

Thunderbolt: Concurrent Smart Contract Execution with Non-blocking Reconfiguration for Sharded DAGs

Junchao Chen Exploratory Systems Lab University of California, Davis jucchen@ucdavis.edu

Lefteris Kokoris-Kogias Mysten Labs kokoris.kogias.eleftherios@gmail.com

Abstract

Sharding has emerged as a critical technique for enhancing blockchain system scalability. However, existing sharding approaches face unique challenges when applied to Directed Acyclic Graph (DAG)-based protocols that integrate expressive smart contract processing. Current solutions predominantly rely on coordination mechanisms like two-phase commit (2PC) and require transaction read/write sets to optimize parallel execution. These requirements introduce two fundamental limitations: (1) additional coordination phases incur latency overhead, and (2) pre-declaration of read/write sets proves impractical for Turingcomplete smart contracts with dynamic access patterns.

This paper presents Thunderbolt, a novel sharding architecture for both single-shard transactions (Single-shard TXs) as well as cross-shard transactions (Cross-shard TXs), and it enables non-blocking reconfiguration to ensure system liveness. Our design introduces four key innovations: First, each replica serves dual roles as a full-shard representative and transaction proposer (shard proposer), employing differentiated execution models: the Execution-Order-Validation (EOV) model for Single-shard TXs and Order-Execution (OE) model for Cross-shard TXs. Second, we develop a DAG-based coordination protocol that establishes deterministic ordering between two transaction types while preserving concurrent execution capabilities. Third, we implement a dynamic concurrency controller that schedules Single-shard TXs without requiring prior knowledge of read/write sets, enabling runtime dependency resolution. Fourth, Thunderbolt introduces a non-blocking shard reconfiguration mechanism to address censorship attacks by featuring frequent shard re-assignment without impeding the construction of DAG nor blocking consensus. This approach maintains continuous DAG construction and consensus progress while preventing persistent adversarial control through periodic shard reassignment. Thunderbolt achieves a $50 \times$ improvement with 64 replicas over serial Tusk execution.

Keywords

Fault-tolerant System, Blockchain, Consensus, Distributed Transaction, Sharding, Reconfiguration, Concurrent Smart Contract

1 Introduction

The emergence of blockchain technology has spurred significant interest in developing resilient systems capable of processing data and transactions under Byzantine conditions, including software Alberto Sonnino Mysten Labs University College London (UCL) alberto@mystenlabs.com

Mohammad Sadoghi Exploratory Systems Lab University of California, Davis msadoghi@ucdavis.edu

errors, hardware failures, network disruptions, and coordinated malicious attacks [7, 10, 30, 36, 43, 63]. These systems enhance reliability and security by enabling collaboration among multiple independent participants [18, 33, 45, 48, 61, 62, 65, 68, 95].

Smart contracts [87, 88], as programmable transaction frameworks embedded in blockchain platforms, empower developers to address real-world challenges through decentralized solutions [17, 54]. However, their adoption is often hindered by execution delays caused by runtime contract code analysis [5]. To overcome this limitation, researchers are actively exploring performance optimization strategies for contract-based blockchain systems.

Several strategies have emerged to improve execution within blockchain systems.

Transaction Processing Models: One practical approach involves enhancing transaction processing. Most blockchain systems adopt the Order-Execute (*OE*) model, where transactions are ordered through consensus before execution [19, 38, 40, 100]. *OE*-based systems often employ deterministic concurrency controls by constructing transaction dependency graphs to optimize parallelism [31, 69, 94, 99]. However, platforms like Hyperledger Fabric [10] utilize the Execute-Order-Validate (*EOV*) model, executing transactions before consensus, to enhance flexibility and Optimistic Concurrency Control (OCC) [53] to improve the concurrency.

Scalability via DAG and Sharding: Another critical area for advancement in blockchain technology is scalability, particularly in supporting parallel execution. Recent advancements leverage Directed Acyclic Graph (DAG)-based consensus protocols to improve scalability. These protocols enable replicas to submit proposals concurrently by building a DAG that links new proposals to historical ones. This architecture has gained significant recognition in the industry due to its robust security features, exceptional scalability, and capability to support smart contracts [11, 12, 26, 50, 51, 58, 73, 81, 84–86]. Complementary to DAG-based approaches, sharding techniques allow parallel transaction processing across each shard, reducing consensus overhead [4, 27, 41, 44, 52, 83, 98, 105, 106].

However, the above approaches do not effectively improve the execution of smart contracts:

Challenge1: Enhancing Transaction Parallelism with-out advanced knowledge. While *OE*-based solutions leverage dependency graphs to optimize parallelism, they require prior knowledge of transaction read/write sets, a constraint incompatible with dynamic smart contracts. Conversely, *EOV*-based approaches face high transaction conflict rates, necessitating advanced conflict resolution algorithms.

EDBT '26, Tampere (Finland)

^{© 2025} Copyright held by the owner/author(s). Published on OpenProceedings.org under ISBN 978-3-98318-102-5, series ISSN 2367-2005. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.

EDBT '26, 24-27 March 2026, Tampere (Finland)



Figure 1: A transaction rescheduling to avoid abortion by moving the WriteA on T_1 after the WriteA on T_2 . They all obtain the correct result based on their operations.

Challenge2: Efficient Cross-Shard Transaction Processing. Existing solutions for cross-shard atomicity, such as relaybased protocols (Sharper [8], BrokerChain [47], and SharDAG [24]) and traditional Two-Phase Commit (2PC) [27, 46], introduce significant delays due to inter-shard coordination. While multishard consensus [8, 70] mitigates 2PC limitations, it sacrifices scalability in large-scale networks with high contention.

The challenges described above raise the question of whether it is possible to design a sharding system that does not depend on understanding the read and write sets of transactions, nor requires additional coordinators to manage cross-shard transactions.

We propose Thunderbolt, an innovative sharding architecture that seamlessly processes both single-shard (Single-shard TXs) and cross-shard transactions (Cross-shard TXs) without centralized coordinators. Thunderbolt effectively integrates the *OE* and *EOV* models, where the *EOV* model enhances parallelism in the execution of Single-shard TXs, while the *OE* model minimizes the abort rate when handling Cross-shard TXs across different shards. Furthermore, the *OE* model ensures a coordinated execution order between these two types of transactions, maintaining the overall correctness of the execution process.

Similar to conventional sharding systems, Thunderbolt organizes transactions into distinct shards to mitigate potential conflicts. In particular, each replica within Thunderbolt corresponds to a single shard and acts as a shard proposer, proposing transactions within that shard. Thunderbolt employs a DAGbased consensus protocol [11, 12, 26, 50, 51, 58, 81, 84–86] to reach an agreement on the execution results provided by each shard proposer.

Inspired by Sui's epoch switching [15], Thunderbolt employs round-robin scheduling [74] to rotate shard proposers periodically, enhancing the system's security and liveness. Proposer rotation is triggered on-demand if a shard fails to propose transactions within a timeout, with seamless DAG transitions preserving protocol continuity.

Thunderbolt also introduces a concurrent executor (*CE*) designed to improve the execution of Single-shard TXs before reaching consensus. Unlike traditional concurrency protocols that primarily manage conflicts based on the order of arrival [3, 32, 89], the *CE* utilizes a non-deterministic ordering system based on the execution run-time states of each transaction. This innovative approach minimizes the abort rate, thereby reducing transaction latency. As illustrated in Figure 1, the *CE* effectively reschedules transactions based on their run-time executions to prevent aborts. For instance, transaction T_2 , which would ordinarily conflict with the write operation of transaction T_1 , can be successfully committed without cancellation.

In summary, this paper makes the following contributions.

• To our knowledge, Thunderbolt is the first sharding consensus mechanism that combines the *OE* and *EOV* models based on DAG-based protocols without requiring any additional coordinators to determine the order between Single-shard TXs and Cross-shard TXs.

- We introduce a new concurrency paradigm that implements a parallel preplay for Single-shard TXs (concurrent consensus execution). Thunderbolt preplays Single-shard TXs followed by parallel verification without needing prior knowledge of the read/write sets.
- Thunderbolt features a non-blocking shard reconfiguration protocol that allows for the rotation of shard assignments without pausing either DAG dissemination or the consensus layer.
- We have implemented a concurrent executor to enhance the parallelism of executing smart contracts without prior knowledge of the read/write sets on Single-shard TXs. The execution engine dynamically arranges transactions based on current assessments to reduce abort rates due to conflicts.
- Our evaluation of Thunderbolt demonstrates an impressive 50× speedup over Tusk [26] with sequential execution using the SmallBank workload on 64 replicas built on Apache ResilientDB (Incubating) [1, 40].

2 Background

Smart Contract. A smart contract is a digital protocol, introduced by Nick Szabo in the mid-1990s [87], designed to facilitate, verify, or enforce the negotiation or performance of a contract automatically. Unlike traditional contracts, which rely on legal systems for enforcement, smart contracts are self-executing and operate on blockchain technology. They are written in code and run on decentralized platforms like Ethereum, ensuring transparency, security, and immutability. Smart contracts eliminate the need for intermediaries, reduce the risk of fraud, and enable trustless transactions between parties. They are widely used in various applications, including decentralized finance (DeFi), supply chain management, and digital identity verification.

These contracts consist of custom functions that operate on user accounts with associated balances. Once deployed to the network, transactions invoking the functions specified in the contract are proposed to execute predefined operations that interact with the user accounts. However, the execution of the contract code occurs within the Ethereum Virtual Machine (EVM) [59], which results in the read and write sets of the contract being indeterminate prior to execution.

DAG-based protocols. Thunderbolt leverages a Directed Acyclic Graph (DAG) structure to address scalability challenges in blockchain systems. This architecture enables efficient transaction proposal mechanisms while maintaining robust security, high throughput, and native support for smart contracts. DAG-based protocols have gained significant traction in the industry due to their ability to decouple transaction dissemination from consensus processes [11, 12, 26, 50, 51, 58, 73, 81, 84–86].

In contrast to traditional linear blockchains, DAG-based protocols allow multiple replicas to propose transactions concurrently. These transactions are constructed into a deterministic DAG structure, ensuring a consistent topological ordering across all honest replicas. Recent advancements in this domain, including Tusk [26, 85], BBCA-Chain [58], Shoal/Shoal++ [11, 84], Mysticeti [13], and Cordial Miners [51], demonstrate how DAGs streamline consensus by separating data propagation from finality mechanisms.

The protocol operates in synchronized rounds, where each DAG vertex in a round consists of two components:

EDBT '26, 24-27 March 2026, Tampere (Finland)



Figure 2: Overview of a DAG-based protocol on Tusk generated by 4 replicas. Each vertice contains both data and a certificate. Each data refers to 2f + 1 certificates from the previous round. Each certificate is generated from its corresponding data, which has received signatures from 2f + 1 replicas. The leader vertice for round r, the solid vertice used to commit the data in its causal history, will be determined before processing round r + 2.

