aligned}  $

Fig.11. TLA<sup>+</sup> module *Samples*.

<sup>⑫</sup>Technically, we need to wrap it in Java first.

to manually check them against each causal consistency variant.

As shown in Table 8, the “bad patterns” based checking algorithms meet their corresponding specifications as expected. It also confirms the satisfaction or violation of the sample histories. This demonstrates, though on test cases of relatively small scales, the correctness of the checking algorithms. Note that it takes much longer to check the history *hb* which consists of two sessions and seven operations directly against the specifications than to use the polynomial “bad patterns” based checking algorithms.

Table 9 shows the time for checking histories *ha*, *hb*, and *hd* against different versions of *CCv* proposed in Subsection 5.3. It demonstrates that each optimization is quite effective in reducing the checking time. Note that *CCv3* is the only version that is feasible for history *hb* with seven operations.

## 5.5 Relating TLA<sup>+</sup> Specification to Jepsen Testing

As summarized in Fig.12, we have two TLA<sup>+</sup> specifications, one for causal consistency variants and the other for “bad patterns” based checking algorithms. We have also a Java implementation of these checking algorithms used in Jepsen testing of MongoDB. Now we explain how they can interact with each other.

On the one hand, utilizing TLC we are able to automatically generate various kinds of histories from the TLA<sup>+</sup> specification of causal consistency variants. Especially, histories satisfying or violating some or all causal consistency variants can be used as test oracles for both the specification and our Java implementation of checking algorithms. They can be used as test oracles for both the specification and our Java implementation of the checking algorithms. On the other hand, it is convenient for MongoDB to generate arbitrarily long histories in real deployment. By checking them against both the TLA<sup>+</sup> specification and our Java implementation of the checking algorithms, we can gain more confidence in our implementation.

## 6 Related Work

### 6.1 Jepsen Testing of MongoDB

The Jepsen team has tested MongoDB concerning its consistency models several times in recent years.

- In 2013, the team tested the election and data replication protocol of MongoDB 2.4.3<sup>[13]</sup>. It showed that acknowledged writes may be lost under network partitions at all consistency levels.

- In 2015, the team tested the single-document consistency of MongoDB 2.6.7<sup>[14]</sup>. It showed that “strictly consistent” reads may see stale versions of documents,

**Table 8.** Model Checking Results on Sample Histories Defined in Fig.11

History	Number of Sessions	Number of Operations	Specification						Checking Algorithm					
			<i>CC</i>		<i>CCv</i>		<i>CM</i>		<i>CCAlg</i>		<i>CCvAlg</i>		<i>CMAlg</i>	
			Result	Time (ms)	Result	Time (ms)	Result	Time (ms)	Result	Time (ms)	Result	Time (ms)	Result	Time (ms)
<i>ha</i>	2	4	✓	1 161	✗	1 155	✓	938	✓	898	✗	802	✓	1 140
<i>hb</i>	2	7	✓	83 089	✓	79 089	✗	82 930	✓	867	✓	990	✗	1 886
<i>hc</i>	2	4	✓	1 073	✗	836	✗	940	✓	950	✗	885	✗	1 321
<i>hd</i>	2	6	✓	2 326	✓	2 318	✓	2 296	✓	945	✓	951	✓	1 166
<i>he</i>	3	6	✗	2 620	✗	3 237	✗	2 673	✗	921	✗	769	✗	868

Note: ✓: satisfaction; ✗: violation.

**Table 9.** Time of Checking Histories Against Different Versions of *CCv*

	<i>ha</i> (4 Operations)	<i>hd</i> (6 Operations)	<i>hb</i> (7 Operations)
<i>CCv</i>	2 927 000 ms	> 24 h	> 24 h
<i>CCv1</i>	2 051 ms	73 020 000 ms	> 24 h
<i>CCv2</i>	1 469 ms	85 000 ms	> 24 h
<i>CCv3</i>	1 161 ms	2 326 ms	83 000 ms

<sup>[13]</sup>Jepsen: MongoDB. <https://aphyr.com/posts/284-call-me-maybe-mongodb>, Oct. 2021.

