

CEP-15: Fast General Purpose Transactions

Elliott Smith, Benedict
benedict@apple.com

Zhang, Tony
nudzhang@umich.edu

Eggleston, Blake
beggleston@apple.com

Andreas, Scott
cscotta@apple.com

Abstract

Modern applications replicate and shard their state to achieve fault tolerance and scalable performance. This presents a coordination problem that modern databases address using leader-based techniques that entail trade-offs: either a scalability bottleneck or weaker isolation. Recent advances in leaderless protocols that claim to address this coordination problem have not yet translated into production systems. This paper outlines distinct performance compromises entailed by existing leaderless protocols in comparison to leader-based approaches. We propose techniques to address these short-comings and describe a new distributed transaction protocol *ACCORD*, integrating these techniques. *ACCORD* is the first protocol to achieve the same steady-state performance as leader-based protocols under important conditions such as contention and failure, while delivering the benefits of leaderless approaches to scaling, transaction isolation and geo-distributed client latency. We propose that this combination of features makes *ACCORD* uniquely suitable for implementing general purpose transactions in Apache Cassandra.

1 Introduction

Modern applications rely upon remote database services to ensure their state is durable and available to clients. To provide these properties, modern databases partition their state into geo-replicated shards. This permits some tolerated combination of failures to coincide without interrupting the service, while ensuring the database may scale to meet user demand. However, a distributed coordination problem is introduced for transaction execution.

Real-world database systems address this by imposing restrictions on functionality or sacrificing performance. Systems that offer transactions using Raft [34] or Multi-Paxos [21] are now common-place [4, 13, 14, 16, 29, 36, 42, 44, 47], but most do not offer cross-shard transactions. These were first introduced by Spanner [8], but required specialised hardware and multiple WAN round-trips. More recently, systems using commodity hardware have begun to catch up: FaunaDB and FoundationDB offer strict-serializable isolation, but order transactions with a global leader process [14, 47]; CockroachDB, YugaByte and DynamoDB avoid this bottleneck, but claim only serializable isolation [6, 40, 44]. Neither group therefore achieves the optimal combination of isolation properties and scalability. Furthermore, being leader-based these systems require additional wide area round-trips for clients that are not co-located with the leader, and for transactions that involve keys whose leaders are not co-located.

Raft and Multi-Paxos confer some important properties though: they may assign their leader role to any healthy process and require only a simple majority of votes, so they may suffer the loss of any minority of replicas and be able to promptly restore their prior steady-state performance. Transactions that share leaders also do not suffer contention penalties, and reads may be performed concurrently - they may even circumvent the leader entirely [23, 31]. Leaderless quorum-based protocols have been proposed [2, 11, 12, 23, 30, 32, 45] that utilise a fast-path to achieve optimal commit latency under low contention, but these have not been used in real systems. We propose that this is in part explained by their unpredictable performance under these same conditions.

In particular, these protocols have fast-path quorums that are disabled by fewer failures than are tolerated overall. For example, Tempo [11] tolerates f failures using $2f + 1$ replicas, but at most one replica may fail before its fast-path is unable to reach decisions. Tapir [45] fares better, with a fast path that survives $\lfloor \frac{f}{2} \rfloor$ failures - but this is half as many as it tolerates overall, and its optimistic concurrency control fails to guarantee forward progress for all transactions.

Protocol	Contention Round-trips	Concurrent Reads	No Livelock	No Aborts	General Purpose
TAPIR	2	✓			✓
Janus	2	✓		✓	
Tempo	2		✓	✓	
ACCORD	1	✓	✓	✓	✓

Table 1: Leaderless protocol properties

Switching to the slow-path imposes a burden on the wider system: each operation now consumes more system resources, as the remaining replicas must process and send more messages. Such additional burden introduces the risk of cascading failure or service degradation. This translates particularly poorly to common failure models. Services that are expected to continue in the face of major incidents such as natural disaster, or the widespread loss of power or connectivity, typically utilise multiple regions for redundancy. Existing leaderless protocols require more regions to predictably accommodate the same number of failures. Since each region incurs an independent risk of failure, the risk of experiencing any failure increases as more regions are added. By requiring more regions to accommodate the same number of failures, these protocols afford less protection overall - as well as any additional fixed costs associated with each region.

In addition, existing leaderless protocols offer a choice of undesirable trade-offs: either commutative commands must be serialized [11] or forward progress is not guaranteed [2, 32, 45] (Table 1). These are both properties that we would rather avoid. Guaranteed forward progress in practice means progress is more predictable, while resource-intensive read-only transactions are common — serializing them introduces a bottleneck that cannot be readily circumvented through scaling.

This paper describes techniques to address these shortcomings of leaderless protocols.

We introduce the idea of flexible [18] fast-path quorums for reaching fast-path consensus, offering protocols the ability to modify their fast-path criteria. We demonstrate that by modifying these quorums in response to failure, we can retain fast-path availability under any number of tolerated failures. This supports any replica layout without steady-state performance penalty versus leader-based approaches.

We additionally show how leaderless protocols that utilise timestamps to order transactions [2, 11, 12] may be combined with topographical knowledge and per-replica message buffers [43] to guarantee single round-trip consensus under tolerated latency and clock skew, regardless of message arrival order.

Finally, we describe ACCORD, a leaderless protocol designed to exploit these techniques while further guaranteeing forward progress and permitting commutative commands to be processed concurrently. We propose that this combination of properties makes ACCORD uniquely suited to providing general purpose distributed transactions.

2 Predictable Leaderless Protocols

There exists a broad literature of leaderless state-machine replication (SMR) protocols [2, 11, 12, 30, 32, 45] that utilise quorums, with a fast-path that can achieve optimal commit latency. These protocols utilise a process that executes a transaction on behalf of a client that we shall refer to as the transaction’s *coordinator*. This process is responsible for seeking votes and advancing the distributed state machine.

We assume a partially synchronous network, that processes fail only by crashing, and that there are no more than f failed processes in any replica set containing r processes where $f = \lfloor \frac{r-1}{2} \rfloor$. We additionally assume a weak failure detector [5], and that each process has a logical clock that is loosely synchronized with real time.

We introduce the concept of *robustness* and *stability* to illustrate some of the compromises entailed by these protocols.

Robustness. Fast-path quorums are said to be collectively *robust to f failures* if, in the presence of f failed processes, there remains at least one such quorum of correct processes.

Stability. A protocol is *stable to f failures* if no additional per-transaction messages must be exchanged in the presence of f failed processes. A protocol is similarly *stable to contention* if no additional messages must be exchanged when conflicting transactions reach processes out of order.

