

Maximal Extractable Value (MEV) Protection on a DAG

Dahlia Malkhi

Chainlink Labs, UK

Pawel Szalachowski

Chainlink Labs, UK

Abstract

Many cryptocurrency platforms are vulnerable to Maximal Extractable Value (MEV) attacks [12], where a malicious consensus leader can inject transactions or change the order of user transactions to maximize its profit.

A promising line of research in MEV mitigation is to enhance the Byzantine fault tolerance (BFT) consensus core of blockchains by new functionalities, like hiding transaction contents, such that malicious parties cannot analyze and exploit them until they are ordered. An orthogonal line of research demonstrates excellent performance for BFT protocols designed around Directed Acyclic Graphs (DAG). They provide high throughput by keeping high network utilization, decoupling transactions' dissemination from their metadata ordering, and encoding consensus logic efficiently over a DAG representing a causal ordering of disseminated messages.

This paper explains how to combine these two advances. It introduces a DAG-based protocol called Fino, that integrates MEV-resistance features into DAG-based BFT without delaying the steady spreading of transactions by the DAG transport and with zero message overhead. The scheme operates without complex secret share verifiability or recoverability, and avoids costly threshold encryption.

2012 ACM Subject Classification Security and privacy → Distributed systems security

Keywords and phrases DAG, MEV, consensus, BFT

Digital Object Identifier 10.4230/OASICS.Tokenomics.2022.6

Acknowledgements We are grateful to Soumya Basu, Christian Cachin, Ari Juels, Mahimna Kelkar, Lefteris Kokoris-Kogias, Oded Naor, Mike Reiter for many comments that helped improve this writeup.

1 Introduction

1.1 MEV

Over the last few years, we have seen exploding interest in cryptocurrency platforms and applications built upon them, like decentralized finance protocols offering censorship-resistant and open access to financial instruments; or non-fungible tokens. Many of these systems are vulnerable to MEV attacks, where a malicious consensus leader can inject transactions or change the order of user transactions to maximize its profit. Thus it is not surprising that at the same time we have witnessed rising phenomena of MEV professionalization, where an entire ecosystem of MEV exploitation, comprising of MEV opportunity “searchers” and collaborating miners, has arisen.

Daian et al. [12] introduced a measure of the “profit that can be made through including, excluding, or re-ordering transactions within blocks”. The original work called the measure miner extractable value, which was later extended by maximal extractable value (MEV) [33] and blockchain extractable value (BEV) [36], to include other forms of attacks, not necessarily performed by miners. At the time of this writing, an “MEV-explore” tool [1] estimates the amount of MEV extracted on Ethereum since the 1st of Jan 2020 to be close to \$700M.



© Dahlia Malkhi and Pawel Szalachowski;
licensed under Creative Commons License CC-BY 4.0

4th International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2022).

Editors: Yackolley Amoussou-Guenou, Aggelos Kiayias, and Marianne Verdier; Article No. 6; pp. 6:1–6:17

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

However, it is safe to assume that the total MEV extracted is much higher, since MEV-explore limits its estimates to only one blockchain, a few protocols, and a limited number of detectable MEV techniques. Although it is difficult to argue that all MEV is “bad” (e.g., market arbitrage can remove market inefficiencies), it usually introduces some negative externalities like: *network congestion*: especially on low-cost chains, MEV actors often try to increase their chances of exploiting an MEV opportunity by sending a lot of redundant transactions, spamming the underlying peer-to-peer network; *chain congestion*: many such transactions finally make it to the chain, making the chain more congested; *higher blockchain costs*: while competing for profitable MEV opportunities, MEV actors bid higher gas prices to prioritize their transactions, which results in overall higher blockchain costs for regular users; *consensus stability*: some on-chain transactions can create such a lucrative MEV opportunity that it may be tempting for miner(s) to create an alternative chain fork with such a transaction extracted by them, which introduces consensus instability risks.

1.2 Blind Order-Fairness on a DAG

A promising way of thwarting some forms of MEV on blockchains is to extend their BFT (Byzantine fault tolerance) consensus core by new properties like Blind Order-Fairness, where no consensus “validator” can learn the content of transactions until they are committed to a total ordering. This notion of “*commit-reveal*” is the key mechanism we focus on in this paper.

An orthogonal line of research concentrates on scaling the BFT consensus of blockchains via a *DAG communication substrate*. A DAG transport can spread yet-unconfirmed transactions in parallel, every message having utility and carrying transactions that eventually become committed to a total ordering. The DAG provides Reliability, Non-equivocation, and Causal-ordering delivery guarantees, enabling a simple and “zero-overhead” consensus ordering on top of it. That is, validators can interpret their DAG locally without exchanging more messages and determine a total ordering of accumulated transactions. Several recent DAG-based BFT systems, including Blockmania [13], Aleph [19], Narwhal [14], DAG-Rider [23], and Bullshark [20], demonstrate excellent performance.

This paper combines these two advances in a solution called Fino, that integrates smoothly two ingredients: A framework for commit-reveal consensus protocols, and a simple DAG-based BFT consensus that leverages the Reliability, Non-Equivocation, and Causality properties of the DAG transport to integrate the commit-reveal framework smoothly and efficiently into consensus.

A commit-reveal framework. The core of the commit-reveal framework of Fino is a standard k -of- n secret sharing approach consisting of two functionalities, Disperse() and Retrieve(). Users first send transactions to validators *encrypted*, so that the consensus protocol commits to an ordering of transactions *blindly*. Disperse() entrusts shares of the secret decryption key for each transaction to validators. After an ordering is committed, Retrieve() opens encrypted transactions by collecting shares from $F + 1$ validators.

Implementing Disperse()/Retrieve() efficiently enough to meet DAG-based consensus throughput is the core challenge Fino addresses. Two straw man approaches are (i) threshold cryptography, but it would incur an order of milliseconds per transaction, (ii) using polynomial secret-sharing and verifying shares during dispersal, also requiring costly cryptography as well as implementing a share-recovery protocol.

The third approach, and the one we recommend using, shares transaction encryption keys using Shamir’s secret-sharing scheme [38], but without secret share verifiability or recoverability (slow and complex), and without threshold encryption (slow). Instead, to

guarantee a unique, deterministic outcome of encrypted transactions, Fino borrows a key insight from DispersedLedger [40] called AVID-M, protecting against share manipulation attacks by verifying that the sharing was correct *post reconstruction*. If the dispersal was incorrect and fails post-reconstruction verification, the transaction is rejected.

Notably, this approach works in microseconds latency and is orders of magnitude faster than threshold encryption.

section 3 describes a fourth approach, a hybrid combining AVID-M with threshold cryptography. Hybrid tackles adversarial conditions where Disperse() partially fails, at the cost of a slower post-reconstruction verification.

