Building Consistent Transactions with Inconsistent Replication

Irene Zhang Naveen Kr. Sharma Adriana Szekeres
Arvind Krishnamurthy Dan R. K. Ports

University of Washington
{iyzhang, naveenks, aaasz, arvind, drkp}@cs.washington.edu

Abstract

Application programmers increasingly prefer distributed storage systems with strong consistency and distributed transactions (e.g., Google’s Spanner) for their strong guarantees and ease of use. Unfortunately, existing transactional storage systems are expensive to use – in part because they require costly replication protocols, like Paxos, for fault tolerance. In this paper, we present a new approach that makes transactional storage systems more affordable: we eliminate consistency from the replication protocol while still providing distributed transactions with strong consistency to applications.

We present TAPIR – the Transactional Application Protocol for Inconsistent Replication – the first transaction protocol to use a novel replication protocol, called inconsistent replication, that provides fault tolerance without consistency. By enforcing strong consistency only in the transaction protocol, TAPIR can commit transactions in a single round-trip and order distributed transactions without centralized coordination. We demonstrate the use of TAPIR in a transactional key-value store, TAPIR-KV. Compared to conventional systems, TAPIR-KV provides better latency and throughput.

1. Introduction

Distributed storage systems provide fault tolerance and availability for large-scale web applications. Increasingly, application programmers prefer systems that support distributed transactions with strong consistency to help them manage application complexity and concurrency in a distributed environment. Several recent systems [4, 11, 17, 22] reflect this

trend, notably Google’s Spanner system [13], which guarantees linearizable transaction ordering.¹

For application programmers, distributed transactional storage with strong consistency comes at a price. These systems commonly use replication for fault-tolerance, and replication protocols with strong consistency, like Paxos, impose a high performance cost, while more efficient, weak consistency protocols fail to provide strong system guarantees.

Significant prior work has addressed improving the performance of transactional storage systems – including systems that optimize for read-only transactions [4, 13], more restrictive transaction models [2, 14, 22], or weaker consistency guarantees [3, 33, 42]. However, none of these systems have addressed both latency *and* throughput for general-purpose, replicated, read-write transactions with strong consistency.

In this paper, we use a new approach to reduce the cost of replicated, read-write transactions and make transactional storage more affordable for programmers. Our key insight is that existing transactional storage systems waste work and performance by incorporating a distributed transaction protocol and a replication protocol that *both* enforce strong consistency. Instead, we show that it is possible to provide distributed transactions with better performance and the same transaction and consistency model using replication with *no* consistency.

To demonstrate our approach, we designed *TAPIR* – the Transactional Application Protocol for Inconsistent Replication. TAPIR uses a new replication technique, called *inconsistent replication* (IR), that provides fault tolerance without consistency. Rather than an ordered operation log, IR presents an *unordered operation set* to applications. Successful operations execute at a majority of the replicas and survive failures, but replicas can execute them in any order. Thus, IR needs no cross-replica coordination or designated leader for operation processing. However, unlike eventual consistency, IR allows applications to enforce higher-level invariants when needed.

Thus, despite IR’s weak consistency guarantees, TAPIR provides *linearizable read-write transactions* and supports globally-consistent reads across the database at a timestamp –

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SOSP’15, October 4–7, 2015, Monterey, CA.

Copyright is held by the owner/author(s).

ACM 978-1-4503-3834-9/15/10.

<http://dx.doi.org/10.1145/2815400.2815404>

¹ Spanner’s linearizable transaction ordering is also referred to as strict serializable isolation or external consistency.

Figure 1: A common architecture for distributed transactional storage systems today. The distributed transaction protocol consists of an atomic commitment protocol, commonly Two-Phase Commit (2PC), and a concurrency control (CC) mechanism. This runs atop a replication (R) protocol, like Paxos.

the same guarantees as Spanner. TAPIR efficiently leverages IR to distribute read-write transactions in a *single round-trip* and order transactions globally across partitions and replicas *with no centralized coordination*.

We implemented TAPIR in a new distributed transactional key-value store called TAPIR-KV, which supports linearizable transactions over a partitioned set of keys. Our experiments found that TAPIR-KV had: (1) 50% lower commit latency and (2) more than $3\times$ better throughput compared to systems using conventional transaction protocols, including an implementation of Spanner’s transaction protocol, and (3) comparable performance to MongoDB [35] and Redis [39], widely-used eventual consistency systems.

This paper makes the following contributions to the design of distributed, replicated transaction systems:

- We define *inconsistent replication*, a new replication technique that provides fault tolerance without consistency.
- We design *TAPIR*, a new distributed transaction protocol that provides strict serializable transactions using inconsistent replication for fault tolerance.
- We build and evaluate TAPIR-KV, a key-value store that combines inconsistent replication and TAPIR to achieve high-performance transactional storage.

2. Over-Coordination in Transaction Systems

Replication protocols have become an important component in distributed storage systems. Modern storage systems commonly partition data into *shards* for scalability and then replicate each shard for fault-tolerance and availability [4, 9, 13, 34]. To support transactions with strong consistency, they must implement both a *distributed transaction protocol* – to ensure atomicity and consistency for transactions across shards – and a *replication protocol* – to ensure transactions are not lost (provided that no more than half of the replicas in each shard fail at once). As shown in Figure 1, these systems typically place the transaction protocol, which combines an atomic commitment protocol and a concurrency control mechanism, on top of the replication protocol (although alternative architectures have also occasionally been proposed [34]).

Distributed transaction protocols assume the availability of an *ordered, fault-tolerant log*. This ordered log abstraction

Figure 2: Example read-write transaction using two-phase commit, viewstamped replication and strict two-phase locking. Availability zones represent either a cluster, datacenter or geographic region. Each *shard* holds a partition of the data stored in the system and has replicas in each zone for fault tolerance. The red, dashed lines represent redundant coordination in the replication layer.

is easily and efficiently implemented with a spinning disk but becomes more complicated and expensive with replication. To enforce the serial ordering of log operations, transactional storage systems must integrate a costly replication protocol with strong consistency (e.g., Paxos [27], Viewstamped Replication [37] or virtual synchrony [7]) rather than a more efficient, weak consistency protocol [24, 40].

The traditional log abstraction imposes a serious performance penalty on replicated transactional storage systems, because it enforces strict serial ordering using expensive distributed coordination *in two places*: the replication protocol enforces a serial ordering of operations across replicas in each shard, while the distributed transaction protocol enforces a serial ordering of transactions across shards. This redundancy impairs latency and throughput for systems that integrate both protocols. The replication protocol must coordinate across replicas on every operation to enforce strong consistency; as a result, it takes *at least two round-trips* to order any read-write transaction. Further, to efficiently order operations, these protocols typically rely on a replica leader, which can introduce a throughput bottleneck to the system.

As an example, Figure 2 shows the redundant coordination required for a single read-write transaction in a system like Spanner. Within the transaction, Read operations go to the shard leaders (which may be in other datacenters), because the operations must be ordered across replicas, even though they are not replicated. To Prepare a transaction for commit, the transaction protocol must coordinate transaction ordering across shards, and then the replication protocol coordinates

the Prepare operation ordering across replicas. As a result, it takes at least two round-trips to commit the transaction.

In the TAPIR and IR design, we eliminate the redundancy of strict serial ordering over the two layers and its associated performance costs. IR is the first replication protocol to provide *pure fault tolerance* without consistency. Instead of an ordered operation log, IR presents the abstraction of an *unordered operation set*. Existing transaction protocols cannot efficiently use IR, so TAPIR is the first transaction protocol designed to provide linearizable transactions on IR.

3. Inconsistent Replication

Inconsistent replication (IR) is an efficient replication protocol designed to be used with a higher-level protocol, like a distributed transaction protocol. IR provides fault-tolerance without enforcing any consistency guarantees of its own. Instead, it allows the higher-level protocol, which we refer to as the *application protocol*, to decide the outcome of conflicting operations and recover those decisions through IR's fault-tolerant, unordered operation set.

3.1 IR Overview

Application protocols invoke operations through IR in one of two modes:

- **inconsistent** – operations can execute in any order. Successful operations persist across failures.
- **consensus** – operations execute in any order, but return a single *consensus result*. Successful operations and their consensus results persist across failures.

inconsistent operations are similar to operations in weak consistency replication protocols: they can execute in different orders at each replica, and the application protocol must resolve conflicts afterwards. In contrast, **consensus** operations allow the application protocol to *decide* the outcome of conflicts (by executing a *decide* function specified by the application protocol) and recover that decision afterwards by ensuring that the chosen result persists across failures as the consensus result. In this way, **consensus** operations can serve as the basic building block for the higher-level guarantees of application protocols. For example, a distributed transaction protocol can decide which of two conflicting transactions will commit, and IR will ensure that decision persists across failures.