Protocol	Fast-Path	Failure		
		Tolerance	Robustness	Stability
TAPIR	$\lceil \frac{3f}{2} \rceil + 1$	f	$\lfloor \frac{f}{2} \rfloor$	$\lfloor \frac{f}{2} \rfloor$
Janus	$2f + 1$	f	0	0
Tempo _f	$2f$	f	1	1
Tempo ₁	$f + 1$	1	1	1
ACCORD	$f + 1$ to $\lceil \frac{3f}{2} \rceil + 1$	f	0 to $\lfloor \frac{f}{2} \rfloor$	f

Table 2: Comparison of failure properties with $2f + 1$ replicas per shard

2.1 Flexible Fast-Path Quorums

Existing protocols [2, 11, 32, 45] have fast-path quorums that are *robust* to fewer failures than are tolerated overall (Table 2). Thus, the protocols are *unstable* in the face of failure: if fast-path quorums become unreachable, remaining processes must exchange additional messages to reach slow-path consensus. Recent work [37] has demonstrated that this trade-off is inherent to leaderless SMR, and that a protocol tolerating f failures cannot have a fast path that is robust to more than $\lfloor \frac{f}{2} \rfloor$.

The principles of Flexible Paxos [18] provide a solution. We declare that fast-path quorums may be reconfigured, but for now leave the mechanism aside. In particular we may reduce the number of quorums, thereby reducing their collective *robustness*. We will demonstrate that in doing so we are able to reduce the size of the remaining quorums, thereby improving *stability*. The trick is that this is a one-sided trade: quorums that contain one or more crashed processes no longer contribute to robustness, so we lose nothing by removing them.

We assume that a protocol utilises three kinds of quorum, or sets of processes that may unanimously reach decisions: fast-path quorums \mathcal{F} for one round-trip consensus, simple quorums \mathcal{Q} for slow-path consensus, and recovery quorums \mathcal{R} for completing interrupted operations. We make no assumptions about the kind or configuration of these quorums but define the set of fast-path quorums to be \mathbb{F} . For recovery to arbitrate between transactions that may each have potentially taken the fast path, it is common to require that any two fast-path quorums must intersect with each other and any recovery quorum at one or more correct replicas [2, 17, 30, 32], i.e., $\mathcal{F} \cap \mathcal{F}' \cap \mathcal{R} \neq \emptyset$. For protocols that follow this design it is possible to guarantee *stability* to all tolerated failures, as we will now demonstrate.

Under maximal tolerated failures there must always be a simple quorum \mathcal{Q}_0 consisting of correct processes that may reach decisions, and thus are always able to assign our fast-path the set containing only \mathcal{Q}_0 . In this case all fast paths use the same quorum and so their intersection is also \mathcal{Q}_0 . For recovery to be possible on the slow-path, all simple and recovery quorums must intersect at one or more correct replicas, so that this set of fast-path quorums maintains our above properties. That is to say, if we fix $\mathbb{F} = \{\mathcal{Q}_0\}$ so that $\mathcal{F} = \mathcal{F}'$ we only require that $\mathcal{F} \cap \mathcal{R} \neq \emptyset$. Since $\mathcal{Q} \cap \mathcal{R} \neq \emptyset$ and $\mathcal{F} = \mathcal{Q}_0$ it must therefore be that $\mathcal{F} \cap \mathcal{R} \neq \emptyset$, and we may reach fast-path consensus under maximal tolerated failures.

This ensures no additional per-transaction messages must be exchanged under any number of tolerated failures, providing optimal *stability to failure*.

2.1.1 Fast-Path Electorates

Most quorum based protocols [2, 11, 21, 30, 34] use simple vote thresholds, where any replica may vote in combination with any other. This is the simplest model to reason about, so we adopt it for ACCORD. We introduce the concept of an *electorate*, namely those replicas whose votes may be counted towards such a vote threshold. Ordinarily all replicas are members of both the simple and fast-path electorates, but in the case of the fast-path we may remove replicas in order to reduce the successful vote threshold. We define a fast-path electorate \mathbb{E} , the membership of which makes fast-path decisions. To maintain the previously outlined recovery properties, namely that the intersection of any recovery quorum with any two fast-path quorums contains at least one correct replica, we derive the vote threshold as follows:

$$|\mathcal{F}| - (|\mathbb{E}| - |\mathcal{F}'|) - f > 0 \quad (1)$$

$$\implies |\mathcal{F}| = \left\lceil \frac{|\mathbb{E}| + f + 1}{2} \right\rceil \quad (2)$$

For example, a replica set of $r = 9$ processes tolerating $f = 4$ failures could choose an electorate of size 9, 7 or 5 conferring fast-path quorums of 7, 6 and 5 respectively. An electorate of 9 could then reach fast-path decisions with fewer than 3 failed processes, at which point the electorate would need to be reduced to either 7 or 5. Under the maximum 4 tolerated failures, all 4 failed processes would need to be excluded from the electorate so that all fast- and slow-path decisions would use the same quorum.

2.2 Timestamp Reorder Buffer

Leaderless consensus protocols can directly process transactions from clients in any region, but this capability introduces the risk of contention: two transactions started at the same time but in different regions may be unable to both reach fast-path consensus. Those replicas closest to each coordinator will witness its corresponding transaction first, so that no fast-path quorum may witness a consistent order of arrival for either transaction. Recent work [43] has demonstrated the potential for loosely synchronized clocks to improve this scenario. Coordinators may assign transactions a future time to be processed on recipient replicas, with the time selected to allow conflicting transactions sent by further-away coordinators to arrive first.

This approach can be refined to guarantee consensus for every transaction under realistic conditions. Firstly, since clocks are not perfectly synchronized messages may still be processed too early. By measuring clock asynchrony, or clock skew we can extend our processing delay to reliably accommodate differences in link latency. Secondly, by utilising a protocol such as Caesar [2] or Tempo [11], that proposes an execution timestamp. Unlike protocols that agree dependencies [30], where every fast-path vote must witness identical histories, timestamp protocols only require that voters witness timestamps in ascending order. Therefore, if the message processing time is derived from the transaction’s proposed execution timestamp we are able to guarantee fast-path consensus for all transactions proposed by a correct and responsive coordinator, using only two weak assumptions: that we can measure both the latency of our network and the error margins of our clocks.

We assume a protocol P that utilises a fast-path for one-round consensus, that its correctness is independent of message arrival order, and that it proposes a timestamp t_0 as a transaction’s intended execution time. We will refer to this round as *PreAccept*. If a fast path quorum witnesses t_0 before any higher timestamp then it must be accepted, and this must uniquely determine the transaction’s execution order.

We define the set \mathbb{C} of all coordinators, $Skew_{Max}$ the maximum instantaneous difference between the logical clocks on any two nodes, and $Latency(C, P)$ as the latency between some coordinator C and some replica P . We finally define $Deadline(t_0, C, P) = t_0 + Skew_{Max} + \max(Latency(C', P) \mid C' \in \mathbb{C}) - Latency(C, P)$. On replica P , receiving a message from coordinator C , this computes the last point in time at which a transaction might arrive from other members of \mathbb{C} with an earlier t_0 .

There are no modifications necessary to protocol P . On receiving a *PreAccept* request, replicas do not immediately process the message, instead buffering it in a queue ordered by t_0 until the computed *Deadline*, at which point the buffer is processed up to and including this message.

Since P ’s safety does not depend on message arrival order, it is unaffected by this change. We assume that our bounds on message latency and clock synchrony are not breached and that processes are responsive. Under these conditions, this generic change guarantees that for any two transactions γ and τ with $t_{0\gamma} < t_{0\tau}$ so that τ could prevent fast-path consensus for γ , it must be that γ is known to \mathbb{E} before τ is processed, and will be processed first. As a result, out-of-order transaction arrival does not result in any additional messages being exchanged.

