

F3B: A Low-Overhead Blockchain Architecture with Per-Transaction Front-Running Protection

Haoqian Zhang ✉

École Polytechnique Fédérale de Lausanne, Switzerland

Louis-Henri Merino ✉

École Polytechnique Fédérale de Lausanne, Switzerland

Ziyan Qu ✉

École Polytechnique Fédérale de Lausanne, Switzerland

Mahsa Bastankhah ✉

École Polytechnique Fédérale de Lausanne, Switzerland

Vero Estrada-Galiñanes ✉

École Polytechnique Fédérale de Lausanne, Switzerland

Bryan Ford ✉

École Polytechnique Fédérale de Lausanne, Switzerland

Abstract

Front-running attacks, which benefit from advanced knowledge of pending transactions, have proliferated in the blockchain space since the emergence of decentralized finance. Front-running causes devastating losses to honest participants and continues to endanger the fairness of the ecosystem. We present Flash Freezing Flash Boys (F3B), a blockchain architecture that addresses front-running attacks by using threshold cryptography. In F3B, a user generates a symmetric key to encrypt their transaction, and once the underlying consensus layer has finalized the transaction, a decentralized secret-management committee reveals this key. F3B mitigates front-running attacks because, before the consensus group finalizes it, an adversary can no longer read the content of a transaction, thus preventing the adversary from benefiting from advanced knowledge of pending transactions. Unlike other mitigation systems, F3B properly ensures that all unfinalized transactions, even with significant delays, remain private by adopting per-transaction protection. Furthermore, F3B addresses front-running at the execution layer; thus, our solution is agnostic to the underlying consensus algorithm and compatible with existing smart contracts. We evaluated F3B on Ethereum with a modified execution layer and found only a negligible (0.026%) increase in transaction latency, specifically due to running threshold decryption with a 128-member secret-management committee after a transaction is finalized; this indicates that F3B is both practical and low-cost.

2012 ACM Subject Classification Security and privacy → Distributed systems security

Keywords and phrases Blockchain, DeFi, Front-running Mitigation

Digital Object Identifier 10.4230/LIPIcs.AFT.2023.3

Related Version *Full Version:* <https://arxiv.org/abs/2205.08529> [65]

Funding This research was supported in part by U.S. Office of Naval Research grant N00014-19-1-2361, the AXA Research Fund, the PAIDIT project funded by ICRC, the IC3-Ethereum Fund, Algorand Centres of Excellence programme managed by Algorand Foundation, and armasuisse Science and Technology. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding sources.

Acknowledgements The authors wish to thank Cristina Basescu, Pasindu Nivanthaka Tennage, Pierluca Borsò-Tan, and Simone Colombo for their extremely helpful comments and suggestions and especially thank Shufan Wang for prototyping F3B on the Ethereum blockchain.



© Haoqian Zhang, Louis-Henri Merino, Ziyan Qu, Mahsa Bastankhah, Vero Estrada-Galiñanes, and Bryan Ford;

licensed under Creative Commons License CC-BY 4.0

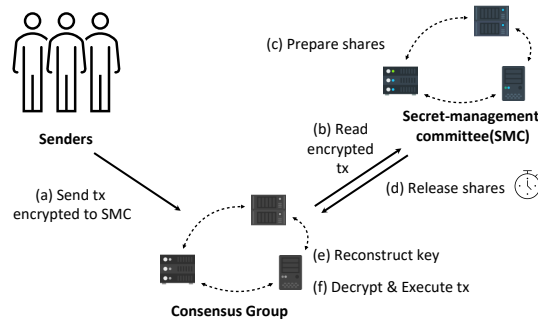
5th Conference on Advances in Financial Technologies (AFT 2023).

Editors: Joseph Bonneau and S. Matthew Weinberg; Article No. 3; pp. 3:1–3:23

Leibniz International Proceedings in Informatics



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



■ **Figure 1** F3B architecture. Senders publish encrypted transactions to the consensus group. The secret-management committee releases the decryption shares once the transactions are no longer pending. Finally, the consensus group reconstruct the key and decrypt and execute the transaction. The secret-management committee and the consensus group can consist of the same set of servers. For clarity in this paper, we logically separate them into two different entities.

1 Introduction

Front-running is the practice of benefiting from the advanced knowledge of pending transactions [20, 7, 1]. Although benefiting some entities involved, this practice puts others at a significant financial disadvantage, making this behavior illegal in traditional markets with established securities regulations [20].

However, the open and pseudonymous nature of blockchain transactions and the difficulties of pursuing miscreants across numerous jurisdictions have made front-running attractive, particularly in decentralized finance (DeFi) [38, 20, 15]. Front-running actors in the blockchain space can read the contents of pending transactions and benefit from them by, *e.g.*, creating their own transactions and positioning them according to the target transaction [4, 15, 20].

Front-running negatively impacts honest DeFi actors and endangers the fairness of this multi-billion market [18]. One estimate suggests that front-running attacks amount to \$280 million in losses for DeFi actors each month [42]. Front-running also threatens the underlying consensus layer’s security by incentivizing unnecessary forks [17, 15].

Despite work addressing front-running, several unmet challenges exist, such as high latency, being restricted to a specific environment, or raising security concerns. Namecoin, an early example of mitigating front-running attacks by having users send a commit and later a reveal transaction, requires two rounds of communication with the underlying blockchain [29]. Submarine further improves Namecoin’s design by hiding the addresses of smart contracts involved, but it induces three rounds of communication to the underlying blockchain [39, 29]. Both approaches induce high latency. Other works have taken a different approach to mitigate front-running attacks by tailoring their solution to a specific application or consensus algorithm [13, 52, 60, 5, 2, 3, 40, 25].

A promising approach is to use threshold encryption, where clients encrypt their transactions to prevent malicious actors from understanding those transactions, as presented in Fairblock [43] and Shutter [57, 58]. However, these schemes require clients to choose a future block to derive the encryption key, which raises security concerns. Suppose a transaction failed to be finalized in the client-chosen block due to, for example, a crypto mania that overwhelms the blockchain network [31] or a deliberate denial of service attack [20]. In this case, the transaction is undesirably revealed (see Section 3.3 for details).

We present Flash Freezing Flash Boys¹ (F3B), a novel blockchain architecture with front-running protection that has a low latency overhead and is compatible with existing consensus algorithms and smart contract implementations. Like Fairblock [43] and Shutter [57, 58], F3B addresses front-running by adopting threshold encryption, but it accomplishes this on a **per-transaction** basis rather than a per-block basis. Rather than selecting an encryption key linked to a future block, clients generate an encryption key for each transaction. This ensures that a transaction remains confidential until the block containing the transaction has received enough confirmations.

As described in Figure 1, F3B’s architecture consists of the following steps: (a) A client encrypts their transaction to a secret-management committee (SMC) and sends their encrypted transaction to the consensus group that operates the underlying blockchain. (b) The SMC reads the encrypted transaction from the underlying blockchain. (c) The SMC prepares the decryption shares for the consensus group. (d) The SMC releases the decryption shares to the consensus group once the underlying blockchain has finalized the transaction. (e) The consensus group reconstructs the key. (f) The consensus group decrypts and executes the transaction. Once the SMC begins to release the decryption shares, malicious actors cannot launch a front-running attack because the transaction is already irreversibly ordered on the blockchain. Although adversaries may attempt to run *speculative* front-running attacks, where they guess the contents of a transaction on metadata information like the sender’s address, these attacks are more likely to fail and can prove to be unprofitable [4]. Nonetheless, we discuss mitigation solutions for these attacks in Section 10.4.

F3B addresses two key practical challenges: (a) mitigating spamming of inexecutable encrypted transactions onto the underlying blockchain, and (b) limiting latency overhead. To mitigate spamming, we introduce a deposit-refund storage fee for storing encrypted transactions, along with the standard execution fee (*e.g.*, gas in Ethereum). To limit the latency overhead, users write only data onto the underlying blockchain once to achieve front-running protection.

We propose two cryptographic threshold schemes that can plug into F3B: TDH2[56] and PVSS [53]. TDH2 enables clients to encrypt their transactions under the same public key of a secret-management committee which is only changeable by time-consuming DKG or resharing protocols. On the other hand, PVSS empowers clients to adopt a different secret-management committee for each transaction but at the cost of the additional preprocessing time for preparing the shares for each transaction.

