



Dissecting Bitcoin and Ethereum Transactions: On the Lack of Transaction Contention and Prioritization Transparency in Blockchains

Johnnatan Messias¹ , Vabuk Pahari¹, Balakrishnan Chandrasekaran²,
Krishna P. Gummadi¹, and Patrick Loiseau³

¹ MPI-SWS, Saarbrücken, Germany

{johnme, vpahari, gummadi}@mpi-sws.org

² Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

b.chandrasekaran@vu.nl

³ Inria, FairPlay Team, Paris, France

patrick.loiseau@inria.fr

Abstract. In permissionless blockchains, transaction issuers include a fee to incentivize miners to include their transactions. To accurately estimate this prioritization fee for a transaction, transaction issuers (or blockchain participants, mjohme@mpi-sws.orgore generally) rely on two fundamental notions of transparency, namely contention and prioritization transparency. Contention transparency implies that participants are aware of every pending transaction that will contend with a given transaction for inclusion. Prioritization transparency states that the participants are aware of the transaction or prioritization fees paid by every such contending transaction. Neither of these notions of transparency holds well today. Private relay networks, for instance, allow users to send transactions privately to miners. Besides, users can offer fees to miners via either direct transfers to miners' wallets or off-chain payments—neither of which are public. In this work, we characterize the lack of contention and prioritization transparency in Bitcoin and Ethereum resulting from such practices. We show that private relay networks are widely used and private transactions are quite prevalent. We show that the lack of transparency facilitates miners to collude and overcharge users who may use these private relay networks despite them offering little to no guarantees on transaction prioritization. The lack of these transparencies in blockchains has crucial implications for transaction issuers as well as the stability of blockchains. Finally, we make our data sets and scripts publicly available.

Keywords: Contention transparency · Prioritization transparency · Private transactions · Bitcoin · Ethereum · MEV

1 Introduction

The rate at which users issue transactions in permissionless blockchains, e.g., Bitcoin [31] and Ethereum [47], is often much higher than the rate at which

© The Author(s) 2024

F. Baldimtsi and C. Cachin (Eds.): FC 2023, LNCS 13951, pp. 221–240, 2024.

https://doi.org/10.1007/978-3-031-47751-5_13

miners can include them in a block [10, 21, 25, 28, 29]. Users typically issue transactions using a wallet software, whose primary functionality is determining an “appropriate” fee for a given transaction. We use the term “fee” to refer generally to the incentive offered by a user to miners for prioritizing the inclusion of their transaction in a block, albeit its exact form may vary, e.g., *fee rate* in Bitcoin and *gas price* in Ethereum. This (prioritization) fee varies, unsurprisingly, as a function of the level of congestion in the blockchain [29] as well as the distribution of fees across available transactions. Inferring either of these is, however, deceptively complicated.

At first glance, these tasks appear straightforward, since every transaction is broadcast to all miners in the blockchain. A user could simply gather all transactions broadcast over time and reconstruct the set of uncommitted transactions available to a miner (i.e., contents of the miner’s Mempool) at any point of time [28]. We refer to this assumption of a public and uniform view (across miners) of all available transactions as *contention transparency*. If contention transparency exists, a user could rank order available transactions by their fee (based on which miners should select transactions for inclusion) and estimate the commit delay of any transaction [29]. Consequently, they could determine the fee that they must pay to guarantee inclusion of their transaction in a given block. We label this assumption that the (prioritization) fee offered by a transaction is only that publicly declared by that transaction as *prioritization transparency*. Neither the contention transparency nor the prioritization transparency, however, holds today in permissionless blockchains.

Lack of Contention Transparency. Not all transactions are publicly broadcast. Users can submit transactions to a subset of miners or mining pools via *private channels* or *relays* that are opaque to the public (i.e., transactions remain private to the relay, until they are committed). Users may also submit their transaction to a specific mining pool that assures them a fast commit time. This paper reveals that such private mining practices (i.e., where transactions are submitted to only a subset of the miners) are becoming commonplace and analyzes the characteristics of these private transactions.

Lack of Prioritization Transparency. The fees offered by a transaction could be substantially more than that publicly declared by it. A transaction could, for instance, privately offer additional fees to a miner to “accelerate” its inclusion in a block. Many such transaction-accelerator (or *front-running as a service (FRaaS)*) platforms exist for Bitcoin [4, 45] and Ethereum [12, 18, 40, 41]. Furthermore, the same transaction could offer different fees to different mining pools (via their relays). The presence of such hidden or dark-fees could fundamentally erode the reliability of any fee prediction: Transaction issuers may end up paying substantially large fees without receiving proportional or any reduction in commit delays. This paper characterizes the prevalence of such dark-fee transactions and analyzes the most popular private relay network available in Ethereum, Flashbots [18]. Furthermore, we conduct active experiments in both Bitcoin and Ethereum to validate our assumptions regarding the prioritization transparency. In addition to showing that transaction fees may not be uniform

across miners, we claim that, given the lack of contention transparency, the lack of prioritization transparency may become more widespread than it is now.

The lack of contention and prioritization transparencies stem from real, non-trivial concerns of transaction issuers. The risk of transactions being front-run by bots [9, 12, 42, 46], for instance, creates the need for transaction privacy. Mining pools that address this need also facilitate, unsurprisingly, off-chain payments via which transaction issuers can (privately) incentivize the miners [4, 29, 45]. We view these developments as natural and logical steps in the evolution of blockchains and back our assertions with empirical observations. We claim, therefore, in contrast to prior work [9, 41], that it is only the opacity of the overall fees issued by a transaction issuer that poses a fundamental threat to the stability of blockchains: Transaction issuers cannot, for instance, precisely infer the fee required to commit their transactions into the next block, and miners can, consequently, overcharge them as the “real” fees are opaque to the rest of the network [46].