- Data Payload: Contains transactions and references to at least 2*f* + 1 certificates from the prior round.
- Certificate: A quorum of 2*f* + 1 cryptographic signatures attesting to the validity of the vertex and its dependencies.

During each round, replicas broadcast their proposed vertices to the network. A vertex becomes certified once 2f + 1 signatures are collected, enabling it to serve as a dependency for new vertices in subsequent rounds. This iterative process ensures liveness while preserving the DAG's causal ordering.

Vertex finalization occurs at fixed intervals, typically every two rounds in Tusk [26] or three rounds in DAG-Rider [50]. A designated leader (selected via round-robin scheduling [74] or distributed randomness [16]) proposes a vertex for commitment. A leader vertex in round r is eligible to be committed during round r + 2 (as in Tusk): 1) The replica must have received at least 2f + 1 vertices from round r + 1, and 2) The leader vertex must be referenced by a minimum of f + 1 vertices in round r + 1.

DAG-based protocols provide the following properties:

- Validity: if an honest replica *R* has a vertice *B* in its local view of the DAG, then *R* also has all the causal history of *B*.
- Consistency: if an honest replica *R* obtains a vertice *B_r* in round *r* from replica *P*, then, eventually, all other honest replicas will have *B_r*.
- Completeness: if two honest replicas have a vertice *B_r* in round *r*, the causal histories of *B_r* are identical.

3 Thunderbolt Overview

Thunderbolt advances smart contract execution efficiency through an innovative sharding architecture augmented by a dynamic shard reconfiguration mechanism. This mechanism counters potential censorship attacks, such as post-execution transaction suppression or biased transaction selection, ensuring network integrity.

Unlike conventional sharding systems that partition replicas into isolated groups governed by separate consensus protocols, Thunderbolt employs a single unified consensus protocol jointly maintained by all replicas. Each replica operates as a shard proposer, processing transactions specific to its assigned shard. To coordinate cross-shard transaction ordering and execution, Thunderbolt leverages a DAG-based consensus protocol, enabling global agreement on transaction validity while preserving shardlevel parallelism via the predetermined leaders underlying the consensus.



Figure 3: The dataflow of Single-shard TXs.

Thunderbolt employs both *EOV* model and *OE* model to address Single-shard TXs and Cross-shard TXs.

- Single-shard TXs (*EOV* Model): Transactions confined to a single shard are processed non-deterministically by a Concurrent Executor (*CE*) within the corresponding shard proposer. After the local preplay, results undergo a consensus across all shards to ensure state consistency.
- Cross-shard TXs (*OE* Model): Cross-shard TXs employ an optimistic concurrency control protocol with deterministic finalization. Atomic commitment is achieved post-consensus, allowing tentative execution optimistically while guaranteeing rollback-free confirmation.

Figure 3 demonstrates an example where two shard proposers propose two single-shard transactions. Further details on Single-shard TXs and Cross-shard TXs are provided in Section 4 and 5.

Thunderbolt also allows the migration of each shard to another replica to avoid a censorship attack, such as dropping transactions.

3.1 System, Threat and Data Model

In this section, we describe the system, threat and the data model. We leave the discussion of the Single-shard TX*s*, the Cross-shard TX*s*, and the shard reconfiguration in Section 4, 5, and 6.

System Model. Thunderbolt is composed of *n* replicas, each of which serves a dual purpose: functioning as a shard and a replica. As a shard, a node maintains a distinct set or partition of data. In its role as a replica, it preserves a copy of the transaction log. Thus, each replica may also be designated as a shard proposer. Furthermore, clients direct transactions to the appropriate shards, and every replica engages in the consensus process to establish a cohesive order for these transactions.

In summary, a node in Thunderbolt assumes three essential roles: 1). It operates as a shard proposer, managing Single-shard TXs within each shard in an independent manner. 2). It acts as a replica that contributes to the consensus process. 3). It serves as a leader that commits Cross-shard TXs in a total order in accordance with the consensus protocol. For clarity in the accompanying illustrations, the terms "replicas" and "shard proposers" may be used interchangeably. We will utilize the term "shard proposers" when delineating shard procedures and will revert to "replicas" in discussions regarding the consensus protocols.

Threat model. We consider a set of *n* replicas (or shards), where at most *f* of these replicas can be faulty and n = 3f + 1. The *f* faulty replicas may exhibit any arbitrary behavior, including Byzantine failures, while the remaining replicas are assumed to be honest and will adhere to the protocol's specifications at all times. Additionally, we assume that clients are not trustworthy and would not expect to send transactions to all of the shards associated with their transactions. The network is expected to be eventually synchronous [29], meaning that messages sent from a replica will eventually arrive within a global stabilization time

(*GST*), which remains unknown to the replicas. Communications between replicas use authenticated point-to-point channels, with messages signed by the sender using a public-private key pair for authentication.

Data model. The data model assumes that each transaction includes a contract code with functions to access data belonging to the sender in the shard. The contract involves two types of operations: <*Read*, K> and <*Write*, K, V>. Here, K represents the key required for access and V is the value that needs to be written to the key K. The contract code is Turing-complete and users could not obtain any information without execution. We also assume that the functions of the contract are idempotent.

Our system is designed with the understanding that data must be partitioned and that each key is assigned a shard ID (*SID*) before it can be utilized. These *SIDs* are predefined and recognized across all shards. They fulfill a dual function: they guide transactions to the appropriate shard proposer and support parallel processing among multiple shards, thereby enhancing overall system efficiency (Section 5). The actual method for partitioning is orthogonal to our work and any existing techniques can be utilized [20–22, 64, 104].

4 Single Shard Transactions

Thunderbolt processes Single-shard TXs through three core components: preplay, execution scheduling, and validation. During each round, a shard proposer initiates the workflow by preplaying a batch of Single-shard TXs. This generates a block containing critical preplay outcomes, which the proposer propagates to other shards via a DAG-based consensus protocol. During consensus, these blocks undergo parallel validation across shards. Once a replica commits a block, it applies the preplay results to its storage.

Preplay. In Thunderbolt, shard proposers play a pivotal role by preplaying transactions to determine their outcomes before block creation. A concurrent executor (*CE*) is employed to preprocess batches of transactions efficiently, producing detailed outputs for each transaction. These outputs include: the read/write sets accessed during execution, the corresponding execution results, and a scheduled execution order (as depicted in Figure 3).

The scheduled order establishes a deterministic serialized sequence, ensuring transaction results remain consistent when executed in the prescribed order. The read/write sets reveal the specific data accessed by each transaction. Crucially, these sets cannot be predetermined and are derived exclusively via the preplay process.

Execution Scheduling. Thunderbolt integrates with any DAGbased dissemination layer that employs a consensus protocol to establish a total block order across replicas. In each round r, a shard proposer R proposes a *CE*-generated block to the DAG, creating a new vertex in the graph. This vertex links to all prior vertices, including those proposed by R in round r - 1. For clarity, "vertices" in the DAG are hereafter referred to as "blocks" in the following sections, implying they have been certified by the protocol.

Validation. When a replica receives data for round r via the DAG (Section 2), Thunderbolt initiates a rigorous validation process to verify the integrity of preplay results within the blocks. Validators construct a local dependency graph using the read/write sets to enable parallel transaction validation rather than



Figure 4: An example illustrates the commitment of each leader in the process of converting Single-shard TXs to Cross-shard TXs on Tusk. In this DAG, Si denotes a Single-shard TX, while $Ci(\{X\})$ represents a Cross-shard TX associated with shards $\{X\}$. The leaders in the odd rounds are selected using round-robin selection. The blocks committed by the same leader have borders of the same color. Si will be converted to Ci if there is any conflict blocks (S10) or could not receive the leader block in time (S24).

sequential checks, optimizing system throughput. Notably, blocks from round r - 1 are validated before those from round r for the same shard proposer.

During re-execution, validators confirm that the computed read sets match the values recorded in the block. A valid dependency graph guarantees consistent read-set results and ensures the final state of each key aligns with the block's declared values. If discrepancies in read-set values are detected, the block is flagged as invalid and discarded. Until blocks are committed, replicas retain preplay results in local storage, either to process Cross-shard TXs (Section 5) or until DAG reconfiguration occurs (Section 6).

5 Cross-shard transactions

Cross-shard TXs involve multiple shards and require consensus to establish a total execution order. This total order ensures that all Cross-shard TXs are executed consistently across the involved shards. However, each shard proposer preplays Single-shard TXs independently and replicates the results; in contrast, the total order for the Cross-shard TXs must be determined first before they can be executed (for instance, preplay optimization cannot be employed for Cross-shard TXs). Therefore, Thunderbolt must coordinate the sequencing between the Cross-shard TXs and the Single-shard TXs to guarantee consistent execution outcomes across all replicas. Prior approaches often rely on coordinators to establish inter-shard order [8, 24, 27, 46, 47, 52, 70, 105], but these introduce communication overhead that degrades performance.

5.1 Rules for Proposals

To address issues, Thunderbolt leverages the DAG's predetermined leaders to enforces a consistent partial order between single-shard and cross-shard transactions via the following rules:

- G1) If a leader *L* commits both a Single-shard TX and a Crossshard TX, the Single-shard TX must be committed first.
- G2) If leader L_i commits Cross-shard TX X in round *i*, any Single-shard TX Y committed by leader L_j in round *j* (where j > i) cannot execute until X is finalized.

It is worth knowing that Leaders are predetermined per round using deterministic methods (e.g., round-robin or global random coins). To enforce these rules, Thunderbolt applies the following proposal rules:

- P1) Cross-shard TXs are submitted directly to the DAG, bypassing the *CE*.
- P2) Leaders committing a batch of transactions must finalize all Single-shard TXs before Cross-shard TXs.

onuru i						
S1	C2 (1,2)		S11	S15	S19	S23
Shard 2	()	⊈\	/	1 W.	/ \X/	Timeout
S2 ₩	S4	58	\$12	S16	S20	S24->C9 (2)
Shard 3	(X	/X	Ύ ₩	7 — X	(/)/	
S3	S5	S9	S13-> Skip	S17	S21	S25
Shard 4		ι	()A	λ	<u>ال</u>	
C1 (3,4)	S6	S10-> Skip	S14-> Skip	S18	S22	S26
Round 1	Round 2	Round 3	Round 4	Round 5	Round 6	Round 7

Figure 5: An example of proposing skipping blocks to restart the preplay of the Single-shard TXs. The Single-shard TXs (S17,S21, and S22), which would be converted to Cross-shard TXs after round 5 in Figure 4, can replay their executions before delivering to the DAG.

- P3) If a shard proposer SL proposes a Single-shard TX X in round r, and the current round's leader L differs from SL, SL must:
 - Wait for *L*'s proposal before preplaying *X*.
 - Convert *X* to a Cross-shard TX if any uncommitted Cross-shard TX *Y* in *L*'s history conflicts with *X*.
 Otherwise, preplay *X* and submit the results.
- P4) If a shard proposer SL proposes a Single-shard TX X in round r, and a prior leader's uncommitted Cross-shard TX Y (in round q < r) conflicts with X, SL converts X to a Cross-shard TX.
- P5) If leader *L* in round *r* commits a Cross-shard TX *X* related to shard A but lacks A's proposal in round r 1, *L* defers committing A and A's subsequent proposals.
- P6) If a shard proposer *SL* proposes a Single-shard TX *X* in round *r* but the leader *L*'s proposal is delayed beyond a timeout, *SL* converts *X* to a Cross-shard TX.

