



Rashnu: Data-Dependent Order-Fairness

Heena Nagda
University of Pennsylvania
hnagda@seas.upenn.edu

Shubhendra Pal Singhal
Georgia Institute of Technology
ssinghal74@gatech.edu

Mohammad Javad Amiri
Stony Brook University
amiri@cs.stonybrook.edu

Boon Thau Loo
University of Pennsylvania
boonloo@seas.upenn.edu

ABSTRACT

Distributed data management systems use state Machine Replication (SMR) to provide fault tolerance. The SMR algorithm enables Byzantine Fault-Tolerant (BFT) protocols to guarantee safety and liveness despite the malicious failure of nodes. However, SMR does not prevent the adversarial manipulation of the order of transactions, where the order assigned by a malicious leader differs from the order in that transactions are received from clients. While *order-fairness* has been recently studied in a few protocols, such protocols rely on synchronized clocks, suffer from liveness issues, or incur significant performance overhead. This paper presents *Rashnu*, a high-performance fair ordering protocol. *Rashnu* is motivated by the fact that fair ordering among two transactions is needed only when both transactions access a shared resource. Based on this observation, we define the notion of *data-dependent order fairness* where replicas capture only the order of data-dependent transactions and the leader uses these orders to propose a dependency graph that represents fair ordering among transactions. Replicas then execute transactions using the dependency graph, resulting in the parallel execution of independent transactions. We implemented a prototype of *Rashnu* where our experimental evaluation reveals the low overhead of providing order-fairness in *Rashnu*.

PVLDB Reference Format:

Heena Nagda, Shubhendra Pal Singhal, Mohammad Javad Amiri, and Boon Thau Loo. *Rashnu: Data-Dependent Order-Fairness*. PVLDB, 17(9): 2335 - 2348, 2024.
doi:10.14778/3665844.3665861

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/HeenaNagda/Order-Fairness>.

1 INTRODUCTION

Distributed systems rely on consensus protocols to provide robustness and high availability [7, 9, 12, 19, 24, 34, 53]. Consensus protocols use the State Machine Replication (SMR) algorithm [45, 57] to ensure that all honest replicas execute transactions in the same order (safety), and all correct transactions are eventually executed

(liveness). Existing consensus protocols typically include a designated leader replica that receives transactions from clients, assigns an order (e.g., a sequence number) to each transaction, which represents the position of the transaction in the final log, and initiates agreement on the order of transaction among all replicas. A malicious leader, however, can control transactions' inclusion and final ordering without violating safety or liveness. Most existing BFT protocols do not prevent such an adversarial order manipulation.

Adversarial manipulation of transactions order is studied in the decentralized finance (DeFi) domain [8, 22, 27, 33, 40, 56, 63] where transaction proposers make profit by including, excluding, or re-ordering transactions within blocks, known as maximal extractable value (MEV) [22]. Consider an exchange transaction to buy a particular asset. A malicious proposer can perform a front-running sandwich attack by placing the buy transaction between two buy and sell transactions (initiated by the malicious proposer) to manipulate asset prices. The proposer buys assets for a lower price to let the victim buy at a higher value and then sells them again, typically at a higher price afterwards. Adversarial manipulation of transactions in Ethereum resulted in extracting more than \$686M in revenue from unsophisticated users (\$1.38B across all EVM powered networks) [17]. Other than profitability, the order of transactions might affect their validity typically when multiple transactions access limited assets, e.g., a ticket booking scenario where the number of available tickets is limited, and through order manipulation, tickets are purchased by users who are not supposed to get one.

Different techniques, such as censorship resistance, random leader election, and threshold encryption, have been proposed to prevent transaction ordering manipulation. Censorship resistance [52] only ensures that correct transactions are eventually ordered, i.e., not censored. However, reordering transactions, e.g., sandwich attacks, is still possible. Similarly, reputation-based systems [5, 21, 42, 47] only detect unfair censorship. The random leader (committee) election or, in general, participation equity, on the other hand, provides opportunities for every replica to propose and commit its transactions, e.g., by becoming the proposer [2, 5, 23, 31, 35, 39, 47, 50, 55, 58, 60]. However, a malicious proposer can still order transactions unfairly in its turn. Finally, using threshold encryption [5, 13, 52, 59], transactions are encrypted, and their content is revealed once their order is fixed. This technique suffers from (1) metadata leakage and (2) collusion attacks between clients and the leader, where the leader becomes aware of a client transaction before ordering it and manipulates its order [37, 38, 43].

Recently, the notion of *order-fairness* is presented to address the manipulation of transaction ordering [14, 37, 38, 43, 44, 62]. Intuitively, order-fairness ensures that if a sufficiently large number of

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 17, No. 9 ISSN 2150-8097.
doi:10.14778/3665844.3665861

replicas (known as γ -fraction) receive a transaction t before another transaction t' , then t must be ordered before t' [38]. To support order-fairness, clients broadcast their transactions to all replicas. In each round, every replica locally orders received transactions according to their received times and sends its ordering to the designated leader of that particular round. The leader then constructs a fair-ordered proposal from the received orderings and initiates consensus on the transaction block.

Existing fair ordering protocols, e.g., Wendy [43, 44], Pompe [62], Aequitas [38], Themis [37] and Quick-order-fairness [14], however, suffer from serious limitations. First, both Wendy [43, 44] and Pompe [62] rely on synchronized clocks between replicas, making these protocols impractical in asynchronous networks. Pompe [62] also determines fair order using the median of timestamps assigned by replicas. The median value, however, can be easily manipulated by faulty replicas. Similarly, in Wendy, if honest nodes' local clocks are far apart, no fairness guarantees can be provided [37]. Second, Aequitas [38], and Quick-order-fairness [14] only guarantee weak liveness and transactions might need to wait an arbitrarily long time before getting committed [37]. Finally, while Themis [37], which bootstraps from HotStuff [61], does not assume synchronized clocks, it suffers from significant performance overhead.

The poor performance of order-fairness protocols is mainly caused by the time required to generate a fair order among *all* transactions. In an asynchronous network, different replicas might receive transactions in different orders, and transactions might be arbitrarily delayed. Byzantine replicas might also send maliciously manipulated local ordering to the leader. Moreover, collecting the local ordering of different replicas might lead to cycles in the final order even when all replicas are honest. Therefore, it becomes time-consuming for the leader to check the order of each pair of transactions in every local ordering and achieve a fair order.

Fair transaction ordering is essential when transactions access a shared resource and manipulating the order gives an unfair advantage to some transactions. However, the execution order of transactions with no dependency does not impact the execution results. This is crucial, especially due to the low workload contention level in many practical applications. Based on this observation, a key insight for our work is that a practical notion of order-fairness can limit the fair ordering only to transactions with some dependencies.

In this paper, we focus on data dependencies among transactions and present the notion of *data-dependent order-fairness* to ensure that if most replicas receive data-dependent transactions in a particular order, the order is preserved in execution.

Using the notion of data-dependent order-fairness, we present a high-performance fair ordering protocol, *Rashnu*¹. In *Rashnu*, each replica, instead of creating a list of received transactions, captures data dependencies between transactions (according to their received order) in the form of a directed acyclic graph (DAG). Upon finishing a round (specified by a threshold, e.g., time window, number of transactions, or both), the replica sends the graph to the leader. The leader then collects the local ordering of different replicas and proposes a global dependency graph representing the fair order of received transactions by initiating consensus on its proposed block.

Considering data dependencies while making the fair ordering of transactions more efficient, brings its own challenges. First, the leader needs to figure out all data dependencies between transactions in all local dependency graphs to extract a fair order. This is challenging because replicas might receive transactions in different orders or even receive different sets of transactions due to the asynchronous nature of the network. As a result, the global view of transactions might be different from the local view of each replica, e.g., while two transactions are independent (no path between them) in the dependency graph of the first replica, there is a path between them in the dependency graph of second replica caused by a third transaction not received by the first replica. Second, replicas need to capture dependencies across different transaction blocks to be able to execute transactions more efficiently. This is because while the execution of some transactions of a block is delayed by some missing transactions, data-independent transactions of the consequent blocks can be executed immediately.

By proposing data-dependent order-fairness, *Rashnu*, on the one hand, reduces the leader latency by capturing the order of only data-dependent transactions and, on the other hand, enables replicas to execute data-independent transactions in parallel. Moreover, *Rashnu* can capture any type of ordering constraint, other than data-dependency, required by applications, e.g., resource constraints. *Rashnu* supports such ordering constraints by enforcing them on the construction of the dependency graph.

Since the fair ordering of transactions is separated from the consensus protocol, *Rashnu* can bootstrap from any leader-based consensus protocol. To be able to compare with existing order-fairness protocols, e.g., Themis [37], we implemented *Rashnu* on HotStuff [61] (used by Themis) as the underlying BFT protocol.

Note that *Rashnu* is mainly proposed for permissioned untrusted environments. However, its underlying techniques can be applied to permissionless blockchain environments as well. In particular, fair ordering techniques can be simply integrated with committee-based permissionless blockchains where a subset of nodes order transactions (through sharding, e.g., Elastico [48] or by separating leader election blocks from transaction ordering blocks [41] where a BFT protocol is used among a subset of miners).

Overall, this paper makes three main contributions.

- Data-dependent order-fairness notion is defined as providing fair ordering only among data-dependent transactions.
- We design *Rashnu*, a high-performance fair-ordering protocol that decouples ordering from consensus and leverages graph-based techniques to achieve order-fairness among data-dependent transactions.
- We implement a prototype of *Rashnu*. Our evaluation results demonstrate the low overhead of *Rashnu* in providing order-fairness. Leveraging parallel execution, *Rashnu* shows even higher throughput compared to its underlying protocol, HotStuff, in compute-intensive workloads.

2 BACKGROUND

A Byzantine fault-tolerant (BFT) protocol runs on a network consisting of a set of replicas that may behave arbitrarily, potentially maliciously. BFT protocols use the State Machine Replication (SMR) algorithm [45, 57] to ensure that honest replicas execute requests

¹Rashnu is the Avestan language name of the Zoroastrian deity of justice.




Figure 1: A Condorcet cycle

in the same order despite the concurrent failure of f Byzantine replicas. SMR BFT protocols need to provide safety and liveness.