This provides *stability to contention*.

2.2.1 Clock Synchrony versus Latency

This technique trades wide area round-trips for an additional latency penalty equal to the bounds on clock synchrony. How useful is this exchange in practice? Many modern data centres have at least one high quality time source such as GPS that is typically combined with a protocol such as NTP or PTP to synchronize commodity time sources residing on the local network, such as those embedded into modern CPUs. In this configuration clock skew may be sub-millisecond [15], so that this penalty will likely be dwarfed by wide-area latency. Importantly any additional latency is uncorrelated with wide-area latency: each locale has its own time source and dissemination mechanism, so that accommodating clock skew represents a fixed global penalty.

3 ACCORD Basic Protocol

ACCORD builds upon a rich history of leaderless state-machine replication (SMR) protocols [2, 11, 12, 32] exploiting the intuitions of EPaxos [30]. This broad approach can be viewed as performing a separate instance of classic Paxos for each transaction to agree its execution order, with a special initial round that permits faster consensus. Recent work has shown that this technique can be easily extended to partial state-machine replication (PSMR) scenarios for improved scalability, with shards replicating only a portion of the global state-machine [11, 32]. ACCORD utilises this approach to agree a timestamp to determine a transaction’s execution order, reducing the incidence of contention. Like Caesar [2], ACCORD uses its consensus acknowledgments to determine the set of conflicting transactions that may execute earlier, however the worst case behaviour is improved so that consensus is reached in at most two round-trips, and live-lock [11] is avoided. Uniquely, this dependency set is not decided consistently during consensus - only a superset of each transaction’s final dependencies is determined, that may be filtered during execution to ensure consistency. This approach achieves an optimal baseline combination of characteristics compared to comparable state-of-the-art protocols, providing a suitable platform for utilising a fast-path electorate and reorder buffer.

As is typical for Paxos [21] derivatives, ballots are used to ensure only the most recent command for a given phase may complete [2, 11, 23, 30]. For brevity we elide the handling of these ballots from the Consensus protocol, and implicitly discard messages if a newer ballot has been witnessed. We additionally discard messages from earlier phases of the protocol.

Transactions are denoted by γ , τ and ν . We say that two transactions γ and τ **conflict** ($\gamma \sim \tau$) if their execution is not commutative, so that either their response or the database state would differ if their execution order were reversed.

Variables declared with a τ subscript are persistent, storing the most recent value for τ on that process. $\rho, \mathbb{E}, \mathcal{Q}, \mathcal{F}$ represent, respectively, a replica-set for a single shard, its fast-path electorate, and a simple and fast-path quorum. By the same token, $\mathbb{P}^\tau, \mathbb{E}^\tau, \mathcal{Q}^\tau, \mathcal{F}^\tau$ denote the replica-sets participating in τ , the union of their fast-path electorates, and a set of responses constituting a simple and fast-path quorum in each replica-set in \mathbb{P}^τ .

3.1 Overview

A client process that wishes to execute a transaction selects a nearby non-replica coordinator C to perform consensus with Algorithm 1 so that an execution timestamp may be decided. Along with this timestamp, the protocol also determines a set of conflicting transactions that might precede our execution. Algorithm 2 is then performed, also by C , so that general-purpose transactions combining cross-shard state may be executed. C steps through several phases across these two protocols: *PreAccept*, *Accept*, *Commit*, *Execute* and *Apply*, all of which must occur in sequence. Only *Accept* may be skipped, if the fast path is taken. *PreAccept* and *Accept* collectively decide the order of execution, which is durably decided prior to entry to *Commit*, where this decision is disseminated to dependent transactions. *Execute* then awaits the execution of those dependencies with lower execution timestamps before computing the result of the transaction. *Apply* persists this result to all replicas. Should the coordinator fail, a weak failure detector invokes the recovery protocol on a replica that has witnessed the transaction, that then takes over the coordinator role from the last successfully completed phase.

In the following sections the consensus and execution protocols will be described in detail, alongside certain simple properties that they maintain. In Section 4 the recovery protocol will be described for finishing interrupted transactions. In Section 3.4 and Section 4 these properties will be used to demonstrate that these protocols in combination maintain the following properties:

Execution consistency. All conflicting transactions are applied in the same order on all replicas.

Execution linearizability. Any transaction γ acknowledged before a conflicting transaction τ has been submitted is executed before τ and applied before τ on all replicas.

Liveness. Any transaction that is known by at least one correct replica will execute eventually and apply on all correct replicas.

3.2 Consensus

ACCORD imposes a total order on conflicting transactions by assigning them each a unique execution timestamp. Timestamps consist of a tuple $(time, seq, id)$, where $time$ is assigned from a per-process monotonically increasing value that is loosely synchronised with the wall clock, seq is a logical time component and id is a unique identifier of the process that created the timestamp. A timestamp sorts by its components in precedence of their declared order. We

use three symbols to represent timestamps $t_0 \leq t \leq T$ where t_0 denotes the timestamp first proposed by *PreAccept*, t a proposed execution timestamp and T_τ the highest timestamp witnessed by a process for τ .

At a high level, a Lamport clock of these timestamps assigns execution times to conflicting transactions, imposing a total order [20]. ACCORD ensures this order is consistent for transactions that have been acknowledged to clients, so that in the event of failure the same total order will be recovered. Initially, a transaction coordinator proposes a timestamp for execution. If a fast-path quorum of replicas unanimously accept and record this timestamp as the most recent, then only this single timestamp may be recovered and it is decided immediately. Otherwise some replicas have responded proposing newer timestamps, so that there are multiple possible Lamport timestamps that may be recovered, and therefore distinct total orders. In this case a slow path using classic Paxos durably agrees which of these possibilities is decided. To execute transactions in timestamp order, replicas also inform coordinators of any conflicting transactions that may take an earlier timestamp. Execution proceeds once these have all committed, and any with an earlier timestamp have executed. This is now described in more detail, alongside certain properties that are maintained.

Algorithm 1 Consensus Protocol

receive τ on coordinator C from client:

- 1: $t_0 \leftarrow (\text{now}, \theta, C)$
- 2: **send** *PreAccept*(τ, t_0) to $\forall p \in \mathbb{E}^\tau$

receive *PreAccept*($\tau, t_{0\tau}$) on p :

- 3: **if** $t_{0\tau} > \max(T_\gamma \mid \gamma \sim \tau)$ **then**
- 4: $t_\tau \leftarrow t_{0\tau}$
- 5: **else**
- 6: $t_\tau \leftarrow \max(T_\gamma \mid \gamma \sim \tau)$
- 7: $t_\tau.(seq, id) \leftarrow (t_\tau.seq + 1, p)$
- 8: **end if**
- 9: $T_\tau \leftarrow t_\tau$
- 10: *PreAccepted* $_\tau \leftarrow \text{true}$
- 11: **reply** *PreAcceptOK*($t_\tau, deps : \{\gamma \mid \gamma \sim \tau \wedge t_{0\gamma} < t_{0\tau}\}$)

receive *PreAcceptOK*($t, deps$) from \mathcal{Q}^τ :

