



# Pushing the Limit: Verified Performance-Optimal Causally-Consistent Database Transactions

Shabnam Ghasemirad<sup>(✉)</sup>, Christoph Sprenger, Si Liu,  
Luca Multazzu, and David Basin

ETH Zurich, Zurich, Switzerland  
shabnam.ghasemirad@inf.ethz.ch

**Abstract.** Modern web services crucially rely on high-performance distributed databases, where concurrent transactions are isolated from each other using concurrency control protocols. Relaxed isolation levels, which permit more complex concurrent behaviors than strong levels like serializability, are used in practice for higher performance and availability. In this paper, we present Eiger-PORT<sup>+</sup>, a concurrency control protocol that achieves a strong form of causal consistency, called TCCv (Transactional Causal Consistency with convergence). We show that Eiger-PORT<sup>+</sup> also provides performance-optimal read transactions in the presence of transactional writes, thus refuting an open conjecture that this is impossible for TCCv. We also deductively verify that Eiger-PORT<sup>+</sup> satisfies this isolation level by refining an abstract model of transactions. This yields the first deductive verification of a complex concurrency control protocol. Furthermore, we conduct a performance evaluation showing Eiger-PORT<sup>+</sup>'s superior performance over the state-of-the-art.

## 1 Introduction

Modern web services are built on top of high-performance database systems operating in partitioned, geo-distributed environments. These systems provide distributed transactions that group the users' read and write requests. To balance data consistency and system performance, databases provide a spectrum of *isolation levels* (I in ACID: Atomicity, Consistency, Isolation, and Durability [46]), defining the degree of separation between concurrent transactions. Isolation is enforced by *concurrency control protocols* (also called *transaction protocols*).

Many applications, such as social networks, opt for weak isolation levels to avoid the performance overhead of stronger levels like serializability [43]. These weaker guarantees allow distributed transactions to remain functional even during network partitions, while still providing useful properties. Notably, *transactional causal consistency* (TCC) represents a successful integration of ideas from the distributed computing and database communities. It extends causal consistency [2, 44]—the strongest consistency level achievable in an always-available system [4]—by incorporating transactional guarantees. The past decade has seen sustained efforts in designing databases supporting performant causally-consistent distributed transactions [35, 3, 11, 39, 38, 47], along with their growing adoption in industry [41, 14, 40]. Nearly all of these systems provide a stronger

variant of TCC, known as TCCv [3,35], that includes *data convergence* requiring views across different clients to eventually converge to the same state.

In this paper, we present a case study on developing and verifying a performance-optimal, causally-consistent, database transaction protocol. Our protocol, Eiger-PORT+, provides TCCv, for which we give a formal proof. Our work faced two challenges. First, it is not a priori clear that such an isolation guarantee is achievable for a performance-optimal protocol. In fact, Lu et al. [38] conjectured that for distributed performance-optimal read-only transactions (PORTs) in the presence of transactional writes, TCC (without convergence) is the strongest achievable isolation level. They presented the Eiger-PORT protocol, which provides this guarantee. In this paper, we constructively refute their conjecture by designing our novel protocol, Eiger-PORT+, which achieves the stronger TCCv guarantee.

Second, transaction protocols are notoriously hard to get right, as witnessed by numerous design-level isolation errors in production databases [24,20,28,21,22], and even in protocols that have undergone pen-and-pencil proofs [42] and model-checking analysis [37]. We thus aim for a full deductive verification of Eiger-PORT+, covering *all* possible behaviors. To our knowledge, the deductive verification of transaction protocols for weak isolation levels, which exhibit complex concurrent behaviors, has not been attempted so far. Previous efforts in this area have focused on simple textbook protocols like two-phase locking achieving serializability or employed model checking, which requires bounding the number of processes and transactions. We address this challenge using our Isabelle/HOL framework [15] built around Xiong et al.’s abstract transaction model [50]. We formalize Eiger-PORT+ and show that it satisfies TCCv using reduction [26] in combination with a refinement of the TCCv instance of the abstract transaction model.

Furthermore, we implement and deploy Eiger-PORT+, along with Eiger-PORT and its precursor Eiger, and conduct a comprehensive performance comparison of these three protocols. Our evaluation demonstrates Eiger-PORT+’s superior performance in terms of system throughput and latency across various scenarios, e.g., with a growing number of clients and servers.

The complete formal development accompanying this paper, including all definitions and proofs, as well as a protocol implementation are available at [17].

**Contributions** Overall, we see our contributions as three-fold:

- **Conjecture refutation.** We formally refute the conjecture that TCC is the strongest achievable isolation level for PORTs in the presence of transactional writes by designing a protocol, Eiger-PORT+, that provably achieves TCCv.
- **Proof of correctness.** We model Eiger-PORT+ in Isabelle/HOL and verify its correctness by showing that its behavior conforms to the TCCv instance of the abstract transaction model [50]. This represents the first complete formal verification of a complex distributed database transaction protocol.
- **Superior performance.** We deploy Eiger-PORT+, along with two state-of-the-art causally-consistent transaction protocols, in a cluster and evaluate their performance. Our experimental results demonstrate Eiger-PORT+’s superior performance, with both lower latency and higher throughput.




Fig. 1: A spectrum of isolation levels.  $A \rightarrow B$  means  $A$  is weaker than  $B$ . RC: read committed [7]; RA: read atomicity [6]; TCC: transactional causal consistency [38], provided by Eiger-PORT [38]; TCCv: TCC with convergence [3,35], offered by Eiger and our Eiger-PORT+; SI: snapshot isolation [7]; (S)SER: (strict) serializability [43]. Protocols supporting PORTs are highlighted in gray.

## 2 Background

### 2.1 Distributed Database Transactions

In a distributed database, vast amounts of data are split up and stored across multiple servers, also called partitions. User requests are submitted as database transactions, initiated by front-end clients. Each client executes the transactions in its own session, where a transaction comprises a sequence of read and/or write operations on data items (or keys) distributed across partitions.

**Isolation levels** Distributed databases offer various isolation levels, differing on how they balance data consistency and system performance. Figure 1 shows a spectrum of practically relevant isolation levels, ranging from weaker ones like Read Committed, through various forms of transactional causality, to stronger guarantees such as Serializability. We briefly explain Read Atomicity and two variants of transactional causality, which are the focus of this work.

**Read Atomicity (RA).** This is also known as *atomic visibility*, requiring that all or none of a transaction’s updates are observed by other transactions. It prohibits *fractured reads* anomalies, such as Carol only observing one direction of a new (bi-directional) friendship between Alice and Bob in a social network.

**Transactional Causal Consistency (TCC).** In addition to RA, this level requires that two causally related transactions appear to all client sessions in the same causal order [2,44]. It prevents *causality violations*, such as Carol observing Bob’s response to Alice’s message without seeing the message itself.

**TCC with Convergence (TCCv).** With TCC, different clients may observe causally unrelated transactions in different orders. TCCv’s *convergence* property prevents this by requiring these clients’ views to converge to the same state [34,3]. For example, this prevents confusion created by Alice and Bob independently posting “Let’s meet at my place” in a road trip planner. In practice, most causally-consistent databases provide convergence.

**Performance-optimal read-only transactions** NOCS [38] is the state-of-the-art impossibility result that captures conflicts between distributed transactions’

performance and their isolation guarantees. NOCS proves that read-only transactions cannot complete with **Non**-blocking communication in **One** round and **Constant**-size metadata, while achieving **Strict** serializability (SSER). At best, three of the four NOCS properties can be satisfied. In particular, protocols satisfying the NOC properties (under isolation levels weaker than SSER) are said to provide *performance-optimal* read-only transactions (PORTs). The NOCS authors [38] introduce the Eiger-PORT protocol providing TCC and PORTs. They also state the following conjecture, which has remained unresolved for four years.