A recent line of work, e.g., Wendy [43, 44], Aequitas [38], Pompe [62], Themis [37] and Quick order-fairness [14], have proposed to add *order-fairness* as the third property that SMR BFT protocols need to guarantee. Order-fairness aims to ensure that the transactions are committed in the same order as they arrived at the network. Order-fairness is parameterized by an order-fairness parameter γ representing the fraction of replicas that receive transactions in a particular order. Wendy [43, 44] and Pompe [62] require replicas to access synchronized local clocks. Pompe further determines the fair order by relying on timestamps assigned by replicas, which can be manipulated by malicious replicas. As a result, we mainly focus on Aequitas [38], Themis [37] and Quick order-fairness [14], which do not consider synchrony assumptions, making them more suitable for asynchronous networks (note that Quick order-fairness [14] provides liveness only when all nodes are honest [37]).

Receive-Order-fairness. In a fair-ordering protocol, each replica needs to individually order transactions locally and send its local order to the leader. Each replica can order transactions based on (1) the timestamp assigned by clients, (2) the propagation time (that can be estimated by measuring network latency), or (3) the received time. Since clients might maliciously assign timestamps to their transactions, replicas cannot rely on the assigned timestamps (unless each client is equipped with trusted hardware). The propagation time also cannot be captured precisely as the network is asynchronous and transactions might be arbitrarily delayed. As a result, Rashnu, similar to existing protocols [14, 37, 38], relies on the transactions receive time.

Condorcet cycles. The notion of (strong) *receive-order-fairness* specifies that if γ fraction of replicas receive a transaction t before another transaction t' , then all honest replicas must order t (strictly) before t' . In an asynchronous network, replicas might receive transactions in different orders. Hence, defining a fair order among all transactions becomes impossible, even if all replicas are honest, as demonstrated by the *Condorcet paradox*. Condorcet paradox states that even if the local order of each individual replica is transitive, there might be situations that lead to non-transitive collective voting preferences. Figure 1 demonstrates the Condorcet paradox with four transactions between four replicas. As can be seen, a majority (3 out of 4) of replicas received t_1 before t_2 (replicas r_1, r_3 and r_4). Similarly t_2 is received before t_3 by a majority of replicas (r_1, r_2 and r_4) and t_3 is received before t_4 by replicas r_1, r_2 and r_3 . However, t_4 is also received by replicas r_2, r_3 and r_4 before t_1 . The collection of these orders results in a cyclic global ordering as generated by the leader (the leader is one of the replicas).

To address this challenge, transactions involved in a Condorcet cycle can be sent to replicas concurrently within the same batch [38]. Specifically, given transactions t and t' where t is received by sufficiently many replicas before t' ; while strong order-fairness

requires the leader to send t before t' , *batch-order-fairness* relaxes this requirement by saying t should be delivered no later than (before or at the same time as) t' . Batch-order-fairness does not specify the order of transactions within a batch and respects a fair order up to this limit. Hence, if a total order among all transactions is required, a deterministic total ordering for transactions in the same cycle, e.g., alphabetical, needs to be specified.

Weak liveness. Batch order-fairness circumvents the Condorcet impossibility results. However, Condorcet cycles can chain together and extend arbitrarily. In this case, using batch order-fairness, the leader waits for a chain to be completed before sending any transaction of the batch to replicas. Hence, liveness might be violated as transactions wait for a long time. To address the weak liveness issue of batch order-fairness, transactions within the same cycle can be delivered contiguously in a set of successive blocks (instead of a single block) where a part of the current cycle can be output later without violating order-fairness as long as no transaction from a later cycle comes before it [37]. Using the *deferred ordering* technique, the leader proposes a partial (incomplete) ordering for some transactions within the block and defers their total ordering to the next consecutive blocks. The total ordering for deferred transactions does not depend on the chaining of Condorcet cycles. As a result, (standard) liveness can be achieved.

3 RASHNU MODEL

Rashnu deploys on a set of n known nodes (replicas) where at most f of them are Byzantine (malicious) at each time. In the Byzantine failure model, faulty replicas may exhibit arbitrary, potentially malicious, behavior. In an asynchronous system, where replicas can fail, no consensus solutions guarantee both safety and liveness (FLP result) [29]. As a result, Rashnu assumes the partially synchronous communication model to circumvent the FLP impossibility. In a partial synchrony model, an unknown global stabilization time (GST) exists, after which messages between honest replicas are received within some known bound Δ . A strong adversary can coordinate malicious replicas and delay communication. However, the adversary is computationally bounded and cannot subvert standard cryptographic assumptions. Replicas are connected with point-to-point bi-directional communication channels, and each client can communicate with any replica. Network links are pairwise authenticated, which guarantees that a malicious replica cannot forge a message from an honest replica. For communication between replicas, we assume the presence of digital signatures and public-key infrastructure (PKI). A collision-resistant hash function $D(\cdot)$ is also used to map a message m to a constant-sized digest $D(m)$.

BFT protocols require the number of replicas $n > 3f$ to guarantee safety with at most f malicious replicas [10, 11, 20, 26, 46]. However, fair ordering of transactions requires larger n . The order-fairness is further parameterized by an order-fairness parameter γ representing the fraction of replicas that receive transactions in a particular order. In Rashnu, n needs to be larger than $\frac{4f}{2\gamma-1}$.

LEMMA 3.1. Given a network consisting of n replicas from which at most f are malicious. The fair ordering of transactions is possible only when $n > \frac{4f}{2\gamma-1}$ where γ is the fraction of replicas (majority or more) that receive transactions in a particular order.

PROOF. (Using quorum size) With n replicas, since f replicas might be faulty, the protocol can rely on a quorum of $n-f$ replicas to generate the final order. Among these $n-f$ replicas, f might be malicious, i.e., f replicas not participating in the quorum are honest slow replicas. As a result, out of the quorum of $n-f$ replicas, only $n-2f$ replicas are guaranteed to be honest. To realize order-fairness if γn replicas receive transactions in a particular order, the final ordering must reflect that order. Since only $n-2f$ replicas within the quorum are guaranteed to be honest, the output of order-fairness must be the same as γn even with $\gamma n - 2f$ replicas broadcasting a particular order. On the other hand, a majority of replicas must agree with the order. Otherwise, we could end up with conflicting fair orders suggested by different sets of γn replicas. More precisely, to ensure that only one of the two different orders $t < t'$ or $t' < t$ is captured between two transactions t and t' , $\gamma n - 2f > \frac{n}{2}$. As a result $n > \frac{4f}{2\gamma-1}$. \square

PROOF. (Using δ -differential validity) The δ -differential validity [30] can also be used to prove the number of required replicas in a fair ordering protocol. Let $c(v)$ denote the number of honest replicas that propose value v . δ -differential validity states that if an honest replica decides v , then every other value v' proposed by another honest replica satisfies $c(v') \leq c(v) - \delta$. Based on this definition, a BFT protocol satisfies δ -differential validity if and only if it never decides a value v' with $c(v') < c(v) - \delta$ where v is the value proposed most often by honest replicas.

In an asynchronous network, δ -differential consensus is achievable only if $\delta \geq 2f$ (i.e., $c(v') < c(v) - 2f$) [30]. In our context, the value v is interpreted as the order of two transactions t and t' . As stated before, the output of fair ordering must be the same even if $\gamma n - 2f$ replicas broadcast the order. As a result, $(1-\gamma)n < (\gamma n - 2f) - 2f$ where $(1-\gamma)n$ is the maximum number of replicas that might propose another order. Hence, $n > \frac{4f}{2\gamma-1}$. \square

The order-fairness parameter γ represents the fraction of replicas that receive transactions in a particular order. The range of possible values for the order-fairness parameter γ can be calculated based on the total number of replicas n .

LEMMA 3.2. Order-fairness parameter γ is in $(\frac{1}{2} + \frac{2f}{n}, 1]$.

PROOF. Intuitively, γ should be larger than $\frac{1}{2}$ to prevent multiple conflicting fair orders. In Rashnu, Since $n > \frac{4f}{2\gamma-1}$, $\gamma > \frac{1}{2} + \frac{2f}{n}$. On the other hand, $\gamma = 1$ is the case where all replicas receive transactions in the same order. As a result, $\frac{1}{2} + \frac{2f}{n} < \gamma \leq 1$. \square

In an asynchronous network, replicas might receive transactions in different orders. Even if all replicas are honest, defining a fair order among all transactions is impossible, as demonstrated by the Condorcet paradox. Condorcet paradox states that even if the local order of all individual replicas is transitive, there might be situations that lead to non-transitive collective voting preferences. Moreover, batching transactions might lead to weak liveness issues when Condorcet cycles are chained together. In Rashnu, inspired by the order deferring technique [37], transactions of the same cycle are proposed contiguously in successive blocks. However, since Rashnu considers dependency only among data-dependent

transactions, cycles are less likely to happen compared to existing fair ordering protocols.

The order of executing transactions matters when transactions compete with each other on the same resources and manipulating the order gives an unfair advantage to some transactions. As a result, our notion of order-fairness limits the fair ordering to data-dependent transactions. Each transaction performs a sequence of reads and writes, each accessing a single record. Rashnu assumes a priori knowledge of transactions' read- and write-set, where the read-set and write-set of transactions are pre-declared or can be obtained from the transactions via static analysis, e.g., all records involved in a transaction, are accessed by their primary keys. Even if that assumption does not hold, the system can employ speculative execution techniques [28] to obtain the read-set and write-set of each transaction. Note that a conservative estimation of read-set and write-set, used by different techniques, while slightly decrease the performance, does not hurt the correctness of Rashnu. Given a transaction t , we use $R(t)$ and $W(t)$ to denote the read-set and write-set of transaction t respectively. Intuitively, two transactions t and t' are data-dependent if they access the same data object and one performs a write operation on the data object.

Definition: (Data-dependent transactions). two transactions t and t' are *data-dependent* if $(R(t) \cap W(t')) \cup (W(t) \cap R(t')) \cup (W(t) \cap W(t')) \neq \emptyset$.

Definition: (Data-dependent order-fairness). Given two data-dependent transactions t and t' . If γ -fraction of replicas receive t before t' , no honest replica outputs t' before t .

Rashnu processes transactions in rounds where each replica collects transactions received from clients and sends a block of transactions to the leader at the end of each round. Note that replicas might not receive the same set of transactions from clients in each round due to network asynchrony.