<sup>[14]</sup>Jepsen: MongoDB stale reads. <https://aphyr.com/posts/322-jepsen-mongodb-stale-reads>, Oct. 2021.

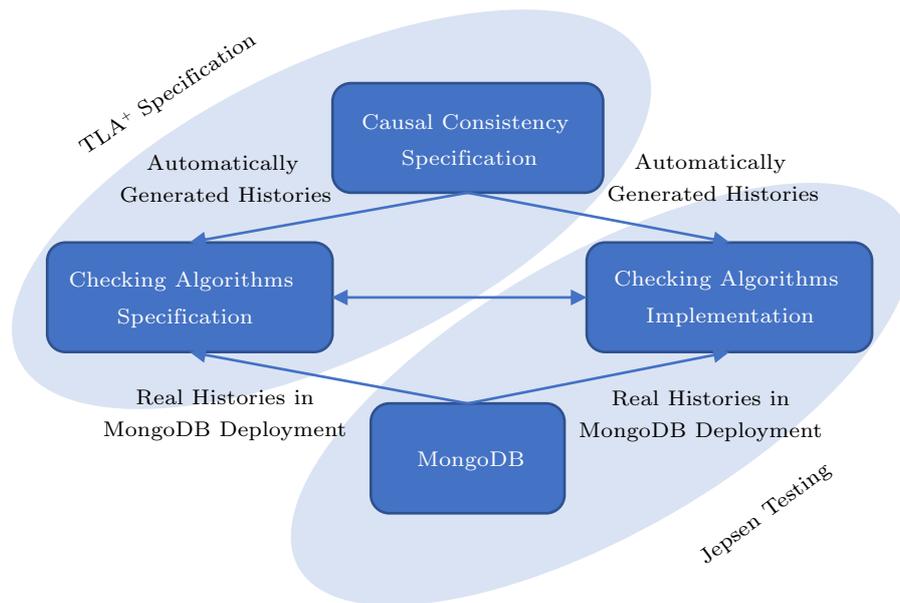


Fig.12. Relating TLA<sup>+</sup> specification to Jepsen testing.

and worse still they may return garbage data that has never been written before.

- In 2017, the team tested the v0 and v1 replication protocols of MongoDB 3.4.0-rc3<sup>⑮</sup>. It showed that the v0 replication protocol may lose the majority-committed documents. The new v1 replication protocol also contains bugs, allowing data loss in all versions up to MongoDB 3.2.11 and 3.4.0-rc4.

- In 2018, the team tested the causal consistency protocol of MongoDB 3.6.4. It showed that in the presence of node failures or network partitions, causal consistency is guaranteed only for reads with `majority readConcern` and writes with `majority writeConcern`. In this paper, we identify several drawbacks of this testing in terms of specification, test case generation, implementation of causal consistency checking algorithms, and testing scenarios. We also propose a more thorough design of Jepsen testing of the causal consistency protocol of MongoDB.

- In 2020, the team tested the transactional consistency models of MongoDB 4.2.6<sup>⑯</sup>. It showed that MongoDB failed to preserve snapshot isolation, even for reads with `majority readConcern` and writes with `majority writeConcern`.

## 6.2 Consistency Checking Problem

Much work has been devoted to the problem of checking whether a given history satisfies a desirable consistency model. Gibbons and Korach<sup>[25]</sup> systematically studied the complexity of the checking problem against strong consistency models, including linearizability<sup>[26]</sup> and sequential consistency<sup>[27]</sup>. Regarding weak consistency models, Wei *et al.*<sup>[28]</sup> addressed the problem of checking PRAM consistency<sup>[29]</sup> over histories of read/write registers. They first proved that for non-differentiated histories, the decision problem is NP-complete, and then proposed a polynomial-time checking algorithm for differentiated histories. Recently, Bouajjani *et al.* addressed the problem of checking causal consistency<sup>[11]</sup>. They considered three well-known variants of causal consistency, namely CC, CM, and CCv. They proved that checking whether a general history of arbitrary replicated objects satisfies CC, CM, or CCv is NP-hard, and that it is NP-complete for histories of read/write registers. Moreover, they proposed polynomial-time algorithms for differentiated histories of read/write registers. In this paper, we fully implement these efficient checking algorithms and utilize them to test the causal consistency protocol of MongoDB.