We summarize our contributions as follows. We characterize the lack of contention transparency in both Bitcoin and Ethereum: We show that the use of private channels or relay networks to submit transactions directly to a subset of miners is becoming widespread. This practice will likely erode prioritization transparency, as transaction issuers may not be able to estimate the appropriate fees, none of which are publicly visible. We characterize the prevalence of such private transactions fees. We found that Flashbots bundles represent 52.11% of all Ethereum blocks. With the lack of prioritization transparency, miners might overcharge users when they send their transactions privately. We also show that Bitcoin miners collude (with an aggregate hashing power of more than 50% of the network’s total hashing power) when including dark-fees transactions. Finally, we release our data sets and the scripts used in our analysis to enable the scientific community to reproduce our results [30].

2 Related Work

There is a rich literature on block rewards as incentives for mining [7, 16, 17, 19, 23, 32, 33, 36, 39, 48]. Recent work also analyzed the implications of relying on transaction fees separately [6] and in conjunction with block rewards [43], as well as the relationship between such incentives and transaction waiting times [10]. These prior work assume that transactions are broadcast to all miners and the fees offered is uniform across miners. None of them acknowledge the issue of transparency.

Basu *et al.* [3] and Lavi *et al.* [25] addressed the inefficiencies in transaction-fee setting mechanisms (i.e., first-price auctions) by proposing alternative mechanisms. They claim that miners might be dishonest, albeit they present no empirical evidence. Siddiqui *et al.* [38] used simulations to show that, if transaction fees are the only incentives, miners will select transactions greedily, thereby increasing the commit times of many transactions. Prior work also analyzed the Ethereum fee (i.e., gas price) mechanism to determine the gas price for a given transaction [1, 26, 27, 44]. The fee estimation and fee-based prioritization schemes in these studies do not take into account dark-fees or private mining.

Many transaction-accelerator, or FRaaS, platforms exist for both Bitcoin [4, 45] and Ethereum [12, 18, 40]. Transaction issuers might resort to such acceleration or off-chain payment channels to hide their true fee from competitors and avoid being front-run [9, 41]. Tim Roughgarden [37] discussed the incentives for off-chain agreements (such as dark-fees) between miners and users for first-price auctions and different deviations of the new Ethereum fee mechanism *EIP-1559 protocol* [5].¹ Roughgarden showed that miners and users cannot strictly increase their joint utility through off-chain payments under EIP-1559 because on-chain bids can be easily replaced by the off-chain bids. However, utility here is only based on the revenue of bidding for block space. The author did not take into account that utility might depend on other factors, such as transaction issuers wanting to keep their actual bids for block space hidden through off-chain payments, which strictly increases their chances of prioritization, as other bidders cannot counter bid, as they are unaware of the bid itself.

Closest to our work are two that analyze private mining. Strehe and Ante [41] investigated *exclusive mining* (or private mining), where transactions issuers and miners collude to include transactions that have been sent through a private network. In this case, the transactions are not publicly disclosed until they have been included in a block; besides, the fees can remain opaque to everyone forever, as such off-chain agreements may use fiat currencies. Weintraub *et al.* [46] measured the popularity of *Flashbots*, the most used private relay network for Ethereum. Our work, in contrast, extensively investigates private transactions in both Bitcoin and Ethereum blockchains. Through active measurements, we empirically show that Bitcoin miners collude and highlight the colluding mining pools. We show that Flashbots bundles are quite prevalent in Ethereum and are mainly used for calling Decentralized Exchanges (DEX) contracts to take advantage of *Maximal Extractable Value (MEV)* opportunities. Finally, we discuss why our findings are still valid after “The Merge”—an Ethereum hard fork deployed on September 15th, 2022 [13, 14].

3 On Contention Transparency

3.1 The Rise of Private Relay Networks

With the lucrative market of Decentralized Finance (DeFi) in Ethereum, today, bots engage in predatory front-running behaviors such as sandwich attacks and transaction-replay attacks [9, 24, 34, 35, 42, 46, 49]. Relay networks help users to counter such attacks: They provide users with a private channel for communicating with miners, who have to prove their identity to participate in the relay. Relay networks help users completely bypass the P2P network: Users send their transactions to the relay network, which in turn relays them to its participant miners. The relay network and its participants claim (a) not to front-run these transactions and (b) to keep them private until they are included in a block [18].

¹ The EIP-1559 went live in the Ethereum’s London hard fork upgrade on August 5th, 2021, at block number 12,965,000.

These transactions, hence, by construction, experience no front-running issues. Relay networks are centralized; if miners misbehave, they may lose their network membership and forfeit their future profits. Multiple relay networks (e.g., bloXroute, Taichi Network [40], and others [8,15]) exist today, but we focus on Flashbots [18], the largest relay network for Ethereum.