4 FAIR TRANSACTION ORDERING

In Rashnu, similar to other fair ordering protocols and even some BFT protocols (e.g., HotStuff [61] and Prime [4]), clients broadcast their transactions to all replicas. Each replica collects a batch of transactions, constructs a local dependency graph for transactions based on the received order, and sends the graph to the leader. The leader then collects all local dependency graphs and generates a global dependency graph that captures fair order among transactions. If the order of two data-dependent transactions can not be determined by the leader, e.g., they are received by an insufficient number of replicas, the leader defers the order to the next blocks.

Figure 2 presents an overview of Rashnu in a simple example. We use this example throughout this section. The example includes four replicas r_1 to r_4 (assuming $f = 1$, n becomes $4f + 1 = 5$ and there are $n - f = 4$ replicas in the quorum) where one of them is the leader and presents two consecutive rounds of the protocol. A set of transactions t_1 to t_{11} are received from clients in these two rounds where each transaction accesses a subset of data objects A, B, and C, as shown in the figure (for simplicity, all data accesses are assumed to be write). This example includes many data dependencies among transactions to capture different corner cases.

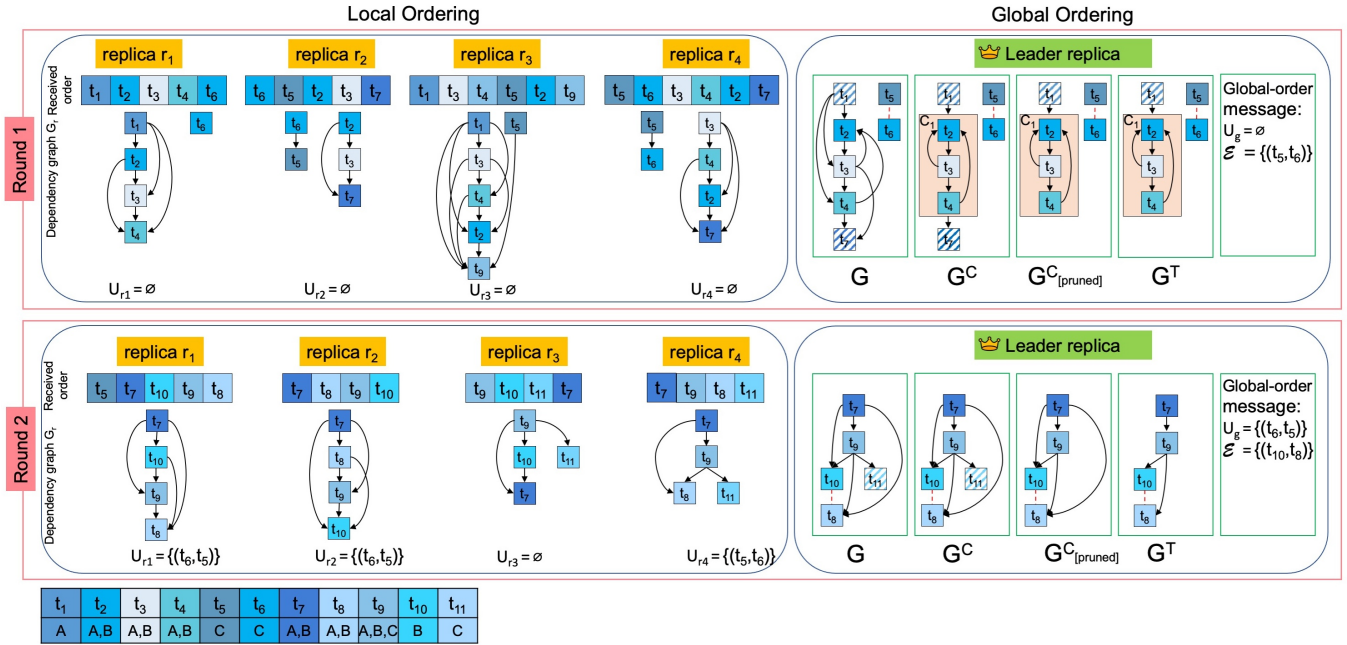


Figure 2: Local and global ordering in Rashnu

In real-world scenarios, however, a large percentage of transactions are typically data-independent.

This section demonstrates how different replicas construct their local ordering and how the leader orders transactions in a fair yet efficient manner. We then discuss the execution of transactions and the correctness of Rashnu.

4.1 Local Ordering

In the local ordering phase, each replica generates a dependency graph for client requests that the replica has received and their order is not determined yet. In each round i , a replica deals with three types of transactions.

- First, old transactions that have been received by the replica in a round j ($j < i$); however, they have not been proposed by a leader (in any round j to $i - 1$) yet. This happens when an insufficient number of replicas receive a transaction in a round. Hence, the leader does not propose it until more replicas receive it (probably, in a later round).
- Second, new transactions that are received by the replica in the current round i ; however, the leader has already proposed them in an earlier round. This is because transactions might be delayed due to network asynchrony, and a replica receives a transaction that has already been received by other replicas, and its order is possibly determined in an earlier round.
- Third, new transactions that are received by the replica in the current round i , and have not been proposed by a leader in any earlier round.

As shown in Algorithm 1, each replica r initiates an empty graph $G_r = (T_r, E_r)$ at the beginning of each round i . The replica first adds all its old transactions (first type) to graph G_r (lines 2-3). The

replica adds an edge (t', t) to the graph G_r for each transaction t if the replica has received t' before t (i.e., $t' < t$) and t and t' are data-dependent (lines 4-6). Transactions t and t' might have been received by the replica in different rounds. Note that if the underlying consensus protocol relies on a stable leader, e.g., PBFT [15], the leader can keep track of the first type of transactions. In this way, replicas do not need to wait for the previous round leader's proposal before sending their local order in the current round. Upon receiving a valid signed request message $m = \langle \text{REQUEST}, t, \tau_c, c \rangle_{\sigma_c}$ from an authorized client c with timestamp τ_c to execute transaction t , the replica checks whether the transaction has already been proposed by the leader in an earlier round (second type). Otherwise (third type), the replica adds vertex t and all its dependencies with existing vertices to the graph G_r (lines 9-13).

If transaction t has already been proposed in an earlier round (second type), its order might not have been determined. When the leader finalizes the global order of transactions (as explained in Section 4.2), if the number of received orders between two data-dependent transactions t and t' is insufficient, the leader can not determine the order and adds a pair (t, t') to a set of *missing pairs* (undirected edges) \mathcal{E} . This set is then used by replicas in the next round to specify the order of missing pairs and complete the previous proposals.

When a replica receives a transaction t that was proposed in an earlier round, the replica checks the set of missing pairs (undirected edges) \mathcal{E} sent by the leader of the previous round to see if t is part of any missing pairs. For any pair $(t', t) \in \mathcal{E}$, if the replica has already received t' , the replica adds (t', t) to the set of *updated local ordering* U_r to specify the order between t' and t . Updated ordering includes a set of edges between transactions of proposals received in previous

Algorithm 1 Local ordering on replica r

Input: (1) a set of incoming transactions in round i ,
(2) the set of missing pairs (i.e., undirected edges) \mathcal{E}

- 1: Initiate an empty graph $G_r = (T_r, E_r)$
- 2: **for** every transaction t that is received in an earlier round but has not been proposed by a leader **do** ▷ First type
- 3: Add transaction (vertex) t to T_r
- 4: **for** every vertex $t' \in T_r$ **do**
- 5: **if** t and t' are data-dependent **then**
- 6: Add (t', t) to E_r
- 7: **while** round i has not been finished **do**
- 8: **for** every incoming transaction t received from clients **do**
- 9: **if** t is not proposed in an earlier round **then** ▷ Third type
- 10: Add a vertex t to G_r
- 11: **for** every vertex $t' \in T_r$ **do**
- 12: **if** t' and t are data-dependent **then**
- 13: Add (t', t) to E_r
- 14: **else if** t is a vertex in an earlier proposal **then** ▷ Second type
- 15: **for** every pair $(t', t) \in \mathcal{E}$ **do**
- 16: **if** t' is already received **then**
- 17: Add (t', t) to U_r
- 18: Send $\langle \langle \text{LOCAL-ORDER}, i, G_r, U_r \rangle_{\sigma_r}, T \rangle$ to the leader

rounds and enables the leader to finalize previous proposals by adding the missing edges.

At the end of round i , each replica r sends a signed local-order message $\langle \langle \text{LOCAL-ORDER}, i, G_r, U_r \rangle_{\sigma_r}, T \rangle$ including the dependency graph G_r and the set of updated local ordering U_r to the leader. The set of received requests T is piggybacked to keep local-order messages small. The local-order messages are then used by the leader in global ordering to prove that this set of requests has been received.

Figure 2 presents the local ordering phase of Rashnu on different replicas in two rounds. Data objects accessed (written) by each transaction are shown at the bottom of the figure. For example, in round 1, replica r_2 receives transactions in this order: $t_6 < t_5 < t_2 < t_3 < t_7$ and generates local dependency graph G_{r_2} by adding edges between data-dependent transactions, e.g., (t_2, t_3) as both write on data objects A and B. As shown, replicas receive different sets of transactions in each round and deal with different cases. For instance, replica r_3 receives transaction t_9 in round 1. However, since it has not been proposed by the leader, r_3 includes t_9 in its graph of round 2 as well (first type). Similarly, while replica r_1 receives transaction t_5 in round 2, the replica does not send t_5 to the leader because it has already been proposed by the leader of round 1 (second type). The set of updated local ordering U_r is empty in round 1 for all replicas because there is no prior proposal to be updated in round 1. However, since the graph of round 1 includes a missing edge ($\mathcal{E} = t_5, t_6$), nodes include the edge in their updated local ordering U_r of round 2 if they have received both transactions. Note that $U_{r_3} = \emptyset$ in round 2, because the node has not received transaction t_6 yet.

4.2 Global Ordering

In the global ordering phase, as demonstrated in Algorithm 2, the leader receives the local-order messages from different replicas and generates the final ordering of transactions. Since f replicas might be faulty and not send their local-order messages, the leader collects a quorum of $n-f$ local-order messages from different replicas (including itself) to generate the global order.

