

Eating Sandwiches: Modular and Lightweight Elimination of Transaction Reordering Attacks

Orestis Alpos   

Institute of Computer Science, University of Bern, Switzerland

Ignacio Amores-Sesar   

Institute of Computer Science, University of Bern, Switzerland

Christian Cachin   

Institute of Computer Science, University of Bern, Switzerland

Michelle Yeo   

IST Austria, Klosterneuburg, Austria

Abstract

Traditional blockchains grant the miner of a block full control not only over which transactions but also their order. This constitutes a major flaw discovered with the introduction of decentralized finance and allows miners to perform MEV attacks. In this paper, we address the issue of sandwich attacks by providing a construction that takes as input a blockchain protocol and outputs a new blockchain protocol with the same security but in which sandwich attacks are not profitable. Furthermore, our protocol is fully decentralized with no trusted third parties or heavy cryptography primitives and carries a linear increase in latency and minimum computation overhead.

2012 ACM Subject Classification Security and privacy → Distributed systems security

Keywords and phrases Consensus, MEV, Byzantine behavior, Rational behavior

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2023.12

Related Version *Full Version*: <https://arxiv.org/abs/2307.02954>

Funding This work has been funded by the Swiss National Science Foundation (SNSF) under grant agreement Nr. 200021_188443 (Advanced Consensus Protocols).

Acknowledgements We would like to thank Krzysztof Pietrzak and Jovana Mičić for useful discussions.

1 Introduction

The field of blockchain protocols has proved to be extremely robust. Since its creation with Bitcoin [45], it had gone through several enhancements such as Ethereum [2] and has seen the appearance of *decentralized finance* (DeFi). With this, some design flaws started to show up. Blockchains would ideally allow users to trade tokens with each other in a secure manner. However, existing designs do not consider users trading tokens of one platform for FIAT currency or tokens of a different platform, arguably one of the major flaws of today's blockchain platforms, *maximal extractable value* (MEV) [26]. Current estimates show that the total volume of MEV since 2020 is around 675M USD [3]. From a social welfare perspective, while MEV is profitable to miners, it presents a serious invisible tax on the users on the blockchain. Indeed the financial losses built up over time could potentially shy away users from the blockchain, and consequently impact the security of the chain.

Sandwich attacks are one of the most common types of MEV [38] accounting for a loss of 174M USD over the span of 33 months [48] for users of Ethereum. Sandwich attacks leverage the miner's ability to *select* and *position* transactions within a block. Consider the simple example of a sequence of transactions that swap one asset X for another asset Y in a



© Orestis Alpos, Ignacio Amores-Sesar, Christian Cachin, and Michelle Yeo; licensed under Creative Commons License CC-BY 4.0

27th International Conference on Principles of Distributed Systems (OPODIS 2023).

Editors: Alysson Bessani, Xavier Défago, Junya Nakamura, Koichi Wada, and Yukiko Yamauchi; Article No. 12; pp. 12:1–12:22



Leibniz International Proceedings in Informatics

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

decentralized exchange where exchange rates are computed automatically based on some function of the number of underlying assets in the pool (e.g., a constant product market maker [7]). Now suppose there is a miner that also wants to swap some units of X for Y . The most favorable position for the miner would be to place their transaction at the start of the sequence, so as to benefit from a lower X -to- Y exchange rate. This approach achieves a simple arbitrage strategy for any sequence of X -to- Y swaps: the miner can insert an X -to- Y exchange at the start of the sequence and use the computed exchange rate to sell, say, k units of X to get units of Y . The miner then front-runs the sequence of X -to- Y swaps, i.e., it inserts its own transaction at the start. To finish off the attack, the miner back-runs the sequence with another transaction of its own that swaps some units of Y to X , i.e., inserts this transaction at the end, and will often obtain more than k units of X . In this way, the miner profits from its insider knowledge and its power to order transactions. Refer to the full version [6] for a detailed description of exchange rate computation and sandwich attacks.

Since any miner of a given block has full control over the transactions added to the block, as well as over the way transactions are ordered, it is straightforward for the miner to launch the above attack. Consequently, this gives miners a lot of power as they control precisely the selection and positioning of transactions with every block they mine. This problem has received broad attention in the practice of DeFi and in the scientific literature.

A classic technique to mitigate this attack is thus to remove the control over the positioning of the transactions in the block from the adversary, whether by using a trusted third party to bundle and order the transactions as in flashbots¹, Eden², or OpenMEV³. Another method works by imposing a fair ordering of the transactions using a consensus algorithm that respects the order in which miners and validators first received the transactions [34]. The classical solutions make the protocol dependent of external factors, besides, affecting efficiency.

In this work, we introduce the *Partitioned and Permuted Protocol*, abbreviated Π^3 , an efficient decentralized algorithm that does not rely on external resources to counter front-running. It renders sandwich attacks unprofitable and can easily be implemented on top of an existing blockchain protocol Π .

Protocol Π^3 determines the final order of transactions in a block B_i , created by a miner M_i , through a *uniformly randomly chosen permutation* Σ_i . To explain the method, let us focus on three transactions in B_i , a victim transaction tx^* submitted by a client, and the front-running and back-running transactions, tx_1 and tx_2 , respectively, created by the miner. Since any relative ordering of these three transactions is equally probable, tx_1 will be ordered before tx_2 with the same probability as tx_2 before tx_1 , hence the miner will profit or make a loss with the same probability. Protocol Π^3 uses a fresh permutation for each block; it is chosen by a set of *leaders*, which are recent miners in the blockchain. We recognize and overcome the following challenges.

First, Σ_i must not be known before creating B_i , otherwise M_i would have the option to use Σ_i^{-1} , the inverse of Σ_i , to initially order the transactions in B_i , so that the final order is the one that benefits M_i . We overcome this by making Σ_i known only *after* M_i has been mined. On the other hand, if Σ_i is chosen after creating B_i , a coalition of leaders would be able to try multiple different permutations and choose the most profitable one – the number of permutations a party can try is only limited by their processing power. For these reasons, we have the leaders *commit* to their contributions to Σ_i before B_i becomes known, producing

¹ <https://www.flashbots.net>

² <https://www.edennetwork.io>

³ <https://openmev.xyz/>

unbiased randomness. To incentivize leaders to open their commitments, our protocol Π^3 employs a delayed reward release mechanism that only releases the payment to leaders when they have generated and opened all commitments.

In some cases, however, performing a sandwich attack might still be more profitable than the block reward, and hence a leader might still choose to not reveal their commitment to bias the resulting permutation. In general, a coalition of k leaders can choose among 2^k permutations out of the $n_t!$ possible ones, where n_t denotes the number of transactions in the block. It turns out that the probability that tx_1 , tx^* , and tx_2 appear in that order in one of the 2^k permutations can be significant for realistic values. Protocol Π^3 mitigates this by *dividing each transaction* into m chunks, which lowers the probability of a profitable permutation in two ways. First, the number of possible permutations is much larger, $(n_t m)!$ instead of $n_t!$. Second, a permutation is now profitable if the majority of chunks of tx_1 appear before the chunks of tx^* , and vice versa for the chunks of tx_2 . As we discuss, the probability of a profitable permutation approaches zero rapidly as the number of chunks m increases. We discuss how to implement the chunking mechanism while preserving transaction integrity and atomicity.

Organization. In this paper, we introduce a construction that takes as input a blockchain protocol Π and produces a new blockchain protocol Π^3 in which sandwich attacks are no longer profitable. We begin by revisiting the concept of *atomic broadcast* [17] and setting the model for the analysis. Secondly, we introduce our construction justifying how miners are incentivized to follow the protocol before moving on to analyzing the construction in detail. Thirdly, we guarantee that the construction does not include any vulnerability to the protocol by showing that Π^3 implements a variant of atomic broadcast if Π does. This part of the analysis is performed in the traditional *Byzantine* model. Fourthly, we consider the *rational* model to show that sandwich attacks are no longer profitable in Π^3 . We consider the dual model of *Byzantine* for the security analysis and *rational* for the analysis of the sandwich attacks because we considered it to be a perfect fit to show that the security of Π^3 is not weakened even against an adversary that obtains nothing for breaking the protocol, as well as, we can assume that any party attempts to extract value from any sandwich attack. In other words, we consider both the security analysis and the analysis of the sandwich attack in the worst scenario possible for the protocol. Lastly, we conclude the paper with an empirical analysis of the protocol under real-life data, as well as an analysis of the additional overhead introduced by our protocol.

2 Related work

The idea to randomize the transaction order within a block is folklore in the blockchain space. It has been explored by Yanai [52] and also implemented in the wild [4]. To the best of our knowledge, we are the first to implement the randomization using on-chain randomness and to provide a security analysis for this model. Additionally, Randomspam [4] also acknowledges that some spamming attacks can occur with randomized transactions, where the attacker aims to insert several low-cost transactions to maximize the probability that some of these transactions are positioned exactly at a profitable transaction. Our work reduces the success probability of these attacks by first chunking each transaction into smaller parts and then permuting all chunked transactions, rendering exact positioning attacks less profitable unless more transactions are added, incurring larger gas costs.

12:4 Eating Sandwiches

A recent line of work [34, 39, 32, 19, 33] formalizes the notion of *fair ordering* of transactions. These protocols ensure, at consensus level, that the final order is consistent with the local order in which transactions are observed by parties. Similarly, the Hashgraph [9] consensus algorithm aims to achieve fairness by having each party locally build a graph with the received transactions. As observed by Kelkar *et al.* [34], a transaction order consistent with the order observed locally for any pair of transactions is not always possible, as Condorcet cycles may be formed. As a result, fair-ordering protocols output a transaction order that is consistent with the view of only some *fraction* of the parties, while some transactions may be output in a batch, i.e., with no order defined among them. Moreover, although order-fairness removes the miner’s control over the order of transactions, it does not eliminate front-running and MEV-attacks: a *rushing* adversary that becomes aware of some tx early enough can broadcast its own tx' and make sure that sufficiently many nodes receive tx' before tx .