3.1.1 IR Application Protocol Interface

Figure 3 summarizes the IR interfaces at clients and replicas. Application protocols invoke operations through a client-side IR library using *InvokeInconsistent* and *InvokeConsensus*, and then IR runs operations using the *ExecInconsistent* and *ExecConsensus* upcalls at the replicas.

If replicas return conflicting/non-matching results for a **consensus** operation, IR allows the application protocol to decide the operation's outcome by invoking the *decide* function – passed in by the application protocol to *InvokeConsensus* – in the client-side library. The *decide* function takes the list of

Figure 3: Summary of IR interfaces and client/replica state.

Figure 4: IR Call Flow.

returned results (the candidate results) and returns a single result, which IR ensures will persist as the *consensus result*. The application protocol can later recover the consensus result to find out its decision to conflicting operations.

Some replicas may miss operations or need to reconcile their state if the consensus result chosen by the application protocol does not match their result. To ensure that IR replicas eventually converge, they periodically *synchronize*. Similar to eventual consistency, IR relies on the application protocol to reconcile inconsistent replicas. On synchronization, a single IR node first upcalls into the application protocol with *Merge*, which takes records from inconsistent replicas and merges them into a *master record* of successful operations and consensus results. Then, IR upcalls into the application protocol with *Sync* at each replica. *Sync* takes the *master record* and reconciles application protocol state to make the replica consistent with the chosen consensus results.

3.1.2 IR Guarantees

We define a *successful operation* to be one that returns to the application protocol. The *operation set* of any IR group includes all successful operations. We define an operation *X* as being *visible* to an operation *Y* if one of the replicas executing *Y* has previously executed *X*. IR ensures the following properties for the operation set:

P1. [Fault tolerance] At any time, every operation in the operation set is in the record of at least one replica in any quorum of $f + 1$ non-failed replicas.

P2. [Visibility] For any two operations in the operation set, at least one is visible to the other.

P3. [Consensus results] At any time, the result returned by a successful **consensus** operation is in the record of at least one replica in any quorum. The only exception is if the consensus result has been explicitly modified by the application protocol through **Merge**, after which the outcome of **Merge** will be recorded instead.

IR ensures guarantees are met for up to f simultaneous failures out of $2f + 1$ replicas² and any number of client failures. Replicas must fail by crashing, without Byzantine behavior. We assume an asynchronous network where messages can be lost or delivered out of order. IR *does not* require synchronous disk writes, ensuring guarantees are maintained even if clients or replicas lose state on failure. IR makes progress (operations will eventually become successful) provided that messages that are repeatedly resent are eventually delivered before the recipients time out.

3.1.3 Application Protocol Example: Fault-Tolerant Lock Server

As an example, we show how to build a simple lock server using IR. The lock server's guarantee is mutual exclusion: a lock cannot be held by two clients at once. We replicate **Lock** as a **consensus** operation and **Unlock** as an **inconsistent** operation. A client application acquires the lock only if **Lock** successfully returns OK as a consensus result.

Since operations can run in any order at the replicas, clients use unique ids (e.g., a tuple of client id and a sequence number) to identify corresponding **Lock** and **Unlock** operations and only call **Unlock** if **Lock** first succeeds. Replicas will therefore be able to later match up **Lock** and **Unlock** operations, regardless of order, and determine the lock's status.

Note that **inconsistent** operations *are not commutative* because they can have side-effects that affect the outcome of **consensus** operations. If an **Unlock** and **Lock** execute in different orders at different replicas, some replicas might have the lock free, while others might not. If replicas return different results to **Lock**, IR invokes the lock server's **decide** function, which returns OK if $f + 1$ replicas returned OK and NO otherwise. IR only invokes **Merge** and **Sync** on recovery, so we defer their discussion until Section 3.2.2.

IR's guarantees ensure correctness for our lock server. P1 ensures that held locks are persistent: a **Lock** operation persists at one or more replicas in any quorum. P2 ensures mutual exclusion: for any two conflicting **Lock** operations, one is visible to the other in any quorum; therefore, IR will never receive $f + 1$ matching OK results, precluding the **decide** function from returning OK. P3 ensures that once the

client application receives OK from a **Lock**, the result will not change. The lock server's **Merge** function will not change it, as we will show later, and IR ensures that the OK will persist in the record of at least one replica out of any quorum.

3.2 IR Protocol

Figure 3 shows the IR state at the clients and replicas. Each IR client keeps an *operation counter*, which, combined with the *client id*, uniquely identifies operations. Each replica keeps an unordered *record* of executed operations and results for **consensus** operations. Replicas add **inconsistent** operations to their record as **TENTATIVE** and then mark them as **FINALIZED** once they execute. **consensus** operations are first marked **TENTATIVE** with the result of locally executing the operation, then **FINALIZED** once the record has the consensus result.

IR uses four sub-protocols – *operation processing*, *replica recovery/synchronization*, *client recovery*, and *group membership change*. Due to space constraints, we describe only the first two here; the third is described in [46] and the last is identical to that of Viewstamped Replication [32].

3.2.1 Operation Processing

We begin by describing IR's normal-case **inconsistent** operation processing protocol without failures:

1. The client sends $\langle \text{PROPOSE}, id, op \rangle$ to all replicas, where id is the operation id and op is the operation.
 2. Each replica writes id and op to its record as **TENTATIVE**, then responds to the client with $\langle \text{REPLY}, id \rangle$.
 3. Once the client receives $f + 1$ responses from replicas (retrying if necessary), it returns to the application protocol and asynchronously sends $\langle \text{FINALIZE}, id \rangle$ to all replicas. (**FINALIZE** can also be piggy-backed on the client's next message.)
 4. On **FINALIZE**, replicas upcall into the application protocol with **ExecInconsistent**(op) and mark op as **FINALIZED**.
- Due to the lack of consistency, IR can successfully complete an **inconsistent** operation with a *single round-trip* to $f + 1$ replicas and no coordination across replicas.

Next, we describe the normal-case **consensus** operation processing protocol, which has both a *fast path* and a *slow path*. IR uses the fast path when it can achieve a *fast quorum* of $\lceil \frac{3}{2}f \rceil + 1$ replicas that *return matching results* to the operation. Similar to Fast Paxos and Speculative Paxos [38], IR requires a fast quorum to ensure that a majority of the replicas in any quorum agrees to the consensus result. This quorum size is necessary to execute operations in a single round trip when using a replica group of size $2f + 1$ [29]; an alternative would be to use quorums of size $2f + 1$ in a system with $3f + 1$ replicas.

When IR cannot achieve a fast quorum, either because replicas did not return enough matching results (e.g., if there are conflicting concurrent operations) or because not enough replicas responded (e.g., if more than $\frac{f}{2}$ are down), then it must take the slow path. We describe both below:

² Using more than $2f + 1$ replicas for f failures is possible but illogical because it requires a larger quorum size with no additional benefit.

1. The client sends $\langle \text{PROPOSE}, id, op \rangle$ to all replicas.
2. Each replica calls into the application protocol with $\text{ExecConsensus}(op)$ and writes id, op , and $result$ to its record as **TENTATIVE**. The replica responds to the client with $\langle \text{REPLY}, id, result \rangle$.
3. If the client receives at least $\lceil \frac{3}{2}f \rceil + 1$ matching $results$ (within a timeout), then it takes the *fast path*: the client returns $result$ to the application protocol and asynchronously sends $\langle \text{FINALIZE}, id, result \rangle$ to all replicas.
4. Otherwise, the client takes the *slow path*: once it receives $f + 1$ responses (retrying if necessary), then it sends $\langle \text{FINALIZE}, id, result \rangle$ to all replicas, where $result$ is obtained from executing the *decide* function.
5. On receiving **FINALIZE**, each replica marks the operation as **FINALIZED**, updating its record if the received $result$ is different, and sends $\langle \text{CONFIRM}, id \rangle$ to the client.
6. On the slow path, the client returns $result$ to the application protocol once it has received $f + 1$ **CONFIRM** responses.

The fast path for **consensus** operations takes a single round trip to $\lceil \frac{3}{2}f \rceil + 1$ replicas, while the slow path requires two round-trips to at least $f + 1$ replicas. Note that IR replicas can execute operations in different orders and *still* return matching responses, so IR can use the fast path without a strict serial ordering of operations across replicas. IR can also run the fast path and slow path in parallel as an optimization.