Flashbots. Flashbots’s users *bundle* one or more transactions in some specific order [18]. Miners are expected to mine the entire bundle (retaining the ordering of transactions within the bundle) and place it at the top of their blocks. The miners receive a fee (paid via a direct transfer to their wallets) for including the bundle in addition to the (traditional) fees associated with the transactions in that bundle. If there are two competing bundles—capturing the same financial opportunity, e.g., liquidations—miners will choose the one with the highest reward (i.e., maximizing financial incentives). The other bundle is *discarded* (since the financial opportunity no longer exists after having been captured by the included bundle), albeit its transactions do *not* expend *any* gas. Therefore, except for a network base fee introduced in EIP-1559, arbitrageurs and liquidators can participate without having any balance in their wallet: If they successfully capture a financial opportunity, they pay the miner from the profit secured and pocket the rest [18]. Flashbots is a *free* to use relay network, and they allow anyone to query whether a transaction used their relay network and the private fees paid to the miner (after it has been committed in a block). We use this publicly available data for analyzing the transactions issued (privately) on Flashbots. Flashbots, however, does not list the discarded bundles (or its transactions): we have access, hence, only to committed transactions.

3.2 Characterizing Private Relay Networks

We gathered all Ethereum blocks mined over a 9-month time period—from September 8th, 2021 to June 30th, 2022—to investigate the behavior of Ethereum mining pools. This data set contains 347,629,393 issued transactions and 1,867,000 blocks (from block number 13,183,000 to 15,049,999). We used miners’ wallet addresses to infer the block owners, but we failed to identify the owners of 46,895 blocks (or 2.51% of the total); we grouped the latter into one category, “Unknown.” Figure 1a shows the distribution of blocks and transactions mined in Ethereum by the top-20 mining pools. We also retrieved 6,937,292 transactions (2% of all issued transactions) contained in 3,284,886 bundles from Flashbots; these are transactions sent privately to miners. 972,911 (52.11%) of blocks in the data set have at least one such Flashbots transaction: *Private transactions are becoming quite common across most of the powerful mining pools in Ethereum.*

Flashbots labels its bundles (and constituent transactions) into one of three categories: (i) *flashbots*, which represent those sent through their private relay; (ii) *rogue*, referring to those delivered to a (Flashbots) miner, but via a different relay network; and (iii) *miner payout*, indicating a bundle containing payouts to users of a mining pool [46]. We find 58.82%, 27.93%, and 13.25% of transactions belonging to the flashbots, miner payout, and rogue categories, respectively. We

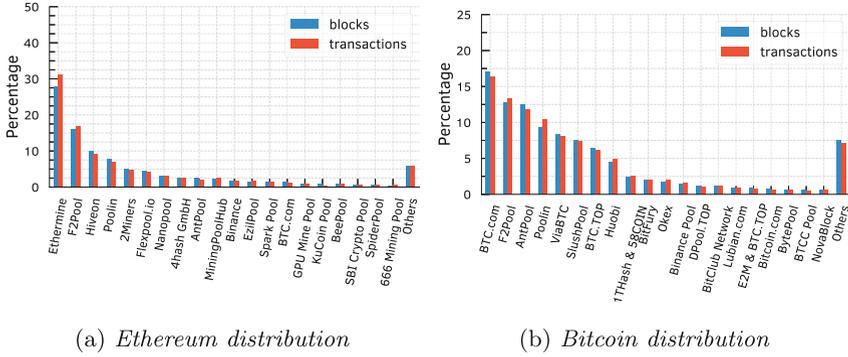


Fig. 1. Blocks mined and transactions confirmed in (a) Ethereum and (b) Bitcoin by the top-20 mining pools; “Others” consolidates the remaining mining pools.

also noticed that 70,260 (1.01%) of all Flashbots transactions failed to execute after inclusion in a block. A small fraction of transactions is, hence, not successfully executed despite using private relays.

Flashbots claims to have $\approx 85\%$ of the total Ethereum hash rate [18]. Per our analyses, however, the majority of the mining pools (47 out of 48—barring EthPool) use Flashbots, accounting for 99.99% of the total Ethereum hash rate, A recent work also corroborates our findings [46].

Some of the most powerful mining pools like Spark Pool² (which cooperates with Taichi Network [40]), Ethermine [15], and F2Pool (part of Eden Network [8]) offer their own relay networks. As these networks allow transaction issuers to send transactions exclusively to a specific miner, we hypothesize that miners would prefer (or prioritize) these transactions to those sent via the public P2P network. Crucially, payments from these private transactions are guaranteed, while those from publicly issued transactions are not—they are available to any miner willing to commit them. *Miners, hence, would likely offer preferential treatment for private transactions.*

3.3 On Preferential Treatment of Private Transactions

We substantiate our hypothesis of preferential treatment for private transactions via an active experiment conducted on September 8th, 2021. We issued 8 transactions, where 4 were sent privately via the Taichi Network, powered by Spark Pool, and 4 through the public Ethereum network (refer Table 4 in Appendix 1).

While running the experiment, we checked if the popular Ethereum blockchain explorers (i.e., Etherscan, Blockchain.com, and Blockchair) observed any of our private transactions; if they did, it would imply that the Taichi Network leaked the transactions to the public. While the public transactions

² Spark Pool suspended their mining services on Sept. 30th, 2021, due to regulatory requirements introduced by Chinese authorities [20].