We define two types of transactions: fixed and pending. Transaction t is fixed if it appears in at least $n-2f$ local-order messages,

Algorithm 2 Global ordering on leader replica π

Input: (1) $n-f$ local-order messages received from different replicas,
(2) list of missing edges \mathcal{E}

- 1: **for** every transaction t in some local-order messages **do**
- 2: **if** t appears in at least $n-2f$ local-order messages **then**
- 3: Label t as fixed
- 4: **else if** t appears in at least $n(1-\gamma) + f + 1$ local-order messages **then**
- 5: Label t as pending
- 6: ▷ Step 1: global dependency graph generation
- 7: Initiate an empty graph $G = (V, E)$
- 8: Initiate updated global ordering dependency list $L_g = \emptyset$
- 9: **for** every fixed or pending transaction t **do**
- 10: add a vertex t to G
- 11: **for** each pair of vertices t and t' **do**
- 12: **if** $w(t, t') > w(t', t) \ \& \ w(t, t') \geq n(1-\gamma) + f + 1$ **then**
- 13: Add (t, t') to E
- 14: **else if** $w(t', t) > w(t, t') \ \& \ w(t', t) \geq n(1-\gamma) + f + 1$ **then**
- 15: Add (t', t) to E
- 16: ▷ Step 2: condensation graph generation
- 17: $G^C = (V^C, E^C) \leftarrow \text{CONDENSATION}(G)$
- 18: ▷ Step 3: graph pruning
- 19: **for** every pending vertex u **do**
- 20: **if** there is no fixed vertex v such that $(u, v) \in E^C$ **then**
- 21: Remove u from V^C
- 22: Remove all edges involving u from E^C
- 23: **for** every pair of data-dependent transactions t and t' with no edges **do**
- 24: **if** t and t' are in two different vertices of V^C **then**
- 25: Add (t, t') to \mathcal{E}
- 26: ▷ Step 4: DAG transitive reduction generation
- 27: $G^T = (V^T, E^T) \leftarrow \text{TRANSITIVEREDUCTION}(G^C)$
- 28: ▷ Step 5: updating previous proposals
- 29: **for** each pair of transactions t and t' in any U_i **do**
- 30: **if** $w(t, t') > w(t', t) \ \& \ w(t, t') > n(1-\gamma) + f + 1$ **then**
- 31: Add (t, t') to U_g
- 32: **else if** $w(t', t) > w(t, t') \ \& \ w(t', t) > n(1-\gamma) + f + 1$ **then**
- 33: Add (t', t) to U_g
- 34: ▷ Step 6: global order proposal
- 35: Send $\langle \text{GLOBAL-ORDER}, G^T, \mathcal{E}, L, U_g, \mathcal{L}_u \rangle_{\sigma_\pi}$ to all replicas

whereas t is pending if at least $n(1-\gamma) + f + 1$ (and less than $n-2f$) local-order messages include t . A fixed transaction has been received by sufficiently many replicas to be included in the final order. Specifically, since at most f local-order messages in the quorum might have been received from faulty replicas, the leader counts on only $n-2f$ messages and if $n-2f$ different replicas receive a transaction, the transaction is ordered safely. On the other hand, pending transactions are not received by enough replicas yet to finalize an order. However, an edge from a pending to a fixed transaction might occur. As a result, we keep such pending transactions in the leader proposal enabling the leader to propose more transactions. We show that this does not violate the order-fairness (Definition 3). In Figure 2, since $\gamma = 1$, $f = 1$ (and $n = 5$), a transaction t is fixed if it appears in 3 or 4 local orderings, e.g., transactions t_2, t_3, t_4, t_5 and t_6 in round 1, while transaction t is pending if only 2 local orderings include t , e.g., transactions t_1 and t_7 in round 1.

We also define a weight function $w : E \mapsto [0, n-f]$ to denote the number of local dependency graphs with a particular edge. Given a set \mathcal{L} of $n-f$ local dependency graphs and two transactions (vertices) t and t' in round i , $w(t, t')$ represents the number of graphs that include an edge from t to t' . Note that for any pair of data-independent transactions t and t' , $w(t, t') = 0$. In Figure 2 and in round 1, $w(t_3, t_4) = 3$ while $w(t_3, t_9) = 1$.

Step 1: global dependency graph generation. Upon receiving a quorum of $n-f$ local-order messages from different replicas in round i , the leader initiates an empty graph $G = (V, E)$ and adds

all fixed and pending transactions to its vertex set (lines 6-9 of Algorithm 2). For each pair of data-dependent transactions t and t' in V , the leader calculates both $w(t, t')$ and $w(t', t)$. If the maximum of $w(t, t')$ and $w(t', t)$ is equal or greater than $n(1 - \gamma) + f + 1$, then the corresponding edge will be added to G (lines 10-14).

The goal of this step is to include as many transactions as possible while making sure that the exclusion of any transaction from the proposal does not violate fairness (Definition 3). In particular, the order-fairness definition states that if γ fraction of replicas receives transaction t before transaction t' , then t' is not ordered before t . When at least $n(1 - \gamma) + f + 1$ replicas propose a particular order $t < t'$ for two transactions t and t' , at most, $n\gamma - 2f - 1$ replicas (from the set of $n - f$ replicas) can propose the reverse order $t' < t$, i.e., $(n - f) - (n(1 - \gamma) + f + 1) = n\gamma - 2f - 1$. This number of replicas is, nevertheless, still lower than the γ fraction of (honest) replicas (as discussed in lemma 3.1, only $n - 2f$ replicas within the quorum of $n - f$ replicas are guaranteed to be honest). As a result, if $n(1 - \gamma) + f + 1$ replicas propose an order $t < t'$, the order can be safely chosen and the order-fairness will not be violated even if all remaining replicas propose the reverse order $t' < t$.

Specifically, when $\gamma = 1$, based on the order-fairness definition, if all honest replicas (i.e., at least $n - 2f$ within a quorum of $n - f$) receive transaction t before t' , then t' should not be ordered before t . As a result, if the protocol observes that $f + 1$ replicas have received t before t' , it can safely order t before t' . This is because it becomes impossible for $n - 2f$ honest replicas within a quorum of $n - f$ to receive t and t' in the reverse order (i.e., $t' < t$), hence, order-fairness can not be violated. Moreover, if neither of the two possible orders satisfies order-fairness, i.e., some ($> f$) received $t < t'$ while others ($> f$) received $t' < t$, the protocol is free to choose one of the orders. As a result, deciding based on $f + 1$ local ordering does not violate order-fairness. In this case, the protocol chooses the order with the higher weight.

In round 1 of Figure 2, the leader adds fixed transactions t_2, t_3, t_4, t_5, t_6 and pending transactions t_1 and t_7 to the graph while transaction t_9 is not added as it is received by only one replica. The leader adds an edge between two transactions if the edge appears in at least $n(1 - \gamma) + f + 1 = 2$ local graphs, e.g., edge (t_2, t_4) is not added because it appears only in G_{r_1} .

Note that malicious replicas might add invalid edges, e.g., between independent transactions, to their graph. However, the leader can detect such edges. Even if the leader does not validate edges, since at most f replicas are Byzantine, no invalid edges will be added to the final graph. Similarly, if the leader is malicious and adds incorrect edges, its malicious behavior can be easily detected by replicas, as the leader must send the $n - f$ local ordering to replicas (as a proof).

Step 2: condensation graph generation. The graph generated by the leader might contain cycles (due to the Condorcet paradox). Rashnu batches transactions involved in a cycle and delivers them to replicas simultaneously. To determine the order of transactions, the final graph must be acyclic. To generate an acyclic graph from a cyclic graph, Rashnu uses the graph condensation technique. Given a graph G , to generate the condensation graph G^C of G , Rashnu first identifies the strongly connected components of G . Each strongly connected component intuitively represents either a

single vertex (transaction) or a cycle in graph G . More formally, a strongly connected component C is a maximal subset of vertices such that any two vertices of this subset are reachable from each other. The condensation of graph G is graph $G^C = (V^C, E^C)$ where each vertex $C \in V^C$ corresponds to a strongly connected component of graph G , and $(C_i, C_j) \in E^C$ if and only if there are two vertices $u \in C_i$ and $v \in C_j$ such that $(u, v) \in E$. Vertex C in G^C is fixed if it includes at least one fixed transaction. Otherwise, C is pending.

As shown in Figure 2, transactions t_2, t_3 and t_4 construct a cycle in graph G of round 1. As a result, the condensation graph G^C consists of 4 vertices (strongly connected components): single-transaction vertices t_1, t_5, t_6 , and t_7 and vertex C_1 consisting of t_2, t_3 and t_4 . The resulting graph G^C is acyclic.

Step 3: graph pruning. Once the condensation graph is generated, the next step is to remove pending transactions with no outgoing path to a fixed transaction. These transactions are removed because they have not been received by a sufficient number of replicas and they do not incorporate in determining the order of a fixed transaction, i.e., we initially add them to the graph because there might be a path from a pending to a fixed transaction, helping in determining the order. Given two data-dependent transactions t and t' where t is fixed, t' is pending, and $(t, t') \in E^C$. Since at least $n(1 - \gamma) + 1$ honest replicas have received t before t' , removing pending transaction t' does not violate fairness.

In Figure 2, transactions t_1 and t_7 are pending in round 1. However, since t_1 has outgoing paths to fixed vertex C_1 , we keep it in the graph; while t_7 is removed from the graph. Removing t_7 enables the next proposer to propose the order of t_7 freely if t_7 appears in a sufficient number of local graphs.

Finally, for every pair of data-dependent transactions t and t' with no edges in between, if t and t' are not in the same vertex of G^C (i.e., they are not part of the same cycle), the leader adds a pair (t, t') to \mathcal{E} . Maintaining the set of missing edges \mathcal{E} is necessary because determining the order of a missing edge might result in a new or extended cycle. Hence, a transaction should not be executed until the order of all its predecessor transactions in the graph is determined. We do not maintain missing edges between transactions of the same cycle because such edges do not contribute to transaction ordering (i.e., the transactions already constructed a cycle). In Figure 2, the edge between t_5 and t_6 in round 1 and the edge between t_8 and t_{10} in round 2 are missing.