3.2.2 Replica Recovery and Synchronization

IR uses a single protocol for recovering failed replicas and running periodic synchronizations. On recovery, we must ensure that the failed replica applies all operations it may have lost or missed in the operation set, so we use the same protocol to periodically bring all replicas up-to-date.

To handle recovery and synchronization, we introduce *view changes* into the IR protocol, similar to Viewstamped Replication (VR) [37]. These maintain IR's correctness guarantees across failures. Each IR view change is run by a leader; leaders coordinate only view changes, *not* operation processing. During a view change, the leader has just one task: to make at least $f + 1$ replicas up-to-date (i.e., they have applied all operations in the operation set) and consistent with each other (i.e., they have applied the same consensus results). IR view changes require a leader because polling inconsistent replicas can lead to conflicting sets of operations and consensus results. Thus, the leader must decide on a *master record* that replicas can then use to synchronize with each other.

To support view changes, each IR replica maintains a current *view*, which consists of the identity of the leader, a list of the replicas in the group, and a (monotonically increasing) *view number* uniquely identifying the view. Each IR replica can be in one of the three states: **NORMAL**, **VIEW-CHANGING** or **RECOVERING**. Replicas process operations only in the **NORMAL** state. We make four additions to IR's operation processing protocol:

1. IR replicas send their current view number in every response to clients. For an operation to be considered successful, the IR client must receive responses with matching view numbers. For **consensus** operations, the view numbers in **REPLY** and **CONFIRM** must match as well. If a client receives responses with different view numbers, it notifies the replicas in the older view.
2. On receiving a message with a view number that is higher than its current view, a replica moves to the **VIEW-CHANGING** state and requests the master record from any replica in the higher view. It replaces its own record with the master record and upcalls into the application protocol with **Sync** before returning to **NORMAL** state.
3. On **PROPOSE**, each replica first checks whether the operation was already **FINALIZED** by a view change. If so, the replica responds with $\langle \text{REPLY}, id, \text{FINALIZED}, v, [result] \rangle$, where v is the replica's current view number and $result$ is the consensus result for **consensus** operations.
4. If the client receives **REPLY** with a **FINALIZED** status for **consensus** operations, it sends $\langle \text{FINALIZE}, id, result \rangle$ with the received $result$ and waits until it receives $f + 1$ **CONFIRM** responses in the same view before returning $result$ to the application protocol.

We sketch the view change protocol here; the full description is available in our TR [46]. The protocol is identical to VR, except that the leader must *merge* records from the latest view, rather than simply taking the longest log from the latest view, to preserve all the guarantees stated in 3.1.2. During synchronization, IR finalizes all **TENTATIVE** operations, relying on the application protocol to decide any consensus results.

Once the leader has received $f + 1$ records, it merges the records from replicas in the latest view into a master record, R , using $\text{IR-MERGE-RECORDS}(records)$ (see Figure 5), where $records$ is the set of received records from replicas in the highest view. IR-MERGE-RECORDS starts by adding all **inconsistent** operations and **consensus** operations marked **FINALIZED** to R and calling **Sync** into the application protocol. These operations must persist in the next view, so we first apply them to the leader, ensuring that they are visible to any operations for which the leader will decide consensus results next in **Merge**. As an example, **Sync** for the lock server matches up all corresponding **Lock** and **Unlock** by id ; if there are unmatched **Locks**, it sets $locked = \text{TRUE}$; otherwise, $locked = \text{FALSE}$.

IR asks the application protocol to decide the consensus result for the remaining **TENTATIVE consensus** operations, which either: (1) have a matching result, which we define as the *majority result*, in at least $\lceil \frac{f}{2} \rceil + 1$ records or (2) do not. IR places these operations in d and u , respectively, and calls $\text{Merge}(d, u)$ into the application protocol, which must return a consensus result for every operation in d and u .

IR must rely on the application protocol to decide consensus results for several reasons. For operations in d , IR cannot tell whether the operation succeeded with the majority result

```

IR-MERGE-RECORDS(records)
1  R, d, u = ∅
2  for ∀op ∈ records
3    if op.type == inconsistent
4      R = R ∪ op
5    elseif op.type == consensus and op.status == FINALIZED
6      R = R ∪ op
7    elseif op.type == consensus and op.status == TENTATIVE
8      if op.result in more than  $\frac{f}{2} + 1$  records
9        d = d ∪ op
10     else
11       u = u ∪ op
12  Sync(R)
13  return R ∪ Merge(d, u)

```

Figure 5: Merge function for the master record. We merge all records from replicas in the latest view, which is always a strict super set of the records from replicas in lower views.

on the fast path, or whether it took the slow path and the application protocol *decide*'d a different result that was later lost. In some cases, it is not safe for IR to keep the majority result because it would violate application protocol invariants. For example, in the lock server, OK could be the majority result if only $\lceil \frac{f}{2} \rceil + 1$ replicas replied OK, but the other replicas might have accepted a conflicting lock request. However, it is also possible that the other replicas *did* respond OK, in which case OK would have been a successful response on the fast-path.

The need to resolve this ambiguity is the reason for the caveat in IR's consensus property (P3) that consensus results can be changed in Merge. Fortunately, the application protocol can ensure that successful consensus results *do not change* in Merge, simply by maintaining the majority results in *d* on Merge *unless they violate invariants*. The merge function for the lock server, therefore, does not change a majority response of OK, *unless* another client holds the lock. In that case, the operation in *d* could not have returned a successful consensus result to the client (either through the fast or the slow path), so it is safe to change its result.

For operations in *u*, IR needs to invoke *decide* but cannot without at least $f + 1$ results, so uses Merge instead. The application protocol can decide consensus results in Merge without $f + 1$ replica results and still preserve IR's visibility property because IR has already applied all of the operations in *R* and *d*, which are the only operations definitely in the operation set, at this point.

The leader adds all operations returned from Merge and their consensus results to *R*, then sends *R* to the other replicas, which call Sync(*R*) into the application protocol and *replace their own records with R*. The view change is complete after at least $f + 1$ replicas have exchanged and merged records and SYNC'd with the master record. A replica can only process requests in the new view (in the NORMAL state) *after* it completes the view change protocol. At this point, any recovering replicas can also be considered recovered. If the leader of the view change does not finish the view change by some timeout, the group will elect a new leader to complete

the protocol by starting a new view change with a larger view number.

3.3 Correctness

We give a brief sketch of correctness for IR, showing why it satisfies the three properties above. For a more complete proof and a TLA+ specification [26], see our technical report [46].

We first show that all three properties hold in the absence of failures and synchronization. Define the set of *persistent operations* to be those operations in the record of at least one replica in any quorum of $f + 1$ non-failed replicas. P1 states that every operation that succeeded is persistent; this is true by quorum intersection because an operation only succeeds once responses are received from $f + 1$ of $2f + 1$ replicas.

For P2, consider any two successful **consensus** operations *X* and *Y*. Each received candidate results from a quorum of $f + 1$ replicas that executed the request. By quorum intersection, there must be one replica that executed both *X* and *Y*; assume without loss of generality that it executed *X* first. Then its candidate result for *Y* reflects the effects of *X*, and *X* is visible to *Y*.

Finally, for P3, a successful **consensus** operation result was obtained either on the fast path, in which case a fast quorum of replicas has the same result marked as TENTATIVE in their records, or on the slow path, where $f + 1$ replicas have the result marked as FINALIZED in their record. In both cases, at least one replica in every quorum will have that result.

Since replicas can lose records on failures and must perform periodic synchronizations, we must also prove that the synchronization/recovery protocol maintains these properties. On synchronization or recovery, the protocol ensures that all persistent operations from the previous view are persistent in the new view. The leader of the new view merges the record of replicas in the highest view from $f + 1$ responses. Any persistent operation appears in one of these records, and the merge procedure ensures it is retained in the master record. It is sufficient to only merge records from replicas in the highest view because, similar to VR, IR's view change protocol ensures that no replicas that participated in the view change will move to a lower view, thus no operations will become persistent in lower views.