- 12: $deps \leftarrow \bigcup \{p.deps \mid p \in \mathcal{Q}^\tau\}$
- 13: **if** $\exists \mathcal{F}^\tau \subseteq \mathcal{Q}^\tau (\forall p \in \mathcal{F}^\tau \cdot p.t = t_0)$ **then**
- 14: **send** *Commit*($\tau, t_0, t_0, deps$) to $\forall p \in \rho \in \mathbb{P}^\tau$
- 15: **go to** Execution Protocol
- 16: **else**
- 17: $t \leftarrow \max(p.t \mid p \in \mathcal{Q}^\tau)$
- 18: **send** *Accept*($\tau, t_0, t, deps$) to $\forall p \in \rho \in \mathbb{P}^\tau$
- 19: **end if**

receive *Accept*($\tau, t_{0\tau}, t_\tau, deps_\tau$):

- 20: $T_\tau \leftarrow \max(t_\tau, T_\tau)$
- 21: *Accepted* $_\tau \leftarrow \text{true}$
- 22: **reply** *AcceptOK*($deps : \{\gamma \mid \gamma \sim \tau \wedge t_{0\gamma} < t_\tau\}$)

receive *AcceptOK*($deps$) on C from \mathcal{Q}^τ :

- 23: $deps \leftarrow \bigcup \{p.deps \mid p \in \mathcal{Q}^\tau\}$
 - 24: **send** *Commit*($\tau, t_0, t, deps$) to $\forall p \in \rho \in \mathbb{P}^\tau$
 - 25: **go to** Execution Protocol
-

Property 3.1 (Timestamp ordering). *For any two conflicting transactions $\gamma \sim \tau$ where γ commits before τ is submitted by a client, $t_\gamma < t_\tau$.*

The coordinator assigns transaction τ a globally unique timestamp t_0 and proposes this to \mathbb{E}^τ , the members of all fast-path electorates participating in τ . This is the lowest possible execution timestamp t that τ may take. On recipient replicas, t_0 is compared with the largest timestamp witnessed for all conflicting transactions $\gamma \sim \tau$. If it is larger, the replica votes in favour of setting $t = t_0$; otherwise a new larger t is proposed. In either case the proposed t is stored by the replica so that it may only vote for higher timestamps for future conflicting transactions.

Each replica additionally includes in its *PreAcceptOK* response the set $deps_\tau$ of conflicting transactions $\gamma \sim \tau$ witnessed by the replica that might take a lower execution timestamp, namely those where $t_{0\gamma} < t_{0\tau}$. Note that any transitive dependency of another $\gamma \in deps_\tau$ where $Committed_\gamma$ may be pruned from $deps_\tau$, as it is durably a transitive dependency of τ .

Property 3.2 (Timestamp consistency). *All processes that commit a transaction do so with the same timestamp.*

If a fast-path quorum \mathcal{F}^τ votes to accept $t = t_0$, this timestamp is durably decided and is committed to all replicas participating in τ . Fast-path quorums and their durability are described in more detail in Section 4.

If insufficient fast-path responses are received, the maximum t from any simple quorum \mathcal{Q}^τ may be taken to impose a valid total order [20]. However there is no single t that can be durably recovered should this coordinator fail. To ensure consistency we require that a recovery coordinator will select the same t that we proceed to *Execute* with. We achieve this with an *Accept* phase that can be viewed as a classic Paxos *Accept* phase with a privileged initial ballot b_0 , similar to Fast Paxos [23]. This phase requires only a simple quorum for all $\rho \in \mathbb{P}^\tau$ before τ may be committed.

In this case every replica includes in its *AcceptOK* response the set of conflicting transactions $\gamma \sim \tau$ witnessed by the replica that might take a lower execution timestamp, namely those where $t_{0\gamma} < t_\tau$.

3.3 Execution

Once t_τ is decided it is logically committed, and disseminated to every shard alongside $deps_\tau$ via *Commit* so that other transactions with an earlier execution timestamp that had witnessed us during consensus may proceed. Note that only t_τ is durably committed, $deps_\tau$ is only recorded for recovery purposes outlined in Section 4. Simultaneously the coordinator issues *Read* messages to at least one process in each $\rho \in \mathbb{P}^\tau$, including the nearest correct process for each shard. In common configurations this will permit a response without any wide-area latency. The *Read* request includes the subset of $deps_\tau$ that interacts with ρ , and replicas wait to answer this message until every such dependency has either been witnessed as committed with a higher execution timestamp $t_\gamma > t_\tau$, or its result has been applied locally. Replicas do not advance further until the coordinator evaluates τ against these responses and submits *Apply* messages, that must in turn be received by replicas before they may process conflicting transactions with a higher timestamp.

Algorithm 2 Execution Protocol

Coordinator C:

```

26: for  $\rho \in \mathbb{P}^\tau$  do
27:    $deps_\rho \leftarrow \{\gamma \mid \gamma \in deps \wedge \rho \in \mathbb{P}^\gamma\}$ 
28:   send Read( $\tau, t, deps_\rho$ ) to some nearby  $p \in \rho$ 
29: end for
receive Commit( $\tau, t_0, t_\tau, deps_\tau$ ):
30:  $Committed_\tau \leftarrow \mathbf{true}$ 
receive Read( $\tau, t_\tau, deps_{\tau, \rho}$ ) on  $p$ :
31: await  $Committed_\gamma \forall \gamma \in deps_{\tau, \rho}$ 
32: await  $Applied_\gamma \forall \gamma \in deps_{\tau, \rho}, t_\gamma < t_\tau$ 
33:  $reads \leftarrow \text{read}(\tau)$ 
34: reply ReadOK( $reads$ )
receive ReadOK( $reads$ ) from each shard:
35:  $result \leftarrow \text{execute}(\tau, reads)$ 
36: send Apply( $\tau, t, deps_{\tau, \rho}, result$ ) to  $\forall p \in \rho \in \mathbb{P}^\tau$ 
37: send  $result$  to client
receive Apply( $\tau, t_\tau, deps_{\tau, \rho}, result_\tau$ ):
38: await  $Committed_\gamma \forall \gamma \in deps_{\tau, \rho}$ 
39: await  $Applied_\gamma \forall \gamma \in deps_{\tau, \rho}, t_\gamma < t_\tau$ 
40:  $\text{apply}(writes, t_\tau)$ 
41:  $Applied_\tau \leftarrow \mathbf{true}$ 

```

Note. The result of execution is persisted at replicas to avoid the following situation: if a transaction involving shards a and b is applied at all correct ρ_a , but is unapplied at all correct ρ_b , we would be unable to re-compute the result

of τ , and transactions that conflict on b would be unable to proceed. By persisting its result, these transactions may consult a to complete τ on b .

3.4 Safety

We now demonstrate that this protocol maintains the atomicity and strict-serializable isolation properties, namely that every transaction is executed by every replica, in the same order, and that to an external observer every transaction occurs in an order consistent with the real-time order of client-visible events.