We implemented a prototype of F3B with post-Merge² Ethereum [23] as the underlying blockchain and Dela [19] as the secret-management committee. We measure the latency overhead by comparing the time it takes to decrypt and execute a transaction with the time it takes just to execute the transaction. Our analysis shows that, with a committee size of 128, the latency overhead is 0.026% and 0.027% for Ethereum under the TDH2 and PVSS respectively; In comparison, Submarine, which also offers per-transaction protection and hides the address of smart contracts as F3B, exhibits a 200% latency overhead, as it requires three rounds of communication with the underlying blockchain [39, 11]. For part of our prototype, we modified Ethereum’s execution layer by adding a new transaction type featuring encryption and delayed execution. By only modifying the execution layer, we

¹ The name *Flash Boys* comes from a popular book revealing this aggressive market-exploiting strategy on Wall Street in 2014 [38].

² The Merge refers to the merge executed on September 15th, 2022, to complete Ethereum’s transition to proof-of-stake consensus.

3:4 Flash Freezing Flash Boys(F3B)

can (a) provide compatibility with various consensus algorithms embedded in Ethereum’s consensus layer, including Proof-of-Work (PoW), Proof-of-Authority (PoA) and the recently added, Proof-of-Stake (PoS) and (b) protect existing smart contracts without requiring any code modifications.

In this paper, our key contributions are as follows:

1. The design of a blockchain architecture with front-running protection that uses threshold encryption on a per-transaction basis, enabling confidentiality for all pending transactions, even if transactions are delayed, while achieving low overhead.
2. The design of two protocols based on TDH2 and PVSS for F3B satisfies various demands and user scenarios.
3. A prototype that, on Ethereum’s execution layer, demonstrates F3B’s ability to be agnostic to (a) the underlying consensus algorithm and (b) to smart contract implementations while achieving low-latency overhead.
4. A systematic evaluation of F3B on post-Merge Ethereum by looking at transaction latency, throughput, and reconfiguration costs.

2 Background

In this section, we present a brief background on blockchain and smart contracts, and we introduce front-running attacks and mitigation strategies.

2.1 Blockchain & Transaction Ordering

A blockchain is an immutable append-only ledger of ordered transactions [44]. However, transactions go through a series of stages before they are finalized – irreversibly ordered – on the blockchain. After a sender creates a transaction, they need to propagate the transaction among the consensus nodes that then place the transaction in a pool of pending transactions, most commonly known as mempool. Notably, these transactions are *not yet irreversibly ordered*, thus opening up the possibility for front-running attacks. Furthermore, under certain probabilistic consensus algorithms, such as PoW or PoS, a transaction inserted onto the blockchain can still be reordered by inducing a fork of the underlying blockchain. Hence, to guarantee irreversible ordering for probabilistic consensus algorithms, a transaction must receive enough block confirmations – the number of blocks succeeding the block containing the transaction [44, 34, 14].

2.2 Smart Contract & Decentralized Exchange

A smart contract is an executable computer program modeled after a contract or an agreement that executes automatically [50]. A natural fit for smart contracts is on top of decentralized fault-tolerant consensus algorithms, such as PBFT-style algorithms, PoW, and PoS, to ensure their execution and integrity [63, 44, 32].

Although Bitcoin uses a form of smart contracts [44], it was not until Ethereum’s introduction that the blockchain space realized Turing-complete smart contracts, the backbone necessary for creating complex decentralized exchanges. To interact with these complex smart contracts, users need to pay *gas*, a pseudo-currency that represents the execution cost by miners [21]. However, the expressiveness of smart contracts comes with significant risks, from inadvertent vulnerabilities to front-running. Front-running is exhibited by the lack of guarantees that the underlying blockchain provides regarding ordering.

2.3 Front-Running Attacks & Mitigation

The practice of front-running involves benefiting from advanced knowledge of pending transactions [20, 7, 1]. In itself, knowledge of pending transactions is harmless, but the ability to act on this information is where the true problem lies. In the context of blockchains, an adversary performs a front-running attack by influencing the order of transactions, provided that transactions in the mempool are entirely in the clear.

Cryptocurrencies suffer from mainly three types of front-running attacks [20]: displacement, insertion, and suppression. *Displacement* is the replacement of a target transaction with a new transaction formulated by the front-running attacker. *Insertion* is the malicious introduction of a new transaction before a target transaction in the finalized transaction ordering. *Suppression* is the long-term or indefinite delaying of a target transaction.

In an ideal world, front-running protection would consist of an *immediate* global ordering of each transaction, as clients broadcast their transactions to prevent attackers from changing their order. In reality, even if all participants were honest, such global ordering is practically impossible due to clock synchronization [16] and consistency problems (*e.g.*, two transactions having the same time). Malicious participants can still carry out front-running attacks, because timings can easily be manipulated.

A more practical solution involves encrypting transactions, thereby preventing the consensus group from knowing the contents of the transactions when ordering them. This solution mitigates front-running attacks as an attacker is hindered from taking advantage of pending *encrypted* transactions.

3 Strawman Protocols

In order to explore the challenges inherent in building a framework, such as F3B, we first examine a couple of promising but inadequate strawman approaches, representative of state-of-the-art proposals [39, 11, 43, 57] but simplified for expository purposes.

3.1 Strawman I: Sender Commit-and-Reveal

The first strawman design has the sender create two transactions: a *commit* and a *reveal* transaction. The commit transaction is simply a commitment (*e.g.*, hash) of the intended reveal transaction, which is simply the typical contents of a transaction that is normally vulnerable to front-running. The sender will propagate the commit transaction and then *wait* until its finality by the consensus group, before releasing the reveal transaction. Once the reveal transaction is propagated, the consensus group proceeds to verify and to execute the transaction, in the execution order that the commit transaction was finalized on the blockchain. Given the finality in the former transaction, the sender is unable to change the contents of the reveal transaction.

This simple strawman protocol mitigates front-running attacks because the commit transaction determines the execution order and the contents of the commit transaction do not expose the contents of the reveal transaction. However, this strawman protocol presents some notable challenges: (a) the sender must remain online to continuously monitor the blockchain to know when to release their reveal transaction, (b) the reveal transaction might be delayed due to a congestion event like the cryptokitties mania [31] or a deliberate denial-of-service (DoS) attack like the Fomo3D incident [20], (c) this approach is subject to output bias, as the consensus nodes or the sender can deliberately choose not to reveal certain transactions during the reveal phase [4], such as only revealing profitable ones and aborting others, and (d) this approach has a significant latency overhead of over 100%, given that the sender must now send two non-overlapping transactions instead of the one standard transaction.

3.2 Strawman II: The Trusted Custodian

A straightforward method for removing the sender from the equation, after sending the commit transaction, is to employ a trusted custodian. After the consensus group finalizes the transaction onto the underlying blockchain, the trusted custodian reveals the transaction's contents.

This strawman protocol mitigates front-running attacks, as the nodes cannot read, before ordering, the contents of the transaction. However, the trusted custodian presents a single point of failure: Consensus nodes cannot decrypt and execute a transaction if the custodian crashes. Instead, by employing a *decentralized* custodian, we can mitigate the single point of failure issues.

3.3 Strawman III: Threshold Encryption with Block Key

The next natural step is to have a decentralized committee that generates a public key for each block, thus enabling a user to encrypt their transaction for a future block. The committee would then release the private key after the block finality. Furthermore, the committee can use identity-based encryption [55] to enable users to derive a future block key based on the block's height.

This strawman protocol seems to mitigate front-running, as the transactions in a block are encrypted until they are finalized in their intended block. However, if an encrypted transaction fails to be included in the specified block, its contents will be revealed shortly thereafter while remaining unfinalized, thus making it vulnerable to front-running. Blockchain networks have repeatedly observed such failures due to congestion, such as cryptokitties manias [31], or well-funded DoS attacks, such as the Fomo3D attack that flooded the Ethereum network with transactions for three minutes [20]. Such an approach can incentivize a consensus node to intentionally produce an empty block by aiming to reveal the pending transactions for that block. Therefore, we require a per-transaction rather than a per-block level of confidentiality, thus ensuring that a transaction is never revealed before it is finalized on the blockchain.

4 System Overview

In this section, we present F3B's system goals, architecture, and models.