Step 4: DAG transitive reduction generation. As an optimization, the generated graph can be simplified by removing the transitive edges. The transitive reduction of a directed graph is another directed graph that has the same reachability relation with the same vertices and as few edges as possible. A transitive reduction of a graph $G^C = (V^C, E^C)$ is graph $G^T = (V^T, E^T)$ where (1) $V^C = V^T$ (vertices are the same) and (2) for each pair of vertices v and u in G^T , there is a path from u to v in G^T if and only if there is a path from u to v in G^C [3]. Since G^C is finite and acyclic, its transitive reduction G^T is unique and is a sub-graph of G^C (i.e., the minimum equivalent graph). In Figure 2 and in round 2, graph G^T has 2 fewer edges compared to G^C . For instance, (t_7, t_{10}) is removed as t_{10} is reachable from t_7 through t_9 .

Step 5: updating previous proposals. The leader also needs to update the previous proposals by adding the missing edges between

Algorithm 3 Order finalization on replica r

Input: a global-order message received from the leader

- ▷ Step 1: Global order validation
- 1: Upon receiving $\langle \text{GLOBAL-ORDER}, G^T, \mathcal{E}, \mathcal{L}, U_g, \mathcal{L}_u \rangle_{\sigma_\pi}$
- 2: Validate G^T , \mathcal{E} and U_g using \mathcal{L} and \mathcal{L}_u
- ▷ Step 2: Establishing consensus
- 3: `HOTSTUFF`(GLOBAL-ORDER)
- ▷ Step 3: Transaction execution
- 4: **for** every edge (t_1, t_2) in U_g where $t_1, t_2 \in \text{Block } B_i$ **do**
- 5: Add (t_1, t_2) to $B_i.G^T$
- 6: **for** every vertex v in $B_i.G^T$ **do**
- 7: **if** B_{i-1} is marked as completed, all predecessors of v in $B_i.G^T$ are executed and there is no missing edge in \mathcal{E} involving v **then**
- 8: **if** v is a single transaction **then**
- 9: Execute transaction v
- 10: **else** ▷ v contains a Condorcet cycle
- 11: Let be v_1, v_2, \dots, v_n be a Hamiltonian cycle of v
- 12: Execute transactions in the specified order
- 13: **if** All transactions of Block $B_i.G^T$ are executed **then**
- 14: Mark block B_i as completed

data-dependent transactions. As explained earlier, each replica r sends a set of updated local ordering dependencies U_r to the leader in its local-order message. When the current leader receives an ordering dependency, i.e., an edge, between two data-dependent transactions t and t' by at least $n(1 - \gamma) + f + 1$ replicas on some previous proposal, the leader adds an edge to its updated global ordering dependencies list U_g . If the leader receives both (t, t') and (t', t) , each from at least $n(1 - \gamma) + f + 1$ replicas, the edge with the highest weight will be added to the list U_g . The leader also removes the edge from the list of missing edges \mathcal{E} . In Figure 2 and in round 2, since replica r_1 and r_2 ($f + 1$ replicas) send (t_6, t_5) in their U_r sets, the leader adds (t_6, t_5) to the list U_g .

Step 6: global order proposal. Once the graph is generated, the leader π multicasts a $\langle \text{GLOBAL-ORDER}, G^T, \mathcal{E}, \mathcal{L}, U_g, \mathcal{L}_u \rangle_{\sigma_\pi}$ message to all replicas. The global-order message includes the dependency graph G^T , the set of missing edges \mathcal{E} , the set \mathcal{L} of $n - f$ local dependency graphs received from different replicas, the updated global ordering dependencies list U_g , and the set \mathcal{L}_u of $n - f$ updated local ordering dependencies received from different replicas. The set \mathcal{L} and \mathcal{L}_u are included to enable replicas to verify the dependency graph and the lists constructed by the leader.

4.3 Order Finalization

The order finalization is performed by all replicas to finalize the order proposed by the leader and execute transactions. As shown in Algorithm 3, order finalization consists of three main steps. First, replicas validate the proposed order. Second, all replicas establish agreement on the proposed order using a BFT protocol, and finally, replicas execute transactions following the proposed order.

Step 1: global order validation. Upon receiving a global-order message from the leader, each replica validates graph G^T , list of missing edges \mathcal{E} and list of updated global ordering dependencies U_g (Algorithm 3, lines 1-2). To validate G^T , the replica ensures that fixed and pending transactions are labeled correctly, and edges are added only if the order is proposed by a sufficient number of replicas. Similarly, the replicas validate both \mathcal{E} and U_g lists by checking the missing edges of the current and previous proposals. If the global order is invalid, (honest) replicas do not participate in the consensus

protocol; hence, the leader will be eventually replaced and the new leader collects all local orders and generates the global order.

Step 2: establishing agreement. Once the global-order message is validated, replicas establish agreement on the proposed order using the utilized consensus protocol. The current deployment of Rashnu uses HotStuff [61] as the underlying BFT protocol, enabling us to compare Rashnu with existing order-fairness protocols, e.g., Themis [37].

Step 3: transaction execution. Once consensus is achieved, each replica updates the previous proposals by adding edges from U_g . Replicas start executing transactions of a block once all predecessor blocks are executed, i.e., marked as completed. A block B_i is marked as completed if its dependency graph $B_i.G^T$ has no missing edges, and all its transactions are executed. In Rashnu, replicas follow the edges in the final dependency graph in executing transactions and can execute data-independent transactions of a block in parallel. Each vertex of a block is a strongly connected component consisting of either a single transaction or a set of transactions that construct a cycle. For each vertex that is a cycle, first, a Hamiltonian cycle is identified. A Hamiltonian cycle is a cycle that visits each vertex exactly once. If the graph includes more than one Hamiltonian cycle, all replicas deterministically use one, e.g., based on transaction ids, and execute transactions of the cycle in that order (lines 6-12).

Note that while Rashnu executes data-independent transactions within a block in parallel, replicas still execute transactions block by block. Specifically, each replica waits for block B_i to be marked as completed before executing any transactions of B_{i+1} . This technique, while simplifying the execution process, could result in unnecessary latency for two main reasons. First, data-independent transactions can be executed in parallel. Hence, when transactions of a current block are being executed, there is no need for data-independent transactions of a successor block to wait. Second, when some edges are missing, no transactions from any successor blocks can be executed, even if the transaction order of a successor block does not depend on the missing edges. To address this issue, Rashnu can capture data dependencies across blocks. While the overhead of capturing data dependencies across blocks might be high in contentious workloads, it results in a significant performance gain in workloads with low to moderate contention.

4.4 Correctness Argument

This section briefly discusses the order-fairness, safety, and liveness of Rashnu. Some arguments are inspired by the correctness arguments of Themis [37].

LEMMA 4.1. If transaction t appears in $n - 2f$ local-order messages (i.e., fixed transaction), t is proposed by an honest leader.

PROOF. The leader includes all fixed and pending transactions in its graph G and only the graph pruning step removes pending vertices (with no outgoing path to a fixed vertex) from the condensation graph. Since a vertex in the condensation graph is pending if it does not contain any fixed transactions, fixed transactions will always be proposed by the leader. The leader needs to send all local orderings to replicas (as proof of construction of G). Hence, if it maliciously excludes a fixed transaction, replicas detect that and do not accept the proposal resulting in replacing the leader. \square

LEMMA 4.2. Given two data-dependent transactions t and t' in a valid leader proposal. The proposal includes either (t, t') or (t', t) .

PROOF. Rashnu assumes a partial synchrony model where at least $n - 2f$ replicas in the quorum (eventually) have sent both transactions to the leader. Since $n - 2f > 2(n(1 - \gamma) + f)$, at least $w(t, t')$ or $w(t', t)$ is equal or greater than $n(1 - \gamma) + f + 1$. Since the leader adds only one edge (the edge with higher weight) even when both $w(t, t')$ are $w(t', t)$ are equal or greater than $n(1 - \gamma) + f + 1$, the final graph includes either (t, t') or (t', t) but not both. \square

LEMMA 4.3. Graph G^T proposed by an honest leader is acyclic.

PROOF. While graph G might contain (Condorcet) cycles, each cycle is part of a vertex (i.e., strongly connected component) in the condensation graph G^C . Since graph pruning and transitive reduction generation steps do not introduce any new edges, the final graph is still acyclic. \square

LEMMA 4.4. Given two order-dependent transactions t and t' received in round i where t is fixed. If a valid leader proposal includes only t then there are at least $n(1 - \gamma) + 1$ honest replicas that have received t before t' .

PROOF. If transaction t' was a fixed transaction, the leader proposal must include it. Hence, t' is not fixed. If t' is not a pending transaction, it has been received by at most $n(1 - \gamma) + f$ replicas. Since t is fixed, it appears in $n - 2f$ local-order messages. As a result, $p = (n - 2f) - (n(1 - \gamma) + f) = \gamma n - 3f$ replicas ordered t before t' . $n > \frac{4f}{2\gamma - 1}$, hence, $p > n(1 - \gamma) + f$, from which at most f replicas might be faulty. Hence, at least $n(1 - \gamma) + 1$ honest replicas received t before t' . If t' is a pending transaction, since it is not included in the leader proposal, there is no path from t' to any fixed transactions that includes t . As a result, either (1) $w(t', t) \leq n(1 - \gamma) + f$ or (2) $n(1 - \gamma) + f + 1 \leq w(t', t) \leq w(t, t')$. The second case implies that at least $n(1 - \gamma) + 1$ honest replicas received t before t' . In the first case, since t is fixed, $w(t, t') \geq (n - 2f) - (n(1 - \gamma) + f) = \gamma n - 3f > n(1 - \gamma) + f$ since $n > \frac{4f}{2\gamma - 1}$. \square

LEMMA 4.5. Given two order-dependent transactions t and t' , if a valid leader proposal includes only t (and t' was not in an earlier proposal) and t' is received before t by γn replicas, t and t' are in the same Condorcet cycle.

PROOF. Since t is in the proposal, it is either fixed or pending. If t is a fixed transaction, there are at least $n(1 - \gamma) + 1$ honest replicas that have received t before t' (lemma 4.4); which contradict the condition (i.e., t' is received before t by γn replicas). Hence, t is pending. Since t is pending there is a path t, t_1, t_2, \dots, t_k from t to some fixed transaction t_k in the proposal. Since t_k is a fixed transaction, based on lemma 4.4, there are at least $n(1 - \gamma) + 1$ honest replicas that have received t_k before t' . Since t' is received before t by γn replicas, $t, t_1, t_2, \dots, t_k, t'$ construct a Condorcet cycle. \square

THEOREM 4.6. Rashnu guarantees data-dependent order-fairness.