Property 3.3 (Dependency safety). *Any coordinator committing τ with t_τ does so with $deps_\tau$ containing all conflicting γ that may be committed with $t_\gamma < t_\tau$.*

Proof. Suppose τ is committed with t_τ and $deps_\tau$, and consider any γ that commits with $t_\gamma < t_\tau$. Each of γ, τ may have committed via the slow- or fast-paths. We illustrate the case where both γ, τ committed via the slow-path; the other cases are similar. By assumption, γ is pre-accepted at some slow-path quorum Q^γ , where each replica in Q^γ pre-accepted γ for some execution timestamp $\leq t_\gamma$. Meanwhile, τ must be accepted at some slow-path quorum Q^τ with t_τ . Then for any replica P in $Q^\gamma \cap Q^\tau$, P must have pre-accepted γ before accepting τ , since otherwise P will not have pre-accepted γ with an execution timestamp less than t_τ . As a result, by the Accept Phase of the protocol, $\gamma \in deps_\tau$. \square

Property 3.4 (Execution order). *For any transactions γ and τ where $\gamma \sim \tau$ and $t_\gamma < t_\tau$, τ executes after γ .*

Proof. By Property 3.3 the coordinator knows all γ that may be committed with $t_\gamma < t_\tau$. By waiting for these to commit we know all γ where $t_\gamma < t_\tau$ and wait for them to apply on a replica of each shard. \square

Property 3.5 (Application order). *For any transaction τ , τ applies at p after all $\gamma \sim \tau$ where, once γ is committed, $t_\gamma < t_\tau$.*

Proof. By Property 3.3 the coordinator knows all γ that may be committed with $t_\gamma < t_\tau$, and waits for them to apply before sending the result of executing τ to p . The subset of $deps_\tau$ that applies on p is included, and p waits for this subset to apply locally before applying τ . \square

Theorem 3.1 (Execution consistency). *For any two conflicting transactions $\gamma \sim \tau$ where γ is applied before τ on some replica, γ will be applied before τ on all replicas.*

Proof. By Properties 3.2, 3.3 and 3.4, $t_\gamma < t_\tau$ and $\gamma \in deps_\tau$. By Properties 3.4 and 3.5 γ executes before τ and applies before τ at all replicas. \square

Theorem 3.2 (Execution linearizability). *For any two conflicting transactions $\gamma \sim \tau$ where τ is proposed by a client after γ has been committed, γ executes before τ , and is applied before τ at all replicas.*

Proof. By Properties 3.1, 3.2 and 3.3, $t_\gamma < t_\tau$ and $\gamma \in deps_\tau$. By Properties 3.4 and 3.5 γ executes before τ and applies before τ at all replicas. \square

4 Recovery Protocol

Algorithm 3 is invoked by a weak failure detector to recover a transaction τ whose coordinator has failed. This protocol contacts a recovery quorum \mathcal{R}^τ that ensures τ is pre-accepted before responding, ensuring the properties of normal execution are maintained against any υ that had not previously committed. For those υ that had already committed we must now demonstrate that we maintain our declared properties while recovering τ . In the case where τ has reached a slow-path decision, any quorum $\mathcal{R} \subseteq \rho$ where $|\mathcal{R}| \geq r - f$ must witness at least one response from the last completed phase of the state machine. The state machine can be picked-up from this point, similarly maintaining the properties of normal execution. We now only require a similar capability for fast-path decisions.

Fast-path recovery operates on a simple intuition: if we may deduce that an incomplete transaction τ either cannot have been committed on the fast path *or* that any υ that may execute after τ has $\tau \in deps_\upsilon$, then we may safely propose $t = t_0$ by the slow path. If υ is committed after the recovery quorum was reached then it must witness τ . Otherwise this is achieved by using the *deps* associated with all *Accepted* or *Committed* υ that supersede τ : if any of these had

Algorithm 3 Recovery Protocol**Coordinator C:**

```

1:  $b \leftarrow$  fresh ballot
2: send  $Recover(b, \tau, t_0)$  to  $\forall p \in \rho \in \mathbb{P}^\tau$ 
receive  $Recover(b, \tau, t_0)$  on  $p$ :
3: if  $b \leq PromisedBallot_\tau$  then
4:   reply  $NACK(b_\tau)$ 
5: else
6:    $PromisedBallot_\tau \leftarrow b$ 
7:    $Accepts \leftarrow \{\gamma \mid \gamma \sim \tau \wedge \tau \notin deps_\gamma \wedge Accepted_\tau\}$ 
8:    $Commits \leftarrow \{\gamma \mid \gamma \sim \tau \wedge \tau \notin deps_\gamma \wedge Committed_\tau\}$ 
9:    $Wait \leftarrow \{\gamma \in Accepts \mid t_{0\gamma} < t_{0\tau} \wedge t_\gamma > t_{0\tau}\}$ 
10:   $Superceding \leftarrow \{\gamma \in Accepts \mid t_{0\gamma} > t_{0\tau}\}$ 
11:     $\cup \{\gamma \in Commits \mid t_\gamma > t_{0\tau}\}$ 
12:  if  $\tau \notin Txns$  then
13:    run Consensus Protocol L3 to L10 ▷ PreAccept
14:  end if
15:  if  $\neg Accepted_\tau \wedge \neg Committed_\tau \wedge \neg Applied_\tau$  then
16:     $deps_\tau \leftarrow \{\gamma \mid \gamma \sim \tau \wedge t_{0\gamma} < t_{0\tau}\}$ 
17:  end if
18:  reply  $RecoverOK(*_\tau, Superceding, Wait)$ 
19: end if
receive  $NACK$  on  $C$ :
20: yield to competing coordinator
receive  $RecoverOK(*, Superceding, Wait)$  from  $p \in \mathcal{R}^\tau$ 
21: if  $\exists p \in \mathcal{R}^\tau (p.Applied_\tau)$  then
22:   send  $response(result)$  to client
23:   send  $Apply(\tau, t_0, p.t, p.deps, result)$  to  $\forall p' \in \rho \in \mathbb{P}$ 
24: else if  $\exists p \in \mathcal{R}^\tau (p.Committed_\tau)$  then
25:   send  $Commit(\tau, t_0, p.t, p.deps) \forall p' \in \rho \in \mathbb{P}^\tau$ 
26:   go to Execution Protocol
27: else if  $\exists p \in \mathcal{R}^\tau (Accepted_\tau)$  then
28:   select  $p$  with highest accepted ballot
29:    $t \leftarrow p.t$ ;  $deps \leftarrow p.deps$ 
30:   go to Consensus Protocol L18 ▷ Accept
31: else
32:    $t \leftarrow t_0$ ;  $deps \leftarrow \bigcup \{p.deps_\tau \mid p \in \mathcal{R}^\tau\}$ 
33:   if  $\exists i \mid \{p \in \mathbb{E} \mid p.t > p.t_0\} > |\mathbb{E}| - |\mathcal{F}_i|$  then
34:      $t \leftarrow \max(p.t \mid p \in \mathcal{R}^\tau)$ 
35:   else if  $\exists p \in \mathcal{R}^\tau (p.Superceding \neq \emptyset)$  then
36:      $t \leftarrow \max(p.t \mid p \in \mathcal{R}^\tau)$ 
37:   else if  $\bigcup \{p.Wait \mid p \in \mathcal{R}^\tau\} \neq \emptyset$  then
38:     await  $Committed_\gamma (\forall \gamma \in \bigcup \{p.Wait \mid p \in \mathcal{R}^\tau\})$ 
39:     restart Recovery Protocol
40:   end if
41:   go to Consensus Protocol L18 ▷ Accept
42: end if

```

not witnessed τ in $deps_{\nu}$, then τ had not reached a fast-path before they did, and since ν supersedes τ it may not reach fast-path consensus after. Conversely, if all have witnessed τ then they will wait for τ to commit so that we may propose any t , including t_0 .