We repeat that all four approaches above are compatible with the Fino framework, though the fastest and simplest one is AVID-M.

DAG-based integration. Fino integrates Disperse()/Retrieve() into a DAG-based consensus without delaying the steady spreading of transactions by the DAG transport and with zero message overhead. Validators interpret their local DAGs to arrive at commit decisions on blind ordering. After an ordering has committed for a batch of transactions, validators invoke Retrieve() and present decryption shares piggybacked on DAG broadcasts. Finally, validators interpret their local DAGs to arrive at unanimous, deterministic outcome of transaction opening.

During periods of synchrony, Fino has a happy path for transaction commit of two DAG latencies, plus a single DAG latency for determining an opening outcome.

2 Background and Preliminaries

2.1 System Model

The Byzantine fault tolerance (BFT) Consensus core of blockchains implements state machine replication (SMR), where non-faulty parties—also known as validators—agree on an order of transactions and execute them consistently. The goal of our work is to extend BFT consensus and build into it protection against MEV attacks which rely on transaction content analysis.

We assume that the system should be resilient to Byzantine faults, i.e., faulty validators can implement any adversarial behavior (like crashing, equivocating, or colluding). We assume a partially synchronous network [18], where asynchronous periods last up to unknown global stabilization time (GST), while in synchronous periods there is a known bound Δ on delays for message delivery. This model requires a BFT threshold $N \geq 3F + 1$, where N and F is the number of all and Byzantine faulty validators, respectively.

2.2 MEV and Blind Order-Fairness

In this paper, we focus on consensus-level MEV mitigation techniques. There are fundamentally two types of MEV-resistant Order-Fairness properties:

Blind Order-Fairness. A principal line of defense against MEV stems from committing to transaction ordering without seeing transaction contents. This notion of MEV resistance, referred to here as Blind Order-Fairness, is used in a recent SoK on Preventing Transaction Reordering by Heimbach and Wattenhofer [21], and is defined as: *“when it is not possible for any party to include or exclude transactions after seeing their contents. Further, it should not be possible for any party to insert their own transaction before any transaction whose contents it already been observed.”*

Time-Based Order-Fairness. Another measure for MEV protection is brought by sending transactions to all BFT parties simultaneously and using the relative arrival order at a majority of the parties to determine the final ordering. In particular, this notion of order fairness ensures that “if sufficiently many parties receive a transaction tx before another tx' , then in the final commit order tx' is not sequenced before tx .”

This defines the transactions order and prevents adversaries that can analyze network traffic and transaction contents from reordering, censoring, and front-/back-running transactions received by Consensus parties. Moreover, Time-Based Order-Fairness protects against a potential collusion between users and BFT leaders/parties because parties explicitly input relative ordering into the protocol. Time-Based Order-Fairness is used in various flavors in several recent works, including Pompe [43], Aequitas [26], Themis [25], “Wendy Grows Up” [28], and “Quick Order Fairness” [11]. We briefly discuss some of those protocols in section 8.

Another notion of fairness found in the literature, that does not provide Order-Fairness, revolves around participation fairness:

Participation Fairness. A different notion of fairness aims to ensure censorship-resistance or stronger notions of participation equity. Participation Fairness guarantees that the committed sequence of transactions includes a certain portion of honest contribution (aka “Chain Quality”). Several BFT protocols address Participation Fairness, including Prime [2], Fairledger [29], HoneyBadger [31]. As mentioned in section 4, some DAG-based BFT protocols like Aleph [19], DAG-Rider [23], Tusk [14], and Bullshark [20] use a layered DAG paradigm. In this approach, Participation Fairness comes essentially for free because every DAG layer must include messages from $2F+1$ participants. It is worth noting that Participation Fairness does not prevent a corrupt party from injecting transactions after it has already observed other transactions, nor a corrupt leader from reordering transactions after reading them, violating both Blind and Time-Based Order-Fairness.

2.3 Threshold Methods

Beside standard cryptographic primitives, like a symmetric encryption scheme or a cryptographic hash function $H()$, we employ primitives specifically designed for hiding and opening transactions.

Secret Sharing. For positive integers k, n , where $k \leq n$, a (k, n) -threshold scheme is a technique where a secret is distributed among n parties, such that it can be reconstructed by any k parties, while any group of $k - 1$ parties cannot learn the secret. The secret S is selected by a trusted third party, called *dealer*, who splits it into n individual *shares* by calling: $s_1, \dots, s_n \leftarrow \text{SS.Split}(S)$. After at least k parties exchange shares, they can combine them and recover the secret by calling $S \leftarrow \text{SS.Combine}(s'_1, \dots, s'_k)$.

Shamir’s Secret Sharing (SSS) [38] is one of the first and most widely used (k, n) -threshold schemes. In SSS, $\text{Split}(S)$ chooses a polynomial f of degree $k - 1$ that hides S at the origin point, and generates n points on the polynomial. Then, any k points allow to interpolate the polynomial and compute its value at $f(0)$.

Threshold Encryption. A (k, n) -threshold encryption is a scheme where messages encrypted under a single public encryption key pk can be decrypted by a private decryption key, shared among n parties – each with its secret key sk_i , where any group of k parties can decrypt a message. The message m is encrypted with an encryption algorithm: $c \leftarrow \text{TE.Enc}(pk, m)$.

We consider schemes where each party, upon receiving a ciphertext c , computes its own *decryption share* by invoking: $ds_i \leftarrow \text{TE.ShareGen}(sk_i, c)$; and any set of k unique decryption shares is enough to recover the plaintext, which we denote $m \leftarrow \text{TE.Dec}(c, ds'_1, \dots, ds'_k)$.

Moreover, some threshold encryption schemes, like shown by Shoup and Gennaro [39], allow parties to verify that a given decryption share corresponds to the ciphertext, which we denote by the $\text{TE.Verify}(c, ds_i)$ function. Note that, without such a verification, it is trivial for a malicious party, providing an incorrect share, to cause decryption failures.

3 The commit-reveal Framework

The core of blind fairness-ordering is a standard k-of-n secret sharing approach. For each transaction tx , a user (“Alice”) picks a secret symmetric key tx-key , and sends tx encrypted with it to validators. Secret sharing allows Alice to share tx-key with validators such that $F + 1$ shares are required to reconstruct tx-key , and no set of F bad validators can open it before it is committed (blindly) to the total order. Honest validators wait to commit to a blind order of transactions first, and only later open them. At the same time, before committing a transaction to the total ordering, validators ensure that they can open it.

More formally, a sharing protocol has two abstract functionalities, $\text{Disperse}(\text{tx})$ and $\text{Retrieve}(\text{tx})$. $\text{Disperse}()$ allows a dealer to disseminate shares of the secret tx-key to validators, with the following guarantees:

Hiding: A set of F bad validators cannot reconstruct tx-key .