<sup>⑮</sup>Jepsen Testing of MongoDB 3.4.0-rc3. <https://jepsen.io/analyses/mongodb-3-4-0-rc3>, Oct. 2021.

<sup>⑯</sup>Jepsen Testing of MongoDB 4.2.6. <https://jepsen.io/analyses/mongodb-4.2.6>, Oct. 2021.

## 7 Conclusions

We proposed a thorough design of Jepsen testing of the causal consistency protocol of MongoDB. It strengthened the official Jepsen testing in 2018 in terms of specification, test case generation, implementation of causal consistency checking algorithms, and testing scenarios. The experimental results confirmed the claim of causal consistency in MongoDB's documentation. We also developed formal specifications of causal consistency and their checking algorithms in TLA+. The model checking results demonstrated the correctness of the checking algorithms and we can gain more confidence in our implementation.

In the future, we will explore the issues discussed in Subsection 5.5 and more intensive experiments are needed. We plan to improve the official Jepsen testing of the transaction protocols of MongoDB 4.2.6. On the other hand, we are also interested in applying formal methods to MongoDB's protocols. Specifically, we will formally specify these protocols in TLA+, verify them using the TLC model checker, and develop mechanical correctness proofs for them using TLAPS.

## References

- [1] Schultz W, Avitabile T, Cabral A. Tunable consistency in MongoDB. *Proc. VLDB Endow.*, 2019, 12(12): 2071-2081. DOI: [10.14778/3352063.3352125](https://doi.org/10.14778/3352063.3352125).
- [2] Tyulenev M, Schwerin A, Kamsky A, Tan R, Cabral A, Mulrow J. Implementation of cluster-wide logical clock and causal consistency in MongoDB. In *Proc. the 2019 International Conference on Management of Data*, June 30-July 5, 2019, pp.636-650. DOI: [10.1145/3299869.3314049](https://doi.org/10.1145/3299869.3314049).
- [3] Abadi D. Consistency tradeoffs in modern distributed database system design: CAP is only part of the story. *IEEE Computer*, 2012, 45(2): 37-42. DOI: [10.1109/MC.2012.33](https://doi.org/10.1109/MC.2012.33).
- [4] Brewer E A. Towards robust distributed systems (abstract). In *Proc. the 19th Annual ACM Symposium on Principles of Distributed Computing*, July 2000, Article No. 7. DOI: [10.1145/343477.343502](https://doi.org/10.1145/343477.343502).
- [5] Gilbert S, Lynch N. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *ACM SIGACT News*, 2002, 33(2): 51-59. DOI: [10.1145/564585.564601](https://doi.org/10.1145/564585.564601).
- [6] Brzezinski J, Sobaniec C, Wawrzyniak D. From session causality to causal consistency. In *Proc. the 12th Euro-micro Conference on Parallel, Distributed and Network-Based Processing*, Feb. 2004, pp.152-158. DOI: [10.1109/EM-PDP.2004.1271440](https://doi.org/10.1109/EM-PDP.2004.1271440).
- [7] Kulkarni S S, Demirbas M, Madappa D, Avva B, Leone M. Logical physical clocks. In *Proc. the 18th International Conference on Principles of Distributed Systems*, Dec. 2014, pp.17-32. DOI: [10.1007/978-3-319-14472-6\\_2](https://doi.org/10.1007/978-3-319-14472-6_2).
- [8] Du J, Iorgulescu C, Roy A, Zwaenepoel W. GentleRain: Cheap and scalable causal consistency with physical clocks. In *Proc. the ACM Symposium on Cloud Computing*, Nov. 2014, Article No. 4. DOI: [10.1145/2670979.2670983](https://doi.org/10.1145/2670979.2670983).