4.1 System Goals

Our system goals, inspired by our strawman protocols, are

- **Front-Running Protection:** prevents entities from practicing front-running.
- **Decentralization:** mitigates a single point of failure or compromise.
- **Confidentiality:** reveals a transaction, only after the underlying consensus layer finalizes it.
- **Compatibility:** remains agnostic to the underlying consensus algorithm and to smart contract implementation.
- **Low-Latency:** exhibits low-latency transaction-processing overhead.

4.2 Architecture Overview

F3B, shown in Figure 1, mitigates front-running attacks by working with a secret-management committee to manage the storage and release of on-chain secrets. Instead of propagating their transactions in cleartext, the sender can now encrypt their transactions and store the

corresponding secret keys with the secret-management committee. Once the transaction is finalized, the secret-management committee releases the secret keys so that consensus nodes of the underlying blockchain can verify and execute transactions. Overall, the state machine replication of the underlying blockchain is achieved in two steps: the first is about the ordering of transactions, and the second is about the execution of transactions. As long as most trustees in the secret-management committee are secure and honest and the key is revealed to the public when appropriate, each consensus node can always maintain the same blockchain state.

F3B encrypts the entire transaction³, such as the smart contract address, inputs, sender's signature, and other metadata, as those information can provide enough information to launch a probabilistic front-running attack, such as the Fomo3D attack [20] or a speculative attack based on the leakage of metadata [4].

4.3 System and Network Model

F3B's architecture consists of three components: *senders* that publish (encrypted) transactions, the *secret-management committee* (SMC) that manages and releases secrets, and the *consensus group* that maintains the underlying blockchain. For the F3B based on the PVSS scheme, the client can choose a different SMC for each transaction. For the F3B based on the THD2 scheme, an SMC has a fixed membership over one epoch. When transiting from one epoch to the next, the SMC can modify its membership under the THD2 scheme with backward secrecy to prevent new trustees from decrypting old transactions without interrupting users' encryption by running a resharing protocol [62].

The secret-management committee and the consensus group can consist of the same set of servers. For clarity in this paper, we logically separate them into two different entities.

For the underlying network, we assume that all honest blockchain nodes and trustees of the SMC are well connected and that their communication channels are synchronous, *i.e.*, if an honest node or trustee broadcasts a message, then all honest nodes and trustees receive the message within a known maximum delay [46].

4.4 Threat Model

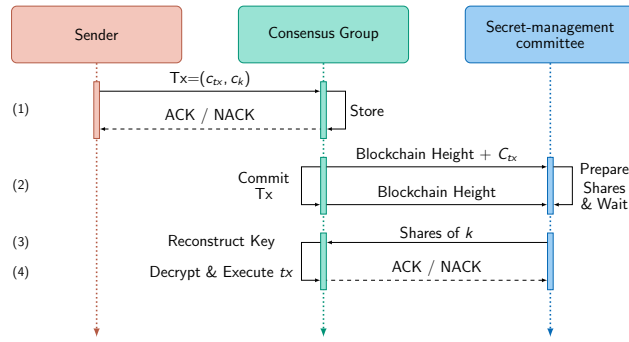
We assume that the adversary is computationally bounded, that the cryptographic primitives we use are secure, and in particular that the Diffie-Hellman problem and its decisional variant are hard. We further assume that all messages are digitally signed and that the consensus nodes and the SMC only process correctly signed messages.

The secret management committee consists of n trustees, where f can fail or behave maliciously. We require $n \geq 2f + 1$ and set the secret-recovery threshold to $t = f + 1$. We assume that the underlying blockchain is secure: *e.g.*, at most f' of $3f' + 1$ validators can fail or misbehave in a PBFT-style or PoS blockchain, or the adversary controls less than 50% computational power in a PoW blockchain. We acknowledge that the security assumptions for the secret management committee and the underlying blockchain might differ, potentially reducing the overall system's security to the least secure subsystem.

We assume that attackers do not launch speculative front-running attacks [4], but we present a discussion on some mitigation strategies for reducing side-channel leakage in Section 10.4.

³ Section 10.4 further discusses how to hide the sender's address.

3:8 Flash Freezing Flash Boys(F3B)



■ **Figure 2** F3B per-transaction protocol steps: (1) Send an encrypted transaction to the underlying blockchain, (2) Prepare shares by trustees while waiting for transaction finality, (3) Reconstruct the key, (4) Execute the transaction.

5 F3B Protocol

In this section, we introduce the F3B’s protocol, starting with some preliminaries, followed by the F3B’s detailed design. The full paper [65] offers a more comprehensive protocol description, and Section 10 introduces some optimizations.

5.1 Preliminaries

In this subsection, we introduce our preliminaries, including our baseline model for the underlying blockchain and the cryptographic primitives used in F3B.

Blockchain Model

To compare F3B’s impact, we model the underlying blockchain to involve a consensus protocol that finalizes transactions into a block that is linked to a previous block. We assume the underlying block’s time as L_b seconds. In PoW and PoS-based blockchains, a transaction is finalized only after a certain number of additional blocks have been added to the chain (also known as block confirmations). Thus, we define that a transaction is finalized after m block confirmations. Therefore, our baseline transaction latency is mL_b .

Shamir’s Secret Sharing

A (t, n) -threshold secret sharing scheme enables a dealer to share a secret s among n trustees such that any group of $t \leq n$ or more trustees can reconstruct s and no group less than t trustees learns any information about s . Whereas a simple secret sharing scheme assumes an honest dealer, verifiable secret sharing (VSS) enables the trustees to verify that the shares they receive are valid [24]. Public verifiable secret sharing (PVSS) further improves VSS to enable a third party to check all shares [53].

Distributed Key Generation (DKG)

DKG is a multi-party (t, n) key-generation process for collectively generating a private-public key pair (sk, pk) , without relying on a single trusted dealer; each trustee i obtains a share sk_i of the secret key sk , and collectively obtains a public key pk [27]. Any client can now use pk to encrypt a secret, and at least t trustees must cooperate to retrieve this secret [56].

5.2 Protocol Outline

We present the outline of F3B protocols with two different threshold cryptographic schemes. Figure 2 presents the protocol outline, and the full paper [65] offers a more comprehensive protocol description.

5.2.1 Protocol based on TDH2

Setup

Before an epoch, the secret-management committee runs a DKG protocol to generate a private key share sk_{smc}^i for each trustee and a collective public key pk_{smc} written onto the underlying blockchain. To offer chosen-ciphertext attack protection and to verify the correctness of secret shares, we utilize the TDH2 cryptosystem [56] containing NIZK proofs.

Per-Transaction Protocol

1. *Write Transaction:* A sender first generates a symmetric key k and encrypts it with pk_{smc} from the underlying blockchain, thus obtaining the resulting ciphertext c_k . Next, the sender creates their signed transaction and symmetrically encrypts it by using k , denoted as $c_{tx} = enc_k(tx)$. Finally, the sender sends (c_{tx}, c_k) to the consensus group who writes the pair onto the blockchain.
2. *Shares Preparation by Trustees:* Once written, each secret-management committee trustee reads c_k from the sender's transaction and prepares their decrypted share of k .
3. *Key Reconstruction:* When the sender's transaction (c_{tx}, c_k) is finalized onto the underlying blockchain (after m block confirmations), each secret-management committee trustee releases their share to the consensus group. The consensus group verifies the decrypted shares and uses them to reconstruct k by Lagrange interpolation of shares when there are at least t valid shares.
4. *Decryption and Execution:* The consensus group finally symmetrically decrypts the transaction $tx = dec_k(c_{tx})$ using k , thus enabling it to execute tx .

Resharing Protocol


To modify a SMC's membership and to offer backward secrecy over epochs, an SMC can periodically run a verifiable resharing protocol [62] to replace certain trustees or redistribute the trustees' private keys. Unlike DKG, resharing keeps the epoch's public key, thus preventing undesirable interruptions of encryption services.

5.2.2 Protocol based on PVSS

Per-Transaction Protocol

0. *Share Preparation By sender:* For every transaction, the sender runs the PVSS protocol [53] to generate an encrypted key share $share_i$ for each trustee, as well as a corresponding NIZK proof and public polynomial commitment. The proof and commitment can be used to verify the correctness of key share and protect against chosen-ciphertext attacks. The sender obtains the symmetric key k from the PVSS protocol.