PROOF. Given two data-dependent transactions t and t' where γn replicas receive t before t' . Four cases can happen in the final ordering. First, t and t' are proposed in the same valid block and in different vertices of the graph G^T (then, following Algorithm 2, t'

is not ordered before t since γn replicas receive t before t'). Second, t and t' are proposed in the same valid block and in the same vertex of the graph G^T (then, t' is not ordered before t – both are part of the same cycle). Third, t' is proposed in a later valid block than t (then, t' is obviously not ordered before t), and fourth, t is proposed in a later valid block than t' (then, based on lemma 4.5, t and t' are in the same Condorcet cycle and t' is not ordered before t). Hence, Rashnu guarantees data-dependent order-fairness. \square

THEOREM 4.7. Rashnu guarantees safety.

PROOF. The safety of Rashnu is a direct consequence of the safety of HotStuff [61], as Rashnu does not modify any phases of the underlying agreement protocol. \square

THEOREM 4.8. Rashnu guarantees liveness.

PROOF. Rashnu considers a partial synchrony model where a correct client transaction t will be (eventually) received by all replicas. As a result, t will appear in at least $n - 2f$ local-order messages (either in the same round or different rounds), becomes a fixed transaction, and is proposed by a leader. The execution of t only depends on missing edges between previously proposed order-dependent transactions and as soon as such edges are added (i.e., corresponding transactions appear in $n - 2f$ local-order messages), t can be executed. However, the order of transaction t does not depend on any transaction that has not been proposed. Hence, in contrast to some other fair ordering protocols like Aequitas [38], Condorcet cycles can not be chained and violate liveness. \square

5 EXPERIMENTAL EVALUATION

Our evaluation has two main goals. First, measuring the overhead of supporting order-fairness in Rashnu compared to its (unfair) underlying BFT protocol, HotStuff [61]. Second, Comparing the performance of Rashnu and the state-of-the-art order-fairness protocol; Themis [37] (the only existing order-fairness protocol that provides the same guarantees as Rashnu without requiring clock synchronization or any optimistic assumption on the honesty of nodes). To this end, we analyze the impact of the following parameters on the performance of HotStuff, Rashnu, and Themis:

- (1) transaction batch size (Section 5.1),
- (2) degrees of contention (Section 5.2),
- (3) network size (Section 5.3),
- (4) order-fairness parameter (Section 5.4),
- (5) compute-intensity of workloads (Section 5.5),
- (6) geo-distribution of replicas (Section 5.6), and
- (7) workload (Section 5.7).

Varying these parameters enables us to simulate many real BFT application workloads [32]. Our Rashnu implementation is bootstrapped from the HotStuff protocol [61]. We use the author's open-source libhotstuff codebase [1] and implemented Rashnu on top of that. We mainly change the HotStuff codebase by enabling the leader to generate a fair order and the replicas to send their local order to the leader and validate the proposed order. Other phases of the HotStuff protocol remain untouched. We have also implemented Themis on top of HotStuff in the same way as Rashnu to enable a fair comparison. We perform most experimental evaluations under the SmallBank benchmark. We initially populate the

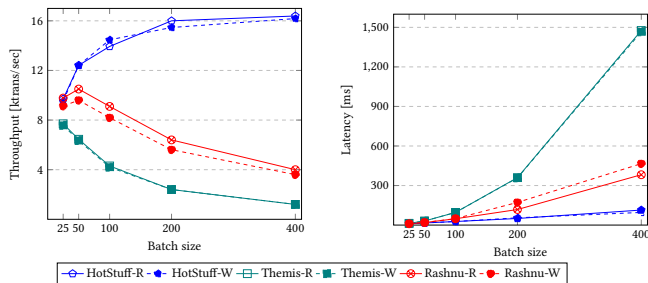


Figure 3: Impact of batch size

system with 10000 records and run each protocol under read-heavy ($P_w = 0.05$) and write-heavy ($P_w = 0.95$) workloads, e.g., Rashnu-R is Rashnu under the read-heavy, and Rashnu-W is Rashnu under the write-heavy workload. To determine the accounts accessed by each transaction, a Zipfian distribution is followed, which can be configured in terms of skewness, e.g., $s = 0$ corresponds to a uniform distribution. We further evaluate Rashnu under different YCSB workloads, as demonstrated in Section 5.7.

We run our experiments on a set of c6220 bare-metal machines on CloudLab [25], each with two Xeon E5-2650v2 processors (8 cores each, 2.6GHz), 64GB RAM and two 1TB SATA 3.5" 7.2K rpm hard drives. These machines are connected by two networks, each with one interface: (1) a 1 Gbps Ethernet control network; (2) a 10 Gbps Ethernet commodity fabric. We report latency and throughput. The results reflect end-to-end measurements from the clients.

5.1 Performance with Different Batch Sizes

In the first set of experiments, we measure the impact of transaction batch size on the performance of different protocols. In this set of experiments, the number of replicas is assumed to be 5 ($4f + 1$ where $f = 1$), the order-fairness parameter $\gamma = 1$, and the account selection follows a uniform distribution. Figure 3 depicts the results for all three protocols. The solid and dashed lines are used for read-heavy and write-heavy workloads, respectively. When blocks are small (block size = 25), Rashnu provides fairness with zero overhead, i.e., Rashnu processes 9675 tps while HotStuff processes 9700 tps, both with 10 ms latency. This is because the cost of generating small dependency graphs is insignificant compared to running consensus among replicas. Increasing the block size, however, results in a gap between HotStuff and Rashnu due to the overhead of fair transaction ordering; generating local dependency graphs, constructing the global dependency graph, and validating the final order. Nevertheless, with block size = 50, Rashnu incurs only 15% throughput overhead; while with the same setting, Themis suffers from 45% throughput overhead. Further increasing the block size makes the gap between HotStuff and Rashnu larger. This is expected because Rashnu must construct larger local and global graphs, requiring checking more and more dependencies. However, compared to Themis, Rashnu shows 233% higher throughput and 74% lower latency with block size 400. This is because, with a larger block size, generating the global order becomes more costly in Themis compared to Rashnu, as Rashnu considers ordering only among data-dependent transactions. The type of workload has a negligible impact on the performance of HotStuff and Themis, as expected. The performance of Rashnu, however, is reduced by 7%




Figure 4: Impact of workload contention

to 10% in write-heavy workloads with different block sizes, compared to read-heavy ones, as Rashnu needs to capture more data dependencies between transactions.

The best batch size could be affected by parameters such as the network size, the geo-distribution, and even the computation power of nodes. Moreover, the choice might differ over time due to the dynamic nature of distributed systems, even for the same protocol and environment. To have a fair comparison, for the remaining set of experiments, we set the batch size to 100 (a number where all protocols demonstrate fairly good performance).

5.2 Varying the Degree of Contention

In the next set of experiments, we study the impact of the workload contention by changing the Zipfian skewness of the smallbank benchmark from $s = 0.01$ (uniform distribution) to $s = 0.5$ and $s = 0.99$ (contentious workload). In this set of experiments, the batch size is 100, $n = 5$, and $\gamma = 1$.

As shown in Figure 4, the performance of HotStuff and Themis is not affected by increasing the workload skewness since they do not construct dependency graphs. Rashnu, however, shows $\sim 10\%$ higher latency (in both read-heavy and write-heavy workloads) and 2% and 10% throughput reduction in read-heavy and write-heavy workloads when we increase the Zipfian skewness from $s = 0.01$ to $s = 0.99$. Overall, Rashnu incurs 22% throughput reduction and 27% higher latency by going from a uniform read-heavy workload to a skewed write-heavy workload. This is the overhead of constructing local and global dependency graphs by replicas and the leader. Note that, even with $s=0.99$ and $P_w=0.95$, Rashnu demonstrates 34% higher throughput and 31% lower latency compared to Themis.

5.3 Performance with Different Network Size

In the third set of experiments, we measure the performance of Rashnu in networks with different sizes, i.e., 5, 21, 41, 61, 81, and 101. We consider request batches of size 100, $\gamma = 1$, and uniform account selection. Since the type of workloads (i.e., read-heavy and write-heavy), does not have a significant impact on the performance of HotStuff and Themis, we only report the results of the read-heavy workload ($P_w = 0.05$) for those two protocols. As depicted in Figure 5, increasing the number of replicas significantly reduces the performance of all three protocols. This is expected because establishing consensus among a large set of replicas is expensive due to the high communication cost.

Interestingly, with more than 20 replicas, Rashnu demonstrates almost the same performance as HotStuff; providing order-fairness




Figure 5: Impact of the network size




Figure 6: Impact of the order-fairness parameter

for free (the same happens for Themis with $n > 55$). This is because the cost of consensus in HotStuff becomes much higher than the overhead of fair ordering by Rashnu and Themis in a large network. This also shows that the low performance of Rashnu and even Themis in large networks is caused by the HotStuff consensus routine and if Rashnu is bootstrapped from a high-performance protocol, it produces better results.

5.4 Varying the Order-fairness Parameter

We next run Themis and Rashnu under different values of the order-fairness parameter γ . Since the network size is a function of the order-fairness parameter (i.e., $n = \frac{4f}{2\gamma-1} + 1$), reducing γ requires a larger n . Specifically, with $f = 1$, we evaluate protocols with $\gamma = 1$ ($n = 5$), $\gamma = 0.9$ ($n = 6$), $\gamma = 0.75$ ($n = 9$), $\gamma = 0.6$ ($n = 21$), and $\gamma = 0.55$ ($n = 41$). In all experiments, the batch size is 100, and the account selection is uniform. The results for Rashnu-R, Rashnu-W, and Themis are shown in Figure 6.

Similar to Figure 5, increasing the number of replicas (resulting from reducing γ) decreases the overall performance. Since the cost of communication dominates the fair ordering overhead, the gap between Rashnu and Themis becomes smaller by increasing γ ; while Rashnu-R demonstrates 75% higher throughput than Themis with $\gamma = 1$ ($n = 5$), it shows only 19% higher throughput with $\gamma = 0.55$ ($n = 41$).