These rules enable Thunderbolt to process Cross-shard TX*s* without blocking shards while maximizing parallelism for singleshard transaction preplay.

EXAMPLE 1. Figure 4 illustrates Thunderbolt's handling of singleshard and cross-shard transactions on Tusk: Single-shard TXs will be executed before Cross-shard TXs under the same leader (Rule P2). S10 is converted to Cross-shard TX C1 due to its dependency on S9's leader history (Rule P3). S13 and S14 become C4 and C5, respectively, because C1 remains uncommitted until round 5 (Rule P4). Leader S18 in round 5 skips C2 and subsequent transactions due to missing S11 (Rule P5). In round 7, S24 converts S4 to Cross-shard TX C9 after failing to receive leader S23's proposal (Rule P6).

5.2 Parallel Execution

When processing Cross-shard TXs, Thunderbolt preserves all sharding metadata for each transaction. Rather than sequential execution, Thunderbolt employs deterministic concurrency control mechanisms, such as QueCC [69], to construct dependency graphs using cross-shard metadata (SID). This enables parallel execution while maintaining consistency across shards.

5.3 Message Failures

In practical deployments, network unreliability may delay message delivery. To mitigate this:

• If a leader *L* cannot include all Single-shard TX*s* linked to a Cross-shard TX*s* due to network delays (e.g., missing shard proposals), *L* bypasses the affected Cross-shard TX and subsequent transactions from the same shard (*C*2 in Figure 4 on Rule P5). This prevents violations of global ordering G2 by excluding incomplete transaction sets. These transactions are later finalized by subsequent leaders.

• If a shard proposer fails to receive the leader's proposal within a round's timeout window, it cannot preplay its Single-shard TX (Rule P6). The proposer instead promotes the transaction to a Cross-shard TX and submits it directly to the DAG (such as *S*24 in Figure 4).

5.4 Preplay Recovering

Under Rule P3, a shard proposer *SL* must convert a Single-shard TX to a Cross-shard TX if it detects conflicting uncommitted Cross-shard TXs in prior leader histories. While this ensures safety, it forfeits the performance benefits of preplay, such as the blocks in shard 3 after round 3 in Figure 4.

To recover preplay, *SL* must verify that all conflicting Crossshard TXs have been finalized by preceding leaders. Since a leader CL_r in round r is finalized within two subsequent rounds (Section 2), *SL* can safely preplay new single-shard transactions once: 1) It receives 2f + 1 certificates in round r + 1, and 2) CL_r is referenced by at least f + 1 blocks in round r + 1. If *SL* identifies any conflicting Cross-shard TXs while proposing Single-shard TXs, *SL* submits skip blocks to the DAG until prior leaders are finalized. For instance, *S*10, *S*13, and *S*14 in Figure 5 are converted to skip blocks. Consequently, subsequent transactions, like *S*17 and *S*22, are reverted to the *EOV* model, restoring preplay capabilities.

6 Shards Reconfiguration

In Byzantine environments, compromised replicas may enable attackers to launch censorship attacks, threatening the integrity of transactions within their assigned shards.

Thunderbolt employs a round-robin selection mechanism [80] to rotate shard proposers if a leader fails to propose transactions for *K* rounds or at fixed intervals of K' rounds (where K' > K). This rotation serves dual purposes:

- Preventing transaction duplication (DDOS [28, 60]): Proposers perform local deduplication to block malicious clients from flooding the system with redundant transactions, a known challenge in DAG-based protocols [12, 15, 85].
- Mitigating censorship: Regular rotation limits the window for a compromised proposer to disrupt operations.

Unlike traditional consensus protocols that depend on notification messages to alter primary replicas, Thunderbolt introduces an innovative mechanism that uses the underlying DAG protocols to facilitate the seamless transition to a new DAG and reconfigure shard proposers. We leverage a round-robin approach to select a new proposer that if the current proposer of shard X is replica R_i , the subsequent proposer of shard X will be $R_{(i \mod n)+1}$.

However, the transmission of blocks to a new proposer may experience delays or omissions due to network issues or the actions of a malicious proposer. If the new proposer for round r is unable to receive the proposal committed in round r - 1 from the previous proposer, the new proposer will stop operations until the block arrives to ensure safety.

To address this challenge, Thunderbolt introduces Shift blocks to reach agreements among shards when a shard reconfiguration should be initiated and switch to a new DAG to process further transactions.

A replica R in a shard broadcasts a Shift block in round r under the following conditions:

- (1) *R* receives no block from a shard proposer after round r K.
- (2) *R* has proposed blocks for at least K' rounds.



Figure 6: An example of DAG reconfiguration. Shard 1 in DAG 1 which is a malicious replica that delays the blocks from round 2, which triggers a reconfiguration to a new DAG (K = 2). Additionally, a further reconfiguration to DAG 2 will take place after the replicas have proposed six rounds (K' = 6).



Figure 7: The architecture of the *CE* consists of a set of executors that execute transactions. The concurrency controller utilizes a dependency graph to determine the order of transactions and their execution results.

- (3) *R* received f + 1 Shift blocks from distinct replicas at round r 1.
- (4) R does not broadcast the Shift block before.

EXAMPLE 2. In Figure 6 where K = 2 and K' = 6, shards 2 and 3 broadcast Shift blocks in round 4 after failing to receive blocks from shard 1 in rounds 2–3. Despite receiving a block in round 4, shard 4 broadcasts a Shift block in round 5 upon receiving two Shift blocks from peers, prioritizing liveness assurance.

Non-blocking Reconfiguration. Since at 2f + 1 honest replicas will commit the same block during the same round (Section 2), we designate the round of the first commit block that includes 2f + 1 Shift blocks as the ending round for the current DAG. Then, each replica will begin a new DAG from the same ending round to ensure the system's safety. For instance, shard 2 will propose block 15 at round 5 after proposing a Shift block at round 4. Finally, block 19 from shard 3 at round 6 is selected as the leader during the consensus process. Subsequently, shard 3 commits all historical blocks from block 19, including the Shift blocks from other shards. At this point, all shards will transition to the new DAG (DAG 2) and start executing transactions within the new shard. Moreover, each replica will propose a Shift block every K' = 6 blocks in the new DAG (DAG 2) to facilitate a transition to the next DAG. This measure is intended to protect against censorship attacks, which may involve dropping transactions, failing to propose blocks, or prioritizing certain transactions over others. Additionally, the non-blocking mechanism provides protection against malicious proposers as the reconfiguration process requires a minimum of 2f + 1 Shift blocks to be effective.

Uncommitted Transactions. Due to the two-round leader commitment latency, transactions uncommitted in the ending round of a DAG are discarded and must be resubmitted. Only transactions from the last two rounds or those excluded by the leader are affected. Clients will automatically retransmit transactions lost due to the reconfiguration.

Censorship Attacks. Malicious replicas may attempt censorship via message drops, transaction rescheduling, DDoS attacks, or proposal halting [28, 60]. As each replica governs an entire shard in Thunderbolt, such attacks can paralyze specific shards.

Time	Transactions	Operations	Dependencies	Execution Order
0	Initial DB	D = 3	ł	{}
1	$T_1:(W, D, 3)$	T_1 writes $D = 3$	$\{T_1\}$	{}
2	<i>T</i> ₂ :(R, <i>D</i> , 3)	$T_2 \text{ reads } D \text{ on } T_1:$ $(D = 3)$	$\{T_1 \rightarrow T_2\}$	{}
3	<i>T</i> ₃ :(R, <i>D</i> , 3)	$T_3 \text{ reads } D \text{ on } T_1:$ $(D = 3)$	$\{\begin{array}{c}T_1 \to T_2 \\ T_1 \to T_3\end{array}\}$	{}
4	<i>T</i> ₃ : Commit	Wait for T_1	$\{\begin{array}{c}T_1 \to T_2 \\ T_1 \to T_3\end{array}\}$	{}
5	<i>T</i> ₁ :(W, <i>D</i> , 5)	T_1 writes $D = 5$. Abort T_2, T_3	$\{T_1\}$	{}
6	$T_3:(\mathbb{R}, D, 5)$ (Re-execute)	$T_3 \text{ reads } D \text{ on } T_1:$ $(D = 5)$	$\{T_1 \rightarrow T_3\}$	{}
7	T_1 : Commit	Commit T_1	$\{T_1 \rightarrow T_3\}$	$\{T_1\}$
8	T ₃ : Commit	Commit T ₃	$\{T_1 \rightarrow T_3\}$	$\{T_1, T_3\}$
9	<i>T</i> ₂ : (W, <i>D</i> , 3)	Invalid and re-execute		
10	$T_2:(\mathbb{R}, D, 5)$ (Re-execute)	$T_2 \text{ reads } D \text{ on } T_1:$ $(D = 5)$	$\{T_2\}$	$\{T_1, T_3\}$
11	<i>T</i> ₂ : (W, <i>D</i> , 2)	T_2 writes $D = 2$	$\{\begin{array}{c}T_1 \to T_2\\T_1 \to T_3\end{array}\}$	$\{T_1, T_3\}$
12	T ₂ : Commit	Commit T ₂	$\{\begin{array}{c}T_1 \to T_2\\T_1 \to T_3\end{array}\}$	$\{T_1, T_3, T_2\}$

Table 1: An example of generating the dependency while executing the transactions $\{T_1, T_2, T_3\}$ that access the data D and determining the execution order.

The reconfiguration mechanism counters this by periodically reassigning shards to new replicas, limiting the impact window of compromised replicas.

7 Concurrent Executor

The concurrent executor (*CE*) is a critical component enabling Thunderbolt to process Single-shard TXs concurrently during the preplay phase. As a nondeterministic concurrency control executor, *CE* generates a serialized execution order, read/write sets, and execution results for transaction batches. These outputs allow any replica to independently verify correctness, even though the *CE*-derived order may differ from the transactions' arrival sequence.

The architecture of *CE* is illustrated in Figure 7, where a set of executors executes the transactions, and a concurrency controller (*CC*) determines the execution results among the transactions.

Transactions progress through a two-phase data flow: an execution phase (operations processing) and a finalization phase (commit/abort decisions). Table 1 exemplifies this workflow using transactions $\{T_1, T_2, T_3\}$.