3:10 Flash Freezing Flash Boys(F3B)



■ **Figure 3** In Ethereum, once they are inserted in the blockchain, the transactions are executed and finalized after receiving m block confirmations. Whereas, in F3B, transactions are encrypted, and their executions are postponed after receiving m block confirmations when the secret-management committee releases the encryption keys. Both scenarios have a similar finality time.

1. *Write Transaction*: A sender first creates the ciphertext c_k with the key shares, NIZK proofs, and commitments generated during share preparation. Next, the sender creates their transaction and symmetrically encrypts it by using the symmetric key k , denoted as $c_{tx} = enc_k(tx)$. Finally, the sender sends (c_{tx}, c_k) to the consensus group who writes the pair onto the blockchain.
2. *Shares Preparation by Trustees*: Same as (2) in 5.2.1.
3. *Key Reconstruction*: Same as (3) in 5.2.1.
4. *Decryption and Execution*: Same as (4) in 5.2.1.

5.3 Overhead Analysis

We analyze both protocols' overheads. Write Transaction (step 1) is identical to sending a transaction to the underlying blockchain. We assume trustees can finish Shares Preparation by Trustees (step 2) within the confirmation time of the tx⁴. Hence, the time for steps 1 and 2 is equivalent to finalizing a transaction on the underlying blockchain and waiting until its finality, which takes mL_b time based on our baseline model (Section 5.1). As in PVSS protocol, the sender can finish Share Preparation By sender (step 0) before having the tx ; thus step 0 does not contribute to the transaction latency. Comparing our protocol with the baseline, Key Reconstruction (step 3) and Decryption and Execution (step 4) are additional steps, and we denote the time of those steps to be L_r .

Figure 3 demonstrates the conceptual difference in finality and execution time between F3B and the baseline. As the secret-management committee releases the secret keys with a delay of m blocks, this introduces an execution delay of m blocks. However, in both cases, to prevent attacks such as double-spending, the recipient should not accept a transaction, until it is finalized. Therefore on a commercial level, F3B is similar to the baseline because it exhibits the finality time of a transaction that is of use to the recipient⁵.

⁴ As we presented in Section 9, confirmation time in Ethereum is much longer than the share preparation time by trustees.

⁵ In F3B, transaction finalization is slower due to the key reconstruction and delayed execution after transaction finality. However, the overhead is negligible compared to finality time, as discussed in Section 9.

5.4 TDH2 and PVSS: Pros and Cons

When applying THD2 and PVSS to F3B, each scheme has some advantages and disadvantages. This subsection offers qualitative comparisons, whereas Section 9 provides quantitative comparisons between the two protocols.

- **Preprocessing:** In TDH2, the secret-management committee needs to do DKG per epoch, whereas in PVSS, the sender needs to prepare shares per transaction.
- **Membership:** In TDH2, the secret-management committee’s membership is fixed per epoch, whereas in PVSS, the sender can choose a different secret-management committee for each transaction, providing the best flexibility.
- **Ciphertext:** TDH2 has a constant ciphertext length, whereas PVSS’s ciphertext grows linearly with the size of the secret-management committee.

In conclusion, no one protocol can completely replace another. System designers need to choose one or both protocols based on their needs and constraints to mitigate front-running attacks.

6 Achieving the System Goals

In this section, we present how F3B achieves the system goals outlined in Section 4.1.

Front-Running Protection: *prevents entities from practicing front-running.* We reason the protection offered by F3B from the definition of front-running: if an adversary cannot benefit from pending transactions, he cannot launch front-running attacks. In F3B, the sole entity that knows the content of a pending transaction is the sender who is financially incentivized to *not* release its contents. The content is revealed only when its transaction is finalized; thus, by definition, the attacker has no means to launch a front-running attack. However, we acknowledge that attackers can use side channels (*e.g.* metadata such as sender’s address and transaction size) of the encrypted transaction to launch *speculative* front-running attacks, as discussed in Section 4.4 and Section 10.4. We present a more comprehensive security analysis discussion in Section 7.

Decentralization: *mitigates a single point of failure or compromise.* Due to the properties of DKG [27], THD2 [56], and PVSS [53], the SMC can handle up to $t - 1$ malicious trustees and up to $n - t$ offline trustees.

Confidentiality: *reveals a transaction, only after the underlying consensus layer finalizes it.* The sender encrypts each transaction with a newly generated symmetric key. The symmetric key is (a) encrypted under the secret-management committee’s public key in TDH2-based protocol, (b) embedded into the encrypted shares in PVSS-based protocol. In both protocols, $f + 1$ trustees are required to retrieve the symmetric key. Per our threat model, only f trustees can behave maliciously; this ensures that the symmetric key cannot be revealed. We outline a more detailed security analysis in Section 7.

Compatibility: *remains agnostic to the underlying consensus algorithm and to smart contract implementation.* F3B requires modifying the execution layer to enable encrypted transactions. However, the consensus layer remains untouched, thus agnostic to the underlying consensus algorithms. Furthermore, F3B does not require to modify smart contract implementations, thus enabling existing smart contracts to benefit from front-running protection automatically.

Low-Latency: *exhibits low-latency transaction-processing overhead.* Similar to the baseline model, F3B requires clients to write only one transaction onto the underlying blockchain. This enables F3B to have a low-latency overhead compared to other front-running protection design that require multiple transactions for the same security guarantees. We present an evaluation of this latency overhead in Section 9.

7 Security Analysis

In this section, we introduce the security analysis of F3B’s protocol.

7.1 Front-Running Protection

From our threat model, we reason about why an attacker can no longer launch front-running attacks with absolute certainty of a financial reward, even with the collaboration of at most f malicious trustees. As we assume that the attacker does not launch speculative attacks based on metadata of the encrypted transactions, the only way the attacker can front-run transactions is by using the plaintext content of the transaction. As the attacker cannot access the content of the transaction before it is finalized on the underlying blockchain, then the attacker cannot benefit from the pending transaction. This prevents front-running attacks (by the definition of front-running). As we assume that the symmetric encryption we use is secure, the attacker cannot decrypt the transaction based on its ciphertext. Due to the properties of TDH2 [56], DKG [27], and PVSS [53] with our threat model, the attacker cannot obtain the private key and/or reconstruct the symmetric key. Recall that the attacker can collude with at most only f trustees, and that $f + 1$ are required to recover or gain information about the symmetric key.

7.2 Replay Attack

We consider a scenario in which an adversary can copy a pending (encrypted) transaction and submit it as their own transaction to reveal the transaction’s contents, before the victim’s transaction is finalized. By revealing the contents of the copied transaction, the attacker can then trivially launch a front-running attack. However, we explain the reason the adversary is unable to benefit from such a strategy.

In the first scenario, the adversary copies the ciphertext c_k and the encrypted transaction c_{tx} from tx_w , then creates a new write transaction tx'_w , digitally signed with their signature. However, even if the adversary’s tx'_w is decrypted and executed before the victim’s transaction tx_w , it effectively results in the blockchain executing tx_w ⁶. This leaves the adversary with no time to front-run their own tx'_w without knowing its contents.

In our second scenario, the adversary instead sends the transaction to a blockchain with smaller m block confirmations. Consider two blockchains b_1 and b_2 whose required number of confirmation blocks are m_1 and m_2 with $m_1 > m_2$. If the adversary changes the label L to L' for the blockchain b_2 instead of blockchain b_1 , the secret-management committee will successfully decrypt the transaction. However, we argue it is hard to form a valid write-transaction with L' by the adversary.

⁶ Note that tx is already a signed transaction; thus, tx_w and tx'_w have the same effect on the blockchain.

For the TDH2 protocol, the adversary would need to generate $e' = H_1(c, u, \bar{u}, w, \bar{w}, L')$ and $f = s + re'$, without knowing the random parameter r and s . Suppose the adversary generates $u = g^r, \bar{u} = \bar{g}^{r'}$ with $r \neq r'$ and $w = g^s, \bar{w} = \bar{g}^{s'}$ with $s \neq s'$. For tx'_w to be valid, we must have $g^f = wu^e$ and $\bar{g}^f = \bar{w}\bar{u}^e$, this implies that $(s - s') + e(r - r') = 0$. As $r \neq r'$, the adversary has only a negligible chance of having tx'_w pass verification.