Another common defense against front-running attacks is the *commit and reveal* technique. The idea is to have a user first commit to a transaction, e.g., by announcing its hash or its encryption, and, once the order is fixed, reveal the actual transaction. However, an adversary can choose not to reveal the transaction, should the final order be non-optimal. Doweck and Eyal [29] employ time-lock puzzle commitments [50], so that a transaction can be brute-force revealed, and protocols such as Unicorn [40] and Bicorn [22] employ verifiable delay functions [15] to mitigate front-running. Whereas they indeed manage to mitigate front running, the main disadvantages of these solutions are threefold: firstly, transactions may be executed much later than submitted, with no concrete upper bound on the revelation time. Secondly, a delay for the time-lock puzzle has to be chosen which matches the network delay and adversary’s computational power. Finally, it is unclear who should spend the computational power to solve the time-lock puzzles, especially in proof of work blockchains where this shifts computational power away from mining.

A different line of work [30, 43, 18, 49, 53] hides the transactions until they are ordered with the help of a *committee*. For instance, transactions may be encrypted with the public key of the committee, so that its members can collaboratively decrypt it. However, this method uses threshold encryption [28] and requires a coordinated setup. Also multi-party computation (MPC) has been used [12, 5, 41] to prevent front-running. MPC protocols used in this setting must be tailor-made so that misbehaving is identified and punished [11, 36]. A disadvantage of the aforementioned techniques is that the validity of a transaction can only be checked after it is revealed. These techniques also rely on strong cryptographic assumptions and coordination within the committee. The protocol presented in this work disincentivizes sandwich attacks without requiring hidden transactions or employing computationally heavy cryptography.

Another widely deployed solution against front-running involves a dedicated *trusted third party*. Flashbots⁴, Eden⁵, and OpenMEV⁶ allow Ethereum users to submit transactions to their services, then order received transactions, and forward them to Ethereum miners. Chainlink’s Fair Sequencing Service [20], in a similar fashion, aims to collect encrypted transactions from users, totally orders them, and then decrypts them. The third-party service may again be run in a distributed way. The drawback with these solutions is that attacks are not eliminated, but trust is delegated to a different set of parties.

⁴ www.flashbots.net

⁵ www.edennetwork.io

⁶ <https://openmev.xyz/>

An orthogonal but complementary line of research is taken by Heimbach and Wattenhofer [31]. Instead of eliminating sandwich attacks, they aim to improve the resilience of ordinary transactions against sandwich attacks by strategically setting their slippage tolerance to reduce the risk of both transaction failure as well as sandwich attacks.

Last but not least, Baum *et al.* [10] and Heimbach and Wattenhofer [31] survey the area of front-running attacks.

3 Model

Notation. For a set X , we denote the set of probability distributions on X by $\mu(X)$. For a probability distribution $\nu \in \mu(X)$, we denote sampling x from X according to ν by $x \leftarrow \nu$.

3.1 Block-based atomic broadcast

Parties broadcast transactions and deliver blocks using the events $\text{bab-broadcast}(tx)$ and $\text{bab-deliver}(b)$, respectively, where block b contains a sequence of transactions $[tx_1, \dots, tx_{n_t}]$. The protocol outputs an additional event $\text{bab-mined}(b, P)$, which signals that block b has been *mined* by party P , where P is defined as the *miner* of b . Notice that $\text{bab-mined}(b, P)$ signals only the creation of a block and not its delivery. In addition to predicate $\text{VT}()$, we also equip our protocol with a predicate $\text{VB}()$ to determine the validity of a block. Moreover, we define a function $\text{FB}()$, which describes how to fill a block: it gets as input a sequence of transactions and any other data required by the protocol and outputs a block. These predicates and function are determined by the higher-level application or protocol.

► **Definition 1.** *A protocol implements block-based atomic broadcast with validity predicates $\text{VT}()$ and $\text{VB}()$ and block-creation function $\text{FB}()$ if it satisfies the following properties, except with negligible probability:*

Validity: *If a correct party invokes a $\text{bab-broadcast}(tx)$, then every correct party eventually outputs $\text{bab-deliver}(b)$, for some block b that contains tx .*

No duplication: *No correct party outputs $\text{bab-deliver}(b)$ for a block b more than once.*

Integrity: *If a correct party outputs $\text{bab-deliver}(b)$, then it has previously output the event $\text{bab-mined}(b, \cdot)$ exactly once.*

Agreement: *If some correct party outputs $\text{bab-deliver}(b)$, then eventually every correct party outputs $\text{bab-deliver}(b)$.*

Total order: *Let b and b' be blocks, and P_i and P_j correct parties that output $\text{bab-deliver}(b)$ and $\text{bab-deliver}(b')$. If P_i delivers b before b' , then P_j also delivers b before b' .*

External validity: *If a correct party outputs $\text{bab-deliver}(b)$, such that $b = [tx_1, \dots, tx_{n_t}]$, then $\text{VB}(b) = \text{TRUE}$ and $\text{VT}(tx_i) = \text{TRUE}$, for $i \in 1, \dots, n_t$. Moreover, if $\text{FB}(tx_1, \dots, tx_{n_t})$ returns b , then $\text{VB}(b) = \text{TRUE}$.*

Fairness: *There exists $C \in \mathbb{N}$ and $\mu \in \mathbb{R}_{>0}$, such that for all $N \geq C$ consecutive delivered blocks, the fraction of the blocks whose miner is correct is at least μ .*

Observe that the properties assure that $\text{bab-mined}(b, P)$ is triggered exactly once for each block b , hence each block has a unique miner. For ease of notation, we define on a block b the fields $b.\text{txs}$, which contains its transactions, and $b.\text{miner}$, which contains its miner. Since blocks are delivered in total order, we can assign them a *height*, a sequence number in their order of delivery, accessible by $b.\text{height}$. Finally, for simplicity we assume that a delivered block allows access to all blocks with smaller height, through an array $b.\text{chain}$. That is, if $b.\text{height} = i$ then $b.\text{chain}[i']$ returns b' , such that $b'.\text{height} = i'$, for all $i' \leq i$.

3.2 Blockchain and network

Blockchain protocols derive their security from different techniques such as *proof of work* (*PoW*) [45], *proof of stake* (*PoS*) [27], *proof of space-time* (*PoST*) [24], or *proof of elapsed time* (*PoET*) [16]. In the remainder of the work, we consider a generic protocol Π that has a probabilistic termination condition, capturing all the model above. Furthermore, we model Π as block-based atomic broadcast.

Parties. Similar to previous works, our protocol does not make explicit use of the number of parties or their identities, and does not require the parties themselves to know this number. We assume an static network of n_p parties. We consider the Byzantine model, where f parties may behave arbitrarily, as well as, the rational model where all parties behave maximizing their utilities.

Transactions & Blocks. A transaction tx contains a set of *inputs*, a set of *outputs*, and a number of digital signatures. Transactions are batched into *blocks*. A block contains a number of transactions, n_t , for simplicity we assume n_t to be constant. A block b may contain parameters specific to protocol Π such as references to previous blocks, but we abstract the logic of accessing them in a field $b.chain$, as explained in the context of Definition 1. We allow conditional execution of transactions across blocks, i.e., a transaction can be executed conditioned on the existence of another transaction in a previous block.

Network. A *diffusion functionality* implements communication among the parties, which is structured into *synchronous* rounds. The functionality keeps a $RECEIVE_i$ string for each party P_i and makes it available to P_i at the start of every round. String $RECEIVE_i$ is used to store all messages P_i receives. When a party P_i instructs the diffusion functionality to *broadcast* a message, we say that P_i has *finished its round* and the functionality tags P_i as finished for this round. The adversary, detailed in Section 5, is allowed to read the string of any party at any moment during the execution and to see any messages broadcast by any party immediately. Furthermore, the adversary can write messages directly and selectively into $RECEIVE_i$ for any P_i , so that only P_i receives the message at the beginning of the next round. This models a *rushing* adversary.

When all non-corrupted parties have finished their round, the diffusion functionality takes all messages that were broadcast by non-corrupted parties in the round and adds them to $RECEIVE_i$ for all parties, this is the reason of the name *synchronous* rounds. Every non-corrupted party communicates changes to its local view at the end of each round. If a non-corrupted party creates a block in round r , the new block is received by all parties by round $r + 1$. Furthermore, even if the adversary causes a block to be received selectively by only some non-corrupted parties in round r , the block is received by all non-corrupted parties by round $r + 2$. The update of the local view also includes the delivery of transactions contained in the blocks that satisfy the conditions to be accepted.

4 Protocol

Our proposed protocol Π^3 (“Partitioned and Permuted Protocol”) contains two modifications to a given underlying blockchain protocol Π in order to prevent sandwich MEV attacks.