7.1 Execution Phase

During the execution phase, the executors access the data within *CC* directly and *CC* maintains a dependency graph to keep track of the relationship between transactions and all the results are stored in the graph directly to avoid accessing the disk IO. The critical characteristic of *CC* is that *CC* only maintains the graph based on the current operations among the transactions without requiring any read/write set knowledge. Furthermore, *CC*

Thunderbolt: Concurrent Smart Contract Execution with Non-blocking Reconfiguration for Sharded DAGs

EDBT '26, 24-27 March 2026, Tampere (Finland)



Figure 8: Dependency Graph on *Thunderbolt*. Edges with the same color represent a dependency graph with a specific key.

mechanism is nondeterministic, which means it can arrange transactions in any order based on their current states. The order of two transactions will not be established until a dependency is created. For instance, in the case of two conflicting transactions, denoted as T_1 and T_2 , which both only write the same key, either transaction may be ordered first. Both execution orders, $[T_1, T_2]$ and $[T_2, T_1]$, are considered valid. A dependency is established based on the commit times of these transactions. If T_1 commits before T_2 , the execution order $[T_1, T_2]$ indicates that T_1 is prioritized over T_2 . Additionally, a third transaction, T_3 , can influence this dependency. A dependency is formed if T_3 reads the value following T_2 's write before T_1 and T_2 commit, resulting in a unique execution order of $[T_1, T_2, T_3]$. As a result, a dependency indicator will be generated for the two conflicting transactions once a read-write conflict arises or when both transactions have been committed.

While receiving an operation from a transaction T sent by the executors identified accessing the key K, denoted as O_k , CCchecks the relationships among the transactions. If T conflicts with other transactions or has been aborted by other transactions, the operation O_k will not be considered valid. For instance, T_2 at time 9 in Table 1 is an example of an invalid operation, as it was aborted by T_1 at time 5 due to its outdated read in D. Transaction Twill be aborted in such cases and require reexecution. Otherwise, if the operation O_k is valid, it will be added to the dependency graph (section 8.1) and obtain the operation result, such as the value V that O_k intends to read. We can obtain the operation results from other transactions directly based on the dependency graph to allow reading uncommitted data, such as T_2 reads D on T_1 at time 2.

7.2 Finalization Phase

During the finalization phase, the executor informs *CC* that the executor has completed all the operations. Then *CC* will update the results to the storage asynchronously once all its dependencies have been committed and assign the execution order to the transactions. If *CC* has terminated the transaction due to conflicts with other transactions, *CC* aborts the transaction and notifies the executor to restart the execution.

8 Dependency Graph in CC

This section explains the dependency graph G, which is central to the CC component. It plays a crucial role in maintaining the causal relationships between transactions during the preplay process in CE. The CC component ensures that the sequential order of execution defined by G is valid.

8.1 Dependency Graph Construction

A Dependency Graph is a graph G(V, E) that plays a crucial role in tracking the causal relationship between transactions in *CC*. Each node $v \in V$ represents a specific transaction. Additionally, each edge $e(u, v, k) \in E$ indicates a connection between two



Figure 9: An example of incorporating operations into the dependency graph and adjusting the graph to ensure correctness. a) T_4 writes A=3. b) T_4 reads A. c) T_4 reads key A on its existing node and retrieves the value from T_3 .

transactions u and v on a key K. This relationship is represented as $u \rightarrow kv$. For example, in Figure 8, transaction T_5 generates an edge $e(T_3, T_5, D)$ from T_3 because T_5 acquires the value 3 of key D from T_3 . Without loss of generality, we have assigned a root node denoted R and added edges $e(R, u, k) \in E$ for each $u \in V$ that accesses the key K but does not have any incoming edge on key K, such as T_7 and T_8 in Figure 8.

If the graph G is acyclic, a sequential order can be established by generating a topological order. It is crucial that every transaction must obtain the same causal order in any topological order from G to ensure consistency. Therefore, G is considered a valid graph only if any sequential order generated from the topological order is a valid serialization order and produces the same results. By following any correct order, all transactions will yield the same execution results. However, because of the nondeterministic characteristic, the results may not be the same as those executed in their arrival order. Each node u maintains all records of the operations triggered by a transaction u, including the resulting values. The sequence of the linking nodes establishes the order of commitment between the two transactions.

Since a transaction is an atomic commitment, we combine the internal operations to simplify the in-node states. However, to trace the conflicts between two nodes, we must retain the first operation if it is a read and the last operation if it is a write, to ensure that the causal relationship is not lost. Thus, we remain at most two operations in the nodes: the first read and the last write.

To help illustrate the algorithm, we define the types of each node depending on the operations it contains on a key:

- A node $v \in V$ is a read node R_v^k if the first operation on key *K* is a read.
- A node $v \in V$ is a write node W_v^k if v contains write operations on key K.
- The root node *R* is a write node.

8.2 Generating New Nodes

This section presents the process of adding operations from a new transaction to the dependency graph G.

CC creates a new node whenever an operation O_k is received from a new transaction *T*. If O_k is a write operation, *T* needs to establish a connection to each casual relation. To avoid pointing to the root and assuming that the earlier transaction will commit first, the non-write nodes *v* on key *K*, which only contains reads, without any outgoing edges (not dependent by other nodes) are selected, and edges e(v, u, k) are added, pointing to *u* (Figure 9 (a)).

On the other hand, when a read operation is performed, CC selects the latest write node u to obtain the latest value or selects the root to read the data value from storage if no write nodes exist. If the write node u is selected, we need to make all other



Figure 10: Cycle of conflicts and cascading aborts. a) T_1 reads *B* and adds a dependency from T_3 following the rules in Section 8.2. b) Cascading aborts from T_1 since T_1 wants to write *A* that breaks the read on T_2 .

write nodes W_v^k contain a path to u to guarantee the correctness for the read after write between u and v. Finally, the operation and its result <Type, Key, Result> will be written into the node u. An example of adding a new read operation on A from T_4 is depicted in Figure 9 (b). T_4 selects T_3 to read to obtain A = 3 and adds an edge from T_3 and a record <R, A, 3> is logged down in the node. Then T_3 will add two edges from T_1 and T_2 , respectively.

8.3 Operations on Existing Nodes

When receiving an operation O_k for the key K from an existing transaction T in G, CC will select the corresponding node u to attach the record. If O_k is a read operation, the result will be directly retrieved if u contains the record for key K. Otherwise, it will proceed with the new node operation as specified in Section 8.2 to choose a previous one to access the value. Figure 9 (c) illustrates an instance where T_4 reads key A as its second operation and retrieves the value from T_3 . If O_k is a write operation, the operation will be appended to the node.

8.4 Conflict Detection

Appending the records to an existing node may lead to transaction conflicts. For instance, a transaction updates the value again but it has been read by another transaction or a dependency cycle is created due to the dependency on another key since we always find the latest write to retrieve the value. Figure 10 (a) depicts a scenario in which T_1 attempts to retrieve the value of *B* from T_3 , which has established a dependency from T_1 due to key *A*, which results in the creation of a dependency cycle. In this case, *CC* will try to read the value from its ancestor, like *B* reads the value from the *Root* Figure 10 (a). If there is still any conflict with other transactions, *CC* will trigger the abort process.

Once conflicts are detected, *CC* triggers an abort process as follows:

(1) If u only contains read operations, abort T itself.

(2) If u contains write operations, cascading abort from T.

In Figure 10 (b), we need to abort T_2 and T_3 since T_1 contains a write operation. However, in Figure 10 (a), we only need to remove T_1 and keep T_3 alive.

9 Thunderbolt Correctness Analysis

In this section, we conduct an analysis of the safety and liveness properties of Thunderbolt. Safety is defined as if two conflicting transactions, T_1 and T_2 , are executed within an honest replica in a specific order O (for instance, $T_1 < T_2$), it is expected that all other honest replicas will also execute T_1 and T_2 in the same order O. Liveness, on the other hand, is characterized by the assurance that client requests will consistently receive a response.

Proof of Safety. We prove the safety of Thunderbolt. We will first analyze the safety **within the same DAG**.

THEOREM 1. In the case of two conflicting transactions, T_1 and T_2 , which may occur as either Single-shard TXs within the same

shard or as Cross-shard TXs, if T_1 is executed prior to T_2 in an honest replica, any other honest replica will also execute T_1 before T_2 .

PROOF. The DAG protocol ensures that the sequence of two blocks proposed by the same proposer aligns with their commit order. Consequently, the order of two Single-shard TXs proposed by the same proposer will be preserved. Furthermore, the order of Cross-shard TXs is established by the consensus, resulting in a consistent global order. This leads to uniform execution of transactions T_1 and T_2 across all honest replicas, irrespective of whether the transactions are Single-shard TXs or Cross-shard TXs.

THEOREM 2. In the scenario where two conflicting transactions occur, where T_1 is a Single-shard TX and T_2 as a Cross-shard TX, if an honest replica executing T_1 prior to T_2 , any other honest replicas will execute T_1 before T_2 as well.

PROOF. Consider a contradiction scenario in which replica R_1 executes T_1 prior to T_2 , while replica R_2 executes T_2 before T_1 . We also suppose leader X commits T_1 and leader Y commits T_2 . Then it can be shown that R_1 is the proposer of T_1 and T_1 is not converted into a cross-shard transaction. Since R_2 executes T_2 ahead of T_1 , the round of leader Y must be less or equal to the round of leader X ($Y \le X$). However, if the round r to T_1 occurs before Y (r < Y), then T_2 cannot be committed (Rule P5). Conversely, if $r \ge Y$, then according to Rule P4, T_1 should be converted into a Cross-shard TX if T_2 remains uncommitted. Thus, it is impossible for R_1 to execute T_1 before T_2 while R_2 executes T_2 before T_1 .

THEOREM 3. When two conflicting transactions, T_1 and T_2 , are executed by honest replicas in the order of $T_1 < T_2$ across two different DAGs, it is ensured that all other honest replicas will execute T_1 and T_2 in the same order within a single DAG.

PROOF. Section 6 illustrates that all honest replicas will transition to the new DAG starting from the same ending round. Refer to the consistency and completeness property of the DAG (Section 2), all honest replicas will execute transactions T_1 and T_2 within the same DAG.

Consequently, we can draw the conclusion that: If two transactions are isolated, they can be executed in any order. Otherwise, all the honest replicas will execute them in the same order. Thus, Thunderbolt holds its safety guarantees.

Proof of Liveness. In an environment where all replicas operate effectively, they will propose blocks within the same DAG. Each block suggested by the respective shard proposers will ultimately be committed. When a malicious replica is identified, the honest replicas will respond by proposing a Shift block. If fewer than 2f + 1 Shift blocks are proposed, the DAG will remain unchanged, and all replicas will continue functioning within the current DAG while proposing new blocks. Conversely, if there are 2f + 1 Shift blocks, all honest replicas will transition to the new DAG within the same round, as detailed in Section 2. After a minimum of 2f + 1 honest replicas successfully relocate to the new DAG, they will be empowered to propose new blocks. Provided that all replicas maintain proper behavior, they will consistently propose blocks within the same DAG, ensuring that each shard proposer's proposed blocks are duly committed.

Thunderbolt: Concurrent Smart Contract Execution with Non-blocking Reconfiguration for Sharded DAGs

10 Correctness of serializability on CC

In this section, we will present the proofs that establish the correctness of *CC*. We will start by defining key concepts related to serializability, followed by a detailed analysis of our findings regarding correctness.