The view change completes by synchronizing at least $f + 1$ replicas with the master record, ensuring P1 continues to hold. This also maintains property P3: a **consensus** operation that took the slow path will appear as FINALIZED in at least one record, and one that took the fast path will appear as TENTATIVE in at least $\lceil \frac{f}{2} \rceil + 1$ records. IR's merge procedure ensures that the results of these operations are maintained, unless the application protocol chooses to change them in its Merge function. Once the operation is FINALIZED at $f + 1$ replicas (by a successful view-change or by a client), the protocol ensures that the consensus result won't be changed and will persist: the operation will always appear as FINALIZED when constructing any subsequent master records and thus

the merging procedure will always Sync it. Finally, P2 is maintained during the view change because the leader first calls Sync with all previously successful operations before calling Merge, thus the previously successful operations will be visible to any operations that the leader finalizes during the view change. The view change then ensures that any finalized operation will continue to be finalized in the record of at least $f + 1$ replicas, and thus be visible to subsequent successful operations.

4. Building Atop IR

IR obtains performance benefits because it offers weak consistency guarantees and relies on application protocols to resolve inconsistencies, similar to eventual consistency protocols such as Dynamo [15] and Bayou [43]. However, unlike eventual consistency systems, which expect applications to resolve conflicts *after they happen*, IR allows application protocols to *prevent conflicts before they happen*. Using **consensus** operations, application protocols can enforce higher-level guarantees (e.g., TAPIR’s linearizable transaction ordering) across replicas despite IR’s weak consistency.

However, building strong guarantees on IR requires careful application protocol design. IR cannot support certain application protocol invariants. Moreover, if misapplied, IR can even provide applications with *worse* performance than a strongly consistent replication protocol. In this section, we discuss the properties that application protocols need to have to correctly and efficiently enforce higher-level guarantees with IR and TAPIR’s techniques for efficiently providing linearizable transactions.

4.1 IR Application Protocol Requirement: *Invariant checks must be performed pairwise.*

Application protocols can enforce certain types of invariants with IR, but not others. IR guarantees that in any pair of **consensus** operations, at least one will be visible to the other (P2). Thus, IR readily supports invariants that can be safely checked by examining *pairs* of operations for conflicts. For example, our lock server example can enforce mutual exclusion. However, application protocols cannot check invariants that require the entire history, because each IR replica may have an incomplete history of operations. For example, tracking bank account balances and allowing withdrawals only if the balance remains positive is problematic because the invariant check must consider the entire history of deposits and withdrawals.

Despite this seemingly restrictive limitation, application protocols can still use IR to enforce useful invariants, including lock-based concurrency control, like Strict Two-Phase Locking (S2PL). As a result, distributed transaction protocols like Spanner [13] or Replicated Commit [34] would work with IR. IR can also support optimistic concurrency control (OCC) [23] because OCC checks are pairwise as well: each committing transaction is checked against every previously

committed transaction, so **consensus** operations suffice to ensure that *at least one replica sees any conflicting transaction* and aborts the transaction being checked.

4.2 IR Application Protocol Requirement: *Application protocols must be able to change consensus operation results.*

Inconsistent replicas could execute **consensus** operations with one result and later find the group agreed to a different consensus result. For example, if the group in our lock server agrees to reject a Lock operation that one replica accepted, the replica must later free the lock, and vice versa. As noted above, the group as a whole continues to enforce mutual exclusion, so these temporary inconsistencies are tolerable and are always resolved by the end of synchronization.

In TAPIR, we take the same approach with distributed transaction protocols. 2PC-based protocols are always prepared to abort transactions, so they can easily accommodate a Prepare result changing from PREPARE-OK to ABORT. If ABORT changes to PREPARE-OK, it might temporarily cause a conflict at the replica, which can be correctly resolved because the group as a whole could not have agreed to PREPARE-OK for two conflicting transactions.

Changing Prepare results does sometimes cause unnecessary aborts. To reduce these, TAPIR introduces two Prepare results in addition to PREPARE-OK and ABORT: ABSTAIN and RETRY. ABSTAIN helps TAPIR distinguish between conflicts with *committed* transactions, which will not abort, and conflicts with *prepared* transactions, which may later abort. Replicas return RETRY if the transaction has a chance of committing later. The client can retry the Prepare *without* re-executing the transaction.

4.3 IR Performance Principle: *Application protocols should not expect operations to execute in the same order.*

To efficiently achieve agreement on consensus results, application protocols should not rely on operation ordering for application ordering. For example, many transaction protocols [4, 19, 22] use Paxos operation ordering to determine transaction ordering. They would perform worse with IR because replicas are unlikely to agree on which transaction should be next in the transaction ordering.

In TAPIR, we use *optimistic timestamp ordering* to ensure that replicas agree on a single transaction ordering despite executing operations in different orders. Like Spanner [13], every committed transaction has a timestamp, and committed transaction timestamps reflect a linearizable ordering. However, TAPIR clients, not servers, propose a timestamp for their transaction; thus, if TAPIR replicas agree to commit a transaction, they have all agreed to the same transaction ordering.

TAPIR replicas use these timestamps to order their transaction logs and multi-versioned stores. Therefore, replicas can execute Commit in different orders but still converge to

the same application state. TAPIR leverages loosely synchronized clocks at the clients for picking transaction timestamps, which improves performance but is not necessary for correctness.

4.4 IR Performance Principle: *Application protocols should use cheaper **inconsistent** operations whenever possible rather than **consensus** operations.*

By concentrating invariant checks in a few operations, application protocols can reduce **consensus** operations and improve their performance. For example, in a transaction protocol, any operation that decides transaction ordering must be a **consensus** operation to ensure that replicas agree to the same transaction ordering. For locking-based transaction protocols, this is any operation that acquires a lock. Thus, every Read and Write must be replicated as a **consensus** operation.

TAPIR improves on this by using optimistic transaction ordering and OCC, which reduces **consensus** operations by concentrating all ordering decisions into a single set of validation checks at the proposed transaction timestamp. These checks execute in Prepare, which is TAPIR’s only **consensus** operation. Commit and Abort are **inconsistent** operations, while Read and Write are not replicated.

5. TAPIR

This section details TAPIR – the Transactional Application Protocol for Inconsistent Replication. As noted, TAPIR is designed to efficiently leverage IR’s weak guarantees to provide high-performance linearizable transactions. Using IR, TAPIR can order a transaction in a *single round-trip* to all replicas in all participant shards without *any centralized coordination*.

TAPIR is designed to be layered atop IR in a replicated, transactional storage system. Together, TAPIR and IR eliminate the redundancy in the replicated transactional system, as shown in Figure 2. As a comparison, Figure 6 shows the coordination required for the same read-write transaction in TAPIR with the following benefits: (1) TAPIR does not have any leaders or centralized coordination, (2) TAPIR Reads always go to the closest replica, and (3) TAPIR Commit takes a single round-trip to the participants in the common case.

5.1 Overview

TAPIR is designed to provide distributed transactions for a scalable storage architecture. This architecture partitions data into shards and replicates each shard across a set of storage servers for availability and fault tolerance. Clients are front-end application servers, located in the same or another datacenter as the storage servers, not end-hosts or user machines. They have access to a directory of storage servers using a service like Chubby [8] or ZooKeeper [20] and directly map data to servers using a technique like consistent hashing [21].

TAPIR provides a general storage and transaction interface for applications via a client-side library. Note that TAPIR is

Figure 6: Example read-write transaction in TAPIR. TAPIR executes the same transaction pictured in Figure 2 with less redundant coordination. Reads go to the closest replica and Prepare takes a single round-trip to all replicas in all shards.

the application protocol for IR; applications using TAPIR do not interact with IR directly.

A TAPIR application Begins a transaction, then executes Reads and Writes during the transaction’s *execution period*. During this period, the application can Abort the transaction. Once it finishes execution, the application Commits the transaction. Once the application calls Commit, it can no longer abort the transaction. The 2PC protocol will run to completion, committing or aborting the transaction based entirely on the decision of the participants. As a result, TAPIR’s 2PC coordinators cannot make commit or abort decisions and do not have to be fault-tolerant. This property allows TAPIR to use clients as 2PC coordinators, as in MDCC [22], to reduce the number of round-trips to storage servers.

TAPIR provides the traditional ACID guarantees with the strictest level of isolation: strict serializability (or linearizability) of committed transactions.

5.2 Protocol

TAPIR provides transaction guarantees using a *transaction processing protocol*, *IR functions*, and a *coordinator recovery protocol*.

Figure 7 shows TAPIR’s interfaces and state at clients and replicas. Replicas keep committed and aborted transactions in a *transaction log* in timestamp order; they also maintain a multi-versioned *data store*, where each version of an object is identified by the timestamp of the transaction that wrote the version. TAPIR replicas serve reads from the versioned data store and maintain the transaction log for synchronization and checkpointing. Like other 2PC-based protocols, each TAPIR

Figure 7: Summary of TAPIR interfaces and client and replica state.

replica also maintains a *prepared list* of transactions that it has agreed to commit.