5.5 Performance with compute intensive workloads

One main advantage of Rashnu is its ability to execute transactions in parallel (following the dependency graph generated in the ordering phase). To demonstrate this benefit, we consider a compute-intensive (busy-wait) workload where executing each transaction




Figure 7: Performance with compute-intensive workloads




Figure 8: Impact of batch size in a distributed setting

takes around $10\mu s$. We repeat experiments of Section 5.2 and change the Zipfian skewness from 0.01 to 0.5 and 0.99. As shown in Figure 7, Rashnu, by executing data-independent transactions in parallel (with eight threads), demonstrates better performance compared to even HotStuff, providing order-fairness for free, i.e., with $s = 0.01$, Rashnu shows 3% higher throughput compared to HotStuff. This demonstrates that the throughput gain of parallel execution outnumbers the overhead caused by dependency graph generation.

5.6 Performance in a Geo-distributed Setting

In this set of experiments, we measure the performance of protocols in an emulated geo-distributed setup. We repeat the first set of experiments (Section 5.1) with an extra 50 ms latency (injected by Linux netem) for sending messages between any pair of replicas. The results are shown in Figure 8.

As expected, the throughput of all three protocols decreases once network latency is added. Interestingly, Rashnu and Themis reach their peak performance on larger block sizes, compared to the local setting (Figure 3). Specifically, while Rashnu demonstrated its best throughput with block size 50 in the local setting, it shows its highest throughput with a block size of 400 in the distributed setting. Similarly, Themis shows its best throughput on the block size of 200.

This demonstrates a trade-off between communication latency and fair ordering. With large latency, the cost of communication becomes higher than the overhead of fair ordering, hence, large block sizes are beneficial. While with low communication latency, the overhead of fair ordering is much higher, thus smaller blocks give better performance. While with a block size of 400, Rashnu incurs 45% throughput reduction compared to HotStuff, its throughput is still 2.5 times the throughput of Themis. In this setting, Rashnu processes transactions with 61% lower latency, compared to Themis.




Figure 9: Performance under YCSB workload

5.7 Performance with YCSB Workloads

In the last set of experiments, we study the performance of different protocols under YCSB workloads. We have chosen 5 different (most relevant) workloads from the YCSB benchmark: A: update heavy workload (A1: R/W = 5/95 and A2: R/W = 50/50), B: read-mostly workload (R/W = 95/5), C: read-only workload (R/W = 100/0), and F: read-modify-write workload. In this set of experiments, the batch size is 100, $n = 5$, $\gamma = 1$, and Zipfian skewness is $s = 0.01$ (uniform distribution).

As shown in Figure 9, Rashnu shows between 14% to 26% lower throughput compared to Hotstuff, while its throughput is still 46% to 61% higher than Themis under different workloads. As expected, Rashnu demonstrates its best performance under the read-only workload (C). Overall, the results are consistent with the results of the SmallBank benchmark.

6 RELATED WORK

Order-fairness has been recently studied by a few protocols, e.g., Wendy [43, 44], Pompe [62], Aequitas [38], Themis [37] and Quick order fairness [14]. Both Wendy and Pompe rely on synchronized clocks between replicas, making these protocols impractical in asynchronous networks. In Wendy, all replicas have access to synchronized local clocks and if all honest replicas receive transaction t before t' , then t is delivered before t' . Pompe, on the other hand, uses a pre-ordering phase and determines the fair order using timestamps assigned by replicas in the pre-ordering phase. Specifically, Pompe orders transaction using their median timestamp. The median timestamp, however, can easily be manipulated by a malicious node that assigns a big timestamp. Moreover, Pompe is vulnerable to censorship [37]. Furthermore, both notions of timed order fairness, used in Wendy, and ordering linearizability, used in Pompe, are strictly weaker than order-fairness studied in Aequitas [38] and Themis [37], as stated in [36]. Aequitas [38] presents the notion of batch-order fairness where all transactions involved in a cycle are delivered to replicas in the same batch. While Aequitas circumvents the Condorcet paradox, it suffers from a liveness issue when Condorcet cycles chain together and extend for an arbitrarily long time. A subsequent study extends the Aequitas approach to permissionless settings [36]. Quick-order-fairness [14] leverages batch-order fairness and also introduces the notion of differential order fairness, inspired by the differential validity notion of consensus. Differential order fairness states that when the number of honest replicas that send transaction t before transaction t' is at least $2f + k$ more than the number of replicas that send t' before t

for some order-fairness parameter $k \geq 0$, the protocol must not deliver t' before t . However, similar to Aequitas, Quick-order-fairness suffers from liveness issues. Moreover, Aequitas and Quick-order-fairness have not been validated by any system implementation. Themis [37], extends Aequitas by addressing its weak liveness issue. Themis introduces the notion of deferred ordering where the actual order of transactions might be deferred to a later proposal, enabling the leader to propose a block without waiting. However, Themis suffers from significant performance overhead, as shown in Section 5, resulting from its complex fair ordering routine. Compared to these protocols, Rashnu addressed both Condorcet cycles and weak liveness and demonstrates high performance.

The fair ordering of transactions has been partially addressed in a few BFT protocols. In Aardvark [18], the leader is monitored to ensure that it does not initiate two new requests from the same client before initiating an old request of another client. Similarly, in PBFT [16], replicas keep the requests in a FIFO queue and only stop the view-change timer when the first request in their queue is executed. Prime [4] introduces a pre-ordering phase where replicas order the received requests locally and share their ordering with each other. In Hashgraph [6], all replicas construct a hashgraph to capture all send and receive events. These protocols, however, do not address challenges like Condorcet cycles.

Fairness has also been used in the domain of consensus with different definition. In permissionless blockchains, e.g., Proof-of-Work, fairness is used to ensure that the mining rewards obtained by different miners are proportional to their relative computational power [2, 47, 49, 51, 54]. Similarly, fairness has been defined as providing opportunities for every replica to propose and commit its requests using fair leader election or fair committee election [2, 5, 31, 39, 47, 55, 60]. However, a malicious leader can still order transactions unfairly in its turn.

7 CONCLUSION

This paper defines the notion of data-dependent order-fairness. We presented a high-performance fair-ordering protocol, Rashnu, that leverages graph-based techniques to achieve order-fairness among data-dependent transactions. Rashnu further utilizes batch ordering and deferred ordering techniques to deal with Condorcet cycles and liveness issues. We implemented a prototype of Rashnu on top of HotStuff and open-sourced its code. Our evaluation demonstrates the efficiency of Rashnu in different scenarios. First, with small batch sizes or in large networks, the overhead of order-fairness in Rashnu is negligible, i.e., Rashnu performs similarly to its underlying consensus protocol HotStuff. Second, Rashnu shows significant performance improvement compared to Themis in different settings, especially in small networks, 27% to 233% throughput improvement on 5 replicas and with varying batch sizes. Finally, in compute-intensive workloads, Rashnu even outperforms HotStuff (by executing transactions in parallel).

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful feedback and suggestions. This work is funded by NSF grants CNS-2104882 and CNS-2107147.