Our first modification involves randomly permuting the transactions in any given block. Note that a naive way of doing so is to use an external oracle (e.g., DRAND [1] or NIST beacon [35]) to generate the randomness which will be applied to a given block. However,

using an external source of randomness relies on strong trust assumptions on the owners of the source, leaves our protocol vulnerable to a single point of failure and it introduces incentive issues between miners of the chain and owners of the external source. To avoid this, our protocol uses miners of the immediately preceding blocks to generate the randomness. These miners, which we refer to as *leaders* for the given block, are in charge of generating random *partial seeds*. These partial seeds are then combined to form a *seed* which will be the input into a PRG to produce a random permutation that is applied to the transactions in the block. To ensure that the permutation is random, we need first to achieve that leaders participate in the generation of random partial seeds and secondly to ensure the partial seeds generated by the leaders are random. That is, the leaders should not commit to the same partial seed each time or collude with other leaders to generate biased partial seeds. To incentivize each leader to participate in the generation of the seed, Π^3 stipulates that they commit to their partial seed and present a valid opening during the commitment opening period, otherwise their reward will be burned. In typical blockchain protocols, the miner of a block receives the block reward immediately. In Π^3 , the miner does not receive the reward until a certain number of additional blocks has been mined. We refer to this as a *waiting phase* and stress that the precise length of the waiting phase is a parameter in our protocol that can be tweaked.

Our second modification is to divide the transfers of each transaction into smaller chunks before permuting the chunked transactions of a block. This modification increases the cardinality of the permutation group in order to reduce the effectiveness of any attack aiming to selectively open partial seeds in order to bias the final permutation.

We stress that our proposed modifications incur minimal computational overhead, since the only possible overhead corresponds to transaction delivery and this aspect is computationally cheap. Thus, the only noticeable impact of our protocol is latency. In the full version [6] we provide an in-depth analysis of the efficiency impact of our proposed modifications.

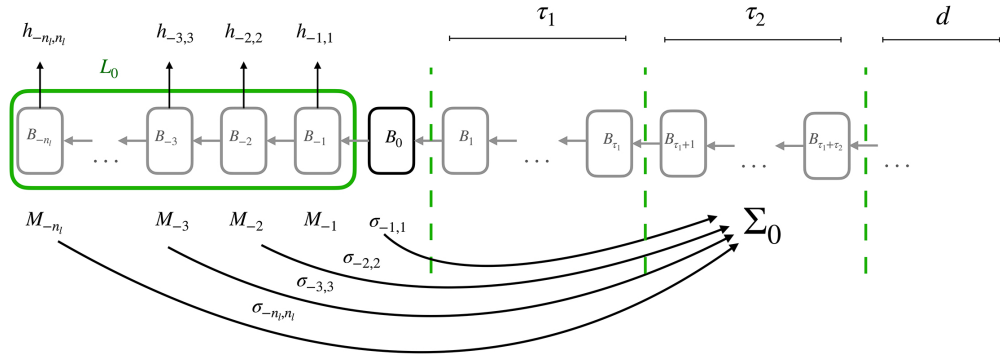
4.1 Permuting transactions

Protocol Π^3 consists of the following four components (see Figure 1): block mining, generation of the random permutation, reward (re)-distribution, and chunking the transactions.

Appending the partial seeds. Let n_ℓ be the size of the leader set for each block. The miner M_i of block B_i is part of the leader set of blocks B_{i+j} , for $j \in [n_\ell]$. M_i must therefore contribute a partial seed $\sigma_{i,j}$ for each of these n_ℓ blocks following B_i . Hence, M_i needs to create n_ℓ random seeds $\sigma_{i,1}, \dots, \sigma_{i,n_\ell}$ and commitments to them, $C(\sigma_{i,1}), \dots, C(\sigma_{i,n_\ell})$. The commitments $C(\sigma_{i,j})$, for $j \in [n_\ell]$, are appended to block B_i , while the seeds $\sigma_{i,j}$ are stored locally by M_i . A block that does not contain n_ℓ commitments is considered invalid.

Looking ahead, we want that any party knowing the committed value can demonstrate it to any other party. Thus, the more standard commitments schemes such as Pedersen commitment [47] are ill-suited. Instead, Π^3 uses a deterministic commitment scheme for committing to permutations, in particular, a collision-resistant cryptographic hash function. When the entropy of the committed values is high enough, then a hash function constitutes a secure commitment scheme. Since the parties commit to a random partial seed, hash functions suffice and yield a cheap commitment scheme.

Opening the commitments. Let $\tau_1, \tau_2 \in \mathbb{N}_{>0}$. Between τ_1 and $\tau_1 + \tau_2$ blocks after the creation of some block B_i , the commitments of the partial permutation to be applied on block B_i must be opened. The miners of these blocks also need to append the openings




■ **Figure 1** We illustrate the routine for the creation of the random permutation σ_0 created by the leader set L_0 and to be applied on block B_0 in Bitcoin. The leaders L_0 for block B_0 is formed by the miners of the n_ℓ blocks before B_0 , marked with the green box. When party M_i mined block B_i , the party generated some random seed $\sigma_{-i,i}$ and included its commitment $h_{-i,i}$ as part of the newly mined block. After block B_0 is mined, the leaders wait for τ_1 blocks before opening the commitments. The commitments must be included in the following τ_2 blocks. Finally the parties wait until every block containing openings are confirmed before delivering block B_0 .

to their blocks, unless a previous block in the chain already contains them (see below for more details). The parameter τ_1 controls the probability of rewriting block B_i after the commitments have been opened. Whereas, parameter τ_2 guarantees that there is enough time for all the honest commitments to be opened and added to some block. Any opening appended a block B_j for $j > i + \tau_1 + \tau_2$ is ignored. We note that specific values of τ_1 and τ_2 might cause our protocol to suffer an increase in latency. We leave these parameters to be specified by the users of our protocol. For the interested reader, we discuss latency-security trade-offs in the full version [6]. The τ_1 blocks created until opening the commitments takes place is known as *silent phase*, whereas the following τ_2 blocks is known as *loud phase*.

A possible way to record the opening of commitments is for the miners that own the commitments to deploy a smart contract that provides a method $open(i, j, \sigma_{i,j})$, where $\sigma_{i,j}$ is a (claimed) opening of the j -th commitment $h_{i,j}$ published in the i -th block B_i . We remark, that the smart contract serves only as proof that an opening to a commitment has been provided, and does not add any functionality to the protocol, so other proof mechanisms can also be considered. The protocol monitors the blockchain for calls to this method. The arguments to each call, as well as the calling party and the block it appears on, are used to determine the final permutations of the blocks and the distribution of the rewards, which we will detail below. We stress that *not* opening a commitment does not impact the progress of protocol, as unopened commitments are ignored.

Deriving the permutation from partial seeds. Let the seed σ_i for block B_i be defined as $\sigma_{i-1,1} \oplus \sigma_{i-2,2} \oplus \dots \oplus \sigma_{i-n_\ell, n_\ell}$. Given the seed σ_i , let $r_i := G(\sigma_i)$, where $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\ell$ is a pseudorandom generator. If at least one of the partial seeds $\sigma_{i,j}$, for $j \in [n_\ell]$, is chosen at random, then σ_i is random as well, and r_i is indistinguishable from a random number [14] without the knowledge of $\sigma_{i,j}$. There are standard algorithms to produce a random permutation from a polynomial number of bits [25].

Incentivizing the behavior. A crucial factor in the security of Π^3 against sandwich MEV attacks is that the permutation used to order transactions within a block should be truly random. Thus, the miners should generate all partial seeds uniformly at random. To



■ **Figure 2** We illustrate the process followed by party M to deliver a block B . We denote by tx_1, \dots, tx_n the n transaction that constitute B . In the first step, party M breaks each transaction tx_i into m transactions tx_i^1, \dots, tx_i^m involving a smaller amount. These smaller transactions are later permuted according to the random permutation Σ . Lastly, party M delivers these small transactions in this new order.

incentivize them to do so, we exploit the fact that all leaders remain in the waiting phase for a period of time, which means that they have not yet received the block rewards and fees for mining their block on the blockchain. Note that the waiting phase is $n_\ell + \tau_1 + \tau_2 + d$ blocks long. This implies that their rewards can be claimed by other miners or burned if a party diverges from the proper execution, according to the rules described below. Consider a partial permutation $\sigma_{i,j}$ committed by miner M_i of block B_i . Recall that $\sigma_{i,j}$ will be applied on block B_{i+j} and that miners can be uniquely identified due to the *bab-mined*() event.

1. Before τ_1 blocks have been appended after block B_{i+j} any other leader of the leader set \mathcal{L}_{i+j} who can append a pre-image of $h_{i,j}$ to the chain can receive the reward and fees corresponding to M_i . This mechanism prevents party M_i from disclosing its commitment before every other leader committed its randomness, thus preventing colluding. A miner whose commitment has been discovered by another leader is excluded from all the leader sets.
2. If the opening of $\sigma_{i,j}$ is not appended to any block, miner M_i loses its reward and fees. This mechanism prevents miners from not opening their commitments. Note that miners are incentivized to include all the valid openings, as discussed below.
3. If any of the previous conditions do not apply, party M_i receives an α fraction of the block reward and fee for $\alpha \in (0, 1)$, which would be paid out the moment M_i leaves the waiting phase. Each miner that appends the opening of M_i 's commitments gets $\frac{(1-\alpha) \cdot w}{n_\ell}$ for each commitment appended.

In the remainder of this work we refer to the block reward and fees as simply *block reward*.

4.2 Chunking transactions

In all commit-and-open schemes, there exists the vulnerability that malicious parties decide to not open their commitments so as to bias the outcome. In our protocol, any coalition of k leaders can choose between 2^k ways to bias the final permutation. If one miner manages to create multiple blocks out of $B_{i-1}, \dots, B_{i-n_\ell}$, this does not even require collusion with others.

The adaptive attacks mounted through withholding can be countered with a simultaneous broadcast abstraction [23], but realizing this is almost impossible in practice [37], especially in the blockchain domain. Alternatively, time-lock puzzles may negate the effect of the

12:10 Eating Sandwiches

delay. But this technique costs computational effort, which may have a negative impact on the environment and possibly also on the protocol. In the particular case of generating a permutation, there is an alternative.