We consider a set of transactions $T = \{T_i\}$ (where $1 \le i \le n$) alongside a predefined commit order *CO*. *CC* generates a sequential order denoted as $SO = [T_1, T_2, ..., T_n]$, producing an outcome represented as $OUT = [OUT_1, OUT_2, ..., OUT_n]$ that each transaction reads and writes some values on some keys. Let *SE* denote one of the possible sequential orders of *SO*, with *OUT'* reflecting its corresponding outcomes. *CC* is deemed serializable if the condition OUT = OUT' holds true.

Definition 4 (Read-Complete). If T_i reads a value updated by T_j in SO, T_i will also read the value updated by T_j in SE. If transaction T_i reads a value updated by a transaction T_j in SO, then T_i will also read the value updated by T_j in SE.

Definition 5 (Write-Complete). If transactions T_i and T_j both write new values to key K, but T_i commits before T_j , which generates an order $T_i < T_j$ in SO, then T_i will also write the values to K before Tj when in SE.

THEOREM 6. CC is considered both Read-Complete and Write-Complete when the dependency graph G is always valid.

PROOF. Firstly, if G is valid, we know that when a read node R_n^k (representing a node u that reads on key K) retrieves a value from a write node W_{μ}^{k} on key K, all the write nodes that are updating values on K either have a direct path to u or a path from v. This guarantees the correctness of read-after-write operations. Consequently, if transaction T_i reads values updated by transaction T_i on key K, where T_i always exists (with the root being a write node), it follows that no other transactions updating values on K will be located between T_i and T_j . Therefore, for any execution order SO generated by G, where T_i reads the value updated by T_i , it can be inferred that T_i will also read the same value from T_i in SE. This demonstrates that CC is Read-Complete. Secondly, since SO is an execution order in CC, if T_i commits before T_j , then T_i will precede T_j in SO. Consequently, T_j will update the values after T_i in SE as well. Therefore, we can conclude that CC is Write-Complete.

THEOREM 7. CC is serializable if CC is both Read-Complete and Write-Complete.

PROOF. Since *CC* is Read-Complete, every transaction T_i that reads the values of certain keys will retrieve the same values regardless of the order generated by *G*. Additionally, since *CC* is Write-Complete, every transaction T_i that updates the values of certain keys will maintain a consistent order among all the orders generated by *G*. Therefore, the transactions will yield the same outcomes in any generated orders *SO* and *SE* (*OUT* = *OUT'*).

11 Concurrency Executor Evaluation

This section evaluates Thunderbolt by assessing its performance on the *CE* and the Thunderbolt protocol. We implement all the baseline comparisons using Apache ResilientDB (Incubating) [1, 40]. Apache ResilientDB is an open-source incubating blockchain project that supports various consensus protocols. It provides a fair comparison of each protocol by offering a highperformance framework. Researchers can focus solely on their protocols without considering system structures such as the network and thread models.

We will begin by comparing *CE* with two baseline protocols: *OCC* [53] and 2PL-No-WAIT [82]. Additionally, we will analyze the performance of Thunderbolt, which is built on Tusk [26], and we will also use Tusk as a baseline for our comparisons. For our input workload, we will utilize SmallBank [2, 6], a benchmarking suite that simulates common asset transfer transactions. This suite is also used to evaluate variant block systems [36, 55, 57, 67, 75, 90, 107].

11.1 Baseline Protocols

We implement OCC [53] and 2PL-No-WAIT [102, 103] to compare the performance against our concurrent executor. We set up our experiments on AWS c5.9xlarge consisting of 36 vCPU, 72GB of DDR3 memory. We use LevelDB as the storage to save the balance of each account.

OCC. Each executor is responsible for locally executing transactions. When an operation within a transaction T requires reading the value of a key K that the executor has not previously accessed during the execution, the executor will retrieve the value from the storage. Each value also contains a version to indicate the time the value was obtained. Any write operation will update the values locally. Upon completion of T, all the updated values will be forwarded to a central verifier. The verifier will cross-check the value versions by comparing them with the current versions in the storage. If there is a mismatch, the commit will be rejected, necessitating the re-execution of T.

2PL-No-WAIT. Each executor performs transactions by directly accessing the storage through a central controller. When an operation within a transaction T requires the read or update of a key K, the controller will lock K to prevent conflicts. If an operation seeks to access K but discovers that another executor has locked it, the executor will release all locks and re-execute T. Upon the completion of T, all the results will be transmitted to storage, and all locks will be released.

11.2 Experiment Setup With Smallbank

SmallBank [2] is a transactional system that comprises six distinct transaction types, five of which are designed to update account balances, while the remaining transaction is a read-only query that retrieves both checking and saving the account details of a user. Our focus is on two types of transactions: SendPayment and GetBalance, which are used to transfer funds between two accounts and retrieve account balances, respectively. Our objective is to evaluate the performance under varying read-write balance workloads. During a SendPayment transaction, the account balances are updated by reading the current balance and then writing the new values back. We created 10,000 accounts and conducted each experiment 50 times to obtain the average outputs.

We evaluated the impact of parallel execution. We measured the performance by uniformly selecting GetBalance with a probability of P_r while SendPayment with $1 - P_r$. We follow a Zipfian distribution to select accounts as transaction parameters and set the Zipfian parameter θ . The value of θ determines the level of account contention, with higher values leading to higher contention. We focus only on data workloads with high contention by setting $\theta = 0.85$.



- Thunderbolt-b500 - Thunderbolt-b300 - OCC-b500 - OCC-b300 - 2PL-No-WAIT-b500 - 2PL-No-WAIT-b500

Figure 11: Evaluation of CE on different numbers of executors.

11.3 Impact from Concurrent Executor

We first evaluate the impact of increasing the number of executors to execute the transactions, then measure the aborts produced by each protocol. We ran two batch sizes *b*300 and *b*500 for each protocol: Thunderbolt-b300, Thunderbolt-b500, OCC-b300, OCC-b500, 2PL-No-WAIT-b300, and 2PL-No-WAIT-b500. We set $P_r = 0.5$ to measure a read-write balanced workflow and $P_r = 0$ on an update-only workflow.

Number of Executors. In the read-write balanced workflow, the results depicted in Figure 11 (a) show that 2PL-No-WAIT protocols with different batch sizes all experience a performance drop when increasing the number of executors beyond 8. However, Thunderbolt and OCC protocols with all the batch sizes obtain their highest throughputs on 12 executors and maintain stable throughput. Thunderbolt-b500 obtained 43*K* TPS while OCC-b500 achieved 35*K* TPS.

In the update-only workflow, the results shown in Figure 11 (b) indicate that OCC and 2PL-No-WAIT stopped increasing earlier in 4 executors (both around 22*K* TPS) while Thunderbolt provides a peek throughput (28*K* TPS) in 12 executors.

These experiments demonstrate that all the protocols do not obtain significant benefits for a large number of executors in a high-competition workflow. However, Thunderbolt can still achieve more parallelism with more executors.

Evaluation of Abort Rates. As we increased the number of executors, we also measured the average number of re-executions for the transactions. The results in Figure 11 indicate that when the number of executors goes beyond 8, all 2PL-NO-WAIT protocols experience a significant increase in the rate of abortions, leading to a drop in throughput from 24k to 18k in the readwrite balanced workflow. While OCC protocols provide a lower rate within the read-write balanced workflow. However, Thunderbolt achieves the lowest abortions, with Thunderbolt-b500 reducing 50% of the abortions from OCC-b500 and 90% from 2PL-NO-WAIT-b500 in all the experiments.

Evaluation of θ . We will now provide an evaluation with various θ values, using a read-write balance workload of $P_r = 0.5$. In particular, we examine high contention workloads where $0.75 \le \theta \le 0.9$, as this range represents the primary focus of our study. The results illustrated in Figure 12 (a and b) demonstrate that at $\theta = 0.75$, both OCC and Thunderbolt show comparable performance. However, as θ increases to 0.9, OCC experiences a



Figure 12: Throughput and average latency within varying θ settings with $P_r = 0.5$ (a and b) and varying P_r settings with $\theta = 0.85$ (c and d).

significant decline in performance, while Thunderbolt continues to achieve higher throughput levels. In contrast, the 2PL-No-WAIT approaches steady throughput, attributable to its locking strategy.

Evaluation of P_r . In this evaluation, we will look at how different values of P_r affect the read and write ratio in the workload, with $\theta = 0.85$. The results in Figure 12(c and d) show that when $P_r = 1$ (all read), all protocols behave similarly. However, the OCC protocol performs slightly better because it allows nonblocking local executions. When conflicts happen, the 2PL-No-WAIT shows a sharp decline in performance. In contrast, both Thunderbolt and OCC also perform worse when we decrease P_r , which leads to more conflicts. On the other hand, as we increase the value of P_r , Thunderbolt achieves better throughput and lower latency than OCC, even when all operations are write-only. All protocols show similar latency at $P_r = 1$, but as P_r increases, the latency for 2PL-No-WAIT rises, while Thunderbolt continues to perform better than OCC.

12 System Evaluation

We conducted evaluations to determine the impact of Thunderbolt built on Tusk. In our evaluation, we compared the performance of Thunderbolt with Tusk, which executes transactions in order after reaching a total order after DAG protocols. We Thunderbolt: Concurrent Smart Contract Execution with Non-blocking Reconfiguration for Sharded DAGs

Number of Replicas

Thunderbolt Thunderbolt-OCC Tusk (tps) 400'000 0 0 0 LAN LAN ٢ 100 Latency 50 16 32 64 16 32 64 Number of Replicas Number of Replicas Throughput (tps) 12,000 0 0 0 WAN WAN ۍ 60 Latency (0 70 00 0 00 00 16 32 64 16 32 64

Figure 13: Throughput and average latency within different replicas within $P_r = 0.5$ in LAN and WAN.

evaluated the impact of different components of the protocol by comparing the results between three systems:

- Thunderbolt: Leverage *CE* + parallel verification.
- Thunderbolt-OCC: Combine OCC + parallel verification.
- Tusk: Utilize the OE model.

Number of Replicas

As the serialized verification will execute the transactions in order to verify the results, we will expect that any DAG-based protocols with serialized verification will provide the same behavior with Tusk. We also leveraged SmallBank as the input workload.

Each replica was configured with *CE* comprising 16 executors to process transaction batches of 500, alongside 16 validators to verify blocks post-consensus. We scaled the system from 8 to 64 replicas. By default, K' was set to a sufficiently large value to disable shard rotation. In the final phase of our evaluation, we examined the impact of varying K' values, which govern the frequency of shard reconfiguration (Section 12).

SmallBank. Throughout the system evaluation, we focus on the SmallBank workload with a read-write balanced scenario ($P_r = 0.5$), where half of the transactions are read-only.

In the smallbank workload, Transaction addresses were selected from a pool of 1000 users with a skew parameter $\theta = 0.85$ to simulate a high-contention workload.