REFERENCES

- [1] 2018. libhotstuff: A general-purpose BFT state machine replication library with modularity and simplicity. <https://github.com/hot-stuff/libhotstuff>.
- [2] Ittai Abraham, Dahlia Malkhi, Karik Nayak, Ling Ren, and Alexander Spiegelman. 2017. Solida: A Blockchain Protocol Based on Reconfigurable Byzantine Consensus. In *Int. Conf. on Principles of Distributed Systems (OPODIS)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [3] Alfred V. Aho, Michael R Garey, and Jeffrey D. Ullman. 1972. The transitive reduction of a directed graph. *SIAM J. Comput.* 1, 2 (1972), 131–137.
- [4] Yair Amir, Brian Coan, Jonathan Kirsch, and John Lane. 2011. Prime: Byzantine replication under attack. *Transactions on Dependable and Secure Computing* 8, 4 (2011), 564–577.
- [5] Avi Asayag, Gad Cohen, Ido Grayevsky, Maya Leshkowitz, Ori Rottenstreich, Ronen Tamari, and David Yakira. 2018. A fair consensus protocol for transaction ordering. In *Int. Conf. on Network Protocols (ICNP)*. IEEE, 55–65.
- [6] Leemon Baird. 2016. The swirls hashgraph consensus algorithm: Fair, fast, byzantine fault tolerance. *Swirls Tech Reports SWIRLDS-TR-2016-01, Tech. Rep* (2016).
- [7] Jason Baker, Chris Bond, James C Corbett, JJ Furman, Andrey Khorlin, James Larson, Jean-Michel Leon, Yawei Li, Alexander Lloyd, and Vadim Yushprakh. 2011. Megastore: Providing scalable, highly available storage for interactive services. In *Conf. on Innovative Data Systems Research (CIDR)*.
- [8] Carsten Baum, James Hsin-yu Chiang, Bernardo David, Tore Kasper Frederiksen, and Lorenzo Gentile. 2021. Sok: Mitigation of front-running in decentralized finance. *Cryptology ePrint Archive* (2021).
- [9] Kenneth P Birman, Thomas A Joseph, Thomas Rauechle, and Amr El Abbadi. 1985. Implementing fault-tolerant distributed objects. *Trans. on Software Engineering* 6 (1985), 502–508.
- [10] Gabriel Bracha. 1984. An asynchronous $(n-1)/3$ -resilient consensus protocol. In *Symposium on Principles of Distributed Computing (PODC)*. ACM, 154–162.
- [11] Gabriel Bracha and Sam Toueg. 1985. Asynchronous consensus and broadcast protocols. *Journal of the ACM (JACM)* 32, 4 (1985), 824–840.
- [12] Nathan Bronson, Zach Amsden, George Cabrera, Prasad Chakka, Peter Dimov, Hui Ding, Jack Ferris, Anthony Giardullo, Sachin Kulkarni, and Harry Li. 2013. TAO: Facebook’s Distributed Data Store for the Social Graph. In *Annual Technical Conf (ATC)*. USENIX Association, 49–60.
- [13] Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. 2001. Secure and efficient asynchronous broadcast protocols. In *Annual Int. Cryptology Conf*. Springer, 524–541.
- [14] Christian Cachin, Jovana Micić, and Nathalie Steinhauer. 2022. Quick Order Fairness. In *Int. Conf. on Financial Cryptography and Data Security (FC)*. Springer, 1–18.
- [15] Miguel Castro and Barbara Liskov. 1999. Practical Byzantine fault tolerance. In *Symposium on Operating Systems Design and Implementation (OSDI)*. USENIX Association, 173–186.
- [16] Miguel Castro and Barbara Liskov. 2002. Practical Byzantine fault tolerance and proactive recovery. *Transactions on Computer Systems (TOCS)* 20, 4 (2002), 398–461.
- [17] Chainlink. 2023. What Is Maximal Extractable Value (MEV)? <https://chain.link/education-hub/maximal-extractable-value-mev>.
- [18] Allen Clement, Edmund L Wong, Lorenzo Alvisi, Michael Dahlin, and Mirco Marchetti. 2009. Making Byzantine Fault Tolerant Systems Tolerate Byzantine Faults. In *Symposium on Networked Systems Design and Implementation (NSDI)*, Vol. 9. USENIX Association, 153–168.
- [19] James C Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, Jeffrey John Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, and Peter Hochschild. 2013. Spanner: Google’s globally distributed database. *Transactions on Computer Systems (TOCS)* 31, 3 (2013), 8.
- [20] Miguel Correia, Nuno Ferreira Neves, and Paulo Verissimo. 2006. From consensus to atomic broadcast: Time-free Byzantine-resistant protocols without signatures. *Comput. J.* 49, 1 (2006), 82–96.
- [21] Tyler Crain, Christopher Natoli, and Vincent Gramoli. 2021. Red Belly: a secure, fair and scalable open blockchain. In *Symposium on Security and Privacy (SP)*. IEEE.
- [22] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *Symposium on Security and Privacy (SP)*. IEEE, 910–927.
- [23] George Danezis, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. 2022. Narwhal and Tusk: a DAG-based mempool and efficient BFT consensus. In *European Conf. on Computer Systems (EuroSys)*. 34–50.
- [24] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Voshall, and Werner Vogels. 2007. Dynamo: Amazon’s highly available key-value store. *Operating Systems Review (OSR)* 41, 6 (2007), 205–220.
- [25] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, et al. 2019. The Design and Operation of {CloudLab}. In *Annual Technical Conf. (ATC)*. USENIX Association, 1–14.
- [26] 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.
- [27] Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. 2019. Sok: Transparent dishonesty: front-running attacks on blockchain. In *Int. Conf. on Financial Cryptography and Data Security (FC)*. Springer, 170–189.
- [28] Jose M Faleiro, Daniel J Abadi, and Joseph M Hellerstein. 2017. High performance transactions via early write visibility. *Proc. of the VLDB Endowment* 10, 5 (2017), 613–624.
- [29] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. 1985. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)* 32, 2 (1985), 374–382.
- [30] Matthias Fitzi and Juan A Garay. 2003. Efficient player-optimal protocols for strong and differential consensus. In *Symposium on Principles of Distributed Computing (PODC)*. 211–220.
- [31] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. 2017. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Symposium on Operating Systems Principles (SOSP)*. ACM, 51–68.
- [32] Vincent Gramoli, Rachid Guerraoui, Andrei Lebedev, Chris Natoli, and Gauthier Voron. 2023. Diablo: A Benchmark Suite for Blockchains. In *European Conf. on Computer Systems (EuroSys)*. ACM.
- [33] Lioba Heimbach and Roger Wattenhofer. 2022. SoK: Preventing Transaction Reordering Manipulations in Decentralized Finance. In *Conf. on Advances in Financial Technologies (AFT)*. ACM, 1–14.
- [34] Robert Kallman, Hideaki Kimura, Jonathan Natkins, Andrew Pavlo, Alexander Rasin, Stanley Zdonik, Evan PC Jones, Samuel Madden, Michael Stonebraker, and Yang Zhang. 2008. H-store: a high-performance, distributed main memory transaction processing system. *Proc. of the VLDB Endowment* 1, 2 (2008), 1496–1499.
- [35] Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. 2021. All you need is dag. In *Symposium on Principles of Distributed Computing (PODC)*. ACM, 165–175.
- [36] Mahimna Kelkar, Soubhik Deb, and Sreeram Kannan. 2022. Order-fair consensus in the permissionless setting. In *ASIA Public-Key Cryptography Workshop*. ACM, 3–14.
- [37] Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. 2023. Themis: Fast, strong order-fairness in byzantine consensus. In *SIGSAC Conf. on Computer and Communications Security (CCS)*. ACM, 475–489.
- [38] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. 2020. Order-fairness for byzantine consensus. In *Annual Int. Cryptology Conf*. Springer, 451–480.
- [39] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynikov. 2017. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual Int. Cryptology Conf*. Springer, 357–388.
- [40] Ariah Klages-Mundt and Andreea Minca. 2019. (In) stability for the blockchain: Deleveraging spirals and stablecoin attacks. *arXiv preprint arXiv:1906.02152* (2019).
- [41] Eleftherios Kokoris Kogias, Philipp Jovanovic, Nicolas Gailly, Ismail Khoffi, Linus Gasser, and Bryan Ford. 2016. Enhancing bitcoin security and performance with strong consistency via collective signing. In *Security Symposium*. USENIX Association, 279–296.
- [42] 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 *Symposium on Security and Privacy (SP)*. IEEE, 583–598.
- [43] Klaus Kursawe. 2020. Wendy, the good little fairness widget: Achieving order fairness for blockchains. In *Conf. on Advances in Financial Technologies (AFT)*. ACM, 25–36.
- [44] Klaus Kursawe. 2021. Wendy Grows Up: More Order Fairness. In *Int. Conf. on Financial Cryptography and Data Security (FC)*. Springer, 191–196.
- [45] Leslie Lamport. 1978. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM* 21, 7 (1978), 558–565.
- [46] Leslie Lamport, Robert Shostak, and Marshall Pease. 1982. The Byzantine generals problem. *Transactions on Programming Languages and Systems (TOPLAS)* 4, 3 (1982), 382–401.
- [47] Kfir Lev-Ari, Alexander Spiegelman, Idit Keidar, and Dahlia Malkhi. 2019. FairLedger: A Fair Blockchain Protocol for Financial Institutions. In *Int. Conf. on Principles of Distributed Systems (OPODIS)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [48] Loi Luu, Viswesh Narayanan, Chaodong Zheng, Kunal Baweja, Seth Gilbert, and Prateek Saxena. 2016. A secure sharding protocol for open blockchains. In *SIGSAC Conf. on Computer and Communications Security (CCS)*. ACM, 17–30.
- [49] Loi Luu, Yaron Velner, Jason Teutsch, and Prateek Saxena. 2017. SmartPool: Practical Decentralized Pooled Mining. In *USENIX Security Symposium*. USENIX, 1409–1426.
- [50] Dahlia Malkhi and Pawel Szalachowski. 2022. Maximal Extractable Value (MEV) Protection on a DAG. *arXiv preprint arXiv:2208.00940* (2022).
- [51] Andrew Miller, Ahmed Kosba, Jonathan Katz, and Elaine Shi. 2015. Nonoutsourcable scratch-off puzzles to discourage bitcoin mining coalitions. In *ACM SIGSAC Conf. on Computer and Communications Security (CCS)*. ACM SIGSAC,

680–691.

- [52] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. 2016. The honey badger of BFT protocols. In *Conf. on Computer and Communications Security (CCS)*. ACM, 31–42.
- [53] Louise E Moser, Peter M Melliar-Smith, Priya Narasimhan, Lauren A Tewksbury, and Vana Kalogeraki. 1999. The Eternal system: An architecture for enterprise applications. In *Int. Enterprise Distributed Object Computing Conf. (EDOC)*. IEEE, 214–222.
- [54] Rafael Pass and Elaine Shi. 2017. Fruitchains: A fair blockchain. In *symposium on Principles of Distributed Computing (PODC)*. ACM, 315–324.
- [55] Rafael Pass and Elaine Shi. 2017. Hybrid Consensus: Efficient Consensus in the Permissionless Model. In *Int. Symposium on Distributed Computing (DISC)*. 6.
- [56] Kaihua Qin, Liyi Zhou, and Arthur Gervais. 2022. Quantifying blockchain extractable value: How dark is the forest?. In *Symposium on Security and Privacy (SP)*. IEEE, 198–214.
- [57] Fred B Schneider. 1990. Implementing fault-tolerant services using the state machine approach: A tutorial. *Computing Surveys (CSUR)* 22, 4 (1990), 299–319.
- [58] Alexander Spiegelman, Neil Girdharan, Alberto Sonnino, and Lefteris Kokoris-Kogias. 2022. Bullshark: Dag bft protocols made practical. In *ACM SIGSAC Conf. on Computer and Communications Security (CCS)*. 2705–2718.
- [59] Chrysoula Stathakopoulou, Signe Rüsçh, Marcus Brandenburger, and Marko Vukolić. 2021. Adding Fairness to Order: Preventing Front-Running Attacks in BFT Protocols using TEEs. In *Int. Symp on Reliable Distributed Systems (SRDS)*. IEEE, 34–45.
- [60] David Yakira, Avi Asayag, Gad Cohen, Ido Grayevsky, Maya Leshkowitz, Ori Rottenstreich, and Ronen Tamari. 2021. Helix: A Fair Blockchain Consensus Protocol Resistant to Ordering Manipulation. *IEEE Transactions on Network and Service Management* 18, 2 (2021), 1584–1597.
- [61] Maofan Yin, Dahlia Malkhi, Michael K Reiter, Guy Golan Gueta, and Ittai Abraham. 2019. HotStuff: BFT consensus with linearity and responsiveness. In *Symposium on Principles of Distributed Computing (PODC)*. ACM, 347–356.
- [62] Yunhao Zhang, Srinath Setty, Qi Chen, Lidong Zhou, and Lorenzo Alvisi. 2020. Byzantine ordered consensus without Byzantine oligarchy. In *Symposium on Operating Systems Design and Implementation (OSDI)*. USENIX Association, 633–649.
- [63] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. 2021. High-frequency trading on decentralized on-chain exchanges. In *Symposium on Security and Privacy (SP)*. IEEE, 428–445.