Binding: After $\text{Disperse}(\text{tx-key})$ completes successfully at an honest validator, any set of honest $F + 1$ validators doing $\text{Retrieve}(\text{tx-key})$ can reconstruct tx-key generating a unique outcome.

Validity: If the dealer is honest, then the outcome from $\text{Retrieve}(\text{tx-key})$ by any $F + 1$ honest validators is tx-key .

It is worth noting that $\text{Disperse}()/\text{Retrieve}()$ does not require certain properties which some use-cases in the literature may need, but not the framework here:

- When $\text{Disperse}()$ completes, it does not guarantee that $\text{Retrieve}()$ can reconstruct an output value that a dealer commits to.
- It doesn't require that individual shares can be recovered.
- It doesn't require being able to use tx-key to derive other values while keeping tx-key itself secret.

We proceed to present four approaches for implementing $\text{Disperse}()/\text{Retrieve}()$. All of them are compatible with the DAG-based consensus protocol in the next section, though the fastest and simplest one, which we recommend using, is AVID-M.

Approach-1: Threshold Cryptography

It is straight-forward to implement $\text{Disperse}/\text{Retrieve}$ using threshold encryption, such that the public encryption $\text{TE.Enc}()$ is known to users and the private decryption $\text{TE.Dec}()$ is shared (at setup time) among validators. To $\text{Disperse}(\text{tx-key})$, the user attaches $\text{TE.Enc}(\text{tx-key})$, the transaction key encrypted with the global threshold key. Once a tx is committed to the total-ordering, $\text{Retrieve}(\text{tx-key})$ is implemented by each validator generating its decryption share via $\text{TE.ShareGen}()$. Validators can try applying $\text{TE.Dec}()$ over $F + 1$ valid decryption shares to decrypt tx .

6:6 Maximal Extractable Value (MEV) Protection on a DAG

Binding stems from two properties. First, a threshold of honest validators can always succeed in generating $F + 1$ valid decryption shares. Second, as indicated in section 2, some threshold cryptography schemes allow verifying that a validator is contributing a correct decryption share, hence by retrieving $F + 1$ valid threshold shares, a unique outcome is guaranteed.

Unfortunately, this method is computationally somewhat heavy: `Disperse()` takes about $300\mu s$ to encrypt using TDH2 on standard hardware, and `Retrieve()` takes about $3500\mu s$ for in a 6-out-of-16 scheme (see section 7).

Approach-2: VSS

Another way for users to `Disperse(tx-key)` is Shamir’s secret sharing (SSS) scheme. To `Disperse(tx-key)`, a user employs `SS.Split(tx-key)` to send individual shares of `tx-key` to each validator. A set of $F + 1$ share holders reveal their shares during `Retrieve()` and combine shares via `SS.Combine()` to reconstruct `tx-key`.

Combining shares is three orders of magnitude faster than threshold crypto and takes a few microseconds in today’s computing environment (see section 7). However, after `Disperse()` completes with $N - F$ validators, “vanilla” `SS.Split()/SS.Combine()` does not guarantee Binding: first, not every set of $F + 1$ honest validators may hold shares from the `Disperse()` phase. Second, retrieving shares from different sets of $F + 1$ validators may result in different outcome from `SS.Combine()`.

VSS schemes support Binding through a share-verification functionality `VSS.Verification(share, commit value)`. Verification allows validators to check that shares are consistent with some committed value S' , such that any set of $F + 1$ shares applied to `SS.Combine()` results in output S' . When a validator receives a share, it should verify the share before acknowledging it. `Disperse()` completes when there are $N - F$ acknowledgements of verifiable shares, certifying that valid shares are held by $F + 1$ honest validators and guaranteeing a unique reconstruction by any subset of $F + 1$. It is possible to add a share-recovery functionality to allow a validator to obtain its individual share prior to `Retrieve()`.

However, despite the vast progress in VSS schemes (see subsection 8.1), share recovery requires linear communication. Additionally, implementing VSS requires non-trivial cryptography.

AVID-M

VSS provides stronger guarantees than necessary to satisfy the Binding requirement of commit-reveal. Specifically, `Disperse()` only needs to guarantee a unique outcome from `Retrieve()`, not success. Borrowing a technique called AVID-M, which was introduced in *DispersedLedger* [40] for reliable information dispersal, validators can completely avoid verifying shares during `Disperse()`. Instead, they verify during `Retrieve()` that the (entire) sharing would have a unique outcome, which is very cheap to do. This works as follows.

To `Disperse(tx)`, a user employs `SS.Split(tx-key)` to send individual shares to validators. This operation requires a deterministic way of splitting secret shares which, in particular, can be based on validator identifiers (e.g., $s_i = P(id_i)$, where P is the polynomial hiding `tx-key`). Additionally, the user combines all shares in a Merkle-tree, certifies the root, and sends with each share a proof of membership, i.e., a Merkle tree path to the root.

When a validator receives a share, it should verify the Merkle tree proof against the certified root before acknowledging the share. `Disperse()` completes when there are $N - F$ acknowledgements, guaranteeing that untampered shares are held by $F + 1$ honest validators.

During `Retrieve()`, $F + 1$ validators reveal individual shares, attaching the Merkle tree path to prove shares have not been tampered with. Note that $F + 1$ shares can be validated by checking only one signature on the root. Validators use `SS.Combine()` to try to reconstruct `tx-key`.

Then they *post-verify* that every subset of $F + 1$ (untampered) shares sent to validators would have the same outcome. To do this, a validator does not need to wait for missing shares nor try combinations of $F + 1$ shares. Moreover, it does not need to communicate with others. A validator simply generates missing shares and re-encodes the Merkle tree. Then, it compares the re-generated Merkle tree with the root certified by the sender. If the comparison fails, the dealer was faulty and the validator rejects `tx`. The unique reconstruction is guaranteed because the post-verification outcome – succeed or fail – becomes fixed upon `Disperse` completion. Each validator arrives at the same outcome no matter which subset of $F + 1$ (untampered) shares are input.

AVID-M takes roughly $50\mu s$ to `SS.split()` and generate Merkle tree commitments, and roughly $50\mu s$ for `Retrieve()`, including post-verification, in a 6-out-of-16 scheme.

Hybrid

Even after `Disperse()` has completed, both VSS and AVID-M may need to interact with a specific set of $F + 1$ honest validators during `Retrieve()`. This implies that the latency of `Retrieve()` is impacted by this specific set of $F + 1$ share holders, and not the speed of the fastest $F + 1$ honest validators. Employing threshold cryptography removes this dependency but incurs a costly computation.

A Hybrid approach combines the benefits of threshold cryptography with AVID-M. In Hybrid, `Retrieve()` works with (fast) secret-sharing during steady-state, and falls back to threshold cryptography, avoiding waiting for specific $F + 1$ validators, during a period of network instability.