Chunking transfers in transactions. Let us assume that a block contains n_{tx} transactions, this means there exist $n_{tx}!$ possible permutations of them. A coalition of k leaders can choose between 2^k possible permutations among the $n_{tx}!$ total permutations. Furthermore, in the simplest case the coalition only aims to order the three transactions that constitute the sandwich attack, thus the fraction of advantageous permutations is $\frac{1}{6}$, the fraction of disadvantageous permutations is $\frac{1}{6}$ and the remaining ones are neutral. If k is big enough, the coalition could still extract enough value to compensate for the lost block rewards of those parties that do not open their commitment.

Therefore we want to increase the size of the permuted space. We assume here that every transaction consists of arbitrary code and a few specialized instructions that are *transfers* of coins or tokens. These may be the *native payment operation* of the blockchain or operations that involve a well-known *standard format for tokens*, which are emulated by smart contracts (such as the ERC-20 standard in Ethereum). We now divide every transfer generated by some transaction into m chunks.

For instance, suppose transaction tx_i consists of Alice paying Bob 1 ETH. In our protocol, each party would locally split tx_i into m chunks tx_i^1, \dots, tx_i^m , consisting of Alice transferring $1/m$ ETH to Bob each. After all transactions are chunked, the permutation will be applied to the larger set of transactions. There exist $(n_{tx}m)!$ permutations and the coalition would need to order the $3m$ chunks that constitute the involved transactions. Furthermore, for a given permutation with some chunks ordered beneficially, there will exist chunks ordered in a disadvantageous way, with overwhelming probability. The coalition needs to optimize the good ordering of some chunks while keeping the bad ordering under control. Obtaining a favorable ordering becomes extremely unlikely as number of chunks m grows (Section 5.2).

Execution of transactions and chunks. Transactions contain arbitrary code whose execution produces an ordered list of transfers, as introduced before. The process of chunking proceeds in two stages.

In the first stage, the party executes the code of all transactions contained in the block serially, in the order determined by the miner; this produces a list of transfers and the corresponding amounts for each transaction. Some transactions may turn out to be invalid, they are removed from the further processing of the block.

In the second stage, for each valid transaction tx a list of m transfer chunks tx^1, \dots, tx^m is produced such that tx^1 contains the *code* executed by tx , and its transfers have their amount set to $1/m$ of the number computed by the code. The transaction chunks tx^2, \dots, tx^m contain *only* the transfers, with their numbers set to $1/m$ of original amounts, but these transactions do not execute further code. If the code executed by tx^1 produces different transfers than in the first stage, the execution of tx^1 is aborted, also the execution of tx^2, \dots, tx^m . We consider two transfers to be the same if they transfer the same amount of coins or tokens from the same source address to the same destination address. Note the blockchain state produced by a transaction in the first stage can differ from the state produced by the same transaction in the actual execution according to the second stage execute; this is the case, for instance, when it interacts with a smart contract.

The block being permuted now contains up to m times as many chunks as the original block contained transactions, each of them transferring $1/m$ the value.

Notice that the permutation is not uniformly random across all choices, but needs to respect that tx^1 appears first within the set of chunks resulting from tx . However, this restriction in the permutation does not constitute any loss of generality since every chunk performs an identical transfer. An adversarial miner can utilize fine-grained conditions such as slippage to additionally control the conditional execution of transactions – and in our case transaction chunks – in a given block. The execution of transactions explained above guarantees atomicity: all chunks are executed or no chunk is executed. In the full version [6] we present an in-depth analysis of how slippage could lead to higher expected revenue, which may also be of independent interest.

4.3 Details

In Algorithm 1 we show the pseudocode for protocol Π^3 , which implements a block-based atomic broadcast (*bab*) primitive. The pseudocode assumes an underlying protocol Π , which is also modeled as a block-based atomic broadcast (*bab*) primitive, as defined in Section 3. The user or high-level application interacts with Π^3 by invoking Π^3 -*broadcast*(tx) events. These are handled by invoking the corresponding Π -*broadcast*(tx) event on the underlying protocol Π (L3-4).

Protocol Π outputs an event *bab-mined*(b, Q) whenever some party Q mines a new block b (L5). For Π^3 , the mining of a new block at height i starts the opening phase for the block at height $i_{\text{open}} = i - \tau_1 - 1$ (L6). Hence, party P loops through the n_ℓ blocks before i_{open} and checks whether it is the miner of each of them (L7-9). If this is the case, P must provide a valid opening to the commitment related to block at height i_{open} . The opening is achieved by a specific type of transaction, for example through a call to a smart contract. In the pseudocode we abstract this into a function *Open*().

Protocol Π outputs an event Π -*deliver*(b) whenever a block b is delivered (L10). According to the analysis of our protocol, this will allow Π^3 to deliver the block $\tau_1 + \tau_2$ positions higher than b , i.e., the block b_{del} at height $i_{\text{del}} = b.\text{height} - \tau_1 - \tau_2$. To this goal, Π^3 first reads the commitments related to b_{del} (L12-13). By construction of Π^3 , a commitment $c_{i,j}$, written on block b_i , is used to order the transactions in block b_{i+j} . Hence, the commitments related to b_{del} have been written on the n_ℓ blocks before b_{del} . Protocol Π^3 then reads the openings to these commitments (L14-17). Again by construction of Π^3 , the openings of the commitments related to b_{del} have been written on the blocks with height $i_{\text{del}} + \tau_1 + 1$ to $i_{\text{del}} + \tau_1 + \tau_2$. For each of these blocks, Π^3 loops through its transactions that contain an opening. L16 then checks whether the opening is for a commitment related to block b_{del} and whether the opening is valid. Protocol Π^3 then calculates the final permutation Σ to be applied to block b (L18-22). As presented in Section 4.1, $\Sigma = \text{PermFromRandBits}(G(\text{seed}))$, where *seed* is the XOR of all valid openings for block b , G is a pseudorandom generator, and *PermFromRandBits* an algorithm that derives a permutation from random bits. The remaining of this block chunks the transactions contained in b (L23-25), applies Σ on the chunked transactions (L26), and swaps the first permuted chunk of each of each transaction with the chunk containing the code (L28). The function *Chunk*() is explained in Section 4.2. Finally, Π^3 delivers block b containing the chunked and permuted transactions through the Π^3 -*deliver*(b) event (L30).

The function *FB*() is an upcall from block-based atomic broadcast. It specifies how a block is filled with transactions and additional data. For simplicity, the pseudocode omits any detail specific to *bab*. It first writes all given transactions on the block, then picks uniformly at random n_ℓ bit-strings of length λ . These are the partial random seeds to be used in the permutation of the following n_ℓ blocks, if the block that is currently being built gets mined and delivered by *bab*. The commitments to these partial seeds are appended on the block.

12:12 Eating Sandwiches

■ **Algorithm 1** Protocol Π^3 . Code for party P .

Implements: Protocol Π^3
Uses: block-based atomic broadcast Π

State:

```

1:    $\sigma[i, j] \leftarrow \perp$ , for all  $i \geq 1, j \in [n_\ell]$ 
2:    $c[i, j] \leftarrow \perp$ , for all  $i \geq 1, j \in [n_\ell]$ 

3: upon event  $\langle \Pi^3\text{-broadcast}, tx \rangle$  do
4:   invoke  $\langle \Pi\text{-broadcast}, tx \rangle$ 

5: upon event  $\langle \Pi\text{-mined}, b, Q \rangle$  do
6:    $i_{\text{open}} \leftarrow b.\text{height} - \tau_1 - 1$ 
7:   for  $i' \in [i_{\text{open}} - n_\ell - 1, i_{\text{open}} - 1]$  do
8:     if  $b.\text{chain}[i'].\text{miner} = P$  then
9:        $\text{Open}(b.\text{chain}[i'].\text{commitments}[i_{\text{open}} - i'])$ 

10: upon event  $\langle \Pi\text{-deliver}, b \rangle$  do
11:    $i_{\text{del}} \leftarrow b.\text{height} - \tau_1 - \tau_2$ 
12:   for  $j \in [n_\ell]$  do                                     // Read commitments for block  $b$ 
13:      $c[i_{\text{del}}, j] \leftarrow b.\text{chain}[i_{\text{del}} - j].\text{commitments}[j]$ 
14:   for  $i' \in [i_{\text{del}} + \tau_1 + 1, i_{\text{del}} + \tau_1 + \tau_2]$  do   // Read the openings for block  $b$ 
15:     for  $tx \in b.\text{chain}[i'].\text{txs}$  such that  $tx = \text{open}(k, l, \sigma)$  do
16:       if  $k + l = i_{\text{del}}$  and  $H(\sigma) = c[i_{\text{del}}, l]$  then
17:          $\sigma[i_{\text{del}}, l] \leftarrow \sigma$ 
18:    $\text{seed} \leftarrow 0^\lambda$ 
19:   for  $j \in [n_\ell]$  do                                     // Compute final permutation for block  $b$ 
20:     if  $\sigma[i_{\text{del}}, j] \neq \perp$  then
21:        $\text{seed} \leftarrow \text{seed} \oplus \sigma[i_{\text{del}}, j]$ 
22:    $\Sigma \leftarrow \text{PermFromRandBits}(G(\text{seed}))$ 
23:    $\text{chunked\_txs} \leftarrow []$ 
24:   for  $tx \in b.\text{txs}$  do                                   // Chunk and permute transactions in block  $b$ 
25:      $\text{chunked\_txs} \leftarrow \text{chunked\_txs} \parallel \text{Chunk}(tx, m)$ 
27:    $\text{chunks} \leftarrow \text{Permute}(\Sigma, \text{chunked\_txs})$ 
28:    $\text{chunks} \leftarrow \text{SwapChunks}(\text{chunks})$  // For each  $tx$ , swap the first chunk in  $\text{chunks}$  with  $tx^1$ 
29:    $b.\text{txs} \leftarrow \text{chunks}$ 
30:   invoke  $\langle \Pi^3\text{-deliver}, b \rangle$ 

31: function  $\text{FB}(txs)$  :
32:    $\text{data} \leftarrow []$ 
33:   for  $tx \in txs$  do
34:      $\text{data} \leftarrow \text{data} \parallel tx$ 
35:   for  $j \in [n_\ell]$  do
36:      $\sigma \xleftarrow{\$} \{0, 1\}^\lambda$ 
37:      $c \leftarrow H(\sigma)$ 
38:      $\text{data} \leftarrow \text{data} \parallel c$ 
39:   return  $\text{data}$ 

40: function  $\text{VB}(b)$  :
41:   if  $(\exists tx \in b.\text{txs} : \neg \text{VT}(tx)) \vee (\exists j \in [n_\ell] : b.\text{commitments}[j] = \perp)$  then
42:     return FALSE
43:   return TRUE