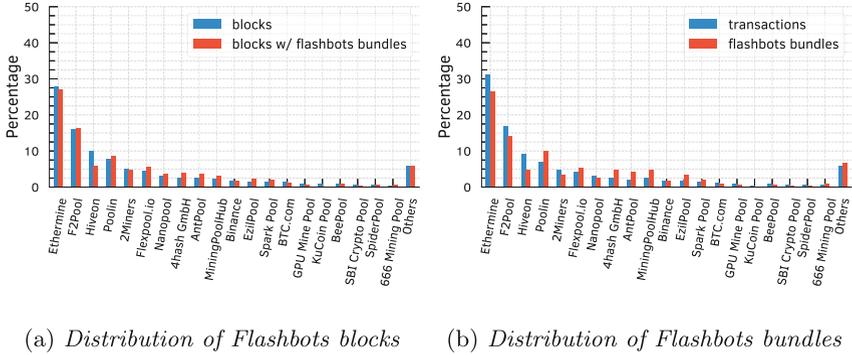


Fig. 2. Distribution of (a) blocks with at least one Flashbots bundle and (b) bundle of transactions per block, per mining pool. Ethermine included 27.05% of all blocks with a Flashbot bundle and 26.63% of all Flashbots bundles, while mining around 28.05% and 31.11% of all blocks and transactions, respectively.

appeared in these blockchain explorers, right after we sent them through the public P2P network, the private transactions were not observed by any of them until the transactions were included in a block. More importantly, our private transactions were *not* flagged by Etherscan (which relies on Flashbots API and more recently on EigenPhi [11]) as private, *even after inclusion in a block*. Measuring the prevalence of private transactions is, hence, challenging; it is likely that our estimates of the volume of private transactions based on such tools represent, hence, a lower bound.

Babel Pool included 2 out of our 4 private transactions. Spark Pool technically supports this mining pool, implying that they “collaborate” in committing private transactions sent over the Taichi network [2]. Our transactions were included, however, in the appropriate position in the block based on their fees. We delve into the prioritization of transactions in the next section.

We also characterize the prevalence of private transactions in Ethereum and indicate that mining pools can each have a distinct set of private transactions in their Mempool. Users, as a result, can no longer rely on the public Mempool alone to estimate their transaction fee. Given the absence of other data, they are highly likely to end up with a false estimate of the “appropriate” transaction fees for their transactions.

4 On Prioritization Transparency

4.1 Prevalence of Transaction Bundling

Flashbots bundles are prevalent in Ethereum (refer Sect. 3.2). Each Flashbots bundle contains at least 1 transaction and at most 631 transactions; on average they contain 2.11 transactions, with a median of 1 and a standard deviation of 6.47. We noticed that Ethermine alone included more than a quarter (26.63%)

of all 3,284,886 bundles (Fig. 2). Also, blocks contain at most 40 bundles, with an average of 3.38, a median of 3, and a standard deviation of 2.64.

Maximal Extractable Value (MEV). Flashbots allows users to bundle together a set of transactions, thereby specifying the order in which they are executed. The bundles can also include public transactions, propagated over the public P2P network. A public transaction that buys a coin on a DEX can, for example, lead to an arbitrage opportunity [35]. A user can include this transaction in a bundle along with one of their own to capture this arbitrage opportunity. The last transaction in the bundle usually pays the miner (based on the profit made) in ether via a direct transfer (i.e., *coinbase transfer*) to their wallet addresses. This essentially means that miners are being offered different prices for mining the same transaction. In other words, miners have a financial incentive for including transactions that are in a bundle at the top of a block, even though the public fee offered through gas price in the transaction data is very low. Hence, each transaction in the bundle has a normal gas price and a *bundle gas price*, which is calculated using the total gas used by all transactions in the bundle and the total miner reward for mining the bundle.

Bundling Public Transactions. To identify bundles with transactions that were probably sent through the public P2P network, we rely on a simple heuristic. Specifically, we focus on transaction bundles of size 2 and 3, and search for transactions that have likely resulted in a publicly sent transaction being bundled. Then, we find bundles issued from different issuers that include a zero and non-zero *max-priority fee*³ transactions. The intuition is that miners have no incentive to include transactions that offer a zero max-priority fee, as they receive no rewards for mining these transactions. Unless they receive extra payment (through Flashbots coinbase transfer). Hence, transactions that have a non-zero max-priority fee were likely sent publicly.

For transaction bundles of size 2, we look for transactions whose issuers are not the same. Furthermore, we look for cases where the first transaction offers a non-zero max-priority fee, with no coinbase transfer to the miner, and the second transaction offers a 0 max-priority fee and a non-zero coinbase transfer.

For transaction bundles of size 3, we look for signs of sandwich attacks [34]. We look for bundles where the first and last transactions have the same issuer, but the second transaction has a different issuer. Additionally, we check that the first and third transactions offer a 0 max-priority fee, meaning that the miner receives no reward from the gas price for mining these transactions. Then, we ensure that the second transaction offers miners a non-zero max-priority fee, while the third offers miners a fee through direct coinbase transfer. This scenario might be a classic sandwich attack, where public transactions are bundled between two private transactions, sent by the same issuer, and the miner gets paid via a coinbase transfer from the third transaction [34].

³ The *max-priority fee* was introduced in EIP-1559 as the unique financial incentive miners get for including publicly announced transactions. The other fees are burned.

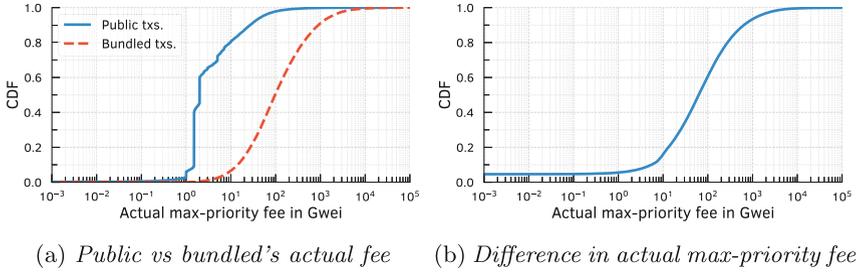


Fig. 3. Diff. between the actual max-priority fee of public transactions and Flashbots bundles; bundles typically offer a larger *effective* fee to the miners.