**Conjecture.** TCC is the strongest isolation level achievable for PORTs in the presence of transactional writes.

Recent studies show that write-heavy workloads involving transactional writes are more prevalent than previously assumed and are expected to become increasingly prominent [52]. This, along with the practical significance of TCCv, motivates our work on refuting the above conjecture and pushing the boundary.

## 2.2 Transition Systems and Refinement

We use *labeled transition systems* (LTSs) to model database protocols and the abstract transaction model. An LTS  $\mathcal{E} = (S, I, \{\xrightarrow{e} \mid e \in E\})$  consists of a set of states  $S$ , a non-empty set of initial states  $I \subseteq S$ , and transition relations  $\xrightarrow{e} \subseteq S \times S$ , one for each event  $e \in E$ . We assume an idling event  $\text{skip} \in E$  with  $s \xrightarrow{\text{skip}} s$ . We often define the relations  $\xrightarrow{e}$  using *guard* predicates  $G_e$  and *update* functions  $U_e$  by  $s \xrightarrow{e} s'$  if and only if  $G_e(s) \wedge s' = U_e(s)$ . A state  $s$  is *reachable* if a sequence of transitions leads from an initial state to  $s$ . We denote the set of reachable states of  $\mathcal{E}$  by  $\text{reach}(\mathcal{E})$ . A set of states  $J$  is an *invariant* if  $\text{reach}(\mathcal{E}) \subseteq J$ .

*Refinement* relates two LTSs  $\mathcal{E}_i = (S_i, I_i, \{\xrightarrow{e}_i \mid e \in E_i\})$ , for  $i \in \{1, 2\}$ . Given *refinement mappings*  $r: S_2 \rightarrow S_1$  and  $\pi: E_2 \rightarrow E_1$  between the LTSs' states and events, we say  $\mathcal{E}_2$  *refines*  $\mathcal{E}_1$ , written  $\mathcal{E}_2 \preceq_{r, \pi} \mathcal{E}_1$ , if (i)  $r(s) \in I_1$  for all  $s \in I_2$  and (ii)  $r(s) \xrightarrow{\pi(e)}_1 r(s')$  whenever  $s \xrightarrow{e}_2 s'$ . Using guards and updates, (ii) reduces to two proof obligations: assuming  $G_e^2(s)$  prove (a)  $G_{\pi(e)}^1(r(s))$  (*guard strengthening*) and (b)  $r(U_e^2(s)) = U_{\pi(e)}^1(r(s))$  (*update correspondence*). Refinement guarantees the inclusion of sets of reachable states (modulo  $r$ ), i.e.,  $r(\text{reach}(\mathcal{E}_2)) \subseteq \text{reach}(\mathcal{E}_1)$ , where  $r$  is applied to each element of  $\text{reach}(\mathcal{E}_2)$ . Refinement proofs often require invariants to strengthen the refinement mapping.

## 2.3 An Abstract Transaction Model

Xiong et al. [50] introduced a centralized operational model for atomic transactions operating on distributed multi-versioned key-value stores (KVSSs), where the database stores data as key-value pairs and each key may be mapped to multiple versions for increased data availability. This model can be instantiated to different isolation guarantees, including RA, TCCv, and SSER (cf. Figure 1). They prove

the equivalence of these model instances to their declarative counterparts based on abstract executions. The model can be used to prove the correctness of both concurrency control protocols and client programs. We have formalized this framework in Isabelle/HOL and extended it with an extensive library of lemmas supporting protocol correctness proofs [15].

Xiong et al.'s model is formulated as an LTS, called the *abstract transaction model*, which abstracts the protocols' distributed collection of KVSs (each representing a shard and/or replica) into a single (centralized) multi-versioned KVS  $\mathcal{K} : \text{key} \rightarrow \text{version list}$  that maps each key to a list of *versions*. Each version  $\mathcal{K}(k, i)$  of a key  $k$  at the list index  $i$  records (i) the value  $v$  stored, (ii) the writer transaction  $t$  that has produced this version, and (iii) the reader set  $T$ , i.e., the set of transactions that have read this version. The pairs  $(t, t')$  for any  $t' \in T$  are called *write-read dependencies*. The relation  $\text{WR}_{\mathcal{K}}$  contains all such pairs. The fact that, in a real, distributed system, each client  $cl$  has a different partial *client view* of  $\mathcal{K}$  is modeled by explicitly representing these views in the model's configurations as mappings  $\mathcal{U}(cl) : \text{key} \rightarrow \text{nat set}$ . This describes, for each key, the set of versions (denoted by list indices) visible to the client. Clients are assumed to process transactions sequentially. The *session order* relation  $\text{SO}$  captures the order of their transactions.

The model assumes the *snapshot property*, ensuring that each transaction reads and writes at most one version of each key. Hence, transactions can be represented by a *fingerprint*  $\mathcal{F} : \text{key} \times \{\text{R}, \text{W}\} \rightarrow \text{value}$ , which maps each key and operation (read or write) to at most one value. It also assumes that views are *atomic*, i.e., clients observe either all or none of a transaction's effects. These properties together establish *atomic visibility*, also called Read Atomicity (RA) (cf. Section 2.1), as the model's baseline isolation guarantee.

The model has two events (plus skip): *commit*, which atomically executes an entire transaction, and *view extension*, which monotonically extends a client's view of the KVS. The commit event's executability depends on the isolation guarantee. Here, we focus on the model's TCCv instantiation, called  $\mathcal{I}_{\text{TCCv}}$ .

**Commit** The commit event's transition relation for TCCv is defined by:

$$\frac{\begin{array}{l} \mathcal{U}(cl) \sqsubseteq u \quad \text{canCommit}_{\text{TCCv}}(\mathcal{K}, u, \mathcal{F}) \quad u \sqsubseteq u' \quad \text{RYW}(\mathcal{K}, \mathcal{K}', u') \\ \text{LWW}(\mathcal{K}, u, \mathcal{F}) \quad \text{wf}(\mathcal{K}, u) \quad \text{wf}(\mathcal{K}', u') \\ t_{sn}^{cl} \in \text{nextTxids}(\mathcal{K}, cl) \quad \mathcal{K}' = \text{UpdateKV}(\mathcal{K}, t_{sn}^{cl}, u, \mathcal{F}) \end{array}}{(\mathcal{K}, \mathcal{U}) \xrightarrow{\text{commit}(cl, sn, u, \mathcal{F})}_{\text{TCCv}} (\mathcal{K}', \mathcal{U}[cl \mapsto u'])}$$

The transition in the conclusion updates the configuration  $(\mathcal{K}, \mathcal{U})$  to the new configuration  $(\mathcal{K}', \mathcal{U}[cl \mapsto u'])$ , where  $\mathcal{K}'$  is the updated KVS and  $\mathcal{U}[cl \mapsto u']$  updates the client  $cl$ 's view to  $u'$ . Both  $\mathcal{K}'$  and  $u'$  are determined by the rule's premises, which act as the event's guards with the following meanings:

- $\mathcal{U}(cl) \sqsubseteq u$ : This condition allows one to extend the client  $cl$ 's current view to a (point-wise) larger one before committing.