The LAN results, as shown in Figure 13, demonstrate that Thunderbolt significantly outperformed Tusk's sequential execution model, achieving a 50x speedup. Specifically, Thunderbolt reached 500*K* TPS with 64 replicas, compared to Tusk's 11*K* TPS. This improvement highlights the benefits of executing transactions in parallel.

We also compared Thunderbolt with Thunderbolt-OCC, which replaces the *CE* with *OCC*. While Thunderbolt-OCC matched Thunderbolt's throughput at 8 replicas, it lagged behind at scale, achieving only 400*K* TPS with 64 replicas. Furthermore, Thunderbolt maintained a low transaction latency of 5 seconds, whereas Tusk's latency soared to 100 seconds under the same conditions.

The WAN results demonstrate similar behavior but with higher latency. The latency gap between Thunderbolt and Tusk become smaller because the WAN latency begins to dominate cost, becoming bottleneck.

Cross-shard Transactions. Next, we evaluated the impact of Cross-shard TXs using 16 replicas. We randomly assigned a percentage P% (0 < $P \le 100$) of transactions to be processed by two shards. Additionally, we assessed the benefits of parallel execution by comparing Thunderbolt-OCC.



Figure 14: Throughput and average latency within different ratios of cross-shard transactions within 16 replicas.



Figure 15: Throughput and average latency within different reconfiguration periods within 8 replicas.

As shown in Figure 14, the performance of both Thunderbolt and Thunderbolt-OCC declined as the percentage P increased. In scenarios with only single-shard transactions (P = 0), both systems achieved 100K TPS. However, when P increased to 8%, Thunderbolt-OCC's throughput dropped to 16K TPS, while Thunderbolt maintained a significantly higher throughput of 64K TPS. Thunderbolt-OCC's performance aligned closely with Tusk, achieving approximately 10K TPS. In contrast, Thunderbolt delivered 19K TPS even when all transactions were cross-shard, demonstrating the advantages of its parallel execution model and non-deterministic ordering on *CE*.

Latency metrics further highlighted Thunderbolt's superiority. Under high-contention workloads, Thunderbolt achieved a transaction latency of 24*s* seconds, while Thunderbolt-OCC's latency was nearly double at 50 seconds.

Reconfiguration Periods. Now, we analyze the performance using different reconfiguration periods K' to transition the shard proposers into a new DAG on 8 replicas. Figure 15 demonstrates that Thunderbolt exhibited lower performance with smaller K'values (80K TPS with K' = 10), attributed to the costly transition between DAGs. Conversely, when K' was increased to over 1000, Thunderbolt demonstrated significantly improved stability, achieving a throughput of 180K TPS. Additionally, the average latency decreased from 1.9s to 1.7s as K' increased from 10 to 5000. Figure 16 also shows the average run time of committing proposals per 100 rounds, that is $\frac{1}{100} \sum (T_{commit(i)} - T_{commit(i-1)})$ where $T_{commit(i)}$ is the time of committing round *i*. We set K' as 300 and it was demonstrated that Thunderbolt will not get stuck during the reconfiguration. The runtime of each round is around 0.07s to 0.1s.

Failures. Finally, we evaluated the impact of replica failures within 16 replicas. We forced f replicas (f = 1 or f = 2) to stop working during the experiments. We randomly designated a percentage P% ($0 < P \le 100$) of the transactions to be processed by two shards.

Figure 17 reveals that Thunderbolt still can provide higher throughputs when some shards stop working. When one replica failed to propose transactions, Thunderbolt

1 (f = 1) obtained 78K TPS working on all Single-shard TXs (P = 0) and 17K TPS on all Cross-shard TXs while when two



Figure 16: Average latency per 100 rounds and reconfigure the shard per 300 rounds.



Figure 17: Throughput and average latency within different ratios of cross-shard transactions within 16 replicas when f (f = 1 or f = 2) replicas failed.

replicas failed to propose transactions, Thunderbolt

2 (f = 2) obtained 66K TPS on all Single-shard TXs and 15K TPS on all Cross-shard TXs. However, the results show that the latency remains stable even if some replicas fail to process, benefiting from the leader rotation of the DAG protocols.

13 Related work

Sharding on DAG-based protocol. Several studies [4, 27, 41, 44, 52, 72, 93, 105] have underscored the necessity of implementing sharding to enhance scalability within blockchain systems. Directed Acyclic Graph (DAG)-based blockchains [11, 12, 14, 25, 26, 26, 50, 50, 51, 56, 58, 73, 81, 84, 85, 85, 86, 96, 97, 101] present a promising alternative for improving concurrent transaction processing through the utilization of the DAG data structure. However, it is noteworthy that only a limited number of approaches offer effective sharding strategies for DAG protocols [8, 24, 27, 47, 70] These strategies often depend on account-based mechanisms to achieve eventual atomicity, employing incentive mechanisms through either the two-phase commit (2PC) protocol or additional coordinating entities.

Execute-Order-Validate. Hyperledger Fabric [10] introduces the Execute-Order-Validate (EOV) framework, which demonstrates high performance primarily in low-contention workloads. To enhance this framework, various techniques have been developed [35, 36, 75, 77, 91, 92], including methods for reordering transactions within a block and improving execution processes. In contrast, Thunderbolt addresses scalability by distributing transactions across multiple shards, thereby enabling each shard proposer to execute transactions in parallel.

Concurrent Execution. Deterministic approaches [31, 69, 99] have been proposed to improve the efficiency of transaction execution by creating a dependency graph. This graph allows for concurrent execution while minimizing conflicts between transactions. Furthermore, segmenting transactions [23, 71, 78, 79] has been introduced as an effective method for reducing these conflicts. CHIRON [66] and BlockSTM [34] offer nondeterministic execution by extracting dependencies from smart contracts or by determining the execution order based on transaction arrival times. Conversely, Thunderbolt adopts a different methodology

that it does not depend on arrival times or read/write set information for transaction processing. Instead, it dynamically assigns execution orders to effectively minimize transaction conflicts.

Concurrent Consensus. A robust and scalable blockchain framework is crucial for the implementation of real-world applications [9, 37, 39, 76]. For example, the PoE [38] model utilizes speculative execution protocols, and protocols, including RCC [40], FlexiTrust [42], and SpotLess [49], incorporate multiple leaders to enhance parallel processing capabilities. However, these protocols currently do not support reconfiguration.

14 Conclusions

We introduce Thunderbolt, an innovative sharding system designed to enhance smart contract execution by integrating the Order-Execute and Execute-Order-Validate models to efficiently handle both single-shard transactions (Single-shard TXs) and cross-shard transactions (Cross-shard TXs). We have developed a concurrent executor that significantly improves the performance of Single-shard TXs without requiring prior knowledge of read/write sets. Additionally, Thunderbolt effectively distributes transactions across multiple shards and utilizes Directed Acyclic Graph (DAG)-based protocols to maintain consistency between single-shard and cross-shard transactions. A key feature of Thunderbolt is its ability to utilize the inherent properties of the DAG to facilitate a non-blocking transition to a new DAG structure. This allows for the rotation of proposers for each shard, helping to prevent malicious activity from any single proposer. Our performance evaluations show that Thunderbolt achieves an impressive speedup of up to 50× compared to the native execution provided by Tusk.

Artifact Availability:

The source code, data, and/or other artifacts have been made available at https://github.com/apache/incubator-resilientdb/tr ee/Thunderbolt.

References

- [1] [n.d.]. Apache ResilientDB (Incubating). https://resilientdb.incubator.apac he.org/
- [2] [n.d.]. smallbank benchmark. http://hstore.cs.brown.edu/documentation/ deployment/benchmarks/smallbank/
- [3] Rakesh Agrawal, Michael J Carey, and Miron Livny. 1987. Concurrency control performance modeling: Alternatives and implications. ACM Transactions on Database Systems (TODS) 12, 4 (1987), 609–654.
- [4] Mustafa Al-Bassam, Alberto Sonnino, Shehar Bano, Dave Hrycyszyn, and George Danezis. 2017. Chainspace: A sharded smart contracts platform. arXiv preprint arXiv:1708.03778 (2017).
- [5] Amjad Aldweesh, Maher Alharby, Maryam Mehrnezhad, and Aad Van Moorsel. 2019. OpBench: A CPU performance benchmark for Ethereum smart contract operation code. In 2019 IEEE International Conference on Blockchain (Blockchain). IEEE, 274–281.
- [6] Mohammad Alomari, Michael Cahill, Alan Fekete, and Uwe Rohm. 2008. The cost of serializability on platforms that use snapshot isolation. In 2008 IEEE 24th International Conference on Data Engineering. IEEE, 576–585.
- [7] Mohammad Javad Amiri, Divyakant Agrawal, and Amr El Abbadi. 2019. Caper: a cross-application permissioned blockchain. Proceedings of the VLDB Endowment 12, 11 (2019), 1385–1398.
- [8] Mohammad Javad Amiri, Divyakant Agrawal, and Amr El Abbadi. 2021. Sharper: Sharding permissioned blockchains over network clusters. In Proceedings of the 2021 international conference on management of data. 76–88.
- [9] Mohammad Javad Amiri, Chenyuan Wu, Divyakant Agrawal, Amr El Abbadi, Boon Thau Loo, and Mohammad Sadoghi. 2022. The Bedrock of BFT: A Unified Platform for BFT Protocol Design and Implementation. CoRR abs/2205.04534 (2022). doi:10.48550/arXiv.2205.04534 arXiv:2205.04534
- [10] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti,

Chrysoula Stathakopoulou, Marko Vukolić, Sharon Weed Cocco, and Jason Yellick. 2018. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. In *Proceedings of the Thirteenth EuroSys Conference*. ACM, 30:1–30:15. doi:10.1145/3190508.3190538