- [9] Akkoorath D D, Tomsic A Z, Bravo M, Li Z, Crain T, Bienuša A, Pregoça N, Shapiro M. Cure: Strong semantics meets high availability and low latency. In *Proc. the 36th International Conference on Distributed Computing Systems*, June 2016, pp.405-414. DOI: [10.1109/ICDCS.2016.98](https://doi.org/10.1109/ICDCS.2016.98).
- [10] Ongaro D, Ousterhout J. In search of an understandable consensus algorithm. In *Proc. the 2014 USENIX Conference on USENIX Annual Technical Conference*, June 2014, pp.305-320.
- [11] Bouajjani A, Enea C, Guerraoui R, Hamza J. On verifying causal consistency. In *Proc. the 44th ACM Symposium on Principles of Programming Languages*, Jan. 2017, pp.626-638. DOI: [10.1145/3009837.3009888](https://doi.org/10.1145/3009837.3009888).
- [12] Burckhardt S. Principles of eventual consistency. *Found. Trends Program. Lang.*, 2014, 1(1/2): 1-150. DOI: [10.1561/25000000011](https://doi.org/10.1561/25000000011).
- [13] Perrin M, Mostéfaoui A, Jard C. Causal consistency: Beyond memory. In *Proc. the 21st ACM Symposium on Principles and Practice of Parallel Programming*, Aug. 2016, Article No. 26. DOI: [10.1145/2851141.2851170](https://doi.org/10.1145/2851141.2851170).
- [14] Ahamad M, Neiger G, Burns J E, Kohli P, Hutto P W. Causal memory: Definitions, implementation, and programming. *Distributed Computing*, 1995, 9(1): 37-49. DOI: [10.1007/BF01784241](https://doi.org/10.1007/BF01784241).
- [15] Lynch N A. Distributed Algorithms. Morgan Kaufmann Publishers Inc., 1996.
- [16] Ouyang H R, Wei H F, Huang Y. Checking causal consistency of MongoDB. In *Proc. the 12th Asia-Pacific Symposium on Internetware*, Nov. 2020, pp.209-216. DOI: [10.1145/3457913.3457928](https://doi.org/10.1145/3457913.3457928).
- [17] Lamport L. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 1978, 21(7): 558-565. DOI: [10.1145/359545.359563](https://doi.org/10.1145/359545.359563).
- [18] Lesani M, Bell C J, Chlipala A. Chapar: Certified causally consistent distributed key-value stores. In *Proc. the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Jan. 2016, pp.357-370. DOI: [10.1145/2837614.2837622](https://doi.org/10.1145/2837614.2837622).
- [19] Lamport L. Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers (1st edition). Addison-Wesley Professional, 2002.
- [20] Lamport L. The temporal logic of actions. *ACM Trans. Program. Lang. Syst.*, 1994, 16(3): 872-923. DOI: [10.1145/177492.177726](https://doi.org/10.1145/177492.177726).
- [21] Wei H F, Tang R Z, Huang Y, Lv J. Jupiter made abstract, and then refined. *Journal of Computer Science and Technology*, 2020, 35(6): 1343-1364. DOI: [10.1007/s11390-020-0516-0](https://doi.org/10.1007/s11390-020-0516-0).
- [22] Yu Y, Manolios P, Lamport L. Model checking TLA+ specifications. In *Proc. the 10th IFIP WG 10.5 Advanced Research Working Conference on Correct Hardware Design and Verification Methods*, Sept. 1999, pp.54-66. DOI: [10.1007/3-540-48153-2\\_6](https://doi.org/10.1007/3-540-48153-2_6).
- [23] Cooper B F, Silberstein A, Tam E, Ramakrishnan R, Sears R. Benchmarking cloud serving systems with YCSB. In *Proc. the 1st ACM Symposium on Cloud Computing*, June 2010, pp.143-154. DOI: [10.1145/1807128.1807152](https://doi.org/10.1145/1807128.1807152).