Each TAPIR client supports one ongoing transaction at a time. In addition to its *client id*, the client stores the state for the ongoing *transaction*, including the *transaction id* and *read and write sets*. The transaction id must be unique, so the client uses a tuple of its client id and *transaction counter*, similar to IR. TAPIR does not require synchronous disk writes at the client or the replicas, as clients do not have to be fault-tolerant and replicas use IR.

5.2.1 Transaction Processing

We begin with TAPIR's protocol for executing transactions.

1. For Write(*key*, *object*), the client buffers *key* and *object* in the write set until commit and returns immediately.
2. For Read(*key*), if *key* is in the transaction's write set, the client returns *object* from the write set. If the transaction has already read *key*, it returns a cached copy. Otherwise, the client sends Read(*key*) to the replica.
3. On receiving Read, the replica returns *object* and *version*, where *object* is the latest version of *key* and *version* is the timestamp of the transaction that wrote that version.
4. On response, the client puts (*key*, *version*) into the transaction's read set and returns *object* to the application.

Once the application calls Commit or Abort, the execution phase finishes. To commit, the TAPIR client coordinates across all *participants* – the shards that are responsible for the keys in the read or write set – to find a single timestamp, consistent with the strict serial order of transactions, to assign the transaction's reads and writes, as follows:

1. The TAPIR client selects a *proposed timestamp*. Proposed timestamps must be unique, so clients use a tuple of their local time and their *client id*.
2. The TAPIR client invokes Prepare(*txn*, *timestamp*) as an IR **consensus** operation, where *timestamp* is the proposed timestamp and *txn* includes the transaction id (*txn.id*) and the transaction read (*txn.read_set*) and write sets (*txn.write_set*). The client invokes Prepare on all participants through IR as a **consensus** operations.

```

TAPIR-OCC-CHECK(txn,timestamp)
1  for  $\forall key, version \in txn.read\_set$ 
2    if  $version < store[key].latest\_version$ 
3      return ABORT
4    elseif  $version < MIN(prepared\_writes[key])$ 
5      return ABSTAIN
6  for  $\forall key \in txn.write\_set$ 
7    if  $timestamp < MAX(PREPARED-READS(key))$ 
8      return RETRY,  $MAX(PREPARED-READS(key))$ 
9    elseif  $timestamp < store[key].latestVersion$ 
10     return RETRY,  $store[key].latestVersion$ 
11  prepared-list[txn.id] = timestamp
12  return PREPARE-OK

```

Figure 8: Validation function for checking for OCC conflicts on Prepare. PREPARED-READS and PREPARED-WRITES get the proposed timestamps for all transactions that the replica has prepared and read or write to *key*, respectively.

3. Each TAPIR replica that receives Prepare (invoked by IR through ExecConsensus) first checks its transaction log for *txn.id*. If found, it returns PREPARE-OK if the transaction committed or ABORT if the transaction aborted.
4. Otherwise, the replica checks if *txn.id* is already in its *prepared list*. If found, it returns PREPARE-OK.
5. Otherwise, the replica runs TAPIR's *OCC validation checks*, which check for conflicts with the transaction's read and write sets at *timestamp*, shown in Figure 8.
6. Once the TAPIR client receives results from all shards, the client sends Commit(*txn*, *timestamp*) if all shards replied PREPARE-OK or Abort(*txn*, *timestamp*) if any shards replied ABORT or ABSTAIN. If any shards replied RETRY, then the client retries with a new proposed timestamp (up to a set limit of retries).
7. On receiving a Commit, the TAPIR replica: (1) commits the transaction to its transaction log, (2) updates its versioned store with *w*, (3) removes the transaction from its prepared list (if it is there), and (4) responds to the client.
8. On receiving a Abort, the TAPIR replica: (1) logs the abort, (2) removes the transaction from its prepared list (if it is there), and (3) responds to the client.

Like other 2PC-based protocols, TAPIR can return the outcome of the transaction to the application as soon as Prepare returns from all shards (in Step 6) and send the Commit operations asynchronously. As a result, using IR, TAPIR can commit a transaction with a *single round-trip* to all replicas in all shards.

5.2.2 IR Support

Because TAPIR's Prepare is an IR **consensus** operation, TAPIR must implement a client-side *decide* function, shown in Figure 9, which merges inconsistent Prepare results from replicas in a shard into a single result. TAPIR-DECIDE is simple: if a majority of the replicas replied PREPARE-OK, then it commits the transaction. This is safe because no conflicting transaction could also get a majority of the replicas to return PREPARE-OK.

```

TAPIR-DECIDE(results)
1  if ABORT ∈ results
2    return ABORT
3  if count(PREPARE-OK, results) ≥ f + 1
4    return PREPARE-OK
5  if count(ABSTAIN, results) ≥ f + 1
6    return ABORT
7  if RETRY ∈ results
8    return RETRY, max(results.retry-timestamp)
9  return ABORT

```

Figure 9: TAPIR’s *decide* function. IR runs this if replicas return different results on Prepare.

```

TAPIR-MERGE(d, u)
1  for ∀op ∈ d ∪ u
2    txn = op.args.txn
3    if txn.id ∈ prepared-list
4      DELETE(prepared-list, txn.id)
5  for op ∈ d
6    txn = op.args.txn
7    timestamp = op.args.timestamp
8    if txn.id ∉ txn-log and op.result == PREPARE-OK
9      R[op].result = TAPIR-OCC-CHECK(txn, timestamp)
10   else
11     R[op].result = op.result
12  for op ∈ u
13    txn = op.args.txn
14    timestamp = op.args.timestamp
15    R[op].result = TAPIR-OCC-CHECK(txn, timestamp)
16  return R

```

Figure 10: TAPIR’s *merge* function. IR runs this function at the leader on synchronization and recovery.

TAPIR also supports Merge, shown in Figure 10, and Sync at replicas. TAPIR-MERGE first removes any prepared transactions from the leader where the Prepare operation is TENTATIVE. This step removes any inconsistencies that the leader may have because it executed a Prepare differently – out-of-order or missed – by the rest of the group.

The next step checks d for any PREPARE-OK results that might have succeeded on the IR fast path and need to be preserved. If the transaction has not committed or aborted already, we re-run TAPIR-OCC-CHECK to check for conflicts with other previously prepared or committed transactions. If the transaction *conflicts*, then we know that its PREPARE-OK did not succeed at a fast quorum, so we can change it to ABORT; otherwise, for correctness, we must preserve the PREPARE-OK because TAPIR may have moved on to the commit phase of 2PC. Further, we know that it is safe to preserve these PREPARE-OK results because, if they conflicted with another transaction, the conflicting transaction *must* have gotten its consensus result on the IR slow path, so if TAPIR-OCC-CHECK did not find a conflict, then the conflicting transaction’s Prepare must not have succeeded.

Finally, for the operations in u , we simply decide a result for each operation and preserve it. We know that the leader is

now consistent with $f + 1$ replicas, so it can make decisions on consensus result for the majority.

TAPIR’s sync function (full description in [46]) runs at the other replicas to reconcile TAPIR state with the master records, correcting missed operations or consensus results where the replica did not agree with the group. It simply applies operations and consensus results to the replica’s state: it logs aborts and commits, and prepares uncommitted transactions where the group responded PREPARE-OK.

5.2.3 Coordinator Recovery

If a client fails while in the process of committing a transaction, TAPIR ensures that the transaction runs to completion (either commits or aborts). Further, the client may have returned the commit or abort to the application, so we must ensure that the client’s commit decision is preserved. For this purpose, TAPIR uses the *cooperative termination protocol* defined by Bernstein [6] for coordinator recovery and used by MDCC [22]. TAPIR designates one of the participant shards as a *backup shard*, the replicas in which can serve as a backup coordinator if the client fails. As observed by MDCC, because coordinators cannot unilaterally abort transactions (i.e., if a client receives $f + 1$ PREPARE-OK responses from each participant, it must commit the transaction), a backup coordinator can safely complete the protocol without blocking. However, we must ensure that no two coordinators for a transaction are active at the same time.

To do so, when a replica in the backup shard notices a coordinator failure, it initiates a Paxos round to elect itself as the new backup coordinator. This replica acts as the proposer; the acceptors are the replicas in the backup shard, and the learners are the participants in the transaction. Once this Paxos round finishes, the participants will only respond to Prepare, Commit, or Abort from the new coordinator, ignoring messages from any other. Thus, the new backup coordinator *prevents* the original coordinator and any previous backup coordinators from subsequently making a commit decision. The complete protocol is described in our technical report [46].