- $\text{canCommit}_{\text{TCCv}}(\mathcal{K}, u, \mathcal{F})$ : This is the central commit condition, which ensures that it is safe to commit a transaction at the TCCv isolation level. It requires that the set of visible transactions  $\text{visTx}(\mathcal{K}, u)$  (i.e., the writers of the versions that the view  $u$  points to) is *closed* under the relation  $\text{SO} \cup \text{WR}_{\mathcal{K}}$ , i.e.,

$$((\text{SO} \cup \text{WR}_{\mathcal{K}})^{-1})^+(\text{visTx}(\mathcal{K}, u)) \subseteq \text{visTx}(\mathcal{K}, u) \cup \text{rdonly}(\mathcal{K}).^1 \quad (1)$$

In other words, following the *causal dependency* relation  $\text{SO} \cup \text{WR}_{\mathcal{K}}$  backwards from visible transactions, we only see visible or read-only transactions.

- $u \sqsubseteq u'$ : This condition captures the *monotonic reads* session guarantee, i.e., the view  $u'$  extends the view  $u$ .
- $\text{RYW}(\mathcal{K}, \mathcal{K}', u')$ : This condition expresses the *read-your-writes* (RYW) session guarantee, stating that each client sees all versions previously written by itself.
- $\text{LWW}(\mathcal{K}, u, \mathcal{F})$ : This captures the *last-write-wins* conflict resolution policy, whereby a client reads each key's latest version in its view.
- $\text{wf}(\mathcal{K}, u)$  and  $\text{wf}(\mathcal{K}, u')$ : This requires that the views  $u$  and  $u'$  are *wellformed*, i.e., they are atomic and contain indices that point to existing versions.
- $t_{sn}^{cl} \in \text{nextTxids}(\mathcal{K}, cl)$ : Transaction identifiers  $t_{sn}^{cl} \in \text{TxID}$  are indexed by the issuing client  $cl$  and a (monotonically increasing) sequence number  $sn$ . This condition obtains a fresh transaction ID  $t_{sn}^{cl}$ , where the sequence number  $sn$  is larger than any of the client  $cl$ 's sequence numbers used in  $\mathcal{K}$ .
- $\mathcal{K}' = \text{UpdateKV}(\mathcal{K}, t_{sn}^{cl}, u, \mathcal{F})$ : The KVS  $\mathcal{K}'$  is obtained from  $\mathcal{K}$  by adding the operations described by the fingerprint  $\mathcal{F}$ : the writes append a new version with the writer ID  $t_{sn}^{cl}$  to the respective key's version list, and the reads add  $t_{sn}^{cl}$  to the respective versions' reader sets.

**View extension** The view extension event for TCCv is defined by the rule:

$$\frac{\mathcal{U}(cl) \sqsubseteq u \quad \text{wf}(\mathcal{K}, u)}{(\mathcal{K}, \mathcal{U}) \xrightarrow{\text{xview}(cl, u)}_{\text{TCCv}} (\mathcal{K}, \mathcal{U}[cl \mapsto u])}$$

and extends a client  $cl$ 's view from  $\mathcal{U}(cl)$  to a wellformed view  $u$ . It abstractly models that additional versions of certain keys become visible to the client.

### 3 Eiger-PORT+: An Overview

Causally-consistent transactions have attracted the attention of both academia and industry in recent years. Eiger [35] is among the first distributed databases providing TCCv. Eiger-PORT [38] improves Eiger's overall performance by optimizing its read-only transactions while sacrificing data convergence, thus allowing diverging client views of concurrent conflicting writes (to the same keys) as in TCC. We show that this sacrifice is unnecessary and design Eiger-PORT+ based on Eiger-PORT. Eiger-PORT+ provides TCCv with both read-only and

<sup>1</sup> This condition differs from the one presented in [50], but is equivalent.

write-only transactions.<sup>2</sup> The key idea of achieving convergent client views is to share with clients the *total order* of versions on a server, which has already been established by uniquely assigned timestamps across versions. In contrast, Eiger-PORT constructs individual, possibly different, orders per client. Eiger-PORT+’s read-only transactions satisfy the NOC properties and are therefore performance-optimal (PORT). This is achieved in the same way as for Eiger-PORT. In particular, both protocols’ read operations use only a fixed number of timestamps as metadata. We now give a high-level description of both Eiger-PORT and Eiger-PORT+, starting with their commonalities and then highlighting their differences.

### 3.1 Timestamps

Both distributed transaction protocols leverage *timestamp-based* concurrency control. The timestamps are based on Lamport clocks [25], which clients and servers maintain and update with each local or communication event. Whenever a transaction commits a new version, the current clock reading is paired with the transaction’s client ID to generate a globally unique commit timestamp. The lexicographic order of these pairs induces a total order on commit timestamps.

Each server maintains a *local safe time* `lst` that corresponds to the minimum of the uncommitted transactions’ timestamps or, if there is none, the maximum committed timestamp for that server. Each client maintains a variable `lst_map`, which maps server IDs to their latest known `lst` value, and a *global safe time* `gst`. The latter is updated to the minimum timestamp in `lst_map` when a read-only transaction starts and acts as the *stable frontier* for that client: all transactions with earlier timestamps are guaranteed to be committed on all servers. Each read sent to a server includes `gst` as a read timestamp, which is used to safely read a committed version with a timestamp lower than the `gst`, or the client’s latest own write (to achieve RYW), if its timestamp is higher than the `gst`.

### 3.2 Read and Write Transactions

Write transactions are similar in both protocols and follow a variant of the two-phase commit (2PC) protocol that always commits [35]. In the *prepare* phase, each timestamped write is sent to the corresponding partition, which adds the write to its local data store as a pending version. In the *commit* phase, each partition sets the version as committed, along with its commit timestamp. However, the two protocols differ in how they handle read transactions in the absence of an own write newer than `gst`. While Eiger-PORT+ always reads the *latest* version below `gst` in this case, Eiger-PORT searches for the latest version below `gst` that either has no write conflicts or is written by a different client. This backward scan is presumably done to maintain read atomicity (RA). However, we show that this scan is unnecessary for RA, can harm performance, and cause client view divergence. Consequently, in addition to providing convergence, Eiger-PORT+ improves performance by eliminating this scan’s overhead.

<sup>2</sup> Eiger-PORT+’s pseudocode, together with its description, is given in [16, Appx. A].






Fig. 2: Alice and Bob reading (shown by arrows) from servers storing X and Y, illustrating convergence in Eiger-PORT+ and lack thereof in Eiger-PORT. Each square represents a version, and its color determines the version's writer, where orange, blue, and gray correspond to Alice, Bob, and other clients respectively.

*Example 1.* Figure 2 illustrates the difference between Eiger-PORT and Eiger-PORT+ in reading versions. In Figure 2a, Alice (orange) and Bob (blue) have written some versions on servers (partitions storing keys) X and Y, and their  $gst$ s are 0. Alice and Bob each perform a transaction reading from servers X and Y. In both protocols, Alice reads  $\{X_1, Y_3\}$  and Bob  $\{X_2, Y_1\}$  according to RYW.