- [24] Bowles J, Caminati M B. A verified algorithm enumerating event structures. In *Proc. the 10th International Conference on Intelligent Computer Mathematics*, July 2017, pp.239-254. DOI: [10.1007/978-3-319-62075-6\\_17](https://doi.org/10.1007/978-3-319-62075-6_17).
- [25] Gibbons P, Korach E. Testing shared memories. *SIAM Journal on Computing*, 1997, 26(4): 1208-1244. DOI: [10.1137/S0097539794279614](https://doi.org/10.1137/S0097539794279614).
- [26] Herlihy M P, Wing J M. Linearizability: A correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 1990, 12(3): 463-492. DOI: [10.1145/78969.78972](https://doi.org/10.1145/78969.78972).
- [27] Attiya H, Welch J L. Sequential consistency versus linearizability. *ACM Trans. Comput. Syst.*, 1994, 12(2): 91-122. DOI: [10.1145/176575.176576](https://doi.org/10.1145/176575.176576).
- [28] Wei H, Huang Y, Cao J, Ma X, Lv J. Verifying Pipelined-RAM consistency over read/write traces of data replicas. *IEEE Transactions on Parallel and Distributed Systems*, 2013, 27(5): 1511-1523. DOI: [10.1109/TPDS.2015.2453985](https://doi.org/10.1109/TPDS.2015.2453985).
- [29] Lipton R J, Sandberg J. PRAM: A scalable shared memory. Technical Report, Department of Computer Science, Princeton University, 1988. <https://www.cs.princeton.edu/research/techreps/TR-180-88>, Aug. 2021.



**Hong-Rong Ouyang** received his B.S. degree in computer science and technology from Nanjing University, Nanjing, in 2020. He is currently a Master student with the Department of Computer Science and Technology and the State Key Laboratory for Novel Software Technology at Nanjing University, Nanjing. His research interests include distributed databases and formal methods.



**Heng-Feng Wei** received his B.S. and Ph.D. degrees in computer science and technology from Nanjing University, Nanjing, in 2009 and 2016, respectively. He is currently a research assistant with Software Institute and the State Key Laboratory for Novel Software Technology at Nanjing University, Nanjing. His research interests include distributed computing and formal methods. He is a member of CCF.



**Hai-Xiang Li** is currently a chief researcher and chief architect of the distributed database system TDSQL at Tencent Inc., Shenzhen. His research interests include distributed computing, cloud database, transaction processing, and query optimization. He is a member of CCF.



**An-Qun Pan** is the technical director of Tencent Billing Platform Department, Shenzhen. He has more than 15 years' experience in the research and development of distributed computing and storage systems. He is currently responsible for the research and development of the distributed database system TDSQL. He is a member of CCF.



**Yu Huang** received his B.S. and Ph.D. degrees in computer science from the University of Science and Technology of China, Hefei, in 2002 and 2007, respectively. He is currently a professor with the Department of Computer Science and Technology and the State Key Laboratory for Novel Software Technology at Nanjing University, Nanjing. His research interests include distributed algorithms, distributed systems, formal methods, and system reliability. He is a member of CCF.

# JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY

Volume 37, Number 1, January 2022

## Content

### Special Section on Software Systems 2021

Preface..... *Tao Xie, Shengchao Qin, Wenhui Zhang, Jun Sun, Lei Bu, and Ge Li* ( 1 )

### Theme: Dependable Software Engineering

Event-Based Semantics of UML 2.X Concurrent Sequence Diagrams for Formal Verification.....  
..... *Inès Mouakher, Fatma Dhaou, and J. Christian Attiogbé* ( 4 )  
DeltaFuzz: Historical Version Information Guided Fuzz Testing.....  
..... *Jia-Ming Zhang, Zhan-Qi Cui, Xiang Chen, Huan-Huan Wu, Li-Wei Zheng, and Jian-Bin Liu* ( 29 )