For the PVSS protocol, the adversary must replace the original generator h with h' derived from $H(L')$. Hence, the adversary has to do the proofs without secrets. The security of PVSS guarantees that they only have a negligible probability of succeeding. Note that the base point has to be random to ensure the security. Using Elligator maps [8] guarantees that the generator h is random.

8 Incentive

F3B must incentivize actors to operate and follow the protocol honestly. In this section, we address the critical incentives that, in F3B, prevent spamming transactions and that deter collaboration among trustees from prematurely revealing transactions.

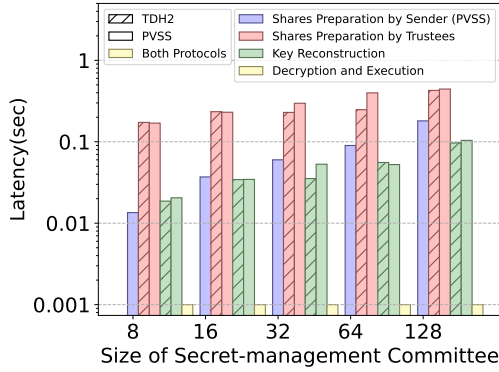
8.1 Spamming Protection

As the consensus group cannot execute encrypted transactions, an adversary could, at a low cost, spam the blockchain with non-executable transactions (*e.g.*, inadequate fees, malformed transactions), thus delaying the finality of honest transactions. To make such an attack costly, we introduce a storage deposit, alongside the traditional execution fee (*e.g.*, gas in Ethereum) and adjustable based on the transaction's size. The underlying blockchain can deduct the storage deposit from the sender's balance, much like paying a transaction fee. Then the blockchain can partially refund the deposit after successful execution by the consensus nodes. This approach imposes a low-cost fee on compliant users and a penalty on those who misbehave.

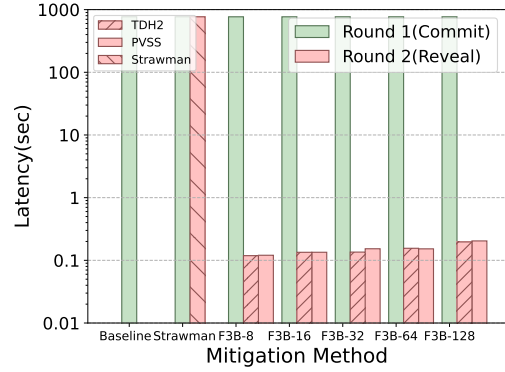
8.2 Operational Incentive

We need a similar incentive structure for the secret-management committees; similar to the way consensus nodes are rewarded for following the blockchain protocol via an execution fee. Whereas SMCs could be rewarded using the execution fee, this fee does not prevent SMC trustees from colluding for their financial gain. For example, an SMC might silently collude with a consensus group by prematurely giving them the decryption shares. Given the difficulty of detecting out-of-band collusion, we need to discourage it from doing so by significantly rewarding anyone who can prove the existence of such collusion.

We propose an incentive structure, where we require each trustee in a secret-management committee to lock an amount c for collateral and, in exchange, they are rewarded proportionally to the staked amount ac for the services they provide. Remind that, based on our threat model (Section 4.4), t trustees must collude altogether to reconstruct a transaction. To maintain security, the potential gain that t trustees benefit from front-running must be less than the potential loss $(1 + a)ct$, which malicious trustees would incur through the slashing protocol described in Section 8.3. Hence, a higher potential loss value $(1 + a)ct$ ensures security for a longer epoch length. If other factors stay consistent, designers can support a longer epoch length by either increasing the collateral requirement (by raising c) or by involving more trustees (by increasing t).



■ **Figure 4** The breakdown latency of each step in F3B by varying the number of secret-management committee trustees from 8 to 128 nodes.



■ **Figure 5** A comparison of the sender commit-and-reveal approach latency with F3B against a baseline modeled in Ethereum. The string “F3B- X ” represents X trustees.

8.3 Slashing Protocol

We need to have a protocol that rewards anyone who can prove a trustee’s or the entire secret-management committee’s misbehavior to discourage the release of the shares prematurely. At the same time, we do not want anyone to accuse a secret-management committee or a specific trustee without repercussions if the SMC or trustee did not actually misbehave.

To accomplish our objective, each trustee of the secret-management committee must stake some amount of cryptocurrency in a smart contract that handles disputes between a defendant (the entire secret-management committee or a particular trustee) and a plaintiff. To start a dispute, the plaintiff will invoke the smart contract with the correct decryption share for a currently pending transaction and their own stake. Suppose the smart contract validates that this is a correct decryption share and that the dispute started before the transaction in question was revealed by the secret-management committee. In this case, the defendant’s stake is forfeited and sent to the plaintiff.

At a protocol level, to prove a correct decryption share in protocol with TDH2, the plaintiff submits $[u_i, e_i, f_i]$ such that $e_i = H_2(u_i, \hat{u}_i, \hat{h}_i)$ where $\hat{u}_i = \frac{u^{f_i}}{u_i^{e_i}}$ and $\hat{h}_i = \frac{g^{f_i}}{h_i^{e_i}}$. In the protocol based on PVSS, the plaintiff submits $[s_i, \pi_{s_i}]$, where π_{s_i} is the NIZK proof that shows $\log_g pk_i = \log_{s_i} \hat{s}_i$. Even if the sender knows s_i , it is impossible to maliciously slash a trustee without the π_{s_i} , which only the corresponding trustee knows.

Deploying such a mechanism would require the smart contract to access the ciphertext of a transaction (*e.g.*, u or \hat{s}_i is necessary to verify the submitted share).

9 Evaluation

We prototype F3B by using post-merge Ethereum [23] as the underlying blockchain and Dela [19] written in Go [28] as the secret-management committee for our evaluation. Remaining consistent with Ethereum’s security assumptions, one epoch length lasts 6.4 minutes and the trustees forming the secret-management committee from validators for a given epoch are randomly selected. We instantiate our cryptographic primitives by using the Edward25519 elliptic curve with 128-bit security supported by Kyber [37], an advanced cryptographic library. We ran experiments on a server with 32GB of memory and 40 (2.1GHz) CPU cores. The network communication delay is simulated to be a fixed 100ms. We further discuss F3B integration with Ethereum in Section 10.

■ **Table 1** Latency Overhead for Ethereum Blockchain.

Confirmations	Latency Overhead varying SMC sizes			
	TDH2		PVSS	
	64	128	64	128
8	0.164%	0.206%	0.160%	0.214%
16	0.082%	0.103%	0.080%	0.107%
32	0.041%	0.052%	0.040%	0.053%
64	0.020%	0.026%	0.020%	0.027%
128	0.010%	0.013%	0.010%	0.013%

■ **Table 2** Storage overhead for two protocols with different Secret-management Committee sizes.

Number of trustees	Storage Overhead (bytes)	
	TDH2 Protocol	PVSS protocol
8	80	792
16	80	1568
32	80	3120
64	80	6224
128	80	12 432

9.1 Latency

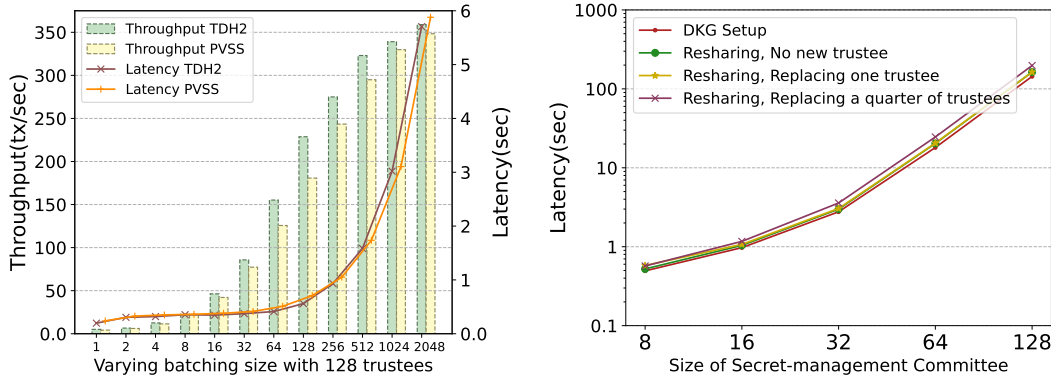
In Figure 4, we present the breakdown latency of each step for both TDH2 and PVSS protocols after a transaction finality while varying the number of SMC trustees from 8 to 128 nodes: (a) shares preparation by trustees, and (b) key reconstruction, and (c) decryption and execution. In addition, we show the time needed for PVSS shares generation by the sender in purple of Figure 4. As discussed in Section 5.3, only (b) and (c) represent the overhead at the per-transaction level.