```

Finally, the predicate $VB()$ specifies that a block is valid if all its transactions are valid, as specified by $VT()$, and if it contains n_ℓ commitments. The predicate $VT()$ is omitted, as its implementation does not affect Π^3 .

5 Analysis

5.1 Security analysis

We model the adversary as an interactive Turing machine (ITM) that corrupts up to t parties at the beginning of the execution. Corrupted parties follow the instructions of the adversary and may diverge arbitrarily from the execution of the protocol. The adversary also has control over the *diffusion functionality*. That is, she can schedule the delivery of messages (within the Δ rounds), as well as read the $RECEIVE_i$ of every party at any moment of the execution and directly write in the $RECEIVE_i$ of any party.

We first show that the security of our construction is derived from the security of the original protocol. Given an execution of protocol Π^3 , we define the equivalent execution in protocol Π as the execution in which every party follows the same steps but the commitment, opening, and randomization of transactions are omitted. We also recall the parameters τ_1 and τ_2 that denote the length (in blocks) of the silent and loud phase respectively.

► **Lemma 2.** *The probability that an adversary can rewrite a block after any honest partial permutations have been opened is negligible in τ_1 .*

Proof. Assume an adversary controlling up to t parties and a block B . We know that if $\tau_1 > d$, protocol Π would deliver block B , thus an adversary cannot revert the chain to modify the order of the transactions stored in B but with negligible probability. ◀

► **Lemma 3.** *The probability that an adversary can rewrite a chain omitting the opening of some honest partial permutation is negligible in τ_2 .*

Proof. The fairness quality of protocol Π states that for any consecutive N blocks, if $N \geq N_0$ the fraction of honest blocks is at least μ . Thus, if $\tau_2 \geq \max\{N_0, \frac{1}{\mu}\}$, there exists at least one honest block containing every opening that is not previously included in the chain. Since ◀

Our construction aims to turn any protocol into a protocol robust against sandwich attacks. However, there might be new vulnerabilities. Intuitively, our construction should not introduce any vulnerability because the only modified aspect is the order in which transactions are delivered. Theorem 5 formalizes this intuition.

► **Remark 4.** Note that every Π -delivered block is also Π^3 -delivered some block after (Line10–30). Note also that every Π^3 -delivered is also Π -delivered. Furthermore, the blocks are delivered in the same order.

► **Theorem 5.** *If protocol Π implements block-based atomic broadcast, then the Partitioned and Permuted Protocol Π^3 implements block-based atomic broadcast.*

Proof. According to Remark 4, the set of Π^3 -delivered blocks is the same as the set of Π -delivered blocks.

Validity. Assume that an honest party Π^3 -broadcasts(tx) transaction tx . The party first Π -broadcasts(tx) (L3–4). The validity property of protocol Π guarantees that eventually a block b containing transaction tx is Π -delivered. According to Remark 4, the honest party eventually Π^3 -delivers a block containing tx and Π^3 satisfies the validity property of block-based atomic broadcast.

No-duplication. Note that Π^3 delivers the same set of blocks as protocol Π , Remark 4. Thus, the no-duplication property of protocol Π^3 is inherited directly from the no-duplication of protocol Π .

Agreement. Consider two honest parties P and Q such that party P Π^3 -delivers block b . Remark 4 guarantees that P also Π -delivers block b . The agreement property of protocol Π ensure that Q eventually Π -delivers block b . Remark 4 guarantees that Q eventually Π^3 -delivers block b . Note that the block b delivered by both P and Q may differ in how the transactions are chunked and permuted. However, Lemmas 2 and 3 guarantee all correct parties agree on the same permutation with all but negligible probability. Hence, we conclude that protocol Π^3 satisfies the agreement property.

Total order. Remark 4 guarantees that the order in which any honest party Π^3 -delivers two block b_1 and b_2 is the same as it Π -delivers them. Thus, the total order property of protocol Π guarantees the total order property of protocol Π^3 .

External validity. This follows from the external validity of Π .

Fairness. According to Remark 4 the same blocks and in the same order are both Π^3 -delivered and Π -delivered. Hence, the fairness property of Π^3 is inherited from the fairness property of protocol Π . \blacktriangleleft

After showing that Π^3 is as secure as the original protocol Π . We turn our attention to analyzing the behavior of Π^3 under sandwich attacks, in the upcoming section.

5.2 Game-theoretic analysis

Here, we aim to show that if we assume all miners are rational, i.e., they prioritize maximizing their own payoff, behaving honestly as according to our protocol Π^3 is a stable strategy.

Strategic games. For $N \in \mathbb{N}$, let $\Gamma = (N, (S_i), (u_i))$ be an N party game where S_i is a finite set of strategies for each party $i \in [N]$. Let $S := S_1 \times \dots \times S_N$ denote the set of outcomes of the game. The utility function of each party i , $u_i : S \rightarrow \mathbb{R}$, gives the payoff of party i given an outcome of Γ . For any party i , a *mixed strategy* s_i is a distribution in $\mu(S_i)$. A *strategy profile* of Γ is $s := s_1 \times \dots \times s_N$ where s_i is a mixed strategy of party i . The *expected utility* of a party i given a mixed strategy profile s is defined as $u_i(s) = \mathbb{E}_{a_1 \leftarrow s_1, \dots, a_N \leftarrow s_N} [u_i(a_1), \dots, u_i(a_N)]$. Finally, we note that if s_i is a Dirac distribution over a single strategy $a_i \in S_i$, we say s_i is a *pure strategy* for party i .

Notation. Let w denote the total reward for mining a block and q the negligible probability that a PPT adversary guesses a correct opening. Recall in Section 4.1 that the total block reward w is split between the miner of the block who gets $\alpha \cdot w$ and the miners that append the correct openings who get $\frac{(1-\alpha) \cdot w}{n_\ell}$ for each correct opening they append. For a given block, we denote by m the number of chunks for *each* transaction in the block, and by λ the utility of the sandwich attack on the block. Specifically, λ refers to the utility of a sandwich attack performed on the original transactions in the order they are in *before* chunking and permuting them. We also denote the optimal sandwich utility by Λ , which is the maximum utility one can get by performing a sandwich attack. Finally, we denote by $\hat{\lambda}_i$ the *average* utility of the sandwich attack taken over all blocks on the chain for a specific miner M_i . This can be computed easily as the transaction mempool is public. We stress that it is important to look at the average sandwich utility for each miner separately and not the average over all miners as the utility a miner can derive from a sandwich attack depends on their available liquidity (i.e., how much assets they can spare to front-run and back-run the transactions).

Quasi-strong ε -Nash Equilibrium. In terms of game theoretic security, we want our protocols to be resilient to deviations of any subset of miners that form a coalition and deviate jointly. The security notion we want to achieve is that of a *quasi-strong ε -Nash Equilibrium* [8, 13, 21]. Let C denote the coalition of players. For any strategy profile s , we denote by $u_C(s)$ the expected utility of the coalition under s . We denote by $u_C(s'_C, s_{-C})$ the expected utility of the coalition when playing according to some other strategy profile s'_C given the other players that are not part of the coalition play according to s .

► **Definition 6** (Quasi-strong ε -Nash Equilibrium). *A quasi-strong ε -Nash Equilibrium is a mixed strategy profile s such that for any other strategy profile s'_C , $u_C(s) \geq u_C(s'_C, s_{-C}) - \varepsilon$ for some $\varepsilon > 0$.*

The notion of a quasi-strong Nash Equilibrium is particularly useful in the context of blockchains as the coalition could potentially be controlled by a single miner with sufficient resources [21]. The notion of an ε -equilibrium is also important in cases where there could be a small incentive (captured by the ε parameter) to deviate from the protocol, and of course the smaller one can make ε , the more meaningful the equilibrium.

Subgame perfection. We also consider games that span several rounds and we model them as extensive-form games (see, e.g., [46] for a formal definition). Extensive form games can be represented as a game tree tx where the non-leaf vertices of the tree are partitioned to sets corresponding to the players. The vertices belonging to each player are further partitioned into information sets I which capture the idea that a player making a move at vertex $x \in I$ is uncertain whether they are making the move from x or some other vertex $x' \in I$. A subgame of an extensive-form game corresponds to a subtree in tx rooted at any non-leaf vertex x that belongs to its own information set, i.e., there are no other vertices that are the set except for x . A strategy profile is a *quasi-strong subgame perfect ε -equilibrium* if it is a quasi-strong ε -Nash equilibrium for all subgames in the extensive-form game.