`Disperse(tx-key)` is implemented in two parts. First, a user applies AVID-M to send validators individual shares and proofs. Second, as a fallback, it sends validators `TE.Enc(tx-key)`. Once `tx` is committed to the total ordering, `Retrieve(tx-key)` has a fast track and a slow track. In the fast track, every validator that holds an AVID-M share reveals it. A validator that doesn't hold an AVID-M share for `tx-key` reveals a threshold key decryption share. In the slow track, validators may give up on waiting for AVID-M shares and reveal threshold key decryption shares, even if they already revealed AVID-M shares.

Post-verification after $F + 1$ valid shares of either kind are revealed checks that **both** the AVID-M and threshold encryption were shared correctly and would have the same outcome. More specifically, a validator both re-encrypts `tx-key` and re-encodes the AVID-M Merkle tree after reconstruction. It compares both outcomes with the sender's. If the comparison fails, the dealer was faulty and the validator rejects `tx`. Binding holds because the post-verification outcome – succeed or fail – becomes fixed upon `Disperse()` completion. Each validator arrives at the same outcome no matter which scheme and what subset of $F + 1$ (untampered) shares are retrieved, thus ensuring Binding.

4 DAG transport

In a DAG-based BFT protocol, validators store messages delivered via reliable and causally ordered broadcast in a local graph. A message inserted into the local DAG has the following guarantees: *Reliability*: there are copies of the message stored on sufficiently many participants, such that eventually, all honest validators can download it. *Non-equivocation*:

6:8 Maximal Extractable Value (MEV) Protection on a DAG

messages by each validator are numbered. If a validator delivers some message as the k 'th from a particular sender, then the message is authenticated by its sender and other validators deliver the same message as the sender k 'th message. *Causal Ordering*: the message carries explicit references to messages which the sender has previously delivered (including its own). Predecessors are delivered locally before the message.

Note that the DAGs at different validators may be slightly different at any moment in time. This is inevitable in a distributed system due to message scheduling. However, a DAG-based Consensus protocol allows each participant to interpret its local DAG, reaching an autonomous conclusion that forms total ordering. Reliability, Non-equivocation and Causal Ordering make the design of such protocols extremely simple as we shall see below.

DAG API. A DAG transport exposes two basic API's, `broadcast()` and `deliver()`.

Broadcast. `broadcast()` is an asynchronous API for a validator to present payload for the DAG transport to be transmitted to all other validators. The main use of DAG messages in a DAG-based Consensus protocol is to pack meta-information on transactions into a block and present the block for broadcast. The DAG transport adds references to previously delivered messages, and broadcasts a message carrying the block and the causal references to all validators.

Deliver. When another validator receives the message carrying the block, it checks whether it needs to retrieve a copy of any of the transactions in the block and in causally preceding messages. Once it obtains a proof-of-availability of all transactions in the block and its causal past, it can acknowledge it. A validator's upcall `deliver(m)` is triggered when sufficiently many acknowledgments for it are gathered, guaranteeing that the message itself, the transactions it refers to, and its entire causal past maintain Reliability, Non-equivocation and Causal Ordering.

It is worth noting that in the protocol discussed in this paper, the Consensus protocol injects meta-information (e.g., complaints), DAG broadcast never stalls or waits for input from it.

Implementing a DAG. There are various ways to implement reliable, non-equivocating and causally-ordered broadcast among $N = 3F + 1$ *validators*, at most F of which are presumed Byzantine faulty and the rest are honest.

Lifetime of a message. A validator packs transaction information and meta-information into a message, adds references to previously delivered messages (including the sender's own preceding messages), and broadcasts the message to all validators.

Echoing. The key mechanism for reliability and non-equivocation is for validators to echo a digest of the first message they receive from a sender with a particular index. When $2F + 1$ echoes are collected, the message can be delivered. There are two ways to echo, one is all-to-all broadcast over authenticated point-to-point channels a la Bracha Broadcast [9]; the other is converge-cast with cryptographic signatures a la Rampart [37] and Cachin et al. [10]. In either case, echoing can be streamlined so the amortized per-message communication is linear, which is anyway the minimum necessary to spread the message.

Layering. Transports are often constructed in a layer-by-layer regime. In this regime, each sender is allowed one message per layer, and a message may refer only to messages in the layer preceding it. Layering is done so as to regulate transmissions and saturate network capacity, and has been demonstrated to be highly effective by various projects [13, 19, 14]. We reiterate that Fino does not require a layered structure.

5 Fino

Fino incorporates MEV protection into a BFT protocol for the partial synchrony model, riding on a DAG transport. BFT validators periodically pack pending encrypted transactions into a batch and use the DAG transport to broadcast them. The key insight for operating commit-reveal on a DAG is that each view must wait until `Retrieve()` completes on transactions `tx` of the previous view.

5.1 The Protocol

Views. The protocol operates in a view-by-view manner. Each view is numbered, as in `view(r)`, and has a designated leader known to everyone.

View change. A validator enters `view(r+1)` when two conditions are met, `viewA` and `viewB`: `viewA` is satisfied when the local DAG of a validator contains $F + 1$ valid votes on `proposal(r)` or $2F + 1$ valid `complaint(r)` on `view(r)`. `viewB` is satisfied when every committed transaction `tx` in the local DAG has $F + 1$ valid shares revealed, hence `Retrieve(tx)` can be completed. Note that `viewB` prevents BFT validators that do not reveal (correct) shares from enabling the protocol to make progress without opening committed transactions.

Proposing. When a leader enters a new `view(r)`, it broadcasts `proposal(r)`.¹ Implicitly, `proposal(r)` suggests to commit to the global ordering of transactions all the messages in the causal history of the proposal. A leader's `proposal(r)` is *valid* if it is well-formatted and is justified in entering `view(r)`.

Voting. When a validator sees a valid leader proposal, it broadcasts `vote(r)`. A validator's `vote(r)` is *valid* if it follows a valid `proposal(r)`.

Committing. Whenever a leader's `proposal(r)` has $F + 1$ valid votes in the local DAG, the proposal and its causal history become *committed*. The commit order induced by a commit decision is described below.

Share revealing. When a validator observes that a transaction `tx` becomes committed, it starts `Retrieve(tx)` and calls `broadcast()` to present its share of the decryption key `tx-key`.