Recall that the overall transaction latency using F3B is $mL_b + L_r$ (Section 5.2). In post-Merge Ethereum, the block time is fixed to 12 seconds, *i.e.*, $L_b = 12$ [9], and, by official standard, a block requires 64 block confirmations (two epochs) to be “finalized”, *i.e.*, $m = 64$ [22].

Figure 5 presents the end-to-end latency comparison between the baseline protocol (Section 5.1), a sender-only commit-and-reveal protocol, as presented in Strawman 1 (Section 3.1), and F3B’s protocol – varying the size of the secret-management committee stated after the string “F3B-”. With the new PoS consensus, finalizing any data in Ethereum requires $mL_b = 64 * 12 = 768$ seconds. The baseline protocol’s total latency is 768 seconds, as it requires only one write to the blockchain. Recall that in the sender-based commit-and-reveal approach (Strawman I), the sender commits a hash to the blockchain, taking 768 seconds, then reveals the transaction in another 768 seconds, totaling 1536 seconds. This results in a 100% latency overhead compared to the baseline, as the two steps must be sequential: the hash must be finalized on the blockchain before the reveal transaction can be propagated. Submarine, a more advanced approach that conceals the smart contract address, requires three sequential transactions. The sender must publish these three transactions in order, with the blockchain finalizing each one before the next one can be sent, suffering a latency delay of $768 * 3 = 2304$ seconds or a 200% latency overhead compared to the baseline [39, 11].

Compared with F3B, the reveal phase (key-reconstruction step) does not require the sender to write any data onto the blockchain. Therefore, we emphasize a significant difference between F3B and other application-based commit-and-reveal approaches, where F3B requires sending only one transaction to the underlying blockchain. Figure 5 shows that our design brings a low-latency overhead of 197ms and 205ms for two protocols, equivalent to 0.026% and 0.027% for Ethereum (relative to the 768 seconds finality time), under an SMC size of 128.

We acknowledge that some Ethereum users may accept a lower confirmation number to accept a transaction, even though Ethereum officially requires 64 blocks [22]. Without loss of generality, we outline different confirmation numbers with F3B’s latency overhead in Table 1.



■ **Figure 6** Performance of the key construction. ■ **Figure 7** The latency cost of DKG setup and three resharing scenarios.

9.2 Throughput

Figure 6 presents the F3B’s throughput results with a secret-management committee consisting of 128 trustees, assuming the underlying blockchain is not the bottleneck. If the keys are individually reconstructed, F3B provides limited throughput due to network transmission overhead incurred from sequential execution. Instead, we can batch keys by reconstructing them concurrently and presenting them in one network transmission. We present this batching effect in Figure 6 by varying the batching size to measure throughput and corresponding latency. By increasing the batching size from 1 to 2048, we can improve throughput from 5 txns/sec to 359 txns/sec with the TDH2 protocol, and from 4 txns/sec to 348 txns/sec with the PVSS protocol. The increased throughput comes with a higher latency cost: With a batching size of 2048, the key reconstruction step of TDH2 now takes 5.71 seconds to process, and the same step of PVSS takes 5.88 seconds; this latency is equivalent to a 0.74% and 0.77% latency overhead over Ethereum. Our results show that F3B provides more than sufficient throughput to support Ethereum (15 tx/sec [51]).

9.3 Reconfiguration in TDH2

Figure 7 demonstrates the cost of reconfiguring a secret-management committee in the TDH2 protocol. Recall that DKG is a one-time setup operation per epoch, bootstrapped during the previous epoch. Our experiment shows that, with a committee size of 128 trustees, DKG takes about 144 seconds; this is about 37.5% of Ethereum’s epoch time (384 seconds). We offer a further discussion about the transition between two epochs in Section 10.1. To provide backward secrecy and dynamic membership, a secret-management committee can run a verifiable resharing protocol [62] within an epoch and, by keeping the public key, without interrupting users’ encryption. Figure 7 illustrates the cost of three scenarios: (a) resharing among the same committee, (b) replacing one trustee, and (c) replacing a quarter of trustees. They all exhibit latency of the same magnitude.

9.4 Storage Overhead

In the TDH2 protocol, as the symmetric key is encrypted with the shared public key, the size of c_k is independent of the number of trustees. We also optimize the original TDH2 protocol to remove the label L from the ciphertext but only insert L in the computation

and verification steps of each party (consensus group, secret-management committee, sender) for protection against replay attacks (Section 7.2). Ultimately, we achieve 80 bytes per transaction of the storage overhead presented in Table 2.

In the PVSS protocol, however, the ciphertext c_k contains encrypted shares, NIZK proofs, and polynomial commitments. The size of the c_k thus approximately grows linearly with the number of trustees, as demonstrated in Table 2. The difference in storage overhead is one of the trade-offs when system designers need to consider using which encryption algorithm in F3B, as discussed in Section 5.4.

10 Discussion

In this section, we discuss some deployment challenges. We leave a detailed analysis for future work.

10.1 Transition of Epoch

Each epoch has its unique public epoch key for a user to encrypt his symmetric key used for encrypting a transaction. However, users will have difficulty choosing the correct epoch key when the time is close to the transition between two epochs. With Fairblock [43] and Shutter [57, 58], undesirable transaction revealing occurs when a user chooses the wrong public key. Whereas, if the transaction is not finalized, F3B never reveals any transaction, regardless of the chosen key; thus offering confidentiality to all unfinalized transactions. We also expect that such an epoch transition is infrequent, compared with a block transition, thus it causes much less trouble to users. In F3B, if a user uses an old epoch key for his encryption, he can safely try again to select the new epoch key. To mitigate the issue even more, the expiring epoch committee can offer some grace period, thus allowing both old and new epoch keys to be valid for a certain period. This significantly reduces the danger of a user choosing an incorrect key.

10.2 Ethereum Gas Fees

Ethereum uses gas fees to cover the cost of executing a transaction and implements a maximum gas limit per block. Incorporating F3B on Ethereum would then require (1) that the gas limit of each transaction to be in cleartext, and (2) that the summation of all transactions' *gas limit* within a block does not exceed the *block gas limit*. This opens the possibility for another type of spamming attack (Section 8.1), where an adversary submits transactions with substantial gas limits, thus leaving little room for other transactions.

Recall that the actual gas used by transactions cannot be determined because the sender encrypts its contents, and that accurately estimating the gas cost of a transaction is particularly difficult due to the uncertainty of the global state when validators process the transaction. One potential approach to mitigating this kind of attack would be to then burn the remaining (unused) gas. However, this approach could be too strict in practice, hence we can instead envision partial refunds: refunding the remaining gas up to a percentage.

10.3 Verifiable Key Propagation

Under our proposed protocol, every consensus node must fetch the shares and run the Lagrange interpolation to reconstruct the key. Would it instead be possible for one of the consensus nodes to reconstruct the symmetric key k from the secret shares and to propagate it to other consensus nodes with a succinct proof?

Therefore, we propose a solution that requires additional storage overhead in exchange for faster verification: Instead of constructing their encrypted transaction as (c_k, c_{tx}) , the sender additionally adds a hash of the symmetric key $h_k = H(k)$ as the third entry, creating the following signed write transaction: $tx_w = [c_k, c_{tx}, h_k]_{sig_{sk_A}}$.

During key reconstruction, after recovering or receiving k , consensus nodes need to verify whether the hash of k is consistent with the one (h_k) published on the ledger. If it is consistent, a consensus node can continue to decrypt the transaction and propagate the key k to others. If it is inconsistent, however, a consensus node must reconstruct the key from decryption shares and publish the shares to the underlying blockchain to slash the sender who provides a wrong h_k .