### Theme: Internetware and Beyond

TOAST: Automated Testing of Object Transformers in Dynamic Software Updates.....  
..... *Ze-Lin Zhao, Di Huang, and Xiao-Xing Ma* ( 50 )  
Characterizing and Detecting Gas-Inefficient Patterns in Smart Contracts.....  
..... *Que-Ping Kong, Zi-Yan Wang, Yuan Huang, Xiang-Ping Chen, Xiao-Cong Zhou, Zi-Bin Zheng, and Gang Huang* ( 67 )  
Simulation Might Change Your Results: A Comparison of Context-Aware System Input Validation in Simulated and Physical  
Environments..... *Jin-Chi Chen, Yi Qin, Hui-Yan Wang, and Chang Xu* ( 83 )  
Meaningful Update and Repair of Markov Decision Processes for Self-Adaptive Systems.....  
..... *Wen-Hua Yang, Min-Xue Pan, Yu Zhou, and Zhi-Qiu Huang* ( 106 )  
Checking Causal Consistency of MongoDB.....  
..... *Hong-Rong Ouyang, Heng-Feng Wei, Hai-Xiang Li, An-Qun Pan, and Yu Huang* ( 128 )  
GridDroid—An Effective and Efficient Approach for Android Repackaging Detection Based on Runtime Graphical User  
Interface..... *Jun Ma, Qing-Wei Sun, Chang Xu, and Xian-Ping Tao* ( 147 )  
Community Smell Occurrence Prediction on Multi-Granularity by Developer-Oriented Features and Process Metrics.....  
..... *Zi-Jie Huang, Zhi-Qing Shao, Gui-Sheng Fan, Hui-Qun Yu, Xing-Guang Yang, and Kang Yang* ( 182 )

### Regular Paper

MacroTrend: A Write-Efficient Cache Algorithm for NVM-Based Read Cache.....  
..... *Ning Bao, Yun-Peng Chai, Xiao Qin, and Chuan-Wen Wang* ( 207 )  
Correlated Differential Privacy of Multiparty Data Release in Machine Learning.....  
..... *Jian-Zhe Zhao, Xing-Wei Wang, Ke-Ming Mao, Chen-Xi Huang, Yu-Kai Su, and Yu-Chen Li* ( 231 )  
On the Discrete-Time Dynamics of Cross-Coupled Hebbian Algorithm.....  
..... *Xiao-Wei Feng, Xiang-Yu Kong, Chuan He, and Dong-Hui Xu* ( 252 )  
A Blockchain-Based Protocol for Malicious Price Discrimination.....  
..... *Li-De Xue, Ya-Jun Liu, Wei Yang, Wei-Lin Chen, and Liu-Sheng Huang* ( 266 )

# JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY

《计算机科学技术学报》

Volume 37 Number 1 2022 (Bimonthly, Started in 1986)

Indexed in: SCIE, Ei, INSPEC, JST, AJ, MR, CA, DBLP

Edited by:

THE EDITORIAL BOARD OF JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY

Zhi-Wei Xu, Editor-in-Chief, P.O. Box 2704, Beijing 100190, P.R. China

Managing Editor: Feng-Di Shu E-mail: [jcst@ict.ac.cn](mailto:jcst@ict.ac.cn) <http://jcst.ict.ac.cn> Tel.: 86-10-62610746

Copyright ©Institute of Computing Technology, Chinese Academy of Sciences 2022  
Sponsored by: Institute of Computing Technology, CAS & China Computer Federation

Supervised by: Chinese Academy of Sciences  
Undertaken by: Institute of Computing Technology, CAS  
Published by: Science Press, Beijing, China  
Printed by: Beijing Baochang Color Printing Co. Ltd

Distributed by:

China: All Local Post Offices

Other Countries: Springer Nature Customer Service Center GmbH, Tiergartenstr. 15, 69121 Heidelberg, Germany

Available Online: <https://link.springer.com/journal/11390>