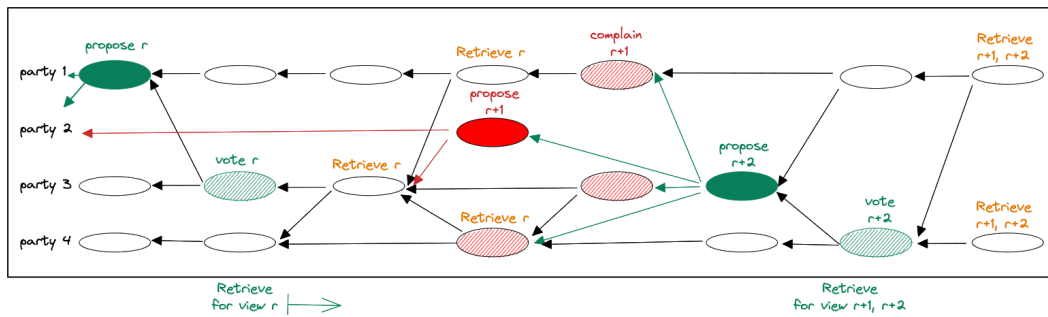
Complaining. If a validator gives up waiting for a commit to happen in `view(r)`, it broadcasts `complaint(r)`. Note, a `vote(r)` by a validator that causally follows a `complaint(r)` by the validator, if exists, is **not** interpreted as a valid vote.

Ordering Commits. When a validator observes that a leader's `proposal(r)` becomes committed, it orders newly committed transactions as follows:

1. Let r' be the highest view $r' < r$ for which `proposal(r')` is in the causal history of `proposal(r)`. `proposal(r')` is recursively ordered.

¹ Recall, `broadcast()` merely presents payload to be transmitted as scheduled by the DAG transport, e.g., piggybacked on messages carrying other transaction info.

6:10 Maximal Extractable Value (MEV) Protection on a DAG



■ **Figure 1** A commit of $\text{proposal}(r)$ is followed by share retrieval. A commit of $\text{proposal}(r+2)$ causes an indirect commit of $\text{proposal}(r+1)$, followed by share revealing of both.

2. The remaining causal predecessors of $\text{proposal}(r)$ which have not yet been ordered are appended to the committed sequence (within this batch, ordering can be done using any deterministic rule to linearize the partial ordering into a total ordering.)

Opening Transactions. When a validator observes that a leader's $\text{proposal}(r)$ becomes committed, it decrypts every committed transaction tx in its causal past that hasn't been decrypted already. That is, let c be the highest view $c < r$ for which $\text{proposal}(r)$ causally follows $F + 1$ valid votes. Transactions in the causal past of $\text{proposal}(c)$ are opened as follows:

1. Let r' be the highest view $r' < c$ for which $\text{proposal}(r')$ is in the causal history of $\text{proposal}(c)$. $\text{proposal}(r')$ is recursively opened.
2. The remaining (committed) transactions tx in $\text{proposal}(c)$'s causal past are opened using $\text{Retrieve}(\text{tx})$ to either produce a key tx-key that decrypts tx or to reject tx .

Figure 1 illustrates a couple of Fino scenarios.

Happy-path scenario. In the first view (view(r)), $\text{proposal}(r)$ becomes committed. The commit sets an ordering for transactions in the causal past of $\text{proposal}(r)$, enabling retrieval of shares for transactions from $\text{proposal}(r)$.

Scenario with a slow leader. A slightly more complex scenario occurs when a view expires because validators do not observe a leader's proposal becoming committed and they broadcast complaints. Figure 1 depicts this happening in view($r+1$). Entering view($r+2$) is enabled by $2F + 1$ complaints about view($r+1$). When $\text{proposal}(r+1)$ itself becomes committed, it indirectly commits $\text{proposal}(r+1)$ as well. Thereafter, validators reveal shares for all pending committed transactions, namely, those in both $\text{proposal}(r+1)$ and $\text{proposal}(r+2)$.

6 Analysis

Fino is minimally integrated with a DAG transport. BFT commit-reveal logic is embedded into the DAG structure simply by broadcasting (more precisely, injecting payloads into broadcasts) in the form of protocol proposals/votes/complaints and revealing shares. Importantly, at no time is DAG broadcast slowed down by the Fino protocol. The reliability and causality properties of the DAG transport make arguing about safety and liveness relatively easy.

6.1 Safety

► **Lemma 1.** *Assume a $\text{proposal}(r)$ ever becomes committed. Denote by r' the minimal view, where $r' > r$, such that $\text{proposal}(r')$ ever becomes committed. Then for every $r \leq q < r'$, $\text{proposal}(q+1)$ causally follows $\text{proposal}(r)$.*

Proof (Sketch). Since $\text{proposal}(r)$ becomes committed, by definition $F + 1$ parties sent votes for it. There are two possibilities for a leader's $\text{proposal}(q+1)$ to be valid. The first is $\text{proposal}(q+1)$ may reference $F + 1$ valid $\text{vote}(q)$ messages. This case occurs only for $q = r$ by the lemma assumption that $r' \geq q + 1$ is the minimal view that ever becomes committed. The second possibility can occur for all $r \leq q < r'$, namely, $\text{proposal}(q+1)$ references $2F + 1$ valid $\text{complaint}(q)$ messages.

In the first case, valid $\text{vote}(r)$ causally follows $\text{proposal}(r)$, and a fortiori $\text{proposal}(r+1)$ causally follows $\text{proposal}(r)$.

In the second case, one of $2F + 1$ $\text{complaint}(q)$ messages is sent by a party who sent a valid $\text{vote}(r)$. By definition, $\text{vote}(r)$ must precede $\text{complaint}(q)$, otherwise it is not considered a (valid) vote. Hence, $\text{proposal}(q+1)$ causally follows $\text{complaint}(q)$ which follows $\text{proposal}(r)$. ◀

► **Lemma 2.** *If ever a $\text{proposal}(r)$ becomes committed, then every valid $\text{proposal}(q)$, where $q > r$, causally follows $\text{proposal}(r)$.*

Proof (Sketch). The proof follows by a simple induction on Lemma 1. ◀

► **Lemma 3.** *If an honest party commits $\text{proposal}(r)$, and another honest party commits $\text{proposal}(r')$, where $r' > r$, then the sequence of transactions committed by $\text{proposal}(r)$ is a prefix of the sequence committed by $\text{proposal}(r')$.*

Proof (Sketch). When $\text{proposal}(r')$ becomes committed, the commit ordering rule is recursively applied to valid proposals in its causal past. By Lemma 2, the (committed) $\text{proposal}(r)$ is a causal predecessor of every valid $\text{proposal}(s)$, for $r < s \leq r'$, and eventually the recursion gets to it. ◀

► **Lemma 4.** *If ever a $\text{proposal}(r)$ becomes committed, then every committed transaction proposed by a valid $\text{proposal}(r')$, where $r' < r$, can be uniquely decrypted.*

Proof (Sketch). When $\text{proposal}(r)$ becomes committed, then by Lemma 3, all the transactions from previously committed proposals are part of the committed transaction history. Moreover, committing $\text{proposal}(r)$ implies the view change (from $r-1$) which occurs only if the condition viewB (see section 5) is met, i.e., every committed transaction tx in the local DAG can complete $\text{Retrieve}(\text{tx})$, hence can be uniquely decrypted (or rejected). ◀

6.2 Liveness

Liveness is guaranteed after GST, i.e., after communication has become synchronous with a known Δ upper bound on transmission delays.