We found 853,394 transactions in 426,697 bundles of length 2, and 1,231,695 transactions in 410,565 bundles of length 3. From those, we found that 110,401 (25.87%) and 37,447 (9.12%) bundles, of lengths 2 and 3, respectively, fit our heuristic. We then calculate the *actual max-priority fee* for these bundles, as the total gas used by all transactions in the bundle divided by the total miner reward (from gas usage and coinbase transfer). Figure 3 shows the price difference miners get for including publicly and bundled transactions. Note that around 40% of transactions differ in the actual max-priority fee by 100 gwei-per-units-of-gas. Flashbots bundles offers much higher gas prices in comparison to the public announced max-priority fee alone.

Towards Liquidations Through Bundling. Lending protocols rely on *over-collateralization* of assets: In order to borrow assets from these protocols, a user has to deposit a collateral of at least 150% of the borrowed amount. To borrow 1 USDC on AAVE, for example, a user would have to collateralize at least 1.5 USDC worth of another asset (e.g., in ETH or BTC). If the ratio of the collateral asset versus the borrowed asset falls below 1.5, the user’s position can be liquidated by any other participant until the ratio stabilizes to 1.5 again. The liquidator then pays back a portion of the user’s debt to receive the collateral asset at a discount. In order to assess an asset’s on-chain value, lending protocols rely on oracle services, e.g., Chainlink Data Feeds. In the case of the two largest lending platforms, AAVE V2 and Compound, for instance, Chainlink provides the price of each asset in ETH and USD, respectively.

We found 16,418 liquidations in AAVE and 6387 liquidations in Compound. Out of these, there were 4863 AAVE liquidations and 2036 Compound liquidations that were sent privately through Flashbots. In AAVE, the three largest collateral assets that were liquidated were WETH (57.58%), LINK (11.84%), and WBTC (8.99%). The debt assets paid for, i.e., the assets borrowed by the users, were USDC (33.77%), USDT (22.27%), DAI (19.39%), and GUSD (5.12%), all of which are stablecoins and account for over 80% of the assets repaid by liquidators. In Compound, the three largest collateral assets that were liquidated were WETH (69.7%), WBTC (10.31%), and UNI (5.5%). The debt assets were

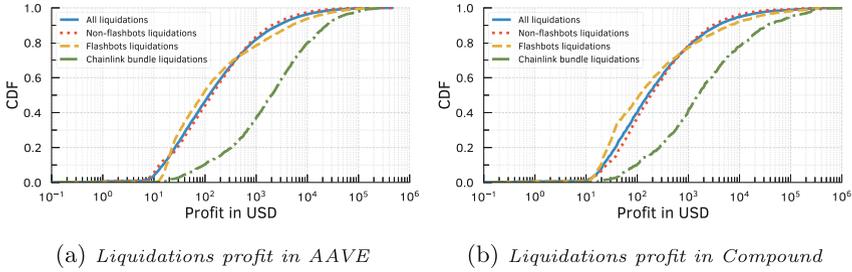


Fig. 4. Profits of liquidators in (a) AAVE and in (b) Compound. Liquidations bundled with Chainlink updates generally provide higher profits.

USDC (38.9%), DAI (30.45%), USDT (23.38%), and TUSD (2.7%), all of which are stablecoins and account for over 90% of the assets repaid by liquidators.

Liquidation with Bundled Oracle Updates. To check the adverse effect of bundling oracle updates, we looked at bundles with Chainlink oracle updates as they are a key part of liquidations. We identified 1165 AAVE liquidations distributed within 1154 bundles (2662 transactions including 1301 oracle updates) that contained at least one oracle update. In Compound, we found 648 liquidations distributed within 641 bundles (1457 transactions including 751 oracle updates) that contained oracle updates. In AAVE, out of 1154 bundles, there were 994 (86.14%) bundles that contained an oracle update followed by a liquidation, and 52 (4.51%) with two oracle updates followed by liquidations. In Compound, out of 641 bundles, there were 548 (85.49%) bundles that contained an oracle update followed by a liquidation, and 39 (6.08%) with two oracle updates followed by liquidations. Out of the total 1813 liquidations in AAVE and Compound we found that only 24 were possible in the previous block. Almost 98.68% of such liquidations were, hence, only possible because of the Chainlink updates in that block.

In order to calculate the profit made by the liquidators, we get the amount of debt that was repaid and the amount of the underlying collateral that was received by the liquidator. We calculate the price of each token at the time of liquidation by looking at the on-chain oracle price from Chainlink at the same block number, where the liquidation took place. For AAVE and Compound, we specifically use the Chainlink on-chain price used by AAVE and Compound in their respective protocols. AAVE uses the price in ETH as a reference for its tokens, whereas Compound’s price oracles are denominated in USD. For AAVE, in order to calculate the profit made by each liquidation, we calculate the profit in ETH, and then multiply the profit by the current Chainlink on-chain price of ETH in USD. Per Fig. 4, liquidations that are bundled with a Chainlink update also have larger profits for liquidators, which implies that the lucrative liquidations are more likely to be bundled together with a Chainlink update.