In more detail, given transactions $\nu \sim \tau$ where ν supersedes τ , any shard ρ that ν and τ have in common, any recovery quorum \mathcal{R}^{τ} , and any subset E of the shard's electorate that may be witnessed by this quorum¹, we must be able to determine either that τ is not pre-accepted or that $\tau \in deps_{\nu}$. To achieve this, as outlined in Section 2.1.1 we ensure that $\mathcal{F}^{\tau} \cap \mathcal{F}^{\nu} \cap E \neq \emptyset$, i.e. that $|\mathcal{F}| = \lceil \frac{|\mathbb{E}|+f+1}{2} \rceil$. Now any τ that did not take the fast path must have fewer than $r - |\mathcal{F}|$ votes for $t = t_0$. However not all τ with this many votes took the fast path, so we must either (1) deduce from other information that they did not, so that we may propose a higher timestamp; or (2) determine that it is safe to propose $t = t_0$ regardless. If τ was committed on the fast-path then proposing $t = t_0$ is correct. If it was not then *Timestamp consistency* and *Timestamp ordering* are unaffected by proposing $t = t_0$, so that we must only demonstrate that *Dependency safety* continues to be maintained in this case.

By definition, if *both* τ and ν may have reached fast-path consensus then $\mathcal{F}^{\tau} \cap \mathcal{F}^{\nu} \neq \emptyset$ so that $\tau \in deps_{\nu}$. By simple induction this maintains *Dependency safety* between all transactions that may have committed on the fast-path.

For any ν committed on the slow path, an *Accept* or *Commit* must be witnessed by \mathcal{R}^{τ} so that we may inspect $deps_{\nu}$. If *Accepted* $_{\nu}$, ν must have reached a simple quorum \mathcal{Q}^{ν} during *PreAccept* before proposing $deps_{\nu}$, so that if $\tau \notin deps_{\nu}$ and $t_{0\nu} > t_{0\tau}$ we know that τ could not have reached fast-path consensus and we may proceed on the slow path. Similarly, if *Committed* $_{\nu}$, ν must have reached \mathcal{Q}^{ν} during *Accept* so we may test $t_{\nu} > t_{0\tau} \wedge \tau \notin deps_{\nu}$. Conversely, if $\tau \in deps_{\nu}$ and ν successfully committed, any execution of ν will know of τ so that *Dependency Safety* is maintained.

This leaves only those transactions γ where $t_{0\gamma} < t_{0\tau}$ that may have reached slow-path consensus with $t_{\gamma} > t_{0\tau}$ but have not been committed. Such transactions may not have witnessed τ without prohibiting τ from reaching fast-path consensus. If we have not determined the correct course of action by other means, we may simply wait for these transactions to commit, or commit them ourselves.

Theorem 4.1 (Liveness). *A transaction τ that is known by at least one healthy replica will always execute eventually and apply on all healthy replicas*

Proof. A healthy coordinator executing the consensus protocol defined in Algorithm 1 always commits τ in a fixed number of steps unless superseded by a recovery coordinator. If the coordinator process is faulty then, given synchrony conditions and weak failure detectors, a recovery coordinator is eventually picked. The recovery protocol defined in Algorithm 3 commits in a fixed number of steps unless there exists an γ where *Accepted* $_{\gamma} \wedge t_{0\gamma} < t_{0\tau} \wedge t_{\gamma} > t_{0\tau}$. Given loosely synchronized clocks there are finite γ , so by simple induction all commands are able to commit eventually. Once committed, execution waits only for those commands $\gamma \in deps_{\tau}$ where $t_{\gamma} < t_{\tau}$, so all commands are able to eventually execute. This result is sent to *Apply* on all healthy replicas. If a healthy process witnesses τ , but does not witness *Apply*, it eventually invokes the recovery protocol. \square

5 Reconfiguring Electorates

A configuration consists of a monotonically increasing *epoch* identifier, the membership of all shards ρ and their fast-path electorates \mathbb{E} . We say that a configuration e_{i+1} has taken effect on ρ once a simple quorum of ρ_{e_i} , the membership of ρ as decided by e_i , is aware of it.

Timestamps are modified to include *epoch* as their first value, so that a transaction τ is coordinated by C with configuration e_1 using $t_{0\tau}.epoch = e_1$. Any $p \in \rho_{e_1}$ with a newer configuration e_2 will respond to *PreAccept* with $t_{\tau}.epoch = e_2$ so that p may not participate in fast-path decision. This alerts C to the new configuration, which is fetched asynchronously. If a fast-path quorum is not reached with ρ_{e_1} , we wait for the new configuration and for \mathcal{Q}_{e_2} *PreAcceptOK* responses - if necessary sending additional *PreAccept* to ρ_{e_2} . The remainder of the Consensus protocol executes for both configurations simultaneously, i.e. so that *Accept* is sent to all $\mathbb{P}^{\tau} = \mathbb{P}_{e_1}^{\tau} \cup \mathbb{P}_{e_2}^{\tau}$, and $\mathcal{Q}^{\tau} = \mathcal{Q}_{e_1}^{\tau} \cup \mathcal{Q}_{e_2}^{\tau}$ must respond with *AcceptOk* before we may proceed to the *Commit* phase. The Execution protocol must wait for $deps_{\tau}$ to apply in $\mathbb{P}_{e_1}^{\tau}$ only, before applying in $\mathbb{P}_{e_2}^{\tau}$ only. The Recovery protocol remains unmodified, depending only on these changes to the Consensus and Execution protocols.

This ensures atomic migration from one configuration to another. No replica that is aware of e_2 may participate in a fast-path decision using e_1 , so if sufficient fast-path responses are received e_2 has not taken effect and τ completes

¹That is, where $t_{0\tau} < t_{0\nu}$, $\rho \in (\mathbb{P}^{\tau} \cap \mathbb{P}^{\nu})$, $\mathbb{E} \subseteq \rho$ and $E = \mathcal{R}^{\tau} \cap \mathbb{E}$

normally. A slow-path decision executes the protocol for both configurations, necessarily informing a majority of ρ of the new configuration, so that it takes effect for transactions that follow τ .

However, by itself this modification does not maintain consistency during recovery. Recall that the simultaneous intersection of any fast-path quorum with any other fast-path quorum and any recovery quorum must include at least one correct process. To ensure this property is maintained, any replica removed from the electorate must broadcast to the new members any transactions it has witnessed but not applied. Until sufficient removed nodes have notified enough new nodes to satisfy this property, the new electorate is forbidden from reaching fast-path decisions by itself. Specifically, newly added members may not participate in a fast-path decision without additional input from removed members of the electorate. These removed members do not participate in any decision, but may be contacted simultaneously to contribute knowledge of old conflicting transactions that were accepted on the fast-path, permitting newly added members' votes to be counted.