After GST, views are semi-synchronized through the DAG. In particular, suppose that after GST, every broadcast by an honest party arrives at all honest parties within DD . Once a $\text{view}(r)$ with an honest leader is entered by the first honest party, within DD all the messages seen by one party are delivered by both the leader and all other honest parties. Hence, within DD , all honest parties enter $\text{view}(r)$ as well. Within two additional DAG latencies, $2 \cdot DD$, the $\text{view}(r)$ proposal and votes from all honest parties are spread to everyone.

6:12 Maximal Extractable Value (MEV) Protection on a DAG

Assuming view timers are set to be at least $3 \cdot DD$, once `view(r)` is entered, a future view will not interrupt a commit. In order to start a future view, its leader must collect either $F + 1$ `vote(r)` messages, hence commit `proposal(r)`; or $2F + 1$ `complaint(r)` expiration messages, which is impossible as argued above.

6.3 Communication complexity

Protocols for the partial synchrony model have unbounded worst case by nature, hence, we concentrate on the costs incurred during steady state when a leader is honest and communication with it is synchronous:

DAG message cost. In order for DAG messages to be delivered reliably, it must implement reliable broadcast. This incurs either a quadratic number of messages carried over authenticated channels, or a quadratic number of signature verifications, per broadcast. In either case, the quadratic cost may be amortized by pipelining, driving it in practice to (almost) linear per message.

Commit message cost. Fino sends $F + 1$ broadcast messages, a proposal and votes, per decision. A decision commits the causal history of the proposal, consisting of (at least) a linear number of messages. Moreover, each message may carry multiple transactions in its payload. As a result, in practice the commit cost is amortized over many transactions.

6.4 Latency

Commit latency. The commit latency in terms of DAG messages is 2, one proposal followed by votes.

Opening latency. During periods of stability, there are no complaints about honest leaders by any honest party. If `tx` is proposed by an honest leader in `view(r)`, it will receive $F + 1$ votes and become committed within two DAG latencies. Within one more DAG latency, every honest party will post a message containing a share for `tx`. Thereafter, whenever $F + 1$ are available, everyone will be able to complete `Retrieve(tx)` with a unique outcome.

6.5 Faulty User

One drawback of the proposed protocol is that it assumes a well-connected and honest user (dealer), able to entrust $F + 1$ honest validators with their shares. If the dealer is malicious or unable to broadcast shares it can lead to the situation where the transport layer buffers a transaction indefinitely without being able to `deliver()` it. The Hybrid approach we presented in section 3 solves this issue by using threshold encryption as a fallback mechanisms (i.e., whenever enough shares cannot be received, parties would proceed with threshold decryption). However, the hybrid construction suffers the slowness of threshold encryption to verify a unique outcome even in faultless executions.

7 Evaluation (prelim)

We implemented the proposed variants of the Fino Disperse and Retrieve functionalities and investigated the computational overhead that these schemes introduce. For secret sharing we implemented SSS [38], while for threshold encryption we implemented the TDH2 scheme [39].

■ **Table 1** Comparison of the implementations' performance.

Scheme	Disperse	ShareGen	ShareVerify	Reconstruct
TDH2-based	311.6 μ s	434.8 μ s	492.5 μ s	763.9 μ s
SSS-AVID-M	52.7 μ s	N/A	2.7 μ s	52.5 μ s
Hybrid opt.	360.3 μ s	N/A	2.7 μ s	361.5 μ s
Hybrid pess.	360.3 μ s	434.8 μ s	492.5 μ s	828.9 μ s

We selected the schemes with the most efficient cryptographic primitives we had access to, i.e., the secret sharing scheme uses the Ed25519 curve [7], while TDH2 uses ristretto255 [15] as the underlying prime-order group. Performance for both schemes is presented in a setting where 6 shares out of 16 are required to recover the plaintext.

The results presented in Table 1 are obtained on an Apple M1 Pro. **Disperse** refers to the overhead on the client-side while **ShareGen** is the operation of deriving a decryption share from the TDH2 ciphertext (it is absent in SSS and the optimistic path of the hybrid scheme). In TDH2, **ShareVerify** verifies if a decryption share matches the ciphertext, while in SSS-based Fino it only checks whether a share belongs to the tree aggregated by the signed Merkle root attached by the client. **Reconstruct** recovers the plaintext from the ciphertext and the number of shares and verifies the outcome integrity (i.e., AVID-M and threshold re-encryption verification, where applicable).

As demonstrated by these micro-measurements, the SSS-AVID-M scheme is the most efficient: in our blind ordering scenario using threshold encryption, each party processing a TDH2 ciphertext would call **ShareGen** once to derive its decryption share, **ShareVerify** $k-1$ times to verify the threshold number of received shares, and **Reconstruct** once to obtain the plaintext. Assuming $k=6$, the total computational overhead for a single transaction to be recovered would take around 3.7ms CPU time. With secret-sharing based Fino, the party would also call **ShareVerify** $k-1$ times and **Reconstruct** once, which requires only 66 μ s CPU time. The hybrid approach introduces an interesting trade-off between its properties and performance. In the optimistic path, it requires around 378 μ s CPU time per transaction, while its fallback introduces negligible overhead to the plain threshold-encryption scheme.

We emphasize that these are micro-benchmarks of the blinding/unblinding of transactions; in the future, we plan to complete a performance evaluation of the entire DAG-based blind-ordering protocol.

8 Related Work and Discussion

Despite the MEV problem being relatively new, there already exist consensus-related systems aiming at solving it. For instance, Flash Freezing Flash Boys (F3B) [42] is a commit-and-reveal architecture, where clients send TDH2-encrypted ciphertexts to be ordered by Consensus, and afterward to be decrypted via shares released by a dedicated secret-management committee. Other closely related techniques and schemes helpful in mitigating MEV are discussed below.

8.1 Verifiable Secret Sharing (VSS)

Although we ultimately end up forgoing verifiability/recoverability of shares, VSS is used in many settings like ours. The overall communication complexity incurred in VSS on the dealer sharing a secret and on a party recovering a share has dramatically improved in recent years. VSS can be implemented inside the asynchronous echo broadcast protocol in $O(n^3)$

communication complexity using Pederson’s original two-dimensional polynomial scheme Non-Interactive Polynomial Commitments [34]. Kate et al. [22] introduce a VSS scheme with $O(n^2)$ communication complexity, utilized for asynchronous VSS by Backes et al. [4]. Basu et al. [5] propose a scheme with linear $O(n)$ communication complexity.