Characterizing Transaction Bundling. To investigate which DEXes protocols are called within Flashbots bundles, we focus on the following contract

calls: 0x Protocol, Balancer, Bancor, Curve, SushiSwap, and Uniswap V1 and V3. In our set of 3,284,886 Flashbots bundles, we find that 2,231,051 (67.92%) unique Flashbots bundles (and 3,076,760 transactions) called at least one of these contracts. Table 1 shows the distribution of the number of transactions and the number of bundles for each of these contracts. We see that Uniswap and SushiSwap are the most bundled DEXes protocols in Flashbots.

4.2 Side Channel (dark-Fee) Payments and Transaction Acceleration

We now focus on the Bitcoin blockchain to study dark-fees transactions.

Prevalence of Transaction Acceleration. Dark-fee transactions (or accelerated transactions) are transactions that offer additional fees to specific mining pools via an opaque and non-public side-channel payment [29]. Messias *et al.* show that in Bitcoin the top 5 mining pools, BTC.com [4], AntPool, ViaBTC [45], F2Pool, and Poolin, deploy transaction acceleration services, which enables users to “accelerate” the confirmation of their transactions by offering mining pools dark-fees [29]. These (dark-)fees are paid in fiat currency through a direct bank transfer or via other crypto coins to the mining pool. They are, therefore, opaque or dark to other participants. Strangely enough, these fees are also non-refundable as the miner receives them regardless of whether they include the transaction in a block or not—a guaranteed payment. The fees paid by the transaction issuer are, furthermore, not made public: only the user and the miner knows the actual fee paid by the transaction inclusion. Since transaction issuers pay the fees off-chain, miners have an incentive for prioritizing these transactions despite the low fee rate offered on-chain. It also implies that the transaction issuer offers a miner a different fee compared to that offered to other miners for including their transaction in a block. Miners do not disclose such private fees paid by issuers. This behavior is different from that of Flashbots in Ethereum: The latter discloses the final dark-fee after the transaction is committed (see Sect. 4.1).

Characterizing Transaction Acceleration. In order to detect accelerated transactions, Messias *et al.* [29] proposed a metric called *signed position prediction error (SPPE)* and *position prediction error (PPE)*. The idea behind these measures is that transactions that have been accelerated through off-chain fees are likely to have been “misplaced” in a block based on the on-chain fee they

Table 1. There are 2,231,051 (67.92%) unique Flashbots bundles, and 3,076,760 (44.35%) transactions, that called the following decentralized exchange contracts in Ethereum: 0x Protocol, Balancer, Bancor, Curve, SushiSwap, Uniswap V1, or V3. Note that a single transaction or bundle might call one or more contracts.

	Balancer	Bancor	Curve v1 & v2	Uniswap v2 & Sushiswap	Uniswap v3	0x Protocol v1, v2 & v3	Total
# of bundles	85,4223.83%	96,1224.31%	53,2962.39%	1,710,98576.69%	1,337,71559.96%	28,7531.29%	2,231,05167.92%
# of transactions	87,8652.86%	99,0403.22%	58,1881.89%	2,533,08482.33%	1,692,48555.01%	29,1000.95%	3,076,76044.35%

offer. Figure 5 shows that the top-6 mining pools in our Bitcoin data set engage in transaction acceleration. Large SPPE values imply that a transaction that should have been included at the bottom is included at the top of the block, confirming acceleration. We rely on this methodology to infer transaction acceleration in Bitcoin and present our data set and findings below.

To identify accelerated transactions, we gathered all Bitcoin blocks mined from Jan. 1st 2018 to Dec. 31st 2020. In total, there are 161,954 blocks from block height 501,951 to 663,904, and 313,575,387 transactions. In Bitcoin, mining pools may indicate their ownership of the block by including a *signature* or *marker* in the *Coinbase* transaction (i.e., the first transaction of every block). We used such markers for identifying the mining pool (owner) of each block following techniques from prior work [22, 29, 36]. We failed to identify, however, the owners of 4911 blocks (approximately 3% of the blocks) and grouped these blocks under the label “Unknown.” Figure 1b shows the distribution of the count of blocks mined and transactions confirmed by the top-20 mining pools. We further removed 65,902,514 (21.02%) *child-pays-for-parent* (CPFP) transactions from our acceleration analyses.

To estimate the prevalence of accelerated transactions in blocks mined by different mining pools, we compute the fraction of blocks mined by the top-15 mining pools, based on their hash rates in our data set (refer to Fig. 1b), that contained transactions with $SPPE \geq 99\%$. Per Fig. 6, we find that many large mining pools such as BTC.com, F2Pool, and ViaBTC are likely including accelerated transactions in a sizeable fraction of their mined blocks, with ViaBTC including it in over 40% of their blocks.

If we consider all mining pools’ transactions with an $SPPE \geq 50\%$ (1,869,043 transactions, in total), from 2018 to 2020, users transferred in total 11,631,217 BTC (or ≈ 223.55 billion USD⁴). The accelerated transactions accounted for 240,226 BTC (or ≈ 4.62 billion USD), corresponding to approximately 2.07%.