Finally, to avoid slow-path consensus during reconfiguration, a new configuration may be agreed to take effect at some future timestamp, allowing it to be disseminated in time for replicas and coordinators to synchronise its application. There will still of course be an interim period during which both fast-paths must be contacted for a fast-path decision to be reached.

6 Related Work

State machine replication (SMR). SMR [38] is a common foundation for building reliable, fault-tolerant systems. Traditional SMR protocols [21–23, 25, 33, 34] rely on a stable leader to efficiently assign a linearizable ordering to client requests. Replicas execute requests in this order, and are able to adopt the leader role should it fail. This provides clients the illusion of a unified service running on a single machine. However, the single leader in these protocols is a bottleneck that limits throughput, and results in increased latency for remote clients in a wide area network.

For better performance in geo-replicated systems, SMR protocols such as Mencius [28] and EPaxos [30] have been proposed. While Mencius permits the leader role to be efficiently rotated between regions to improve client latency, if any replica is faulty progress stalls until this is detected and recovered. EPaxos instead discards the concept of a leader, permitting any replica to coordinate a command at the same time as any other. Dependency graphs and topological sorting ensure a consistent order of execution for non-commutable commands. Recently EPaxos has been demonstrated to suffer degraded performance under contended workloads [2, 12, 43]. Two approaches to this problem have been outlined: loosely synchronized clocks have been employed to synchronize the order in which replicas process conflicting commands [43], and alternative protocols such as Caesar [2] and Atlas [12] have been proposed that agree an execution timestamp, permitting replicas to reach consensus without witnessing identical histories, so that conflicts may be reduced. ACCORD combines both techniques, proposing an execution timestamp and processing messages in timestamp order. This combination provides ACCORD stronger *stability* properties. Additionally, ACCORD has a worst case execution path of two wide area round-trips instead of Caesar's three, ACCORD does not order non-commutative commands in comparison to Atlas.

Separately, SMR approaches have been proposed that vary failure handling properties. Vertical Paxos [24] demonstrated consensus with an optimal two replicas, by permitting the membership to be reconfigured between rounds. Flexible Paxos [18] demonstrated that quorums need not be homogenous, and that these quorums may be reconfigured between rounds, affording similar advantages to Vertical Paxos. Fast Flexible Paxos [17], Atlas [12] and Tempo [11] exploit Flexible Paxos' non-homogenous quorums to trade failure tolerance for steady-state performance. ACCORD demonstrates that optimal failure tolerance may be maintained alongside optimal *stability to failure* by permitting the set of quorums that may reach fast-path decisions to be reconfigured, and that by reducing the number of such quorums we may reach fast-path decisions under any tolerated failures.

Transactional systems. Transactional storage systems typically partition data into multiple shards for scalability, and replicate each shard for fault tolerance. Transactions that span across shards hence rely on concurrency control mechanisms to achieve varying degrees of consistency. For performance reasons, however, many systems [3, 7, 10, 19, 27, 39, 46] offer isolation levels short of strict serializability, or accept scalability bottlenecks [14, 47]. ACCORD differs by featuring scalable strict serializable transactions.

Among systems that do offer strict serializable transactions, Spanner [8], CLOCC [1, 26], and Granola [9] build a transaction layer on top of a replication layer, leading to increased latency [45]. Calvin [41] and SLOG [35] similarly rely on a distinct Paxos layer to order multi-shard transactions. Tapir [32] is the first to address the over-coordination

inherent to such architectures, offering strict serializable multi-shard transactions in a single wide area round-trip. Likewise, Janus [32] uses an EPaxos-like approach to provide multi-shard transactions, offering better performance than Tapir under contended workloads, albeit with limited stored procedure semantics and no stability to failure. More recently, Tempo [11] improves upon Janus by leveraging timestamps instead of explicit dependencies, at a cost of ordering non-commutative transactions.

Compared to prior systems that achieve strict serializable multi-shard transactions, ACCORD achieves optimal performance by utilizing real time timestamps and a message reorder buffer. Unlike Tempo, ACCORD does not rely on additional periodic broadcast mechanisms for timestamp stability, and commutative commands do not interfere. Importantly, ACCORD addresses the poor fast-path stability of existing systems by introducing a configurable fast-path electorate. This provides optimal failure tolerance and consistent performance under any number of tolerated failures. ACCORD is the first leaderless protocol that is sufficiently stable for practical use in a large scale industrial database system. Finally, to the best of our knowledge, no commercial or open-source database systems offer strict serializable transactions across regions in a single wide area round-trip.

References

- [1] Atul Adya, Robert Gruber, Barbara Liskov, and Umesh Maheshwari. Efficient optimistic concurrency control using loosely synchronized clocks. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data*, SIGMOD '95, page 23–34, New York, NY, USA, 1995. Association for Computing Machinery.
- [2] Balaji Arun, Sebastiano Peluso, Roberto Palmieri, Giuliano Losa, and Binoy Ravindran. Speeding up consensus by chasing fast decisions. In *47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 04 2017.
- [3] Nathan Bronson, Zach Amsden, George Cabrera, Prasad Chakka, Peter Dimov, Hui Ding, Jack Ferris, Anthony Giardullo, Sachin Kulkarni, Harry Li, Mark Marchukov, Dmitri Petrov, Lovro Puzar, Yee Jiun Song, and Venkat Venkataramani. TAO: Facebook’s distributed data store for the social graph. In *2013 USENIX Annual Technical Conference (USENIX ATC 13)*, pages 49–60, San Jose, CA, June 2013. USENIX Association.
- [4] Apache Cassandra. <https://cassandra.apache.org/>.
- [5] Tushar Deepak Chandra, Vassos Hadzilacos, and Sam Toueg. The weakest failure detector for solving consensus. *J. ACM*, 43(4):685–722, July 1996.
- [6] CockroachDB. <https://www.cockroachlabs.com>.
- [7] Brian F. Cooper, Raghu Ramakrishnan, Utkarsh Srivastava, Adam Silberstein, Philip Bohannon, Hans-Arno Jacobsen, Nick Puz, Daniel Weaver, and Ramana Yerneni. Pnuts: Yahoo!’s hosted data serving platform. *Proc. VLDB Endow.*, 1(2):1277–1288, August 2008.
- [8] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, JJ Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Yasushi Saito, Michal Szymaniak, Christopher Taylor, Ruth Wang, and Dale Woodford. Spanner: Google’s globally-distributed database. In *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, pages 261–264, Hollywood, CA, October 2012. USENIX Association.
- [9] James Cowling and Barbara Liskov. Granola: Low-overhead distributed transaction coordination. In *Proceedings of the 2012 USENIX Annual Technical Conference*, Boston, MA, USA, June 2012. USENIX.
- [10] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Voshall, and Werner Vogels. Dynamo: Amazon’s highly available key-value store. 41(6):205–220, 2007.
- [11] Vitor Enes, Carlos Baquero, Alexey Gotsman, and Pierre Sutra. *Efficient Replication via Timestamp Stability*, page 178–193. Association for Computing Machinery, New York, NY, USA, 2021.