The induced game. Let us divide our protocol into epochs: each epoch is designed around a given block say B_i and begins with the generation of random seeds for B_i and ends with appending the openings for the committed random seeds for B_i (i.e., block $B_{i+\tau_1+\tau_2}$). We define the underlying game Γ induced by any given epoch of our protocol Π^3 . Γ is a $(\tau_2 + 1)$ -round extensive form game played by $n_\ell + \tau_2$ parties (n_ℓ leaders comprising the leader set L_i for any block B_i and the τ_2 miners that mine the blocks $B_{i+\tau_1+1} \dots B_{i+\tau_1+\tau_2}$). Note that although we have $\binom{N}{\tau_2}$ sets of τ_2 miners to choose from (where N is the total number of miners in the chain) to be the miners of the blocks $B_{i+\tau_1+1} \dots B_{i+\tau_1+\tau_2}$, we can simply fix any set of τ_2 miners together with L_i to be the parties of Γ as we assume all miners are rational and so the analysis of the utilities of any set of τ_2 miners will be the same in expectation. We use A to denote the set of all miners in τ_2 . In what follows, we assume an arbitrary but fixed ordering of the miners in A . Round 1 of Γ consists of only the parties in L_i performing actions, namely picking a random seed and committing to it. In rounds $2, \dots, \tau_2 + 1$ of Γ , each member of L_i can act by choosing to open their commitment or not. However, the moment a member of L_i opens its commitment in a given round, they lose the chance to open their commitment in any subsequent round. Only one miner from A and according to the imposed ordering acts in each round from round 2 to $\tau_2 + 1$ of Γ . The choice of actions of the miner in any of these rounds are the subsets of the set of existing commitment openings (from members of L_i) to append to their block. Finally we note that the $L_i \cap A$ is not necessarily empty and thus miners in the intersection can choose to open and append their commitment in the same round.

12:16 Eating Sandwiches

Let us define the honest strategy profile as the profile in which all members of L_i choose to generate a random seed in round 1 of Γ , all members of L_i open their commitments at round τ_2 (i.e., at block $B_{i+\tau_1+\tau_2-1}$), and each member of A appends all existing opened commitments that appear in the previous round. We denote the honest strategy profile by s . The security notion we want to achieve for our protocol is a quasi-strong subgame perfect ε -equilibrium (refer to Definition 6). Looking ahead, we will also prove that ε can be made arbitrarily small by increasing the number m of chunks.

► **Lemma 7.** *The expected utility of an honest leader is at least $(1 - q)^{n_\ell} \alpha w$.*

Proof. The expected utility for a user following the honest strategy comes from the sum of the block reward, the expected utility from the ordering of any of their transactions within the block, and appending valid openings of committed seeds (if any) to their blocks. The expected utility from the ordering of transactions is 0 due to symmetry: each possible order is equally likely, for each order that gives some positive utility, there exists a different order producing the same negative utility. The expected utility from the block reward is $(1 - q)^{n_\ell} \alpha w$. Thus, the total expected utility of an honest miner is at least $(1 - q)^{n_\ell} \alpha w$. ◀

We outline and analyze two broad classes of deviations or attacks any coalition can attempt in this setting. The first class happens at round 1 of Γ where the members of the coalition commit to previously agreed seeds to produce a specific permutation of the transactions. The coalition then behaves honestly from round 2 to $\tau_2 + 1$ of Γ . We call this attack the *chosen permutation attack* and denote this attack strategy by s_{CP} . In the second class, the coalition behaves honestly at round 1 of Γ , but deviates from round 2 onwards where some members selectively withhold opening or appending commitments to bias the final permutation. We call this attack the *biased permutation attack*, and denote it by s_{BP} .

Chosen permutation attack. Before we describe and analyze the chosen permutation attack (for say a block B_i), we first show that a necessary condition for the attack to be successful, that is, the coalition's desired permutation happens almost surely, is that at least all n_ℓ leaders in L_i have to be involved in the coalition (members of A can also be involved in the coalition, however as we will show this will simply increase the cost). To do so, we let S denote the set of permutations over the list of transactions and their chunks, and we define what we mean by a protocol Π_{perm} (involving n parties) outputs random a permutation in S by the following indistinguishability game called *random permutation indistinguishability* played between a PPT adversary, a challenger, and a protocol Π_{perm} . First, the adversary corrupts up to $n - 1$ parties. The adversary has access to the corrupted parties' transcripts. Then, the challenger samples σ_0 uniformly at random from S , and sets σ_1 to be the output of Π_{perm} . After that, the challenger flips a random bit b and sends σ_b to the adversary. The game ends with the adversary outputting a bit b' . If $b' = b$, the adversary wins the game. We say a protocol Π_{perm} outputs a random permutation if the the adversary wins the above game with probability $\frac{1}{2} + \varepsilon$ for some negligible ε . Let us define the output of a single round of Π^3 as the random permutation that is generated from the seeds generated from all leaders in the round according to the algorithm described in Section 4.1. The following lemma states that as long as a single leader is honest, the output of Π^3 is pseudorandom.

► **Lemma 8.** *An adversary that corrupts at most $n_\ell - 1$ leaders in a single round of Π^3 can only win the random permutation indistinguishability game with negligible probability.*

Proof. The proof follows in the same way as introduced by M. Blum [14], with the addition of the PRG. ◀

Lemma 8 implies that launching the chosen permutation attack and thus choosing to deviate at round 1 of Γ comes with an implicit cost: either a single miner has to mine n_ℓ blocks in a row so the miner single-handedly forms the coalition, or *all* leaders in L_i have to be coordinated into playing according to a predefined strategy.

► **Lemma 9.** *Given the underlying blockchain is secure, the expected utility of the single miner when playing according to s_{CP} is at most $\frac{\lambda}{2^{n_\ell}}$ more than the expected utility of following the honest strategy.*

Proof. Since the underlying blockchain is secure, a necessary condition is that a single miner cannot own more than $\frac{1}{2}$ of the total amount of resources owned by all miners of the protocol. Thus, the probability of mining n_ℓ blocks in a row is strictly less than $\frac{1}{2^{n_\ell}}$. This means that the expected utility under the attack strategy $u_C(s_{CP}) < \frac{\lambda}{2^{n_\ell}} + n_\ell \alpha w$, which is at most $\frac{\lambda}{2^{n_\ell}}$ larger than the expected utility under the honest strategy which is $u_C(s) = n_\ell \alpha w$. ◀

The attack strategy of a coalition composed by more than one miner is more complex compared to the case where there is a single miner, as the coalition needs to ensure its members coordinate strategies. First, the coalition works with the miner of block B_i to select and fix a permutation generated by a specific PRG seed σ_i . Then, the coalition secret shares σ_i among its members⁷. After that, the coalition sets up some punishment scheme to penalize members that do not reveal their partial seeds⁸. Finally, the coalition commits and reveals these partial seeds in accordance to the protocol Π^3 . Let \mathcal{C} denote the expected cost of coordinating the whole chosen permutation attack for the coalition. For this attack to succeed, the expected coordination cost has to be smaller than the expected profit λ .

► **Lemma 10.** *The chosen permutation attack fails to be profitable compared to the honest strategy if $\mathcal{C} > \lambda$.*

Proof. From Lemma 7, the expected revenue of an honest miner is $(1 - q)^{n_\ell} \alpha w$, thus the expected revenue of the coalition when following the honest strategy is $u_C(s) = n_\ell \cdot (1 - q)^{n_\ell} \alpha w$. The expected revenue for the chosen permutation attack strategy is $u_C(s_{CP}) = n_\ell \cdot (1 - q)^{n_\ell} \alpha w + \lambda - \mathcal{C}$. Thus, assuming $\mathcal{C} > \lambda$, and since the expected revenue from a mixed strategy is a convex combination of the revenues of the honest and attack strategies, the pure honest strategy gives a strictly larger expected payoff compared to any mixed strategy. ◀

► **Remark 11.** Computing, or even estimating, the coordination cost is non-trivial as it consists of several dimensions and also depends on a myriad of factors and assumptions. A few notable costs are, firstly, timing costs. The coalition has to convince and coordinate all the leaders to agree on a permutation and also commit and reveal them during a short interval of d blocks. This involves the cost of securely communicating with all the leaders and also the computational cost involved in setting up the secret sharing scheme. A second factor is the choice of the initial order of transactions, which the coalition would have to also agree on with the miner of the attacked block. Picking transactions greedily would be the simplest choice as finding the optimal set of transactions from the mempool is NP-hard [42]. Finally, the coalition has to set up a punishment scheme to penalize members that do not

⁷ This not only prevents members from knowing the partial seeds of other members and hence stealing their block reward, but also additionally safeguards the partial seeds of the members against the miner of block B_i who cannot generate a partial seed of their block and hence has nothing to lose.

⁸ This ensures that every member will reveal their partial seeds and the permutation will be generated properly.

reveal their permutations. If we ignore the cost of setting up such a scheme, this can be implemented using a deposit scheme with the size of the deposit at least the value of the expected additional per user profit from the sandwich attack [51]. This implies an opportunity cost at least linear in $\frac{\lambda}{n_\ell}$, as well as the assumption that each member has at least $\frac{\lambda}{n_\ell}$ to spare to participate in the attack. Additionally, we note that the coalition could extend to miners from A which are outside the leader set L_i . However, since these miners do not contribute to generating the random seeds, they simply add to the communication cost of the coalition. Finally, we note that the coordination cannot be planned in advance due to the unpredictability of the block mining procedure.