5.3 Correctness

To prove correctness, we show that TAPIR maintains the following properties³ given up to f failures in each replica group and any number of client failures:

- **Isolation.** There exists a global linearizable ordering of committed transactions.
- **Atomicity.** If a transaction commits at any participating shard, it commits at them all.
- **Durability.** Committed transactions stay committed, maintaining the original linearizable order.

We give a brief sketch of how these properties are maintained despite failures. Our technical report [46] contains a more

³ We do not prove database consistency, as it depends on application invariants; however, strict serializability is sufficient to enforce consistency.

complete proof, as well as a machine-checked TLA+ [26] specification for TAPIR.

5.3.1 Isolation

We execute Prepare through IR as a **consensus** operation. IR’s visibility property guarantees that, given any two Prepare operations for conflicting transactions, at least one must execute before the second at a common replica. The second Prepare would not return PREPARE-OK from its TAPIR-OCC-CHECK: IR will not achieve a fast quorum of matching PREPARE-OK responses, and TAPIR-DECIDE will not receive enough PREPARE-OK responses to return PREPARE-OK because it requires at least $f + 1$ matching PREPARE-OK. Thus, one of the two transactions will abort. IR ensures that the non-PREPARE-OK result will persist as long as TAPIR does not change the result in Merge (P3). TAPIR-MERGE never changes a result that is not PREPARE-OK to PREPARE-OK, ensuring the unsuccessful transaction will never be able to succeed.

5.3.2 Atomicity

If a transaction commits at any participating shard, the TAPIR client must have received a successful PREPARE-OK from every participating shard on Prepare. Barring failures, it will ensure that Commit eventually executes successfully at every participant. TAPIR replicas always execute Commit, even if they did not prepare the transaction, so Commit will eventually commit the transaction at every participant if it executes at one participant.

If the coordinator fails and the transaction has already committed at a shard, the backup coordinator will see the commit during the polling phase of the cooperative termination protocol and ensure that Commit eventually executes successfully at the other shards. If no shard has executed Commit yet, the Paxos round will ensure that all participants stop responding to the original coordinator and take their commit decision from the new backup coordinator. By induction, if the new backup coordinator sends Commit to any shard, it, or another backup coordinator, will see it and ensure that Commit eventually executes at all participants.

5.3.3 Durability

For all committed transactions, either the client or a backup coordinator will eventually execute Commit successfully as an IR **inconsistent** operation. IR guarantees that the Commit will never be lost (P1) and every replica will eventually execute or synchronize it. On Commit, TAPIR replicas use the transaction timestamp included in Commit to order the transaction in their log, regardless of when they execute it, thus maintaining the original linearizable ordering.

5.4 TAPIR Extensions

The extended version of this paper [46] describes a number of extensions to the TAPIR protocol. These include a protocol for globally-consistent read-only transactions at a timestamp,

and optimizations to support environments with high clock skew, to reduce the quorum size when durable storage is available, and to accept more transactions out of order by relaxing TAPIR’s guarantees to non-strict serializability.

6. Evaluation

In this section, our experiments demonstrate the following:

- TAPIR provides better latency *and* throughput than conventional transaction protocols in both the datacenter and wide-area environments.
- TAPIR’s abort rate scales similarly to other OCC-based transaction protocols as contention increases.
- Clock synchronization sufficient for TAPIR’s needs is widely available in both datacenter and wide-area environments.
- TAPIR provides performance comparable to systems with weak consistency guarantees and no transactions.

6.1 Experimental Setup

We ran our experiments on Google Compute Engine [18] (GCE) with VMs spread across 3 geographical regions – US, Europe and Asia – and placed in different availability zones within each geographical region. Each server has a virtualized, single core 2.6 GHz Intel Xeon, 8 GB of RAM and 1 Gb NIC.

6.1.1 Testbed Measurements

As TAPIR’s performance depends on clock synchronization and round-trip times, we first present latency and clock skew measurements of our test environment. As clock skew increases, TAPIR’s latency increases and throughput decreases because clients may have to retry more Prepare operations. It is important to note that TAPIR’s performance depends on the *actual* clock skew, not a worst-case bound like Spanner [13].

We measured the clock skew by sending a ping message with timestamps taken on either end. We calculate skew by comparing the timestamp taken at the destination to the one taken at the source plus half the round-trip time (assuming that network latency is symmetric). The average RTT between US-Europe was 110 ms; US-Asia was 165 ms; Europe-Asia was 260 ms. We found the clock skew to be low, averaging between 0.1 ms and 3.4 ms, demonstrating the feasibility of synchronizing clocks in the wide area. However, there was a long tail to the clock skew, with the worst case clock skew being around 27 ms – making it significant that TAPIR’s performance depends on actual rather than worst-case clock skew. As our measurements show, the skew in this environment is low enough to achieve good performance.

6.1.2 Implementation

We implemented TAPIR in a transactional key-value storage system, called TAPIR-KV. Our prototype consists of 9094 lines of C++ code, not including the testing framework.

We also built two comparison systems. The first, OCC-STORE, is a “standard” implementation of 2PC and OCC,

Table 1: Transaction profile for Retwis workload.

Transaction Type	# gets	# puts	workload %
Add User	1	3	5%
Follow/Unfollow	2	2	15%
Post Tweet	3	5	30%
Load Timeline	rand(1,10)	0	50%

combined with an implementation of Multi-Paxos [27]. Like TAPIR, OCC-STORE accumulates a read and write set with read versions at the client during execution and then runs 2PC with OCC checks to commit the transaction. OCC-STORE uses a centralized timestamp server to generate transaction timestamps, which we use to version data in the multi-versioned storage system. We verified that this timestamp server was not a bottleneck in our experiments.

Our second system, LOCK-STORE is based on the Spanner protocol [13]. Like Spanner, it uses 2PC with S2PL and Multi-Paxos. The client acquires read locks during execution at the Multi-Paxos leaders and buffers writes. On Prepare, the leader replicates these locks and acquires write locks. We use loosely synchronized clocks at the leaders to pick transaction timestamps, from which the coordinator chooses the largest as the commit timestamp. We use the client as the coordinator, rather than one of the Multi-Paxos leaders in a participant shard, for a more fair comparison with TAPIR-KV. Lacking access to TrueTime, we set the TrueTime error bound to 0, eliminating the need to wait out clock uncertainty and thereby giving the benefit to this protocol.

6.1.3 Workload

We use two workloads for our experiments. We first test using a synthetic workload based on the Retwis application [30]. Retwis is an open-source Twitter clone designed to use the Redis key-value storage system [39]. Retwis has a number of Twitter functions (e.g., add user, post tweet, get timeline, follow user) that perform Puts and Gets on Redis. We treat each function as a transaction, and generate a synthetic workload based on the Retwis functions as shown in Table 1.

Our second experimental workload is YCSB+T [16], an extension of YCSB [12] – a commonly-used benchmark for key-value storage systems. YCSB+T wraps database operations inside simple transactions such as read, insert or read-modify-write. However, we use our Retwis benchmark for many experiments because it is more sophisticated: transactions are more complex – each touches 2.5 shards on average – and longer – each executes 4-10 operations.

6.2 Single Datacenter Experiments

We begin by presenting TAPIR-KV’s performance within a single datacenter. We deploy TAPIR-KV and the comparison systems over 10 shards, all in the US geographic region, with 3 replicas for each shard in different availability zones. We populate the systems with 10 million keys and make

Figure 11: Average Retwis transaction Latency (Zipf coefficient 0.75) versus throughput within a datacenter.

Figure 12: Average wide-area latency for Retwis transactions, with leader located in the US and client in US, Europe or Asia.

transaction requests with a Zipf distribution (coefficient 0.75) using an increasing number of closed-loop clients.

Figure 11 shows the average latency for a transaction in our Retwis workload at different throughputs. At low offered load, TAPIR-KV has lower latency because it is able to commit transactions in a single round-trip to all replicas, whereas the other systems need two; its commit latency is thus reduced by 50%. However, Retwis transactions are relatively long, so the difference in *transaction* latency is relatively small.

Compared to the other systems, TAPIR-KV is able to provide roughly $3\times$ the peak throughput, which stems directly from IR’s weak guarantees: it has no leader and does not require cross-replica coordination. Even with moderately high contention (Zipf coefficient 0.75), TAPIR-KV replicas are able to inconsistently execute operations and still agree on ordering for transactions at a high rate.