10.4 Metadata Leakage

In our architecture, adversaries can only observe encrypted transactions until they are finalized, thus preventing the revelation of transaction contents to launch front-running attacks. Nevertheless, to launch speculative attacks, adversaries can rely on side channels such as transaction metadata. Concretely, as the sender needs to pay the storage fee (Section 8.1) for publishing an encrypted transaction to the underlying blockchain, this leaks the sender's address. Knowledge of the sender's address can help in launching a front-running attack because an adversary might be able, based on the sender's history, to predict the sender's behavior. To prevent this second-order front-running attack, a sender can use a different address to pay for the storage fee. The underlying blockchain can also offer anonymous payment to users, such as Zerocash [49] or a mixing service [66], to further hide the payment address. Another side-channel leakage is the size of the encrypted transaction or the time the transaction is propagated. A possible remedy for mitigating metadata leakage is PURBs [45].

10.5 Key Storage and Node Catchup

In our protocol, if a new node wants to join the consensus group, it cannot execute the historical transactions to catch up, unless it obtains all decryption keys. The secret-management committee or consensus group can store these keys independently from the blockchain, but this requires them to maintain an additional immutable ledger. As consensus nodes already maintain one immutable storage medium, namely the underlying blockchain, the keys can be stored on this medium as metadata; and the blockchain rule can require storing valid keys when producing blocks.

However, this optimization brings about a timing issue, *i.e.*, When should the blockchain require the consensus group to store keys in a block? From our protocol, the transaction is finalized at block height n and revealed at block height $n + m$, thus making the earliest block to write the key at block height $n + m + 1$. With respect to the latest block height to write the key, there is much more flexibility and we need to consider the balance between the delay tolerance for all consensus nodes to retrieve the key and the time that consensus nodes must retain the key. Assuming that the key reconstruction step takes up to δ block times, the key should be written in or before the block $n + m + \delta$.

Although this setup would work well for a blockchain with fixed block time, care must be taken for blockchains where block time is probabilistic as the key might not have been replicated to all consensus nodes at block height $n + m + \delta$, thus some artificial delay for new blocks could be induced.

11 Related Work

Namecoin is a decentralized name service and an early work on front-running protection using a commit-and-reveal design [29]. In Namecoin, a user first broadcasts a salted hash of their name and then, after finality, broadcasts the actual name. Our first strawman protocol (Section 3.1) is based on Namecoin.

After Namecoin, Eskandari et al. [20] systematized front-running attacks on the blockchain by presenting three types of front-running attacks: displacement, insertion, and suppression. Daian et al. [15] also quantified front-running attacks from an economic point of view, determining that front-running attacks can also pose a security risk to the underlying consensus layer by incentivizing unnecessary forks driven by the maximal extractable value (MEV).

Many previous works explore the idea of applying threshold cryptography on blockchain. Virtual ASICs use threshold encryption to implement an all-or-nothing broadcast in the blockchain layer [26]. Sikka [59], Ferveo (Anoma) [5], Schmid [52], Dahlia [40], and Helix [3] apply threshold encryption to mitigate front-running but only present discussions with specific consensus algorithms. Fairblock [43] and Shutter [57, 58] enable encrypted transactions on a per-block basis, but if an encrypted transaction fails to be included in the sender-chosen block, then the transaction would be revealed; our Strawman III design (Section 3.3) is based on their approach.

Calypso is a framework that enables on-chain secrets by adopting threshold encryption governed by a secret-management committee [33]. Calypso allows ciphertexts to be stored on the blockchain and collectively decrypted by trustees according to a predefined policy. F3B leverages Calypso to specifically mitigate front-running attacks and extends its functionality to release the transaction contents once finalized automatically. F3B adopts per-transaction encryption, thus protecting all unfinalized transactions from front-running attacks, even if the transactions are delayed.

Other works adopt different approaches to mitigate front-running. A series of recent studies focus on fair ordering [30, 35, 36], but they cannot prevent an adversary with a rapid network connection [4]. Wendy explores the possibility of combining fair ordering with commit-and-reveal [36] but is in need of quantitative overhead analysis. Submarine is an application-layer front-running protection approach that extends a commit-and-reveal design to prevent leakage of the smart contract address. However, it presents a high latency overhead by requiring senders to have three rounds of communication with the underlying blockchain [39, 11].

Some works adopt time-lock puzzles [48] to blind transactions. For example, the injective protocol [12] uses a verifiable delay function [10] to achieve a proof-of-elapsed-time. However, an open challenge remains to link the time-lock puzzle parameters to an actual real-world delay [4].

Finally, works such as MEV-SGX [41], Tesseract [6], Secret Network [54], and Fairy [60] use a trusted execution environment [64] to mitigate front-running. Nevertheless, these approaches use a centralized component that is then subject to a single point of failure or compromise [47, 61].

12 Conclusion

In this paper, we have introduced F3B, a novel blockchain architecture that addresses front-running attacks with TDH2 and PVSS as threshold encryption protocols on a per-transaction basis. Our evaluation of F3B demonstrates that F3B is agnostic to consensus algorithms

and to existing smart-contract implementations. We have also shown that F3B meets the necessary throughput while presenting a low-latency overhead, thus fitting with Ethereum. Given that the deployment of F3B would require modifications to a blockchain's execution layer, F3B, in return, would also provide a substantial benefit: the F3B-deployed blockchain would now, by default, contain standard front-running protection for all applications in need at once without requiring any modifications to smart contracts themselves.

References

- 1 Nasdaq: Front running. <https://www.nasdaq.com/glossary/f/front-running>, 2018. Accessed: 2022-04-17.
- 2 Prashant Ankalkoti and SG Santhosh. A relative study on bitcoin mining. *Imperial Journal of Interdisciplinary Research (IJIR)*, 3(5):1757–1761, 2017.
- 3 Avi Asayag, Gad Cohen, Ido Grayevsky, Maya Leshkowitz, Ori Rottenstreich, Ronen Tamari, and David Yakira. Helix: A scalable and fair consensus algorithm resistant to ordering manipulation. *Cryptology ePrint Archive*, 2018.
- 4 Carsten Baum, James Hsin-yu Chiang, Bernardo David, Tore Kasper Frederiksen, and Lorenzo Gentile. Sok: Mitigation of front-running in decentralized finance. *Cryptology ePrint Archive*, 2021.
- 5 Joseph Bebel and Dev Ojha. Ferveo: Threshold decryption for mempool privacy in bft networks. *Cryptology ePrint Archive*, 2022.
- 6 Iddo Bentov, Yan Ji, Fan Zhang, Lorenz Breidenbach, Philip Daian, and Ari Juels. Tesseract: Real-time cryptocurrency exchange using trusted hardware. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 1521–1538, 2019.
- 7 Dan Bernhardt and Bart Taub. Front-running dynamics. *Journal of Economic Theory*, 138(1):288–296, 2008.
- 8 Daniel J Bernstein, Mike Hamburg, Anna Krasnova, and Tanja Lange. Elligator: elliptic-curve points indistinguishable from uniform random strings. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 967–980, 2013.
- 9 Blocks, 2022. Accessed: 2022-10-03. URL: <https://ethereum.org/en/developers/docs/blocks>.
- 10 Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *Annual international cryptology conference*, pages 757–788. Springer, 2018.
- 11 Lorenz Breidenbach, Phil Daian, Florian Tramèr, and Ari Juels. Enter the hydra: Towards principled bug bounties and {Exploit-Resistant} smart contracts. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 1335–1352, 2018.
- 12 Eric Chen and Albert Chon. Injective protocol: A collision resistant decentralized exchange protocol [White paper], 2018. URL: <https://coinpare.io/whitepaper/injective-protocol.pdf>.
- 13 Michele Ciampi, Muhammad Ishaq, Malik Magdon-Ismail, Rafail Ostrovsky, and Vassilis Zikas. Fairmm: A fast and frontrunning-resistant crypto market-maker. *Cryptology ePrint Archive*, 2021.
- 14 Coinbase. Coinbase confirmations, 2022(?). Accessed: 2022-03-03. URL: <https://help.coinbase.com/en/coinbase/getting-started/crypto-education/glossary/confirmations>.
- 15 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 *2020 IEEE Symposium on Security and Privacy (SP)*, pages 910–927. IEEE, 2020.
- 16 N Neha Dalwadi and C Mamta Padole. Comparative study of clock synchronization algorithms in distributed systems. *Advances in Computational Sciences and Technology*, 10(6):1941–1952, 2017.