Biased permutation attack. The intuition behind this attack is that any coalition that controls $k \leq n_\ell$ commitments can choose to select the ones to open or append, which allows the coalition to choose among 2^k possible permutations in order to bias the final ordering. This can be achieved in two situations: either k out of n_ℓ leaders of L_i form a coalition and decide which of their commitments to open, or some subset of miners in the loud phase (of size say \tilde{k}) form a coalition and end up controlling k openings, let $\kappa := \min\{k, \tilde{k}\}$. Unlike in the case of the chosen permutation attack, it suffices consider the case where we have a *single* miner that happens to either occupy k leader positions among the group of leaders L_i or mine the \tilde{k} blocks that belong to the coalition in the loud phase. This is because the case where a coalition of distinct miners that collude only adds additional coordination cost. The probability that any such coalition gains any additional utility by performing the biased permutation attack compared to the honest strategy can be upper-bounded. Let *revenue* denote the utility the coalition would gain from performing the biased permutation attack.

► **Lemma 12.** *The probability that a coalition of κ members performing the biased permutation attack achieves utility of at least $\kappa w > 0$ is $\text{P}[\text{revenue} \geq \kappa w] \leq 1 - (1 - e^{-\frac{2m\kappa w}{\lambda}})^{2^k}$.*

Proof. Given a random permutation and a sandwich attack with original utility λ (utility if the order of the transactions were not randomized), denote by $\{X_i(\sigma)\}_{i=1}^m$ the utility produced by chunk i . The sum of these random variables $X(\sigma) = \sum_{i=1}^m X_i(\sigma)$ represents the total utility of a sandwich attack (after chunking and permuting). X takes values in $[-\lambda, \lambda]$, thus the variables $\{X_i(\sigma)\}$ take values in $[-\frac{\lambda}{m}, \frac{\lambda}{m}]$, are equally distributed and are independent. We define the random variables $Y_i(\sigma) = X_i(\sigma) + \frac{\lambda}{m} \in [0, \frac{2\lambda}{m}]$, and $Y(\sigma) = \sum_{i=1}^m Y_i(\sigma) \in [0, 2\lambda]$. Using lemma 7, $\text{E}[Y_i(\sigma)] = \frac{\lambda}{m}$ and $\text{E}[Y(\sigma)] = \lambda$. Applying Chernoff's bound [44] to Y ,

$$\text{P}[Y(\sigma) \geq (1 + \delta)\text{E}[Y(\sigma)]] \leq e^{-\frac{2\delta^2 \text{E}[Y(\sigma)]^2}{m(\frac{\lambda}{m})^2}} = e^{-2m\delta^2} \quad (1)$$

for $\delta > 0$. We can rewrite Equation 1 as follows:

$$\text{P}[\text{revenue}(\sigma) \geq \delta\lambda] = \text{P}[\text{revenue}(\sigma) + \lambda \geq (1 + \delta)\lambda] = \text{P}[Y(\sigma) \geq (1 + \delta)\text{E}[Y(\sigma)]] \leq e^{-2m\delta^2}.$$

Using the law of total probability we obtain that $\text{P}[\text{revenue}(\sigma) \leq \delta\lambda] \geq 1 - e^{-2m\delta^2}$. Considering the maximum over the 2^k possible permutations σ and $\delta = \frac{\kappa w}{\lambda}$ we conclude that

$$\text{P}[\text{revenue} \geq \kappa w] = 1 - \text{P}[\text{revenue} \leq \kappa w] = 1 - \text{P}[\text{revenue}(\sigma) \leq \kappa w] \leq 1 - (1 - e^{-\frac{2m\kappa w}{\lambda}})^{2^k}. \blacktriangleleft$$

► **Lemma 13.** *The probability that a coalition of κ members has positive additional utility is: $\text{P}[\text{revenue} \geq 0] \leq \max_{k' \leq \kappa} \left\{ 1 - (1 - e^{-\frac{2mk'w}{\lambda}})^{2^k} \right\}$.*

Proof. Lemma 12 states a bound for the probability that a coalition of κ parties has a utility of at least $\kappa w > 0$, the penalty for not opening κ commitments. Thus, the general case for a coalition aiming to maximize profit is the maximum over $k' \leq \kappa$. \blacktriangleleft

Recall that Λ is the maximal utility and let $p_{k,\lambda}$ denote $\max_{k' \leq \kappa} \{1 - (1 - e^{-\frac{2m(1-q)^{n_\ell} k' w}}{\lambda})^{2^k}\}$. Then, the expected additional utility from the biased permutation attack of a single miner controlling k leaders is no greater than $p_{k,\lambda}\Lambda$.

Lemmas 10,12 and 13 allow us to prove our main theorem.

► **Theorem 14.** *Suppose $\mathcal{C} > \lambda$, then the honest strategy $s = ((\text{random seed})_{i=1}^{n_\ell}, (\text{open})_{i=1}^{n_\ell})$ is a quasi-strong subgame perfect ε -equilibrium in Γ for $\varepsilon = \max\{\frac{\lambda}{2^{n_\ell}}, p_{k,\lambda}\Lambda\}$.*

Proof. We first observe that the expected utility of a coalition that mixes both the chosen and biased permutation attack strategies is no greater than the expected utility of a coalition that performs the chosen permutation attack with a different chosen permutation that accounts for the biasing of the permutation in the second round of Γ . Hence, it suffices to analyze the expected utility of the coalition when implementing either of these strategies, i.e., deviating at round 1 of Γ or from rounds 2 onwards.

We first analyze the expected utility of a coalition when implementing the chosen permutation attack, which occurs at round 1 or Γ . Since we assume $\mathcal{C} > \lambda$, from Lemma 9 and Lemma 10, we see that any additional expected payoff of any coalition that deviates only at round 1 of Γ by implementing the chosen permutation attack compared to the expected revenue of behaving honestly is at most $\frac{\lambda}{2^{n_\ell}}$.

Now we analyze the expected utility of a coalition when implementing the biased permutation attack. From Lemmas 12 and 13, we see that the strategy that implements the biased permutation attack across all of rounds 2 to $\tau_2 + 1$ of Γ only gives at most $p_{k,\lambda}\Lambda$ more payoff in expectation compared to following the honest strategy s in these rounds.

As such, if we set $\varepsilon = \max\{\frac{\lambda}{2^{n_\ell}}, p_{k,\lambda}\Lambda\}$ to be the largest difference in additional expected revenues between both strategies, we see that $s = ((\text{random seed})_{i=1}^{n_\ell}, (\text{open})_{i=1}^{n_\ell})$ is a quasi-strong ε -subgame perfect equilibrium of Γ . ◀

► **Remark 15.** Recall that ε bounds the additional expected utility an adversary can gain by deviating from the honest strategy profile s . The security of our protocol therefore improves as $\varepsilon = \max\{\frac{\lambda}{2^{n_\ell}}, p_{k,\lambda}\Lambda\}$ decreases. We observe that the first component $\frac{\lambda}{2^{n_\ell}}$ goes to 0 exponentially as the size of the leader set n_ℓ increases. As for the second component $p_{k,\lambda}$, we conduct an empirical analysis of sandwich attacks on Ethereum, see full version [6], to estimate $p_{k,\lambda}$ and we show that this value approaches zero as the number of chunks m increases.

6 Conclusion

In this paper we introduced a new construction that can be implemented on top of any blockchain protocol with three main properties. First, the construction does not add any vulnerability to the old protocol, i.e., the security properties remain unchanged. Secondly, performing sandwich attacks in the new protocol is no longer profitable. Thirdly, the construction incurs in minimal overhead with the exception of a minor increase in the latency of the protocol. Our empirical study of sandwich attacks on the Ethereum blockchain also validates the design principles behind our protocol, demonstrating that our protocol can be easily implemented to mitigate sandwich MEV attacks on the Ethereum blockchain.

References

- 1 Distributed randomness beacon, 2023. URL: <https://drand.love/>.
- 2 Ethereum, 2023. URL: <https://ethereum.org/en/whitepaper/>.