As new versions are added to the servers, the  $gst$ s advance to higher timestamp values. Assume that versions  $Y_1$  to  $Y_4$  are conflicting writes, i.e., the transactions writing  $Y_2$  to  $Y_4$  had already started when  $Y_1$  was committed. Alice and Bob then again read keys X and Y. Bob behaves the same in both protocols (cf. Figures 2b and 2c): He reads the last committed versions below its  $gst$ ,  $\{X_3, Y_3\}$ , as he has no writes above its  $gst$  and the read versions are written by other clients. However, Alice's behavior differs in the two protocols. Given Alice's new  $gst$ , in Eiger-PORT she reads  $\{X_5, Y_1\}$  (cf. Figure 2b), while in Eiger-PORT+ she reads  $\{X_5, Y_4\}$  (cf. Figure 2c). As  $X_5$  is her newest own write on X, this is the same for both protocols. But Alice has no newer own writes to Y and the latest committed version,  $Y_4$ , is Alice's write and has write conflicts. Thus, Eiger-PORT performs a scan to find the latest version written by a *different* client below Alice's  $gst$ , i.e.,  $Y_1$ , while Eiger-PORT+ reads  $Y_4$ , irrespective of its writer or write conflicts.

Hence, in Eiger-PORT, Alice reads  $Y_3$  before  $Y_1$ , while Bob reads  $Y_1$  before  $Y_3$ , which results in diverging client views. This behaviour is allowed by TCC where convergence is not required. In contrast, in Eiger-PORT+, Bob reads  $Y_1$  before  $Y_3$  and Alice reads  $Y_3$  before  $Y_4$ , both agreeing with the same convergent order, i.e., the versions' total order established on the servers ( $Y_1 < Y_3 < Y_4$ ).



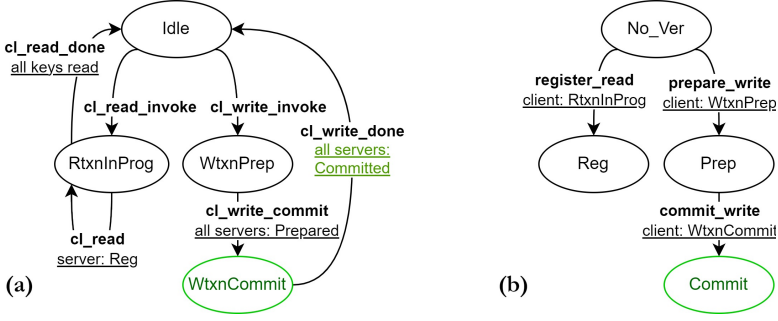


Fig. 3: Eiger-PORT+: state diagrams of (a) a client's `cl_state` and (b) a server's `svr_state` for a given transaction.

## 4 Formal Modeling and Verification

We formally model Eiger-PORT+ (Section 4.1), and use our proof technique (Section 4.2) to verify its TCCv isolation guarantee (Sections 4.3 and 4.4).

### 4.1 Formalizing Eiger-PORT+

To model the protocol, we consider a distributed KVS with one transaction coordinator per client and several servers that handle the clients' transactions. For simplicity, we integrate the coordinator into the client and assume that each server manages one key. We also assume clients execute transactions sequentially.

We formalize Eiger-PORT+ as an LTS. Its states are the protocol's global configurations, consisting of the clients' and the servers' local configurations. As depicted in Figure 3, these local configurations include control states indicating a protocol execution's progress. Each event changes either one client's or one server's configuration and advances its respective control state, ensuring that the clients and servers are independent components with interleaved events. We allow these components to directly access each other's local configurations to exchange information. This is a standard abstraction in protocol modeling, which can later be refined into explicit message-passing communication.

We next define our LTS model's configurations and describe the sequences of events associated with read and write transactions in more detail. We slightly deviate from the Isabelle syntax to stay closer to standard mathematical notation.

**Configurations** We model the client, server, and global configurations as records in Isabelle/HOL (Figure 4). The global configuration contains the client (`cls`) and server (`svrs`) local configurations, whose types are parameterized by the type `'v` of values, and three history variables, which we discuss in Section 4.3.

Besides the global safe time, `gst`, and the `lst_map`, already discussed above, the client configuration consists of its state, `cl_state`, a transaction sequence number, `cl_sn`, and the client's (Lamport) clock, `cl_clock`. The state is described by the type `txn_state`, which has four constructors for idle (`Idle`), read

```

— transaction state
datatype 'v txn_state =
  Idle |
  RtxnInProg ts (key set) (key → 'v) |
  WtxnPrep (key → 'v) |
  WtxnCommit ts (key → 'v)

— version state
datatype 'v ver_state =
  No_Ver |
  Reg |
  Prep ts ts 'v |
  Commit ts ts ts 'v (txid → ts × ts)

— client configuration
record 'v cl_conf =
  cl_state : 'v txn_state
  cl_sn : sqn
  cl_clock : ts
  gst : ts
  lst_map : key → ts

— server configuration
record 'v svr_conf =
  svr_state : txid → 'v ver_state
  svr_clock : ts
  lst : ts

— global configuration
record 'v global_conf =
  cls : cl_id → 'v cl_conf
  svrs : key → 'v svr_conf
  rtxn_rts : txid → ts
  wtxn_cts : txid → ts
  cts_order : key → txid list

```

*— three history variables*

Fig. 4: Eiger-PORT+: the client, server, and global configurations.

transaction in progress (*RtxnInProg*), write transaction prepare (*WtxnPrep*), and write transaction commit (*WtxnCommit*). The latter three states include a key-value map describing the values (to be) read or written.

A server's configuration consists of a function mapping each transaction ID to a version state (*ver\_state*), the server's (Lamport) clock, and its local safe time (*lst*). A version state may either be idle (*No\_Ver*), registered read (*Reg*) for a read transaction, or prepared (*Prep*) or committed (*Commit*) for a write transaction. The latter two states include timestamps in their parameters and the commit state includes a readermap to record information about the transactions reading this version (similar to the abstract model's reader sets).

*Notation.* To improve readability, we sometimes omit the projections *cls* and *svrs*, writing, e.g., (*gst s cl*) for (*gst (cls s cl)*).

**Read-only transactions** proceed as follows (cf. Figure 3). The *cl\_read\_invoke* event of a client starts a read-only transaction and transitions from *Idle* to (*RtxnInProg clk keys ∅*) state, where *clk* is the client's current clock reading, *keys* is the (finite and non-empty) set of keys to be read, and *∅* is the empty key-value mapping, where the subsequently read values will be recorded. As mentioned, this event also updates the client's global safe time, *gst*, to the minimum of the servers' local safe times stored in *lst\_map*. As a result, more up-to-date versions of certain keys may become visible to the client.

Once a client has invoked a read, the involved servers (in *keys*) follow with a *register\_read* event, where they transition from *No\_Ver* to *Reg* state and access the client's *gst* to determine the latest own write newer than the *gst* (if any), or the latest transaction with a commit timestamp  $cts \leq gst$ . This transaction is

```

1 definition cl_write_commit c1 kv_map cts sn u clk s s'  $\longleftrightarrow$ 
2   — guards:
3   cl_state s c1 = WtxnPrep kv_map  $\wedge$ 
4   ( $\forall k \in \text{dom kv\_map. is\_prepared (svr\_state s k (Tn sn c1))} \wedge$ 
5   cts = Max {get_prepared_ts (svr_state s k (Tn sn c1)) |
6     k  $\in$  dom kv_map}  $\wedge$ 
7   sn = cl_sn s c1  $\wedge$ 
8   u = views_of s c1  $\wedge$ 
9   clk = cts + 1  $\wedge$ 
10  — updates (unmentioned variables remain unchanged):
11  cl_state s' c1 = WtxnCommit cts kv_map  $\wedge$ 
12  cl_clock s' c1 = clk  $\wedge$ 
13  wtxn_cts s' = (wtxn_cts s)(Tn sn c1  $\mapsto$  cts)  $\wedge$ 
14  cts_order s' = extend_cts_order s (Tn sn c1) cts kv_map

```

Fig. 5: Eiger-PORT+'s client commit event.

recorded in the read version's readmap along with the current `lst` and updated server clock. This information is then accessed in the client's subsequent `cl_read` event, reading the version's value and updating its own clock and `lst_map`. When the client has read all requested values, i.e.,  $\text{dom kv\_map} = \text{keys}$  holds for its state (`RtxnProg clk keys kv_map`), the event `cl_read_done` brings it back to `Idle`.