6.3 Wide-Area Latency

For wide-area experiments, we placed one replica from each shard in each geographic region. For systems with leader-based replication, we fix the leader’s location in the US and move the client between the US, Europe and Asia. Figure 12 gives the average latency for Retwis transactions using the same workload as in previous section.

When the client shares a datacenter with the leader, the comparison systems are faster than TAPIR-KV because TAPIR-KV must wait for responses from all replicas, which takes 160 ms to Asia, while OCC-STORE and LOCK-STORE can

Figure 13: Abort rates at varying Zipf co-efficients with a constant load of 5,000 transactions/second in a single datacenter.

commit with a round-trip to the local leader and one other replica, which is 115 ms to Europe.

When the leader is in a different datacenter, LOCK-STORE suffers because it must go to the leader on Read for locks, which takes up to 160 ms from Asia to the US, while OCC-STORE can go to a local replica on Read like TAPIR-KV. In our setup TAPIR-KV takes longer to Commit, as it has to contact the *furthest* replica, and the RTT between Europe and Asia is more expensive than two round-trips between US to Europe (likely because Google’s traffic goes through the US). In fact, in this setup, IR’s slow path, at two RTT to the two closest replicas, is *faster* than its fast path, at one RTT to the furthest replica. We do not implement the optimization of running the fast and slow paths in parallel, which could provide better latency in this case.

6.4 Abort and Retry Rates

TAPIR is an optimistic protocol, so transactions can abort due to conflicts, as in other OCC systems. Moreover, TAPIR transactions can also be forced to abort or retry when conflicting timestamps are chosen due to clock skew. We measure the abort rate of TAPIR-KV compared to OCC-STORE, a conventional OCC design, for varying levels of contention (varying Zipf coefficients). These experiments run in a single region with replicas in three availability zones. We supply a constant load of 5,000 transactions/second.

With a uniform distribution, both TAPIR-KV and OCC-STORE have very low abort rates: 0.005% and 0.04%, respectively. Figure 13 gives the abort rate for Zipf co-efficients from 0.5 to 1.0. At lower Zipf co-efficients, TAPIR-KV has abort rates that are roughly an order of magnitude lower than OCC-STORE. TAPIR’s lower commit latency and use of optimistic timestamp ordering reduce the time between Prepare and Commit or Abort to a single round-trip, making transactions less likely to abort.

Under heavy contention (Zipf coefficient 0.95), both TAPIR-KV and OCC-STORE have moderately high abort rates: 36% and 40%, respectively, comparable to other OCC-based systems like MDCC [22]. These aborts are primarily due to the most popular keys being accessed very frequently. For these workloads, locking-based systems like LOCK-STORE

Figure 14: Comparison with weakly consistent storage systems.

would make better progress; however, clients would have to wait for extended periods to acquire locks.

TAPIR rarely needs to retry transactions due to clock skew. Even at moderate contention rates, and with simulated clock skew of up to 50 ms, we saw less than 1% TAPIR retries and negligible increase in abort rates, demonstrating that commodity clock synchronization infrastructure is sufficient.

6.5 Comparison with Weakly Consistent Systems

We also compare TAPIR-KV with three widely-used eventually consistent storage systems, MongoDB [35], Cassandra [25], and Redis [39]. For these experiments, we used YCSB+T [16], with a single shard with 3 replicas and 1 million keys. MongoDB and Redis support master-slave replication; we set them to use synchronous replication by setting WriteConcern to REPLICAS_SAFE in MongoDB and the WAIT command [41] for Redis. Cassandra uses REPLICATION_FACTOR=2 to store copies of each item at any 2 replicas.

Figure 14 demonstrates that the latency and throughput of TAPIR-KV is comparable to these systems. We do not claim this to be an entirely fair comparison; these systems have features that TAPIR-KV does not. At the same time, the other systems do not support distributed transactions – in some cases, not even single-node transactions – while TAPIR-KV runs a distributed transaction protocol that ensures strict serializability. Despite this, TAPIR-KV’s performance remains competitive: it outperforms MongoDB, and has throughput within a factor of 2 of Cassandra and Redis, demonstrating that strongly consistent distributed transactions are not incompatible with high performance.

7. Related Work

Inconsistent replication shares the same principle as past work on commutativity, causal consistency and eventual consistency: operations that do not require ordering are more efficient. TAPIR leverages IR’s weak guarantees, in combination with optimistic timestamp ordering and optimistic concurrency control, to provide semantics similar to past work on distributed transaction protocols but with both lower latency and higher throughput.

Table 2: Comparison of read-write transaction protocols in replicated transactional storage systems.

Transaction System	Replication Protocol	Read Latency	Commit Latency	Msg At Bottleneck	Isolation Level	Transaction Model
Spanner [13]	Multi-Paxos [27]	2 (leader)	4	$2n + \text{reads}$	Strict Serializable	Interactive
MDCC [22]	Gen. Paxos [28]	2 (any)	3	$2n$	Read-Committed	Interactive
Repl. Commit [34]	Paxos [27]	$2n$	4	2	Serializable	Interactive
CLOCC [1, 31]	VR [37]	2 (any)	4	$2n$	Serializable	Interactive
Lynx [47]	Chain Repl. [45]	–	$2n$	2	Serializable	Stored procedure
TAPIR	IR	2 (to any)	2	2	Strict Serializable	Interactive

7.1 Replication

Transactional storage systems currently rely on strict consistency protocols, like Paxos [27] and VR [37]. These protocols enforce a strict serial ordering of operations and no divergence of replicas. In contrast, IR is more closely related to eventually consistent replication protocols, like Bayou [43], Dynamo [15] and others [24, 25, 40]. The key difference is that applications resolve conflicts after they happen with eventually consistent protocols, whereas IR **consensus** operations allow applications to decide conflicts and recover that decision later. As a result, applications can enforce higher-level guarantees (e.g., mutual exclusion, strict serializability) that they cannot with eventual consistency.

IR is also related to replication protocols that avoid coordination for *commutative operations* (e.g., Generalized Paxos [28], EPaxos [36]). These protocols are more general than IR because they do not require application invariants to be pairwise. For example, EPaxos could support invariants on bank account balances, while IR cannot. However, these protocols consider two operations to commute if their order does not matter when applied to *any* state, whereas IR requires only that they produce the same results *in a particular execution*. This is a form of state-dependent commutativity similar to SIM-commutativity [10]. As a result, in the example from Section 3.1.3, EPaxos would consider any operations on the same lock to conflict, whereas IR would allow two unsuccessful Lock operations to the same lock to commute.

7.2 Distributed Transactions

A technique similar to optimistic timestamp ordering was first explored by Thomas [44], while CLOCC [1] was the first to combine it with loosely synchronized clocks. We extend Thomas’s algorithm to: (1) support multiple shards, (2) eliminate synchronous disk writes, and (3) ensure availability across coordinator failures. Spanner [13] and Granola [14] are two recent systems that use loosely synchronized clocks to improve performance for read-only transactions and independent transactions, respectively. TAPIR’s use of loosely synchronized clocks differs from Spanner’s in two key ways: (1) TAPIR depends on clock synchronization only for performance, not correctness, and (2) TAPIR’s performance is tied to the *actual* clock skew, not TrueTime’s worst-case estimated bound. Spanner’s approach for read-only transactions com-

plements TAPIR’s high-performance read-write transactions, and the two could be easily combined.

CLOCC and Granola were both combined with VR [31] to replace synchronous disk writes with in-memory replication. These combinations still suffer from the same redundancy – enforcing ordering both at the distributed transaction and replication level – that we discussed in Section 2. Other layered protocols, like the examples shown in Table 2, have a similar performance limitation.

Some previous work included in Table 2 improves throughput (e.g., Warp [17], Transaction Chains [47], Tango [5]), while others improve performance for read-only transactions (e.g., MegaStore [4], Spanner [13]) or other limited transaction types (e.g., Sinfonia’s mini-transactions [2], Granola’s independent transactions [14], Lynx’s transaction chains [47], and MDCC’s commutative transactions [22]) or weaker consistency guarantees [33, 42]. In comparison, TAPIR is the first transaction protocol to provide better performance (both throughput and latency) for general-purpose, read-write transactions using replication.