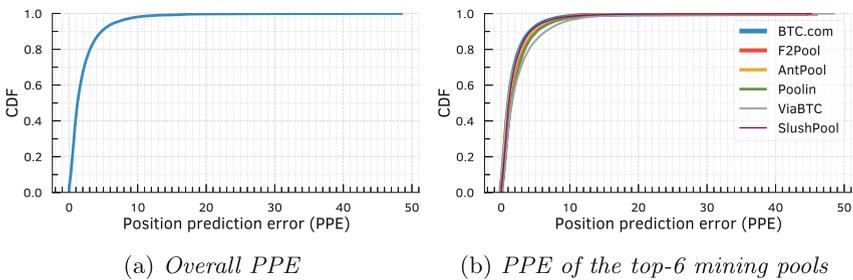


Fig. 5. Bitcoin position prediction error (PPE). (a) There are 160,962 blocks with non-CPFP txs; 80% of all blocks has PPE less than 3.06% (mean is 2.09% and std. deviation is 2.75.). (b) PPEs of top-6 mining pools per their normalized hash rate, showing that all large mining pools engage in transaction acceleration.

⁴ Based on the Bitcoin exchange rate on October 19th 2022, 1 BTC = 19,219.90 USD.

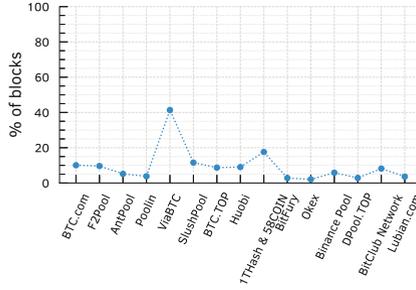


Fig. 6. Blocks with accelerated transactions (with SPPE \geq 99%) are quite common among the top 15 mining pools. In Bitcoin, the mining pools with a high percentage of such blocks are ViaBTC (41.36%), 1THash & 58COIN (17.58%), SlushPool (11.58%), BTC.com (10.03%), and F2Pool (9.63%).

Aggregated Power of Colluding Miners. In order to check the impact of transactions acceleration services on commit time of transaction, we ran active real-world experiments. Specifically, we paid ViaBTC [45] to accelerate selected transactions (see Table 5 in Appendix 2) during periods of high congestion between November 26th and December 1st, 2020. From 10 Mempool snapshots during this period, we selected transactions that offered a very low fee-rate (i.e., 1–2 sat-per-byte) for acceleration. To keep our acceleration costs low, we selected transactions with the smallest size (which was 110 bytes) within this set. For each of the 10 snapshots, we had multiple transactions with such low fee-rates and small size, for a total of 212 transactions across all the snapshots. We randomly selected one transaction from each snapshot (i.e., 10 transactions) and paid ViaBTC 205 EUR to accelerate them.

Table 2. Accelerated transactions have fewer delays and are included at the top of the block, i.e., at higher positions compared to non-accelerated transactions.

<i>metrics</i>	<i>delay in # of blocks</i>		<i>perc. position in a block</i>	
	<i>acc.</i>	<i>non-acc.</i>	<i>acc.</i>	<i>non-acc.</i>
minimum	1	9	0.07	17.47
25-perc	1	148	0.08	75.88
median	2	191	0.09	87.92
75-perc	2	247	0.20	95.00
maximum	3	326	4.39	99.95
average	1.8	198.5	0.79	84.46

We then compare the priority with which the accelerated transactions and the 202 (= 212 – 10) non-accelerated transactions with similar fee rates and sizes were included in the Bitcoin blockchain. The impact of acceleration was

strikingly apparent as shown in Table 2. All 10 accelerated transactions were included within 1–3 blocks after their acceleration, with an average delay of 1.8 blocks. In contrast, the minimum delay for the 202 non-accelerated transactions of comparable fee-rates and sizes was 9 blocks, with an average delay of 198.5 blocks. Interestingly, 38 of the non-accelerated transactions were yet to be included in the blockchain by December 4th, 2020. Similarly, the accelerated transactions were included in top 0.07–4.39 percentile positions, with an average 0.79 percentile position, while the non-accelerated transactions were included in the beyond top 17.47–99.95 percentile positions, with an average 84.46 percentile position. From the above observations, it is clear that the transactions we accelerated were included with high priority, meaning Bitcoin mining pools take off-chain fees into account when prioritizing transactions.

Although, we accelerated our transactions using ViaBTC mining pool, our 10 transactions were included by 5 different mining pools, namely F2Pool, AntPool, Binance, Huobi, and ViaBTC. As we accelerated transaction during time of high congestion in Bitcoin, no mining pool would have included a transaction offering 1–2 sat-per-byte, unless they were accelerated. Since we only paid the ViaBTC mining pool, this implies that ViaBTC is colluding with other mining pools to accelerate transactions that offer off-chain fees. Except for Binance, all these colluding pools rank amongst the top-8 mining pools in terms of their hash rates at the time of our experiments. Table 3 shows the individual as well as the combined hash rates of these 5 colluding mining pools over the last day, last week, and last month before the conclusion of our experiment on December 1st, 2020. The most striking and the most worrisome fact is that the combined hash rates of these colluding mining pools exceeds 55% of the total Bitcoin hash rate. Additionally, if mining pools are colluding to include accelerated transactions, then they might also potentially collude in malicious ways.

Table 3. If we rank the miners who confirmed the accelerated transactions based on their daily, weekly, and monthly hash rate power, at the time these experiments were conducted, the combined hash power of these mining pools exceeds 55% of the Bitcoin’s total hashing power.