**Write-only transactions** are initiated by the `cl_write_invoke` event, in which the client transitions from `Idle` to the state (`WtxnPrep kv_map`). The key-value map `kv_map` describes the keys and associated values to be written and corresponds to the transaction's (write-only) fingerprint (Section 2.3). Once all servers have followed into their prepared state, the client can execute the `cl_write_commit` event, which is defined in Figure 5.

This event has eight parameters: the client ID `c1`, the key-value map `kv_map`, the commit timestamp `cts`, which is the maximum of the involved server's prepared timestamps for the current transaction (`Tn sn c1`) (lines 5-6), the client's current sequence number `sn` (line 7), the abstract view `u` (line 8, see Section 4.3), the updated Lamport clock `clk` (line 9), and the global states `s` and `s'` before and after the event. As  $\text{cts} > \text{cl\_clock}$  always holds here, there is no need to take their maximum to determine `clk`. The guards at lines 3 and 4 require that the client is in the prepared state and that all involved servers have followed into their own prepared state. The client's state is updated to (`WtxnCommit cts kv_map`) (line 11) and the clock is updated (line 12). We will discuss the history variable updates at lines 13-14 in Section 4.3.

After the client's commit event, the involved servers commit the transaction on their side using `commit_write` events. When all servers have committed, the client executes the `cl_write_done` event to return to its idle state.

## 4.2 Proof Technique

We now discuss our protocol verification technique based on refinement, and a technique for commuting independent events required to complement refinement.

The main goal of our verification is to prove the following result, stating that all of Eiger-PORT+'s reachable states are allowed by the abstract model  $\mathcal{I}_{\text{TCCV}}$ :

$$r_{\text{EPP}}(\text{reach EPP}) \subseteq \text{reach } \mathcal{I}_{\text{TCCV}}. \quad (2)$$

Recall from Section 2.2, that this would follow from a proof of  $\text{EPP} \preceq_{r_{\text{EPP}}, \pi_{\text{EPP}}} \mathcal{I}_{\text{TCCV}}$ , for suitable refinement mappings  $r_{\text{EPP}}$  and  $\pi_{\text{EPP}}$  on protocol states and events. However, such a direct proof would fail for the following reason.

Eiger-PORT+ uses timestamps to define an order on versions and to identify “safe-to-read” versions. Hence, to ensure that clients always read the latest version in their view, the refinement mapping must reconstruct the version lists of the abstract KVS in the order of their commit timestamps. Otherwise, the proof of the abstract guard LWW (Section 2.3) will fail. However, the *execution order* of commits and the *order of the associated commit timestamps* may not coincide. We call such commits *inverted commits*. Having inverted commits in an execution may require *inserting* a key's new version to its version list rather than *appending* it. Since the abstract model only ever appends new versions at the end of the version lists, the refinement proof alone would fail for executions with inverted commits.

To address this problem, we introduce an extra proof step to reorder the inverted commits before the refinement. To this end, we define a modified protocol model  $\widehat{\text{EPP}}$  that restricts transaction commits to those that do not introduce any inverted commits. We decompose the proof of (2) into the following two steps:

$$\text{reach EPP} = \text{reach } \widehat{\text{EPP}}, \quad (3)$$

$$r_{\text{EPP}}(\text{reach } \widehat{\text{EPP}}) \subseteq \text{reach } \mathcal{I}_{\text{TCCV}}. \quad (4)$$

We prove (4) by the refinement  $\widehat{\text{EPP}} \preceq_{r_{\text{EPP}}, \pi_{\text{EPP}}} \mathcal{I}_{\text{TCCV}}$ , which works for the restricted model. To prove (3), we use a proof technique based on Lipton's reduction method [26] to successively reorder inverted commits in executions by commuting causally independent events, while preserving the executions' final state.

### 4.3 Refinement Mapping

For the refinement proof, we need to find the refinement mappings  $\pi_{\text{EPP}}$  and  $r_{\text{EPP}}$ . To define  $\pi_{\text{EPP}}$ , we must identify protocol events that refine the abstract commit event and the abstract view extension event. The protocol events `cl_read_done` and `cl_write_commit` refine the abstract commit event since reads and writes are guaranteed to commit after these events are executed. The event `cl_read_invoke` refines the abstract view extension event, as this event updates a client's `gst` and thus extends its view. All other events refine `skip`.

To define  $r_{\text{EPP}}$ , we must reconstruct an abstract configuration from Eiger-PORT+'s protocol configuration  $s$ , i.e.,

$$r_{\text{EPP}} s = (\text{kvs\_of } s, \text{views\_of } s),$$

where the function  $(\text{kvs\_of } s)$  reconstructs the abstract KVS and  $(\text{views\_of } s)$  reconstructs the abstract client views. To help define these components, we add

*history variables* to the global configuration. We now describe these history variables and then the functions `kvs_of` and `views_of`.

**History variables** The global configuration includes three history variables. The variables `rtxn_rts` and `wtxn_cts` serve as shortcuts to respectively map transaction IDs directly to the read timestamp (`gst`) and commit timestamps of the corresponding read-only and write-only transactions. These variables get updated in the corresponding commit events. The variable `cts_order` maps each key to a list of client-committed write-only transactions, *ordered by their commit timestamps*. The client commit event extends `cts_order` by *inserting* the committed transaction’s ID at the position corresponding to its commit timestamp into the transaction ID list of each key written by the transaction (line 13 in Figure 5). This variable is used to facilitate the reconstruction of the abstract KVS.

**Abstract KVS** The function `kvs_of` reconstructs the abstract KVS from the `cts_order` history variable by mapping each key’s list of client-committed transactions to an abstract version list. For a given key `k` and transaction `t` in the list, we extract each version’s value and readerset from server `k`’s (prepared or committed) state for `t`. Since the abstract model always reads the latest versions in a client’s view, as expressed by the guard `LWW` of the abstract commit event (see Section 2.3), the `cts_order` must be sorted by commit timestamp.

**Abstract views** We define the function `views_of`, reconstructing the abstract views from Eiger-PORT+’s model configurations in two steps. We first construct a function `get_view`, where  $(\text{get\_view } s \text{ cl } k)$  denotes the set of client-committed transactions `t`, whose commit timestamp is less than or equal to the client’s `gst`, i.e.,  $(\text{wtxn\_cts } s \text{ t}) \leq (\text{gst } s \text{ cl})$  or that are the client `cl`’s own transactions (for `RYW`). Second, we use `cts_order` to map the transactions IDs in the range of `get_view` to their positions in the `cts_order`, which correspond to indices into the abstract version lists.

#### 4.4 Correctness: Eiger-PORT+ satisfies $\text{TCCv}$

We can now state our main result of this section.

**Theorem 1 (Correctness of Eiger-PORT+).** *The Eiger-PORT+ model EPP satisfies  $\text{TCCv}$ , i.e.,  $r_{\text{EPP}}(\text{reach EPP}) \subseteq \text{reach } \mathcal{I}_{\text{TCCv}}$ .*

As described in Section 4.2, we combine refinement and reduction in this proof. We devote the remainder of this subsection to sketching both parts of our proof, stated as Lemmas 1 and 2 below, followed by describing the invariants used in these proofs.