- 3 MEV over time, 2023. URL: <https://explore.flashbots.net/>.
- 4 Random ordering of equally-priced transactions incentivises competitive spam, 2023. URL: <https://github.com/ethereum/go-ethereum/issues/21350>.
- 5 Ittai Abraham, Benny Pinkas, and Avishay Yanai. Blinder - scalable, robust anonymous committed broadcast. In *CCS*, pages 1233–1252. ACM, 2020. doi:10.1145/3372297.3417261.
- 6 Orestis Alpos, Ignacio Amores-Sesar, Christian Cachin, and Michelle Yeo. Eating sandwiches: Modular and lightweight elimination of transaction reordering attacks. *CoRR*, abs/2307.02954, 2023. arXiv:2307.02954, doi:10.48550/ARXIV.2307.02954.
- 7 Guillermo Angeris and Tarun Chitra. Improved price oracles: Constant function market makers. In *AFT*, pages 80–91. ACM, 2020. doi:10.1145/3419614.3423251.
- 8 Robert J. Aumann. Acceptable points in general cooperative n-person games. In Albert William Tucker and Robert Duncan Luce, editors, *Contributions to the Theory of Games (AM-40), Volume IV*, pages 287–324. Princeton University Press, Princeton, 1959. doi:10.1515/9781400882168-018.
- 9 Leemon Baird and Atul Luykx. The hashgraph protocol: Efficient asynchronous BFT for high-throughput distributed ledgers. In *COINS*, pages 1–7. IEEE, 2020. doi:10.1109/COINS49042.2020.9191430.
- 10 Carsten Baum, James Hsin-yu Chiang, Bernardo David, Tore Kasper Frederiksen, and Lorenzo Gentile. Sok: Mitigation of front-running in decentralized finance. *IACR Cryptol. ePrint Arch.*, page 1628, 2021. URL: <https://eprint.iacr.org/2021/1628>.
- 11 Carsten Baum, Bernardo David, and Rafael Dowsley. Insured MPC: efficient secure computation with financial penalties. In *Financial Cryptography*, volume 12059 of *Lecture Notes in Computer Science*, pages 404–420. Springer, 2020. doi:10.1007/978-3-030-51280-4_22.
- 12 Carsten Baum, Bernardo David, and Tore Kasper Frederiksen. P2DEX: privacy-preserving decentralized cryptocurrency exchange. In *ACNS (1)*, volume 12726 of *Lecture Notes in Computer Science*, pages 163–194. Springer, 2021. doi:10.1007/978-3-030-78372-3_7.
- 13 B. Douglas Bernheim, Bezalel Peleg, and Michael D Whinston. Coalition-proof nash equilibria I. Concepts. *Journal of Economic Theory*, 42(1):1–12, 1987. doi:10.1016/0022-0531(87)90099-8.
- 14 Manuel Blum. Coin flipping by telephone a protocol for solving impossible problems. *SIGACT News*, 15(1):23–27, jan 1983. doi:10.1145/1008908.1008911.
- 15 Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *CRYPTO (1)*, volume 10991 of *Lecture Notes in Computer Science*, pages 757–788. Springer, 2018. doi:10.1007/978-3-319-96884-1_25.
- 16 Mic Bowman, Debajyoti Das, Avradip Mandal, and Hart Montgomery. On elapsed time consensus protocols. In *INDOCRYPT*, volume 13143 of *Lecture Notes in Computer Science*, pages 559–583. Springer, 2021. doi:10.1007/978-3-030-92518-5_25.
- 17 Christian Cachin, Rachid Guerraoui, and Luís E. T. Rodrigues. *Introduction to Reliable and Secure Distributed Programming (2. ed.)*. Springer, 2011. doi:10.1007/978-3-642-15260-3.
- 18 Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient asynchronous broadcast protocols. In *CRYPTO*, volume 2139 of *Lecture Notes in Computer Science*, pages 524–541. Springer, 2001. doi:10.1007/3-540-44647-8_31.
- 19 Christian Cachin, Jovana Micic, Nathalie Steinhauer, and Luca Zanolini. Quick order fairness. In *Financial Cryptography*, volume 13411 of *Lecture Notes in Computer Science*, pages 316–333. Springer, 2022. doi:10.1007/978-3-031-18283-9_15.
- 20 Chainlink Labs. Chainlink 2.0: Next steps in the evolution of decentralized oracle networks. Whitepaper, 2021. URL: <https://research.chain.link/whitepaper-v2.pdf>.
- 21 Krishnendu Chatterjee, Amir Kafshdar Goharshady, and Arash Pourdamghani. Probabilistic smart contracts: Secure randomness on the blockchain. In *IEEE ICBC*, pages 403–412. IEEE, 2019. doi:10.1109/BLOC.2019.8751326.

- 22 Kevin Choi, Arasu Arun, Nirvan Tyagi, and Joseph Bonneau. Bicorn: An optimistically efficient distributed randomness beacon. *IACR Cryptol. ePrint Arch.*, page 221, 2023. URL: <https://eprint.iacr.org/2023/221>.
- 23 Benny Chor, Shafi Goldwasser, Silvio Micali, and Baruch Awerbuch. Verifiable secret sharing and achieving simultaneity in the presence of faults (extended abstract). In *FOCS*, pages 383–395. IEEE Computer Society, 1985. doi:10.1109/SFCS.1985.64.
- 24 Bram Cohen and Krzysztof Pietrzak. The chia network blockchain. Whitepaper, 2019. URL: <https://www.chia.net/wp-content/uploads/2022/07/ChiaGreenPaper.pdf>.
- 25 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, 3rd Edition*. MIT Press, 2009. URL: <http://mitpress.mit.edu/books/introduction-algorithms>.
- 26 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 Symposium on Security and Privacy*, pages 910–927. IEEE, 2020. doi:10.1109/SP40000.2020.00040.
- 27 Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *EUROCRYPT (2)*, volume 10821 of *Lecture Notes in Computer Science*, pages 66–98. Springer, 2018. doi:10.1007/978-3-319-78375-8_3.
- 28 Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In *CRYPTO*, volume 435 of *Lecture Notes in Computer Science*, pages 307–315. Springer, 1989. doi:10.1007/0-387-34805-0_28.
- 29 Yael Doweck and Ittay Eyal. Multi-party timed commitments. *CoRR*, abs/2005.04883, 2020. arXiv:2005.04883.
- 30 Sisi Duan, Michael K. Reiter, and Haibin Zhang. Secure causal atomic broadcast, revisited. In *DSN*, pages 61–72. IEEE Computer Society, 2017. doi:10.1109/DSN.2017.64.
- 31 Lioba Heimbach and Roger Wattenhofer. Eliminating sandwich attacks with the help of game theory. In *AsiaCCS*, pages 153–167. ACM, 2022. doi:10.1145/3488932.3517390.
- 32 Mahimna Kelkar, Soubhik Deb, and Sreeram Kannan. Order-fair consensus in the permissionless setting. In *APKC@AsiaCCS*, pages 3–14. ACM, 2022. doi:10.1145/3494105.3526239.
- 33 Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. Themis: Fast, strong order-fairness in byzantine consensus. *IACR Cryptol. ePrint Arch.*, page 1465, 2021. URL: <https://eprint.iacr.org/2021/1465>.
- 34 Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 451–480. Springer, 2020. doi:10.1007/978-3-030-56877-1_16.
- 35 John Kelsey, Luís T. A. N. Brandão, Rene Peralta, and Harold Booth. Nistir 8213. a reference for randomness beacons: Format and protocol version 2, 2011.
- 36 Aggelos Kiayias, Hong-Sheng Zhou, and Vassilis Zikas. Fair and robust multi-party computation using a global transaction ledger. In *EUROCRYPT (2)*, volume 9666 of *Lecture Notes in Computer Science*, pages 705–734. Springer, 2016. doi:10.1007/978-3-662-49896-5_25.
- 37 Gillat Kol and Moni Naor. Games for exchanging information. In *STOC*, pages 423–432. ACM, 2008. doi:10.1145/1374376.1374437.
- 38 Kshitij Kulkarni, Theo Diamandis, and Tarun Chitra. Towards a theory of maximal extractable value I: constant function market makers. *CoRR*, abs/2207.11835, 2022. arXiv:2207.11835, doi:10.48550/ARXIV.2207.11835.
- 39 Klaus Kursawe. Wendy, the good little fairness widget: Achieving order fairness for blockchains. In *AFT*, pages 25–36. ACM, 2020. doi:10.1145/3419614.3423263.
- 40 Arjen K. Lenstra and Benjamin Wesolowski. Trustworthy public randomness with sloth, unicorn, and trx. *Int. J. Appl. Cryptogr.*, 3(4):330–343, 2017. doi:10.1504/IJACT.2017.10010315.

- 41 Donghang Lu, Thomas Yurek, Samarth Kulshreshtha, Rahul Govind, Aniket Kate, and Andrew Miller. Honeybadgermpc and asynchromix: Practical asynchronous MPC and its application to anonymous communication. In *CCS*, pages 887–903. ACM, 2019. doi:10.1145/3319535.3354238.
- 42 Mohsen Alambardar Meybodi, Amir Kafshdar Goharshady, Mohammad Reza Hooshmandasl, and Ali Shakiba. Optimal mining: Maximizing bitcoin miners’ revenues from transaction fees. In *Blockchain*, pages 266–273. IEEE, 2022. doi:10.1109/BLOCKCHAIN55522.2022.00044.
- 43 Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of BFT protocols. In *CCS*, pages 31–42. ACM, 2016. doi:10.1145/2976749.2978399.
- 44 Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005. doi:10.1017/CB09780511813603.
- 45 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Whitepaper, 2009. URL: <http://bitcoin.org/bitcoin.pdf>.
- 46 Martin J. Osborne and Ariel Rubinstein. *A course in game theory*. The MIT Press, Cambridge, USA, 1994. electronic edition.
- 47 Torben P. Pedersen. Cps, certificate practice statement. In *Encyclopedia of Cryptography and Security*. Springer, 2005. doi:10.1007/0-387-23483-7_83.
- 48 Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? In *IEEE Symposium on Security and Privacy*, pages 198–214. IEEE, 2022. doi:10.1109/SP46214.2022.9833734.
- 49 Michael K. Reiter and Kenneth P. Birman. How to securely replicate services. *ACM Trans. Program. Lang. Syst.*, 16(3):986–1009, 1994. doi:10.1145/177492.177745.
- 50 R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, Massachusetts Institute of Technology, USA, 1996.
- 51 Nikolaž I. Schwartzbach. Deposit schemes for incentivizing behavior in finite games of perfect information. *CoRR*, abs/2107.08748, 2021. arXiv:2107.08748.
- 52 Avishay Yanai. Blinderswap: MEV meets MPC. https://www.youtube.com/watch?v=KQ4xK79YkFE&ab_channel=IC3InitiativeforCryptocurrenciesandContracts, 2021. Accessed 03/08/23.
- 53 Haoqian Zhang, Louis-Henri Merino, Vero Estrada-Galiñanes, and Bryan Ford. F3B: A low-latency commit-and-reveal architecture to mitigate blockchain front-running. *CoRR*, abs/2205.08529, 2022. doi:10.48550/ARXIV.2205.08529.