- [11] Balaji Arun, Zekun Li, Florian Suri-Payer, Sourav Das, and Alexander Spiegelman. 2024. Shoal++: High Throughput DAG BFT Can Be Fast! arXiv preprint arXiv:2405.20488 (2024).
- [12] Kushal Babel, Andrey Chursin, George Danezis, Lefteris Kokoris-Kogias, and Alberto Sonnino. 2023. Mysticeti: Low-latency dag consensus with fast commit path. arXiv preprint arXiv:2310.14821 (2023).
- [13] Kushal Babel, Andrey Chursin, George Danezis, Lefteris Kokoris-Kogias, and Alberto Sonnino. 2023. Mysticeti: Low-Latency DAG Consensus with Fast Commit Path. arXiv preprint arXiv:2310.14821 (2023).
- [14] Leemon Baird. 2016. The swirlds hashgraph consensus algorithm: Fair, fast, byzantine fault tolerance. Swirlds Tech Reports SWIRLDS-TR-2016-01, Tech. Rep 34 (2016), 9–11.
- [15] Same Blackshear, Andrey Chursin, George Danezis, Anastasios Kichidis, Lefteris Kokoris-Kogias, Xun Li, Mark Logan, Ashok Menon, Todd Nowacki, Alberto Sonnino, et al. 2023. Sui lutris: A blockchain combining broadcast and consensus. arXiv preprint arXiv:2310.18042 (2023).
- [16] Dan Boneh, Ben Lynn, and Hovav Shacham. 2001. Short signatures from the Weil pairing. In International conference on the theory and application of cryptology and information security. Springer, 514-532.
- [17] Vitalik Buterin et al. 2014. A next-generation smart contract and decentralized application platform. *white paper* 3, 37 (2014), 2–1.
- [18] Michael Casey, Jonah Crane, Gary Gensler, Simon Johnson, and Neha Narula. 2018. The impact of blockchain technology on finance: A catalyst for change. (2018).
- [19] Miguel Castro and Barbara Liskov. 2002. Practical Byzantine Fault Tolerance and Proactive Recovery. ACM Trans. Comput. Syst. 20, 4 (2002), 398–461. doi:10.1145/571637.571640
- [20] Rick Cattell. 2011. Scalable SQL and NoSQL data stores. Acm Sigmod Record 39, 4 (2011), 12–27.
- [21] Ronnie Chaiken, Bob Jenkins, Per-Åke Larson, Bill Ramsey, Darren Shakib, Simon Weaver, and Jingren Zhou. 2008. Scope: easy and efficient parallel processing of massive data sets. *Proceedings of the VLDB Endowment* 1, 2 (2008), 1265–1276.
- [22] Xueying Chen and Min-ge Xie. 2014. A split-and-conquer approach for analysis of extraordinarily large data. *Statistica Sinica* (2014), 1655–1684.
- [23] Zhihao Chen, Haizhen Zhuo, Quanqing Xu, Xiaodong Qi, Chengyu Zhu, Zhao Zhang, Cheqing Jin, Aoying Zhou, Ying Yan, and Hui Zhang. 2021. SChain: a scalable consortium blockchain exploiting intra-and inter-block concurrency. *Proceedings of the VLDB Endowment* 14, 12 (2021), 2799–2802.
- [24] Feng Cheng, Jiang Xiao, Cunyang Liu, Shijie Zhang, Yifan Zhou, Bo Li, Baochun Li, and Hai Jin. 2024. Shardag: Scaling dag-based blockchains via adaptive sharding. In 2024 IEEE 40th International Conference on Data Engineering (ICDE). IEEE, 2068–2081.
- [25] Anton Churyumov. 2016. Byteball: A decentralized system for storage and transfer of value. URL https://byteball. org/Byteball. pdf (2016), 11.
- [26] George Danezis, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. 2022. Narwhal and Tusk: A DAG-Based Mempool and Efficient BFT Consensus. In Proceedings of the Seventeenth European Conference on Computer Systems. Association for Computing Machinery, New York, NY, USA, 34–50. doi:10.1145/3492321.3519594
- [27] Hung Dang, Tien Tuan Anh Dinh, Dumitrel Loghin, Ee-Chien Chang, Qian Lin, and Beng Chin Ooi. 2019. Towards scaling blockchain systems via sharding. In Proceedings of the 2019 international conference on management of data. 123–140.
- [28] Christos Douligeris and Aikaterini Mitrokotsa. 2004. DDoS attacks and defense mechanisms: classification and state-of-the-art. *Computer networks* 44, 5 (2004), 643–666.
- [29] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. 1988. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)* 35, 2 (1988), 288–323.
- [30] Muhammad El-Hindi, Carsten Binnig, Arvind Arasu, Donald Kossmann, and Ravi Ramamurthy. 2019. BlockchainDB: A shared database on blockchains. Proceedings of the VLDB Endowment 12, 11 (2019), 1597–1609.
- [31] Jose M Faleiro, Daniel J Abadi, and Joseph M Hellerstein. 2017. High performance transactions via early write visibility. *Proceedings of the VLDB Endowment* 10, 5 (2017).
- [32] Peter Franaszek and John T Robinson. 1985. Limitations of concurrency in transaction processing. ACM Transactions on Database Systems (TODS) 10, 1 (1985), 1–28.
- [33] Lan Ge, Christopher Brewster, Jacco Spek, Anton Smeenk, Jan Top, Frans Van Diepen, Bob Klaase, Conny Graumans, and Marieke de Ruyter de Wildt. 2017. Blockchain for agriculture and food: Findings from the pilot study. Number 2017-112. Wageningen Economic Research.
- [34] Rati Gelashvili, Alexander Spiegelman, Zhuolun Xiang, George Danezis, Zekun Li, Dahlia Malkhi, Yu Xia, and Runtian Zhou. 2023. Block-stm: Scaling blockchain execution by turning ordering curse to a performance blessing. In Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming. 232–244.
- [35] Christian Gorenflo, Lukasz Golab, and Srinivasan Keshav. 2020. XOX Fabric: A hybrid approach to blockchain transaction execution. In 2020 IEEE

International Conference on Blockchain and Cryptocurrency (ICBC). IEEE, 1–9.

- [36] Christian Gorenflo, Stephen Lee, Lukasz Golab, and Srinivasan Keshav. 2020. FastFabric: Scaling hyperledger fabric to 20 000 transactions per second. *International Journal of Network Management* 30, 5 (2020), e2099.
- [37] Suyash Gupta, Mohammad Javad Amiri, and Mohammad Sadoghi. 2023. Chemistry behind Agreement. In 13th Conference on Innovative Data Systems Research, CIDR 2023, Amsterdam, The Netherlands, January 8-11, 2023. www.cidrdb.org.
- [38] Suyash Gupta, Jelle Hellings, Sajjad Rahnama, and Mohammad Sadoghi. 2021. Proof-of-Execution: Reaching Consensus through Fault-Tolerant Speculation. In Proceedings of the 24th International Conference on Extending Database Technology.
- [39] Suyash Gupta, Jelle Hellings, and Mohammad Sadoghi. 2021. Fault-Tolerant Distributed Transactions on Blockchain. Morgan & Claypool Publishers. (2021).
- [40] Suyash Gupta, Jelle Hellings, and Mohammad Sadoghi. 2021. RCC: Resilient Concurrent Consensus for High-Throughput Secure Transaction Processing. In 37th IEEE International Conference on Data Engineering, ICDE 2021, Chania, Greece, April 19-22, 2021. IEEE, 1392–1403. doi:10.1109/ICDE51399.2021.00124
- [41] Suyash Gupta, Sajjad Rahnama, Jelle Hellings, and Mohammad Sadoghi. 2020. ResilientDB: Global Scale Resilient Blockchain Fabric. *Proc. VLDB Endow.* 13, 6 (2020), 868–883. doi:10.14778/3380750.3380757
- [42] Suyash Gupta, Sajjad Rahnama, Shubham Pandey, Natacha Crooks, and Mohammad Sadoghi. 2023. Dissecting BFT Consensus: In Trusted Components we Trust!. In Proceedings of the Eighteenth European Conference on Computer Systems, EuroSys 2023, Rome, Italy, May 8-12, 2023, Giuseppe Antonio Di Luna, Leonardo Querzoni, Alexandra Fedorova, and Dushyanth Narayanan (Eds.). ACM, 521–539.
- [43] Suyash Gupta, Sajjad Rahnama, and Mohammad Sadoghi. 2020. Permissioned blockchain through the looking glass: Architectural and implementation lessons learned. In 2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS). IEEE, 754–764.
- [44] Jelle Hellings and Mohammad Sadoghi. 2023. ByShard: sharding in a Byzantine environment. VLDB J. 32, 6 (2023), 1343-1367.
- [45] Maurice Herlihy. 2019. Blockchains from a distributed computing perspective. Commun. ACM 62, 2 (2019), 78–85.
- [46] Zicong Hong, Song Guo, and Peng Li. 2022. Scaling blockchain via layered sharding. *IEEE Journal on Selected Areas in Communications* 40, 12 (2022), 3575–3588.
- [47] Huawei Huang, Xiaowen Peng, Jianzhou Zhan, Shenyang Zhang, Yue Lin, Zibin Zheng, and Song Guo. 2022. Brokerchain: A cross-shard blockchain protocol for account/balance-based state sharding. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 1968–1977.
- [48] Maged N Kamel Boulos, James T Wilson, and Kevin A Clauson. 2018. Geospatial blockchain: promises, challenges, and scenarios in health and healthcare. 10 pages.
- [49] Dakai Kang, Sajjad Rahnama, Jelle Hellings, and Mohammad Sadoghi. 2024. SpotLess: Concurrent Rotational Consensus Made Practical through Rapid View Synchronization. In 40th IEEE International Conference on Data Engineering, ICDE 2024, Utrecht, Netherlands, May 13-17, 2024. IEEE.
- [50] Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. 2021. All you need is dag. In Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing. 165–175.
- [51] Idit Keidar, Oded Naor, Ouri Poupko, and Ehud Shapiro. 2022. Cordial miners: Fast and efficient consensus for every eventuality. arXiv preprint arXiv:2205.09174 (2022).
- [52] Eleftherios Kokoris-Kogias, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Ewa Syta, and Bryan Ford. 2018. Omniledger: A secure, scale-out, decentralized ledger via sharding. In 2018 IEEE symposium on security and privacy (SP). IEEE, 583–598.
- [53] Hsiang-Tsung Kung and John T Robinson. 1981. On optimistic methods for concurrency control. ACM Transactions on Database Systems (TODS) 6, 2 (1981), 213-226.
- [54] Satpal Singh Kushwaha, Sandeep Joshi, Dilbag Singh, Manjit Kaur, and Heung-No Lee. 2022. Ethereum smart contract analysis tools: A systematic review. *IEEE Access* 10 (2022), 57037–57062.
- [55] Ziliang Lai, Chris Liu, and Eric Lo. 2023. When private blockchain meets deterministic database. Proceedings of the ACM on Management of Data 1, 1 (2023), 1–28.
- [56] Chenxin Li, Peilun Li, Dong Zhou, Zhe Yang, Ming Wu, Guang Yang, Wei Xu, Fan Long, and Andrew Chi-Chih Yao. 2020. A decentralized blockchain with high throughput and fast confirmation. In 2020 {USENIX} Annual Technical Conference ({USENIX} {ATC} 20). 515–528.
- [57] Mingxuan Li, Yazhe Wang, Shuai Ma, Chao Liu, Dongdong Huo, Yu Wang, and Zhen Xu. 2023. Auto-tuning with reinforcement learning for permissioned blockchain systems. *Proceedings of the VLDB Endowment* 16, 5 (2023), 1000– 1012.
- [58] Dahlia Malkhi, Chrysoula Stathakopoulou, and Maofan Yin. 2023. BBCA-CHAIN: One-Message, Low Latency BFT Consensus on a DAG. arXiv preprint arXiv:2310.06335 (2023).
- [59] Microsoft. [n. d.]. eEVM. https://github.com/microsoft/eEVM
- [60] Jelena Mirkovic and Peter Reiher. 2004. A taxonomy of DDoS attack and DDoS defense mechanisms. ACM SIGCOMM Computer Communication Review 34, 2 (2004), 39–53.