A related notion is robust secret sharing (RSS) introduced by Krawczyk [27] and later revised by Bellare and Rogaway [6]. RSS allows recoverability in case of incorrect (not just missing) shares, however, these schemes assume an honest dealer. Duan et al. [17] propose ARSS extending RSS to the asynchronous setting.

8.2 Time-Based Order-Fairness

Blind Order-Fairness is achieved by deterministically ordering encrypted transactions and then, after the order is final, decrypting them. The deterministic order can be enhanced by sophisticated ordering logic present in other protocols. In particular, Fino can be extended to provide Time-Based Fairness additionally ensuring that the received transactions are not only unreadable by parties, but also their relative order cannot be influenced by malicious parties (the transaction order would be defined by the time of its ciphertext arrival).

For instance, Pompē [43] proposes a property called *Ordering Linearizability*: “if all correct parties timestamp transactions tx , tx' such that tx' has a lower timestamp than tx by everyone, then tx' is ordered before tx .” It implements the property based on an observation that if parties exchange transactions associated with their receiving timestamps, then for each transaction its median timestamp, computed out of $2F+1$ timestamps collected, is between the minimum and maximum timestamps of honest parties. Fino can be easily extended by the Linearizability property offered by Pompē and the final protocol is similar to the Fino with Blind Order-Fairness (see above) with only one modification. Namely, every time a new batch of transactions becomes committed, parties independently sort transactions by their aggregate (median) timestamps.

More generally, Fino can easily incorporate other Time-based Fairness ordering logic. Note that in Fino, the ordering of transactions is determined on encrypted transactions, but time ordering information should be open. The share revealing, share collection, and unique decryption following a committed ordering are the same as presented previously. The final protocol offers much stronger properties since it not only hides payloads of unordered transactions from parties, but also prevents parties from reordering received transactions.

One form of Time-based Order-Fairness introduced in Aequitas [26] is *Batch-Order Fairness*: “if sufficiently many (at least $\frac{1}{2}$ of the) parties receive a transaction tx before another transaction tx' , then no honest party can deliver tx in a block after tx' , ” and a protocol achieving it. Other forms of Time-Based Order-Fairness which may be used in Fino include “Wendy Grows Up” [28], which introduced *Timed Relative Fairness*, “if there is a time t such that all honest parties saw (according to their local clock) tx before t and tx' after t , then tx must be scheduled before tx' , ” and “Quick Order Fairness” [11], which defined *Differential-Order Fairness*, “when the number of correct parties that broadcast tx before tx' exceeds the number that broadcast tx' before tx by more than $2F + \kappa$, for some $\kappa \geq 0$, then the protocol must not deliver tx' before tx (but they may be delivered together).”

Themis [25] is a protocol realizing *Batch-Order Fairness*, where parties do not rely on timestamps (as in Pompē) but only on their relative transaction orders reported. Themis can also be integrated with Fino, however, to make it compatible this design requires some modifications to Fin’s underlying DAG protocol. More concretely, Themis assumes that the fraction of bad parties cannot be one quarter, i.e., F out of $4F+1$. A leader makes a proposal based on $3F+1$ out of $4F+1$ transaction orderings (each reported by a distinct party). Therefore, we would need to modify the DAG transport so that parties reference $3F+1$ preceding messages (rather than $2F+1$).

8.3 DAG-based BFT

There are many known BFT solutions for partial synchrony, and more specifically, several recent solutions that ride on a DAG [14, 20, 24]. When constructing Fino, we wanted to build MEV protection into a simple DAG-based BFT solution, described below. Notwithstanding, we remark that Fino’s MEV protection can possibly be incorporated into other DAG-riding BFT solutions.

We borrowed a subprotocol from Bullshark [20] that deals with partial synchrony, modifying it so that DAG transmissions would never stall waiting for BFT protocol steps or timers. Fino is different from the borrowed Bullshark component in that BFT protocol steps (e.g., view changes, proposals, votes, and complaints) are injected into the DAG at any time, independent of DAG layers. It uses `broadcast()` (defined in section 2) to present these steps as payloads to the DAG transport, completely asynchronously, while normal DAG transmissions continue. A hallmark of the DAG-riding approach, which Fino preserves while adding Blind Order-Fairness, is zero message overhead.

Narwhal [14] is a recent DAG transport that has a layer-by-layer structure, each layer having at most one message per sender and referring to $2F+1$ messages in the preceding layer. A similarly layered DAG construction appears earlier in Aleph [19], but it does not underscore the separation of transaction dissemination from DAG messages. Narwhal-HS is a BFT Consensus protocol within [14] based on HotStuff [41] for the partial synchrony model, in which Narwhal is used as a “mempool”. In order to drive Consensus decisions, Narwhal-HS adds messages outside Narwhal, using the DAG only for spreading transactions.

DAG-Rider [23] and Tusk [14] build randomized BFT Consensus for the asynchronous model riding on Narwhal. These protocols are zero message overhead over the DAG, not exchanging any messages outside the Narwhal protocol. Both are structured with purpose-built DAG layers grouped into “waves” of 4 (2) layers each. Narwhal waits for the Consensus protocol to inject input value every wave, though in practice, this does not delay the DAG materially.

Bullshark [20] builds BFT Consensus riding on Narwhal for the partial synchrony model. It is designed with 8-layer waves driving commit, each layer purpose-built to serve a different step in the protocol. Bullshark is a “zero message overhead” protocol over the DAG, however, due to a rigid wave-by-wave structure, the DAG is modified to wait for Bullshark timers/steps to insert transactions into the DAG. In particular, if leader(s) of a wave are faulty or slow, some DAG layers wait to fill until consensus timers expire.

More generally, since the eighties, causally ordered reliable broadcast has been utilized in forming distributed Consensus protocols, e.g., [8, 35, 30, 32, 3, 16].

The notion of the secure, causal, reliable broadcast was introduced by Reiter and Birman [37], and later refined by Cachin [10] and Duan et al. [17]. This primitive was utilized in a variety of BFT replicated systems, but not necessarily in the form of zero message overhead protocols riding on a DAG. Several of these pre blockchain-era BFT protocols are DAG based, notably Total [32] and ToTo [16], both of which are BFT solutions for the asynchronous model.

9 Conclusions

In this paper, we show how DAG-based BFT protocols can be enhanced to mitigate MEV – arguably one of the main threats to the success of cryptocurrencies. We investigate the design space of achieving Blind Order-Fairness, and we propose an approach that focuses on practicality. We present preliminary correctness and performance results which indicate that our scheme is secure and efficient. In the future, we plan to extend our analysis and experiments, and investigate other cryptographic tools which can improve our scheme.