-
- [12] Vitor Enes, Carlos Baquero, Tuanir Rezende, Alexey Gotsman, Matthieu Perrin, and Pierre Sutra. State-machine replication for planet-scale systems. In *EuroSys '20: Proceedings of the Fifteenth European Conference on Computer Systems*, pages 1–15, 04 2020.
- [13] etcd. <https://etcd.io>.
- [14] FaunaDB. <https://fauna.com>.
- [15] Yilong Geng, Shiyu Liu, Zi Yin, Ashish Naik, Balaji Prabhakar, Mendel Rosenblum, and Amin Vahdat. Exploiting a natural network effect for scalable, fine-grained clock synchronization. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 81–94, Renton, WA, April 2018. USENIX Association.
- [16] Hazelcast. <https://hazelcast.com>.
- [17] Heidi Howard, Aleksey Charapko, and Richard Mortier. Fast flexible paxos: Relaxing quorum intersection for fast paxos. In *International Conference on Distributed Computing and Networking 2021, ICDCN '21*, page 186–190, New York, NY, USA, 2021. Association for Computing Machinery.
- [18] Heidi Howard, Dahlia Malkhi, and Alexander Spiegelman. Flexible paxos: Quorum intersection revisited. *CoRR*, abs/1608.06696, 2016.
- [19] Tim Kraska, Gene Pang, Michael J. Franklin, Samuel Madden, and Alan Fekete. Mdcc: Multi-data center consistency. In *Proceedings of the 8th ACM European Conference on Computer Systems*, EuroSys '13, page 113–126, New York, NY, USA, 2013. Association for Computing Machinery.
- [20] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, July 1978.
- [21] Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, May 1998.
- [22] Leslie Lamport. Generalized consensus and paxos. Technical Report MSR-TR-2005-33, March 2005.
- [23] Leslie Lamport. Fast paxos. *Distributed Computing*, 19:79–103, October 2006.
- [24] Leslie Lamport, Dahlia Malkhi, and Lidong Zhou. Vertical paxos and primary-backup replication. Technical Report MSR-TR-2009-63, May 2009.
- [25] Leslie Lamport and Mike Massa. Cheap paxos. In *International Conference on Dependable Systems and Networks (DSN 2004)*, June 2004.
- [26] Barbara Liskov, Miguel Castro, Liuba Shrira, and Atul Adya. Providing persistent objects in distributed systems. In *Proceedings of the 13th European Conference on Object-Oriented Programming, ECOOP '99*, page 230–257, Berlin, Heidelberg, 1999. Springer-Verlag.
- [27] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Don't settle for eventual: Scalable causal consistency for wide-area storage with cops. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles, SOSP '11*, page 401–416, New York, NY, USA, 2011. Association for Computing Machinery.
- [28] Yanhua Mao, Flavio P. Junqueira, and Keith Marzullo. Mencius: Building efficient replicated state machines for wans. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation, OSDI'08*, page 369–384, USA, 2008. USENIX Association.
- [29] MongoDB. <https://www.mongodb.com>.
- [30] Iulian Moraru, David G. Andersen, and Michael Kaminsky. There is more consensus in egalitarian parliaments. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles, SOSP '13*, page 358–372, New York, NY, USA, 2013. Association for Computing Machinery.

-
- [31] Iulian Moraru, David G. Andersen, and Michael Kaminsky. Paxos quorum leases: Fast reads without sacrificing writes. In *Proceedings of the ACM Symposium on Cloud Computing*, SOCC '14, page 1–13, New York, NY, USA, 2014. Association for Computing Machinery.
- [32] Shuai Mu, Lamont Nelson, Wyatt Lloyd, and Jinyang Li. Consolidating concurrency control and consensus for commits under conflicts. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 517–532, Savannah, GA, November 2016. USENIX Association.
- [33] Brian M. Oki and Barbara H. Liskov. Viewstamped replication: A new primary copy method to support highly-available distributed systems. In *Proceedings of the Seventh Annual ACM Symposium on Principles of Distributed Computing*, PODC '88, page 8–17, New York, NY, USA, 1988. Association for Computing Machinery.
- [34] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, pages 305–319, Philadelphia, PA, June 2014. USENIX Association.
- [35] Kun Ren, Dennis Li, and Daniel J. Abadi. Slog: Serializable, low-latency, geo-replicated transactions. *Proc. VLDB Endow.*, 12(11):1747–1761, July 2019.
- [36] RethinkDB. <https://rethinkdb.com>.
- [37] Tuanir França Rezende and P. Sutra. Leaderless state-machine replication: Specification, properties, limits (extended version). *ArXiv*, abs/2008.02512, 2020.
- [38] Fred B. Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Comput. Surv.*, 22(4):299–319, December 1990.
- [39] Yair Sovran, Russell Power, Marcos K. Aguilera, and Jinyang Li. Transactional storage for geo-replicated systems. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, SOSP '11, page 385–400, New York, NY, USA, 2011. Association for Computing Machinery.
- [40] Doug Terry. Transactions and scalability in cloud databases - why can't we have both? In *Proceedings of the 18th International Workshop on High Performance Transaction Systems (HPTS)*, November 2019.
- [41] Alexander Thomson, Thaddeus Diamond, Shu-Chun Weng, Kun Ren, Philip Shao, and Daniel J. Abadi. Calvin: Fast distributed transactions for partitioned database systems. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, SIGMOD '12, page 1–12, New York, NY, USA, 2012. Association for Computing Machinery.
- [42] TiKV. <https://tikv.org>.
- [43] Sarah Tollman, Seo Jin Park, and John Ousterhout. Epaxos revisited. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 613–632. USENIX Association, April 2021.
- [44] YugaByte. <https://www.yugabyte.com>.
- [45] Irene Zhang, Naveen Kr. Sharma, Adriana Szekeres, Arvind Krishnamurthy, and Dan R. K. Ports. Building consistent transactions with inconsistent replication. *ACM Trans. Comput. Syst.*, 35(4), December 2018.
- [46] Yang Zhang, Russell Power, Siyuan Zhou, Yair Sovran, Marcos K. Aguilera, and Jinyang Li. Transaction chains: Achieving serializability with low latency in geo-distributed storage systems. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, SOSP '13, page 276–291, New York, NY, USA, 2013. Association for Computing Machinery.
- [47] Jingyu Zhou, Meng Xu, Alexander Shraer, Bala Namasivayam, Alex Miller, Evan Tschannen, Steve Atherton, Andrew J. Beamon, Rusty Sears, John Leach, Dave Rosenthal, Xin Dong, Will Wilson, Ben Collins, David Scherer, Alec Grieser, Young Liu, Alvin Moore, Bhaskar Muppuna, Xiaoge Su, and Vishesh Yadav. Foundationdb: A distributed unbundled transactional key value store. In *Proceedings of the 2021 International Conference on Management of Data*, SIGMOD/PODS '21, page 2653–2666, New York, NY, USA, 2021. Association for Computing Machinery.