EDBT '26, 24-27 March 2026, Tampere (Finland)

- [61] Satoshi Nakamoto. 2009. Bitcoin: A Peer-to-Peer Electronic Cash System. https://bitcoin.org/bitcoin.pdf
- [62] Arvind Narayanan and Jeremy Clark. 2017. Bitcoin's academic pedigree. Commun. ACM 60, 12 (2017), 36–45.
- [63] Senthil Nathan, Chander Govindarajan, Adarsh Saraf, Manish Sethi, and Praveen Jayachandran. 2019. Blockchain meets database: Design and implementation of a blockchain relational database. arXiv preprint arXiv:1903.01919 (2019).
- [64] Shamkant Navathe, Stefano Ceri, Gio Wiederhold, and Jinglie Dou. 1984. Vertical partitioning algorithms for database design. ACM Transactions on Database Systems (TODS) 9, 4 (1984), 680–710.
- [65] Faisal Nawab and Mohammad Sadoghi. 2019. Blockplane: A global-scale byzantizing middleware. In 2019 IEEE 35th International Conference on Data Engineering (ICDE). IEEE, 124–135.
- [66] Ray Neiheiser, Arman Babaei, Giannis Alexopoulos, Marios Kogias, and Eleftherios Kokoris Kogias. 2024. CHIRON: Accelerating Node Synchronization without Security Trade-offs in Distributed Ledgers. arXiv preprint arXiv:2401.14278 (2024).
- [67] Zeshun Peng, Yanfeng Zhang, Qian Xu, Haixu Liu, Yuxiao Gao, Xiaohua Li, and Ge Yu. 2022. Neuchain: a fast permissioned blockchain system with deterministic ordering. *Proceedings of the VLDB Endowment* 15, 11 (2022), 2585–2598.
- [68] Michael Pisa and Matt Juden. 2017. Blockchain and economic development: Hype vs. reality. Center for global development policy paper 107, 150 (2017), 1–42.
- [69] Thamir M Qadah and Mohammad Sadoghi. 2018. Quecc: A queue-oriented, control-free concurrency architecture. In Proceedings of the 19th International Middleware Conference. 13–25.
- [70] Naina Qi, Yong Yuan, and Fei-Yue Wang. 2022. DAG-BLOCK: A novel architecture for scaling blockchain-enabled cryptocurrencies. *IEEE Transactions* on Computational Social Systems 11, 1 (2022), 378–388.
- [71] Dai Qin, Angela Demke Brown, and Ashvin Goel. 2021. Caracal: Contention management with deterministic concurrency control. In Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles. 180–194.
- [72] Sajjad Rahnama, Suyash Gupta, Rohan Sogani, Dhruv Krishnan, and Mohammad Sadoghi. 2022. RingBFT: Resilient Consensus over Sharded Ring Topology. In Proceedings of the 25th International Conference on Extending Database Technology. OpenProceedings.org, 2:298–2:311. doi:10.48786/edbt.2022.17
- [73] Mayank Raikwar, Nikita Polyanskii, and Sebastian Müller. 2024. SoK: DAGbased Consensus Protocols. In 2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). IEEE, 1–18.
- [74] Rasmus V Rasmussen and Michael A Trick. 2008. Round robin scheduling-a survey. European Journal of Operational Research 188, 3 (2008), 617–636.
- [75] Pingcheng Ruan, Dumitrel Loghin, Quang-Trung Ta, Meihui Zhang, Gang Chen, and Beng Chin Ooi. 2020. A transactional perspective on execute-ordervalidate blockchains. In Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data. 543–557.
- [76] Mohammad Sadoghi and Spyros Blanas. 2019. Transaction Processing on Modern Hardware. Morgan & Claypool Publishers. doi:10.2200/S00896ED1V 01Y201901DTM058
- [77] Ankur Sharma, Felix Martin Schuhknecht, Divya Agrawal, and Jens Dittrich. 2019. Blurring the lines between blockchains and database systems: the case of hyperledger fabric. In *Proceedings of the 2019 International Conference on Management of Data*. 105–122.
- [78] Dennis Shasha, Francois Llirbat, Eric Simon, and Patrick Valduriez. 1995. Transaction chopping: Algorithms and performance studies. ACM Transactions on Database Systems (TODS) 20, 3 (1995), 325–363.
- [79] Dennis Shasha, Eric Simon, and Patrick Valduriez. 1992. Simple rational guidance for chopping up transactions. In Proceedings of the 1992 ACM SIGMOD International Conference on management of Data. 298–307.
- [80] Madhavapeddi Shreedhar and George Varghese. 1995. Efficient fair queueing using deficit round robin. In Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication. 231– 242.
- [81] Nibesh Shrestha, Rohan Shrothrium, Aniket Kate, and Kartik Nayak. 2024. Sailfish: Towards Improving Latency of DAG-based BFT. *Cryptology ePrint Archive* (2024).
- [82] Eljas Soisalon-Soininen and Tatu Ylönen. 1995. Partial strictness in two-phase locking. In International Conference on Database Theory. Springer, 139–147.
- [83] Alberto Sonnino, Shehar Bano, Mustafa Al-Bassam, and George Danezis. 2020. Replay attacks and defenses against cross-shard consensus in sharded distributed ledgers. In 2020 IEEE European Symposium on Security and Privacy (EuroS&P). IEEE. 294–308.
- [84] Alexander Spiegelman, Balaji Aurn, Rati Gelashvili, and Zekun Li. 2023. Shoal: Improving dag-bft latency and robustness. arXiv preprint arXiv:2306.03058 (2023).
- [85] Alexander Spiegelman, Neil Giridharan, Alberto Sonnino, and Lefteris Kokoris-Kogias. 2022. Bullshark: Dag bft protocols made practical. In Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. 2705–2718.
- [86] Chrysoula Stathakopoulou, Michael Wei, Maofan Yin, Hongbo Zhang, and Dahlia Malkhi. 2023. BBCA-LEDGER: High Throughput Consensus meets Low Latency. arXiv preprint arXiv:2306.14757 (2023).

- [87] Nick Szabo. 1997. Formalizing and securing relationships on public networks. *First monday* (1997).
- [88] N. Szabo. 1997. The Idea of Smart Contracts. http://www.fon.hum.uva.nl/r ob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2 006/szabo.best.vwh.net/idea.html
- [89] Yong Chiang Tay, Nathan Goodman, and Rajan Suri. 1985. Locking performance in centralized databases. ACM Transactions on Database Systems (TODS) 10, 4 (1985), 415–462.
- [90] Parth Thakkar and Senthilnathan Natarajan. 2020. Scaling hyperledger fabric using pipelined execution and sparse peers. arXiv preprint arXiv:2003.05113 (2020).
- [91] Parth Thakkar and Senthilnathan Natarajan. 2021. Scaling blockchains using pipelined execution and sparse peers. In Proceedings of the ACM Symposium on Cloud Computing. 489–502.
- [92] Parth Thakkar, Senthil Nathan, and Balaji Viswanathan. 2018. Performance benchmarking and optimizing hyperledger fabric blockchain platform. In 2018 IEEE 26th international symposium on modeling, analysis, and simulation of computer and telecommunication systems (MASCOTS). IEEE, 264–276.
- [93] Gang Wang, Zhijie Jerry Shi, Mark Nixon, and Song Han. 2019. Sok: Sharding on blockchain. In Proceedings of the 1st ACM Conference on Advances in Financial Technologies. 41-61.
- [94] Zhaoguo Wang, Shuai Mu, Yang Cui, Han Yi, Haibo Chen, and Jinyang Li. 2016. Scaling multicore databases via constrained parallel execution. In Proceedings of the 2016 International Conference on Management of Data. 1643–1658.
- [95] Gavin Wood et al. 2014. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper* 151, 2014 (2014), 1–32.
 [96] Jiang Xiao, Shijie Zhang, Zhiwei Zhang, Bo Li, Xiaohai Dai, and Hai Jin.
- [96] Jiang Xiao, Shijie Zhang, Zhiwei Zhang, Bo Li, Xiaohai Dai, and Hai Jin. 2022. Nezha: Exploiting concurrency for transaction processing in dagbased blockchains. In 2022 IEEE 42nd International Conference on Distributed Computing Systems (ICDCS). IEEE, 269–279.
- [97] Jie Xu, Yingying Cheng, Cong Wang, and Xiaohua Jia. 2021. Occam: A secure and adaptive scaling scheme for permissionless blockchain. In 2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS). IEEE, 618–628.
- [98] Di Yang, Chengnian Long, Han Xu, and Shaoliang Peng. 2020. A review on scalability of blockchain. In Proceedings of the 2020 2nd International Conference on Blockchain Technology. 1–6.
- [99] Chang Yao, Divyakant Agrawal, Gang Chen, Qian Lin, Beng Chin Ooi, Weng-Fai Wong, and Meihui Zhang. 2016. Exploiting single-threaded model in multi-core in-memory systems. *IEEE Transactions on Knowledge and Data Engineering* 28, 10 (2016), 2635–2650.
- [100] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham. 2019. HotStuff: BFT Consensus with Linearity and Responsiveness. In Proceedings of the ACM Symposium on Principles of Distributed Computing. ACM, 347–356. doi:10.1145/3293611.3331591
- [101] Haifeng Yu, Ivica Nikolić, Ruomu Hou, and Prateek Saxena. 2020. Ohie: Blockchain scaling made simple. In 2020 IEEE Symposium on Security and Privacy (SP). IEEE, 90–105.
- [102] Xiangyao Yu, George Bezerra, Andrew Pavlo, Srinivas Devadas, and Michael Stonebraker. 2014. Staring into the abyss: An evaluation of concurrency control with one thousand cores. (2014).
- [103] Xiangyao Yu, Andrew Pavlo, Daniel Sanchez, and Srinivas Devadas. 2016. Tictoc: Time traveling optimistic concurrency control. In Proceedings of the 2016 International Conference on Management of Data. 1629–1642.
- [104] Matei Zaharia, Reynold S Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J Franklin, et al. 2016. Apache spark: a unified engine for big data processing. *Commun. ACM* 59, 11 (2016), 56–65.
- [105] Mahdi Zamani, Mahnush Movahedi, and Mariana Raykova. 2018. Rapidchain: Scaling blockchain via full sharding. In Proceedings of the 2018 ACM SIGSAC conference on computer and communications security. 931–948.
- [106] Luyi Zhang, Yujue Wang, Yong Ding, Hai Liang, Changsong Yang, and Chunhai Li. 2023. Sharding Technologies in Blockchain: Basics, State of the Art, and Challenges. In International Conference on Blockchain and Trustworthy Systems. Springer, 242–255.
- [107] Yan Zhang, Jia Kang Wang, and Ying Jie Han. 2023. CChain: a high throughput blockchain system. In Second International Conference on Digital Society and Intelligent Systems (DSINS 2022), Vol. 12599. SPIE, 129–136.