8. Conclusion

This paper demonstrates that it is possible to build distributed transactions with better performance and strong consistency semantics by building on a replication protocol with *no* consistency. We present inconsistent replication, a new replication protocol that provides fault tolerance without consistency, and TAPIR, a new distributed transaction protocol that provides linearizable transactions using IR. We combined IR and TAPIR in TAPIR-KV, a distributed transactional key-value storage system. Our experiments demonstrate that TAPIR-KV lowers commit latency by 50% and increases throughput by $3\times$ relative to conventional transactional storage systems. In many cases, it matches the performance of weakly-consistent systems while providing much stronger guarantees.

Acknowledgements

We thank Neha Narula and Xi Wang for early feedback on the paper. This work was supported by the National Science Foundation under grants CNS-0963754, CNS-1217597, CNS-1318396, CNS-1420703, and CNS-1518702, by NSF GRFP and IBM Ph.D. fellowships, and by Google. We also appreciate the support of our local zoo tapers, Ulan and Bintang.

References

- [1] A. Adya, R. Gruber, B. Liskov, and U. Maheshwari. Efficient optimistic concurrency control using loosely synchronized clocks. *Proc. of SIGMOD*, 1995.
- [2] M. K. Aguilera, A. Merchant, M. Shah, A. Veitch, and C. Karamanolis. Sinfonia: a new paradigm for building scalable distributed systems. In *Proc. of SOSP*, 2007.
- [3] P. Bailis, A. Davidson, A. Fekete, A. Ghodsi, J. M. Hellerstein, and I. Stoica. Highly available transactions: Virtues and limitations. In *Proc. of VLDB*, 2014.
- [4] J. Baker, C. Bond, J. Corbett, J. Furman, A. Khorlin, J. Larson, J.-M. Léon, Y. Li, A. Lloyd, and V. Yushprakh. Megastore: Providing scalable, highly available storage for interactive services. In *Proc. of CIDR*, 2011.
- [5] M. Balakrishnan, D. Malkhi, T. Wobber, M. Wu, V. Prabhakaran, M. Wei, J. D. Davis, S. Rao, T. Zou, and A. Zuck. Tango: Distributed data structures over a shared log. In *Proc. of SOSP*, 2013.
- [6] P. A. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison Wesley, 1987.
- [7] K. Birman and T. A. Joseph. Exploiting virtual synchrony in distributed systems. In *Proc. of SOSP*, 1987.
- [8] M. Burrows. The Chubby lock service for loosely-coupled distributed systems. In *Proc. of OSDI*, 2006.
- [9] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. *ACM Transactions on Computer Systems*, 2008.
- [10] A. T. Clements, M. F. Kaashoek, N. Zeldovich, R. T. Morris, and E. Kohler. The scalable commutativity rule: Designing scalable software for multicore processors. In *Proc. of SOSP*, 2013.
- [11] B. F. Cooper, R. Ramakrishnan, U. Srivastava, A. Silberstein, P. Bohannon, H.-A. Jacobsen, N. Puz, D. Weaver, and R. Yerneni. Pnuts: Yahoo!’s hosted data serving platform. *Proc. of VLDB*, 2008.
- [12] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. Benchmarking cloud serving systems with YCSB. In *Proc. of SOCC*, 2010.
- [13] J. C. Corbett et al. Spanner: Google’s globally-distributed database. In *Proc. of OSDI*, 2012.
- [14] J. Cowling and B. Liskov. Granola: low-overhead distributed transaction coordination. In *Proc. of USENIX ATC*, 2012.
- [15] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels. Dynamo: Amazon’s highly available key-value store. In *Proc. of SOSP*, 2007.
- [16] A. Dey, A. Fekete, R. Nambiar, and U. Rohm. YCSB+T: Benchmarking web-scale transactional databases. In *Proc. of ICDEW*, 2014.
- [17] R. Escriva, B. Wong, and E. G. Sirer. Warp: Multi-key transactions for key-value stores. Technical report, Cornell, Nov 2013.
- [18] Google Compute Engine. <https://cloud.google.com/products/compute-engine/>.
- [19] J. Gray and L. Lamport. Consensus on transaction commit. *ACM Transactions on Database Systems*, 2006.
- [20] P. Hunt, M. Konar, F. P. Junqueira, and B. Reed. Zookeeper: Wait-free coordination for internet-scale systems. In *Proc. of USENIX ATC*, 2010.
- [21] D. Karger, E. Lehman, T. Leighton, R. Panigrahy, M. Levine, and D. Lewin. Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the world wide web. In *Proc. of STOC*, 1997.
- [22] T. Kraska, G. Pang, M. J. Franklin, S. Madden, and A. Fekete. MDCC: multi-data center consistency. In *Proc. of EuroSys*, 2013.
- [23] H.-T. Kung and J. T. Robinson. On optimistic methods for concurrency control. *ACM Transactions on Database Systems*, 1981.
- [24] R. Ladin, B. Liskov, L. Shriram, and S. Ghemawat. Providing high availability using lazy replication. *ACM Transactions on Computer Systems*, 1992.
- [25] A. Lakshman and P. Malik. Cassandra: a decentralized structured storage system. *ACM SIGOPS Operating Systems Review*, 2010.
- [26] L. Lamport. The temporal logic of actions. *ACM Trans. Prog. Lang. Syst.*, 1994.
- [27] L. Lamport. Paxos made simple. *ACM SIGACT News*, 2001.
- [28] L. Lamport. Generalized consensus and Paxos. Technical Report 2005-33, Microsoft Research, 2005.
- [29] L. Lamport. Lower bounds for asynchronous consensus. *Distributed Computing*, 19(2):104–125, Oct. 2006.
- [30] C. Leau. Spring Data Redis – Retwis-J, 2013. <http://docs.spring.io/spring-data/data-keyvalue/examples/retwisj/current/>.
- [31] B. Liskov, M. Castro, L. Shriram, and A. Adya. Providing persistent objects in distributed systems. In *Proc. of ECOOP*, 1999.
- [32] B. Liskov and J. Cowling. Viewstamped replication revisited, 2012.
- [33] W. Lloyd, M. J. Freedman, M. Kaminsky, and D. G. Andersen. Don’t Settle for Eventual: Scalable Causal Consistency for Wide-area Storage with COPS. In *Proc. of SOSP*, 2011.
- [34] H. Mahmoud, F. Nawab, A. Pucher, D. Agrawal, and A. E. Abbadi. Low-latency multi-datacenter databases using replicated commit. *Proc. of VLDB*, 2013.
- [35] MongoDB: A open-source document database, 2013. <http://www.mongodb.org/>.
- [36] I. Moraru, D. G. Andersen, and M. Kaminsky. There is more consensus in egalitarian parliaments. In *Proc. of SOSP*, 2013.
- [37] B. M. Oki and B. H. Liskov. Viewstamped replication: A new primary copy method to support highly-available distributed systems. In *Proc. of PODC*, 1988.
- [38] D. R. K. Ports, J. Li, V. Liu, N. K. Sharma, and A. Krishnamurthy. Designing distributed systems using approximate synchrony in data center networks. In *Proc. of NSDI*, 2015.

- [39] Redis: Open source data structure server, 2013. <http://redis.io/>.
- [40] Y. Saito and M. Shapiro. Optimistic replication. *ACM Computing Surveys*, 2005.
- [41] S. Sanfilippo. WAIT: synchronous replication for Redis. <http://antirez.com/news/66>, Dec. 2013.
- [42] Y. Sovran, R. Power, M. K. Aguilera, and J. Li. Transactional storage for geo-replicated systems. In *Proc. of SOSP*, 2011.
- [43] D. B. Terry, M. M. Theimer, K. Petersen, A. J. Demers, M. J. Spreitzer, and C. H. Hauser. Managing update conflicts in bayou, a weakly connected replicated storage system. In *Proc. of SOSP*, 1995.
- [44] R. H. Thomas. A majority consensus approach to concurrency control for multiple copy databases. *ACM Transactions on Database Systems*, 4(2):180–209, June 1979.
- [45] R. van Renesse and F. B. Schneider. Chain replication for supporting high throughput and availability. In *Proc. of OSDI*, 2004.
- [46] I. Zhang, N. K. Sharma, A. Szekeres, A. Krishnamurthy, and D. R. K. Ports. Building consistent transactions with inconsistent replication (extended version). Technical Report 2014-12-01 v2, University of Washington CSE, Sept. 2015. Available at <http://syslab.cs.washington.edu/papers/tapir-tr-v2.pdf>.
- [47] Y. Zhang, R. Power, S. Zhou, Y. Sovran, M. K. Aguilera, and J. Li. Transaction chains: achieving serializability with low latency in geo-distributed storage systems. In *Proc. of SOSP*, 2013.