- 17 dapp.org. Uniswap v2 audit report, 2020. Accessed: 2022-01-22. URL: <https://dapp.org.uk/reports/uniswapv2.html>.
- 18 Defi pulse. URL: <https://www.defipulse.com/?time=All>.
- 19 Dedis ledger architecture. URL: <https://github.com/dedis/de1a>.
- 20 Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. Sok: Transparent dishonesty: Front-running attacks on blockchain. In *Financial Cryptography Workshops*, volume 11599 of *Lecture Notes in Computer Science*, pages 170–189. Springer, 2019.
- 21 Gas and fees, 2022. Accessed: 2022-10-03. URL: <https://ethereum.org/en/developers/docs/gas/>.
- 22 Gasper, 2022. Accessed: 2022-10-03. URL: <https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/gasper/>.
- 23 The merge, 2022. Accessed: 2022-10-03. URL: <https://ethereum.org/en/upgrades/merge/>.
- 24 Paul Feldman. A practical scheme for non-interactive verifiable secret sharing. In *28th Annual Symposium on Foundations of Computer Science (sfcs 1987)*, pages 427–438. IEEE, 1987.
- 25 Flashbots protect. URL: <https://docs.flashbots.net/flashbots-protect/overview>.
- 26 Chaya Ganesh, Claudio Orlandi, Daniel Tschudi, and Aviv Zohar. Virtual asics: generalized proof-of-stake mining in cryptocurrencies. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, pages 173–191. Springer, 2021.
- 27 Rosario Gennaro, Stanisław Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 295–310. Springer, 1999.
- 28 Go. The go programming language, 2009. URL: <https://go.dev>.
- 29 Harry A Kalodner, Miles Carlsten, Paul Ellenbogen, Joseph Bonneau, and Arvind Narayanan. An empirical study of namecoin and lessons for decentralized namespace design. In *WEIS*. Citeseer, 2015.
- 30 Mahimna Kelkar, Soubhik Deb, and Sreeram Kannan. Order-fair consensus in the permissionless setting. *Cryptology ePrint Archive*, 2021.
- 31 Olga Kharif. Cryptokitties mania overwhelms ethereum network’s processing, 2017. URL: <https://www.bloomberg.com/news/articles/2017-12-04/cryptokitties-quickly-becomes-most-widely-used-ethereum-app>.
- 32 Eleftherios Kokoris Kogias, Philipp Jovanovic, Nicolas Gailly, Ismail Khoffi, Linus Gasser, and Bryan Ford. Enhancing bitcoin security and performance with strong consistency via collective signing. In *25th usenix security symposium (usenix security 16)*, pages 279–296, 2016.
- 33 Eleftherios Kokoris-Kogias, Enis Ceyhun Alp, Linus Gasser, Philipp Jovanovic, Ewa Syta, and Bryan Ford. Calypso: Private data management for decentralized ledgers. *Cryptology ePrint Archive*, 2018.
- 34 Kraken. Cryptocurrency deposit processing times, 2022(?). Accessed: 2022-03-03. URL: <https://support.kraken.com/hc/en-us/articles/203325283->.
- 35 Klaus Kursawe. Wendy, the good little fairness widget: Achieving order fairness for blockchains. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 25–36, 2020.
- 36 Klaus Kursawe. Wendy grows up: More order fairness. In *International Conference on Financial Cryptography and Data Security*, pages 191–196. Springer, 2021.
- 37 <https://github.com/dedis/kyber>The Kyber Cryptography Library, 2010 – 2022.
- 38 Michael Lewis. *Flash boys: a Wall Street revolt*. WW Norton & Company, 2014.
- 39 LibSubmarine. Defeat front-running on ethereum, 2017(?). Accessed: 2022-01-24. URL: <https://libsubmarine.org>.
- 40 Dahlia Malkhi and Pawel Szalachowski. Maximal extractable value (mev) protection on a dag. *arXiv preprint arXiv:2208.00940*, 2022.
- 41 Mev-sgx: A sealed bid mev auction design, 2021. URL: <https://ethresear.ch/t/mev-sgx-a-sealed-bid-mev-auction-design/9677>.

3:22 Flash Freezing Flash Boys(F3B)

- 42 Edvardas Mikalaukas. 280 million stolen per month from crypto transactions, 2021. Accessed: 2022-02-16. URL: <https://cybernews.com/crypto/flash-boys-2-0-front-runners-draining-280-million-per-month-from-crypto-transactions>.
- 43 Peyman Momeni. Fairblock: Preventing blockchain front-running with minimal overheads. Master's thesis, University of Waterloo, 2022.
- 44 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260, 2008.
- 45 Kirill Nikitin, Ludovic Barman, Wouter Lueks, Matthew Underwood, Jean-Pierre Hubaux, and Bryan Ford. Reducing metadata leakage from encrypted files and communication with purbs. *Proceedings on Privacy Enhancing Technologies*, 2019(4):6–33, 2019.
- 46 Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2017.
- 47 Hany Ragab, Alyssa Milburn, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. Crosstalk: Speculative data leaks across cores are real. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1852–1867. IEEE, 2021.
- 48 Ronald L Rivest, Adi Shamir, and David A Wagner. Time-lock puzzles and timed-release crypto. Technical report, Massachusetts Institute of Technology. Laboratory for Computer Science, 1996.
- 49 Eli Ben Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE symposium on security and privacy*, pages 459–474. IEEE, 2014.
- 50 Alexander Savelyev. Contract law 2.0: ‘smart’ contracts as the beginning of the end of classic contract law. *Information & communications technology law*, 26(2):116–134, 2017.
- 51 Markus Schäffer, Monika di Angelo, and Gernot Salzer. Performance and scalability of private ethereum blockchains. In *International Conference on Business Process Management*. Springer, 2019.
- 52 Noah Schmid. Secure causal atomic broadcast, 2021. URL: <https://crypto.unibe.ch/archive/theses/2021.bsc.noah.schmid.pdf>.
- 53 Berry Schoenmakers. A simple publicly verifiable secret sharing scheme and its application to electronic voting. In *Annual International Cryptology Conference*, pages 148–164. Springer, 1999.
- 54 Secret markets: Front running prevention for automated market makers, 2020. URL: <https://scrt.network/blog/secret-markets-front-running-prevention>.
- 55 Adi Shamir. Identity-based cryptosystems and signature schemes. In *Workshop on the theory and application of cryptographic techniques*, pages 47–53. Springer, 1984.
- 56 Victor Shoup and Rosario Gennaro. Securing threshold cryptosystems against chosen ciphertext attack. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 1–16. Springer, 1998.
- 57 Combating front-running and malicious mev using threshold cryptography. URL: <https://blog.shutter.network>.
- 58 Shutterized beacon chain. URL: <https://ethresear.ch/t/shutterized-beacon-chain/12249>.
- 59 URL: <https://sikka.tech>.
- 60 Chrysoula Stathakopoulou, Signe Rüsçh, Marcus Brandenburger, and Marko Vukolić. Adding fairness to order: Preventing front-running attacks in bft protocols using tees. In *2021 40th International Symposium on Reliable Distributed Systems (SRDS)*, pages 34–45. IEEE, 2021.
- 61 Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F Wenisch, Yuval Yarom, and Raoul Strackx. Foreshadow: Extracting the keys to the intel {SGX} kingdom with transient {Out-of-Order} execution. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 991–1008, 2018.

- 62 Theodore M Wong, Chenxi Wang, and Jeannette M Wing. Verifiable secret redistribution for archive systems. In *First International IEEE Security in Storage Workshop, 2002. Proceedings.*, pages 94–105. IEEE, 2002.
- 63 Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.
- 64 Bin Cedric Xing, Mark Shanahan, and Rebekah Leslie-Hurd. Intel® software guard extensions (intel® SGX) software support for dynamic memory allocation inside an enclave. *Proceedings of the Hardware and Architectural Support for Security and Privacy 2016*, pages 1–9, 2016.
- 65 Haoqian Zhang, Louis-Henri Merino, Vero Estrada-Galinanes, and Bryan Ford. F3b: A low-latency commit-and-reveal architecture to mitigate blockchain front-running. *arXiv preprint arXiv:2205.08529*, 2022.
- 66 Jan Henrik Ziegeldorf, Roman Matzutt, Martin Henze, Fred Grossmann, and Klaus Wehrle. Secure and anonymous decentralized bitcoin mixing. *Future Generation Computer Systems*, 80:448–466, 2018.