**Restricted model and reduction proof** We define the restricted model  $\widehat{\text{EPP}}$  by adding a guard to the event `cl_write_commit`, which requires that the unique commit timestamp  $(\text{cts}, \text{cl})$  of the client’s transaction is greater than any commit timestamp of a transaction in the `cts_order`. This ensures that the client commit only appends, but does not insert, the new transaction into the `cts_order`. For this model, we prove the following lemma.

**Lemma 1.**  $\text{reach EPP} = \text{reach } \widehat{\text{EPP}}$ .

*Proof (sketch).* By construction of  $\widehat{\text{EPP}}$ , the inclusion “ $\supseteq$ ” is easily shown by a refinement. For the inclusion “ $\subseteq$ ”, consider any execution  $e$  of EPP ending in some state  $s$ . We prove by reduction that we can reorder all inverted commits in  $e$  of EPP, while preserving its final state  $s$ . We first show that the relevant events of two transactions with inverted commit timestamps are pairwise causally independent. We then prove that adjacent causally independent events in  $e$  can be commuted. Using a measure function on executions, we show that this process terminates in an execution  $\hat{e}$  without inverted commits ending in state  $s$ , thus an execution of  $\widehat{\text{EPP}}$ . Hence, any state reachable in EPP is also reachable in  $\widehat{\text{EPP}}$ .  $\square$

**Refinement proof** Next, we establish a refinement between the restricted model  $\widehat{\text{EPP}}$  and the abstract model  $\mathcal{I}_{\text{TCCv}}$  instantiated to  $\text{TCCv}$ , using the refinement mapping defined in Section 4.3.

**Lemma 2.**  $r_{\text{EPP}}(\text{reach } \widehat{\text{EPP}}) \subseteq \text{reach } \mathcal{I}_{\text{TCCv}}$ .

*Proof (sketch).* We show guard strengthening and update correspondence for every event of  $\widehat{\text{EPP}}$ . This is easy for most events, which refine skip. The interesting cases are the read invoke event, which refines the abstract view extension event, and the client commit and read done events, refining the abstract commit event.

We focus here on the client commit event. The update correspondence proof relies on the absence of inverted commits in  $\widehat{\text{EPP}}$  and thus client commits appending versions to KVS version lists. For guard strengthening, we must show that all guards of the abstract commit event (cf. Section 2.3) are implied by the concrete guards. We discuss the most interesting ones. View atomicity (part of view wellformedness) holds by construction, since all versions of a transaction have the same **cts** and the abstracted view includes all transactions with a **cts** below the client’s **gst** (and also its own writes). Similarly, we prove that **LWW** holds, i.e., a client reads the latest version in its view. To show that **canCommit** holds, we prove an invariant stating that the clients’ views remain closed under  $\text{SO} \cup \text{WR}_{\mathcal{K}}$ . This proof in turn requires several invariants about timestamps.  $\square$

**Invariants and lemmas** Our proofs rely on numerous invariants and lemmas. We present the most important ones categorized as follows.

- **Freshness of transaction IDs:** The clients’ current transaction ID is fresh, i.e., does not occur in the KVS until the commit.
- **Past and future transactions:** stating that the respective client and servers are in particular start states (e.g., **Idle**) or end states (e.g., **Commit**).
- **Views:** These invariants include view wellformedness, view closedness (for  $\text{canCommit}_{\text{TCCv}}$ ), and session guarantees (monotonic reads and RYW).
- **Timestamps:** This category includes lemmas showing the monotonic increase of timestamps and the following invariant for any client  $\text{cl}$  and server  $\text{k}$ :

$$\text{gst } s \text{ cl} < \text{lst\_map } s \text{ cl } k < \text{lst } s \text{ k} < \text{svr\_clock } s \text{ k}.$$

This invariant states that the following timestamps are in a strictly increasing order: client  $c_1$ 's global safe time, client  $c_1$ 's entry for server  $k$  in its map of local safe times, server  $k$ 's local safe time, and  $k$ 's local clock value.

- **Client commit order:** The `cts_order` history variable is sorted by commit timestamps and contains only client-committed (and distinct) transactions.

Note that the first three categories are generic and many of their invariants directly imply related guards needed in the abstract commit event's refinement. The last two categories can easily be adapted to other timestamp-based protocols.

## 5 Deployment and Evaluation

We have implemented and deployed our Eiger-PORT+ protocol on a cluster for a comprehensive performance evaluation [17]. Eiger-PORT+ pushes the limit of the state-of-the-art, with *superior performance* and a *stronger isolation guarantee*.

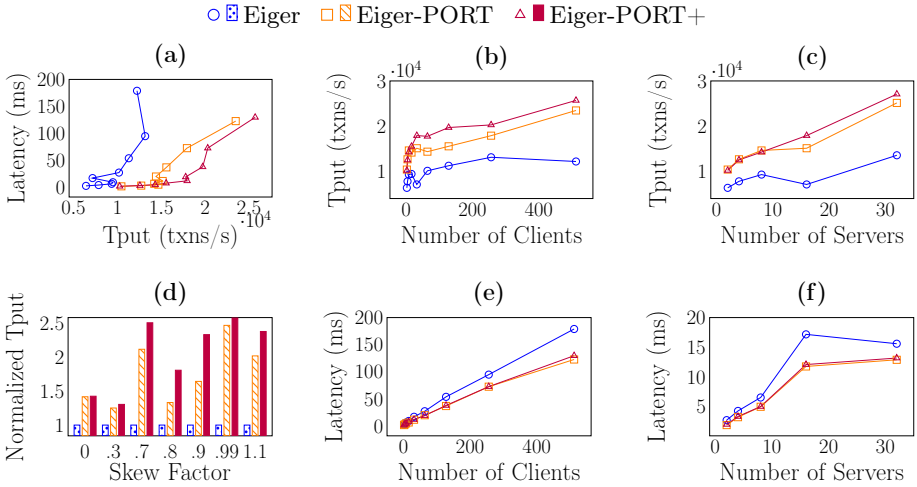


Fig. 6: Performance comparison among the Eiger-family protocols.

**Deployment** We implement Eiger-PORT+, along with Eiger and Eiger-PORT, in the same codebase, each consisting of around 12 kLoC in Java. We use Eiger-PORT's workload generator with default parameters of 32 threads per client, 1 million keys, 90% read proportion, and the Zipfian key-access distribution with a skew factor of 0.8. We deploy these three protocols on a CloudLab [13] cluster of machines, each with 2.4 GHz Quad-Core Xeon CPU and 12 GB RAM. By default, we use eight servers to partition the database and eight client machines to load the servers. We plot each data point using the average over five 60-second trials.

**Evaluation** Overall, Eiger-PORT+ is highly performant and superior to both competing protocols (Figure 6a). In particular, compared to the performance-optimal transaction protocol Eiger-PORT providing TCC, Eiger-PORT+ exhibits



*higher* throughput with a *stronger* isolation guarantee including convergence. Eiger-PORT+ also scales well with an increasing number of clients (Figure 6b) and servers (Figure 6c), with up to 1.8x (resp. 2.5x) throughput improvement over Eiger-PORT (resp. Eiger). In addition, despite varying skews, Eiger-PORT+'s throughput consistently surpasses that of its competitors (Figure 6d). This improvement becomes more pronounced with highly skewed workloads (or larger skew factors). This is because higher skewness results in increased concurrency, which would trigger additional server-side computation in Eiger-PORT and more rounds of communication in Eiger. Notably, despite its higher throughput and stronger isolation guarantee, Eiger-PORT+ demonstrates similar latency to Eiger-PORT (Figures 6e and 6f), which has been proven to be latency-optimal [36].