References

- 1 Mev-explore v1. <https://explore.flashbots.net/>, 2022. Accessed: 2022-07-18.
- 2 Yair Amir, Brian Coan, Jonathan Kirsch, and John Lane. Prime: Byzantine replication under attack. *IEEE TDSC*, 2010.
- 3 Yair Amir, Danny Dolev, Shlomo Kramer, and Dalia Malki. *Transis: A communication sub-system for high availability*. Hebrew University of Jerusalem. Leibniz Center for Research in Computer . . . , 1991.
- 4 Michael Backes, Amit Datta, and Aniket Kate. Asynchronous computational vss with reduced communication complexity. In *CT-RSA*. Springer, 2013.
- 5 Soumya Basu, Alin Tomescu, Ittai Abraham, Dahlia Malkhi, Michael K Reiter, and Emin Gün Sirer. Efficient verifiable secret sharing with share recovery in bft protocols. In *ACM CCS*, 2019.
- 6 Mihir Bellare and Phillip Rogaway. Robust computational secret sharing and a unified account of classical secret-sharing goals. In *ACM CCS*, 2007.
- 7 Daniel J Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. *Journal of cryptographic engineering*, 2012.
- 8 Ken Birman and Thomas Joseph. Exploiting virtual synchrony in distributed systems. In *ACM SOSP*, 1987.
- 9 Gabriel Bracha. Asynchronous byzantine agreement protocols. *Information and Computation*, 1987.
- 10 Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient asynchronous broadcast protocols. In *CRYPTO*. Springer, 2001.
- 11 Christian Cachin, Jovana Mićić, and Nathalie Steinhauer. Quick order fairness. *arXiv preprint*, 2021. [arXiv:2112.06615](https://arxiv.org/abs/2112.06615).
- 12 Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *IEEE SP*. IEEE, 2020.
- 13 George Danezis and David Hryczyn. Blockmania: from block dags to consensus. *arXiv preprint*, 2018. [arXiv:1809.01620](https://arxiv.org/abs/1809.01620).
- 14 George Danezis, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. Narwhal and tusk: a dag-based mempool and efficient bft consensus. In *EuroSys*, 2022.
- 15 Henry de Valence, Jack Grigg, George Tankersley, Filippo Valsorda, and Isis Lovecruft. The ristretto255 group. IETF CFRG Internet Draft, 2020.
- 16 Danny Dolev, Shlomo Kramer, and Dalia Malki. Early delivery totally ordered multicast in asynchronous environments. In *FTCS*. IEEE, 1993.
- 17 Sisi Duan, Michael K Reiter, and Haibin Zhang. Secure causal atomic broadcast, revisited. In *IEEE/IFIP DSN*. IEEE, 2017.
- 18 Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 1988.
- 19 Adam Gagol and Michał Świątek. Aleph: A leaderless, asynchronous, byzantine fault tolerant consensus protocol. *arXiv preprint*, 2018. [arXiv:1810.05256](https://arxiv.org/abs/1810.05256).
- 20 Neil Girdharan, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. Bullshark: Dag bft protocols made practical. *arXiv preprint*, 2022. [arXiv:2201.05677](https://arxiv.org/abs/2201.05677).
- 21 Lioba Heimbach and Roger Wattenhofer. Sok: Preventing transaction reordering manipulations in decentralized finance. *arXiv preprint*, 2022. [arXiv:2203.11520](https://arxiv.org/abs/2203.11520).
- 22 Aniket Kate, Gregory M Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In *ASIACRYPT*. Springer, 2010.
- 23 Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is dag. In *ACM PODC*, 2021.
- 24 Idit Keidar, Oded Naor, and Ehud Shapiro. Cordial miners: A family of simple, efficient and self-contained consensus protocols for every eventuality. *arXiv preprint*, 2022. [arXiv:2205.09174](https://arxiv.org/abs/2205.09174).

- 25 Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. Themis: Fast, strong order-fairness in byzantine consensus. *Cryptology ePrint Archive*, 2021.
- 26 Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. In *CRYPTO*. Springer, 2020.
- 27 Hugo Krawczyk. Secret sharing made short. In *CRYPTO*. Springer, 1993.
- 28 Klaus Kursawe. Wendy grows up: More order fairness. In *Financial Crypto*. Springer, 2021.
- 29 Kfir Lev-Ari, Alexander Spiegelman, Idit Keidar, and Dahlia Malkhi. Fairledger: A fair blockchain protocol for financial institutions. *arXiv preprint*, 2019. [arXiv:1906.03819](https://arxiv.org/abs/1906.03819).
- 30 Peter M Melliar-Smith, Louise E. Moser, and Vivek Agrawala. Broadcast protocols for distributed systems. *IEEE TPDS*, 1990.
- 31 Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of bft protocols. In *ACM CCS*, 2016.
- 32 Louise E Moser and Peter M Melliar-Smith. Byzantine-resistant total ordering algorithms. *Information and Computation*, 1999.
- 33 Alexandre Obadia, Alejo Salles, Lakshman Sankar, Tarun Chitra, Vaibhav Chellani, and Philip Daian. Unity is strength: A formalization of cross-domain maximal extractable value. *arXiv preprint*, 2021. [arXiv:2112.01472](https://arxiv.org/abs/2112.01472).
- 34 Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *CRYPTO*. Springer, 1991.
- 35 Larry L Peterson, Nick C Buchholz, and Richard D Schlichting. Preserving and using context information in interprocess communication. *ACM TOCS*, 1989.
- 36 Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? *arXiv preprint*, 2021. [arXiv:2101.05511](https://arxiv.org/abs/2101.05511).
- 37 Michael K Reiter and Kenneth P Birman. How to securely replicate services. *ACM TOPLAS*, 1994.
- 38 Adi Shamir. How to share a secret. *Communications of the ACM*, 1979.
- 39 Victor Shoup and Rosario Gennaro. Securing threshold cryptosystems against chosen ciphertext attack. In *EUROCRYPT*. Springer, 1998.
- 40 Lei Yang, Seo Jin Park, Mohammad Alizadeh, Sreeram Kannan, and David Tse. DispersedLedger: High-throughput byzantine consensus on variable bandwidth networks. In *USENIX NSDI*, 2022.
- 41 Maofan Yin, Dahlia Malkhi, Michael K Reiter, Guy Golan Gueta, and Ittai Abraham. Hotstuff: Bft consensus with linearity and responsiveness. In *ACM PODC*, 2019.
- 42 Haoqian Zhang, Louis-Henri Merino, Vero Estrada-Galinanes, and Bryan Ford. Flash freezing flash boys: Countering blockchain front-running. In *DINPS*, 2022.
- 43 Yunhao Zhang, Srinath Setty, Qi Chen, Lidong Zhou, and Lorenzo Alvisi. Byzantine ordered consensus without byzantine oligarchy. In *USENIX OSDI*, 2020.