## 6 Related Work

We compare our work with other efforts on verifying transaction protocols, distinguishing them based on testing, model checking, and deductive verification.

**Testing** Previous work has devised various testers either for specific isolation levels, such as SI [20] and SER [49], or for a range of levels [19,22,24,28]. The underlying techniques are usually based on a characterization of anomalies, e.g., specified using dependency graphs [1] or axioms [8,9]. In contrast to our work, testing can only verify individual protocol executions and therefore can easily miss rarely occurring isolation bugs. On the other hand, while we are verifying a protocol model, testing can be done on the actual implementation.

**Model checking** The Maude model checker has been used to verify the RAMP and LORA protocols for RA [30,27], the Walter protocol for (parallel) SI [32], the ROLA protocol for Update Atomicity [31], and the MegaStore and P-Store protocols for SER [18,42]. Maude has also been used to verify various non-transactional consistency properties of the Cassandra key-value store [33]. Using the TLA+ model checker, the Azure CosmosDB has been verified against several non-transactional consistency properties [5] and TiDB against SI [45]. This model checker has also been used to verify some properties of concurrency control protocols other than isolation guarantees [23,53]. For model checking to be feasible, one must usually impose certain bounds (e.g., on the number of processes and transactions). In contrast, our work provides fully general verification results, which hold for arbitrary protocol executions.

**Deductive verification** Xiong et al. [50] also combine reduction and refinement of their abstract transaction model to prove that the COPS protocol [34] (with read-only transactions but single writes) satisfies TCCv and the Clock-SI protocol [12] satisfies SI. However, while being general, pen-and-paper proofs of such complex protocols are error-prone. Previous work on mechanized deductive verification, to the best of our knowledge, only covers either non-transactional consistency properties or serializability of textbook protocols. Chapar is a framework for verifying causal consistency of non-transactional KVSs. PVS and Event-B have been applied to verify (S)SER of the two-phase locking (2PL) protocol [10,51]. In

contrast, we have designed a new protocol, Eiger-PORT+, which is substantially more complex than 2PL, and we have verified that it satisfies TCCv.

## 7 Conclusion

We have designed Eiger-PORT+, a novel, causally-consistent, database transaction protocol, and formally verified its isolation guarantee of TCCv in Isabelle/HOL. In particular, TCCv was previously conjectured to be incompatible with transactional writes in the presence of performance-optimal read-only transactions. We have formally refuted this conjecture by our protocol design and its verification. Moreover, this case study represents the first complete formal verification of a complex distributed database transaction protocol. Our verification effort, excluding the verification framework, amounts to 10.3k lines of Isabelle/HOL code, composed of 0.7 kLoC for the model and 9.6 kLoC for the proof, which required 108 invariants. In addition, we have conducted a comprehensive evaluation, demonstrating Eiger-PORT+'s superior performance over two state-of-the-art protocols. We believe our protocol is an attractive choice for database applications opting for TCCv.

We see several avenues for future work. First, to facilitate formal protocol modeling and correctness proofs, we will develop an abstract distributed protocol model as an intermediate refinement step. This model will capture structure common to protocols and factor out recurring parts of correctness proofs. Second, we intend to support additional protocol features, for example, open-loop clients, which optimize transactional writes by immediately starting a new transaction, once they commit the previous one. An interesting case study in this context would be Eiger-NOC2 [29], a recent successor of Eiger-PORT+ that improves its performance by using open-loop clients and other features. However, this protocol has not yet been formally verified to provide TCCv. Third, we envision verifying our implementation and connecting it to our protocol verification results, possibly following the Igloo methodology [48].

**Acknowledgements.** We thank the anonymous reviewers for their valuable feedback. This research is supported by an ETH Zurich Career Seed Award and the Swiss National Science Foundation project 200021-231862 “Formal Verification of Isolation Guarantees in Database Systems”.

## References

1. Adya, A.: Weak consistency: a generalized theory and optimistic implementations for distributed transactions. Ph.D. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science (1999)
2. Ahamad, M., Neiger, G., Burns, J.E., Kohli, P., Hutto, P.W.: Causal memory: Definitions, implementation, and programming. *Distributed Comput.* **9**(1), 37–49 (1995)

3. Akkoorath, D.D., Tomsic, A.Z., Bravo, M., Li, Z., Crain, T., Bieniusa, A., Preguiça, N.M., Shapiro, M.: Cure: Strong semantics meets high availability and low latency. In: ICDCS 2016. pp. 405–414. IEEE Computer Society (2016)
4. Attiya, H., Ellen, F., Morrison, A.: Limitations of highly-available eventually-consistent data stores. In: PODC 2015. pp. 385–394. ACM (2015)
5. Azure: azure-cosmos-tla. <https://github.com/Azure/azure-cosmos-tla> (2022)
6. Bailis, P., Fekete, A., Ghodsi, A., Hellerstein, J.M., Stoica, I.: Scalable atomic visibility with ramp transactions. *ACM Transactions on Database Systems (TODS)* **41**(3), 1–45 (2016)
7. Berenson, H., Bernstein, P., Gray, J., Melton, J., O’Neil, E., O’Neil, P.: A critique of ANSI SQL isolation levels. *ACM SIGMOD Record* **24**(2), 1–10 (1995)
8. Biswas, R., Enea, C.: On the complexity of checking transactional consistency. *Proc. ACM Program. Lang.* **3**(OOPSLA), 165:1–165:28 (2019)
9. Cerone, A., Gotsman, A.: Analysing snapshot isolation. *J. ACM* **65**(2), 11:1–11:41 (2018)
10. Chklyae, D., Hooman, J., van der Stok, P.: Mechanical verification of transaction processing systems. In: ICFEM 2000. pp. 89–100. IEEE Computer Society (2000)
11. Didona, D., Guerraoui, R., Wang, J., Zwaenepoel, W.: Causal consistency and latency optimality: Friend or foe? *Proc. VLDB Endow.* **11**(11), 1618–1632 (2018)
12. Du, J., Elnikety, S., Zwaenepoel, W.: Clock-si: Snapshot isolation for partitioned data stores using loosely synchronized clocks. In: SRDS ’13. pp. 173–184. IEEE Computer Society (2013)
13. Duplyakin, D., Ricci, R., Maricq, A., Wong, G., Duerig, J., Eide, E., Stoller, L., Hibler, M., Johnson, D., Webb, K., Akella, A., Wang, K., Ricart, G., Landweber, L., Elliott, C., Zink, M., Cecchet, E., Kar, S., Mishra, P.: The design and operation of CloudLab. In: USENIX ATC’19. pp. 1–14 (Jul 2019)
14. ElectricSQL: <https://electric-sql.com/> (2024)
15. Ghasemirad, S., Liu, S., Sprenger, C., Liu, S., Multazzua, L., Basin, D.: VerIso: Verifiable isolation guarantees for database transactions. *Proc. VLDB Endow.* **18** (2025), To appear.
16. Ghasemirad, S., Sprenger, C., Liu, S., Multazzu, L., Basin, D.: Pushing the limit: Verified performance-optimal causally-consistent database transactions. *CoRR abs/2411.07049* (2025). <https://doi.org/10.48550/ARXIV.2411.07049>
17. Ghasemirad, S., Sprenger, C., Liu, S., Multazzu, L., Basin, D.: Pushing the limit: Verified performance-optimal causally-consistent database transactions (artifacts) (Jan 2025). <https://doi.org/10.5281/zenodo.14622073>
18. Grov, J., Ölveczky, P.C.: Formal modeling and analysis of Google’s Megastore in Real-Time Maude. In: Specification, Algebra, and Software - Essays Dedicated to Kokichi Futatsugi. LNCS, vol. 8373, pp. 494–519. Springer (2014)
19. Gu, L., Liu, S., Xing, T., Wei, H., Chen, Y., Basin, D.: IsoVista: Black-box checking database isolation guarantees. *Proc. VLDB Endow.* **17**(12) (2024)
20. Huang, K., Liu, S., Chen, Z., Wei, H., Basin, D., Li, H., Pan, A.: Efficient black-box checking of snapshot isolation in databases. *Proc. VLDB Endow.* **16**(6), 1264–1276 (2023)
21. Jepsen: Jepsen Analyses (2024), <https://jepsen.io/analyses>
22. Jiang, Z.M., Liu, S., Rigger, M., Su, Z.: Detecting transactional bugs in database engines via graph-based oracle construction. In: OSDI’23. USENIX Association (2023)
23. Katsarakis, A., Ma, Y., Tan, Z., Bainbridge, A., Balkwill, M., Dragojevic, A., Grot, B., Radunovic, B., Zhang, Y.: Zeus: locality-aware distributed transactions. In: EuroSys ’21. pp. 145–161. ACM (2021)

24. Kingsbury, K., Alvaro, P.: Elle: Inferring isolation anomalies from experimental observations. *Proc. VLDB Endow.* **14**(3), 268–280 (2020)
25. Lamport, L.: Time, clocks, and the ordering of events in a distributed system. *Commun. ACM* **21**(7), 558–565 (jul 1978)
26. Lipton, R.J.: Reduction: A method of proving properties of parallel programs. *Commun. ACM* **18**(12), 717–721 (1975)
27. Liu, S.: All in one: Design, verification, and implementation of SNOW-optimal read atomic transactions. *ACM Trans. Softw. Eng. Methodol.* **31**(3) (Mar 2022)
28. Liu, S., Gu, L., Wei, H., Basin, D.: Plume: Efficient and complete black-box checking of weak isolation levels. *Proc. ACM Program. Lang.* **8**(OOPSLA2) (2024)
29. Liu, S., Multazzu, L., Wei, H., Basin, D.: NOC-NOC: Towards Performance-optimal Distributed Transactions. *Proc. ACM Manag. Data* **2**(1) (mar 2024)
30. Liu, S., Ölveczky, P.C., Rahman, M.R., Ganhotra, J., Gupta, I., Meseguer, J.: Formal modeling and analysis of RAMP transaction systems. In: *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, 2016. ACM (2016)
31. Liu, S., Ölveczky, P.C., Wang, Q., Gupta, I., Meseguer, J.: Read atomic transactions with prevention of lost updates: ROLA and its formal analysis. *Formal Aspects Comput.* **31**(5), 503–540 (2019)
32. Liu, S., Ölveczky, P.C., Wang, Q., Meseguer, J.: Formal modeling and analysis of the Walter transactional data store. In: *WRLA '18. LNCS*, vol. 11152, pp. 136–152. Springer (2018)
33. Liu, S., Rahman, M.R., Skeirik, S., Gupta, I., Meseguer, J.: Formal modeling and analysis of Cassandra in Maude. In: *ICFEM '14. LNCS*, vol. 8829, pp. 332–347. Springer (2014)
34. Lloyd, W., Freedman, M.J., Kaminsky, M., Andersen, D.G.: Don't settle for eventual: scalable causal consistency for wide-area storage with COPS. In: *SOSP 2011*. pp. 401–416. ACM (2011)
35. Lloyd, W., Freedman, M.J., Kaminsky, M., Andersen, D.G.: Stronger semantics for low-latency geo-replicated storage. In: *NSDI 2013*. pp. 313–328. USENIX Association (2013)
36. Lu, H., Hodsdon, C., Ngo, K., Mu, S., Lloyd, W.: The SNOW theorem and latency-optimal read-only transactions. In: *OSDI 2016*. pp. 135–150. USENIX Association (2016)
37. Lu, H., Mu, S., Sen, S., Lloyd, W.: NCC: Natural concurrency control for strictly serializable datastores by avoiding the Timestamp-Inversion pitfall. In: *OSDI '23*. pp. 305–323. USENIX Association (2023)
38. Lu, H., Sen, S., Lloyd, W.: Performance-optimal read-only transactions. In: *OSDI 2020*. pp. 333–349. USENIX Association (2020)
39. Mehdi, S.A., Little, C., Crooks, N., Alvisi, L., Bronson, N., Lloyd, W.: I can't believe it's not causal! scalable causal consistency with no slowdown cascades. In: *NSDI 2017*. pp. 453–468. USENIX Association (2017)
40. Microsoft: Azure CosmosDB DB. <https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels> (2024)
41. Neo4j: <https://neo4j.com/> (2024)
42. Ölveczky, P.C.: Formalizing and validating the P-Store replicated data store in Maude. In: *WADT '16. LNCS*, vol. 10644, pp. 189–207. Springer (2016)
43. Papadimitriou, C.H.: The serializability of concurrent database updates. *Journal of the ACM (JACM)* **26**(4), 631–653 (1979)
44. Perrin, M., Mostefaoui, A., Jard, C.: Causal consistency: Beyond memory. *SIGPLAN Not.* **51**(8), 26:1–26:12 (Feb 2016)

45. PingCAP: tla-plus. <https://github.com/pingcap/tla-plus> (2022)
46. Silberschatz, A., Korth, H.F., Sudarshan, S.: Database System Concepts, Seventh Edition. McGraw-Hill Book Company (2020), <https://www.db-book.com/>
47. Spirovska, K., Didona, D., Zwaenepoel, W.: Optimistic causal consistency for geo-replicated key-value stores. *IEEE Trans. Parallel Distributed Syst.* **32**(3), 527–542 (2021)
48. Sprenger, C., Klenze, T., Eilers, M., Wolf, F., Müller, P., Clochard, M., Basin, D.: Igloo: Soundly linking compositional refinement and separation logic for distributed systems verification. In: *ACM Program. Lang.* 4, OOPSLA, Article 152 (2020)
49. Tan, C., Zhao, C., Mu, S., Walfish, M.: Cobra: Making transactional key-value stores verifiably serializable. In: *14th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2020*. pp. 63–80. USENIX Association (2020)
50. Xiong, S., Cerone, A., Raad, A., Gardner, P.: Data consistency in transactional storage systems: A centralised semantics. In: *34th European Conference on Object-Oriented Programming, ECOOP 2020. LIPIcs*, vol. 166, pp. 21:1–21:31. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
51. Yadav, D., Butler, M.J.: Rigorous design of fault-tolerant transactions for replicated database systems using Event B. In: *Rigorous Development of Complex Fault-Tolerant Systems [FP6 IST-511599 RODIN project]*. LNCS, vol. 4157, pp. 343–363. Springer (2006)
52. Yang, J., Yue, Y., Rashmi, K.V.: A large-scale analysis of hundreds of in-memory key-value cache clusters at Twitter. *ACM Trans. Storage* **17**(3), 17:1–17:35 (2021)
53. Zhang, I., Sharma, N.K., Szekeres, A., Krishnamurthy, A., Ports, D.R.K.: Building consistent transactions with inconsistent replication. In: *SOSP '15*. p. 263–278. ACM (2015)

**Open Access.** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution, and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