<i>Mining Pool</i>	<i>Hash-rate</i>		
	<i>last 24h</i>	<i>last week</i>	<i>last month</i>
F2Pool	19.9%	18.7%	19.9%
AntPool	12.5%	10.6%	10.2%
Binance	9.6%	10.3%	10.0%
Huobi	8.1%	9.3%	9.8%
ViaBTC	5.1%	7.1%	7.7%
<i>Total</i>	55.2%	56%	57.6%

Furthermore, due to the lack of transparency into their queue, miners can charge higher prices for their acceleration services when colluding. It means that they can overcharge the transaction issuers for including their transactions.

5 Concluding Discussion

In this section, we discuss the implications of our findings regarding the lack of transparency in transaction contention and prioritization. We also argue why our findings and implications would be relevant even in the face of recent changes to blockchain protocols, e.g., Ethereum Improvement Protocol (EIP) 1559 and the Ethereum Paris Network Upgrade (a.k.a. the Merge).

Implications for Publicly Mined Transactions. Most wallet software and crypto-exchanges today rely on reconstructing the current public Mempool state in order to suggest a suitable fee to transaction issuers. With the lack of contention and prioritization transparency, transaction issuers can no longer accurately recreate the current Mempool state for different miners. Consequently, they cannot reliably estimate the fees transactions need to pay for their desired prioritization. Worse, as the fraction of privately mined and accelerated transactions keeps rising, the transaction fees will become less (reliably) predictable in the future.

Implications for Privately Mined Transactions. The problem of reliable fee estimation for a desired level of prioritization is even worse for privately mined transactions that are announced on private relay networks. When transaction issuers announce on a private relay network today, they are often unsure what fraction of total hash rate is controlled by the miners listening to the private relay network. It is important to estimate the hash rate controlled by private mining pools to estimate the commit (waiting) times for transactions. Furthermore, transaction issuers on private relay networks are completely blind to other competing transactions. This opacity allows miners offering private mining and transaction acceleration services to overcharge and demand exorbitant fees to commit transactions. For example, in the Ethereum blockchain, users are observed to be overcharged by miners for having their transactions confirmed with high priority through Flashbots bundles [46].

Relevance of Findings in Light of EIP-1559 and the Ethereum Merge. Our observations about the lack of transparency and their implications are fundamental to the current blockchain architectures and hold both before and after the recent major improvements to blockchains, e.g., EIP-1559 and the Ethereum Merge. While EIP-1559 attempts to improve the estimation of transaction fees that need to be offered, it does not address the problems associated with the lack of transaction contention and prioritization transparency. Similarly, after the Ethereum Merge, *validators* that stake a certain amount of ETH rather than *miners* would be responsible for selecting and validating transactions to include in the next block [13]. Our observations about private mining would still hold for private validation and the implications would still be valid after the Merge.

In conclusion, our work shows that with private mining and accelerated transactions, the promise of the public decentralized blockchain does not hold. Firstly, mining pools with combined hash rates of over 50% are colluding with each other, showing a centralization in the system. Then, they can also censor certain transactions, breaking the ethos of decentralized public blockchains with no central authorities. Second, it breaks the assumption that all activities in the blockchain are transparent. Although this is true for transactions included in the blockchain, prioritization of transactions is becoming more opaque with the rise of private mining and off-chain fees. Hence, we make the case that to fulfill the transparency promise of public blockchains, prioritization of transactions should be transparent as well. Third, with private mining in Ethereum, Flashbots is increasingly being used for malicious and predatory activities such as sandwich attacks, which essentially levies a tax on users interacting with financial institutions on the blockchain (e.g., in DEX). These concerns need to be addressed if public blockchains are going to live up to their promises.

Acknowledgments. This research was supported in part by a European Research Council (ERC) Advanced Grant “Foundations for Fair Social Computing”, funded under the European Union’s Horizon 2020 Framework Programme (grant agreement no. 789373). It was also supported by MIAI @ Grenoble Alpes (ANR-19-P3IA-0003) and by the French National Research Agency under grant ANR-20-CE23-0007.

Appendix 1 Ethereum Private Transaction Experiment

We conducted 4 active experiments where we issued 8 Ethereum transactions; half issued publicly and the other half privately through a private-channel network known as Taichi Network [40]. Table 4 summarizes the transactions in our experiment. Spark Pool and Babel Pool included all private transactions (2 transactions each) sent directly to these miners through Taichi Network.

Table 4. We conducted 4 active experiments in Ethereum by simultaneously accelerating transactions privately and publicly via Taichi Network. Private transactions were included only by Spark Pool and Babel Pool. If we rank these mining pools according to their hash-rate, they account for 27.72% of the total Ethereum hash-rate.

#	type	tx hash	block number	miner	tx. position per # of txs.	block delay (in blocks)	fee paid (in Ether)	base fee (Gwei)	max fee (Gwei)	max priority fee (Gwei)	gas price (Gwei)	block timestamp in UTC
1	public	bbe88e...a4f000	13,183,516	Nanopool	305/336	1	0.00190489	88.98082939	116.52835749	1.72836605	90.70919543	2021-09-08 06:39:18
	private	c46b75...ead538	13,183,520	Babel Pool	29/39	5	0.00225209	105.51391459	120.56586232	1.72836605	107.24228063	2021-09-08 06:40:29
2	public	6d994f...c1aadd	13,183,561	Binance	209/213	2	0.00244137	114.95482846	137.64014705	1.30100683	116.25583529	2021-09-08 06:49:26
	private	a4d4ae...42eb05	13,183,565	Spark Pool	294/296	6	0.00240978	113.45059961	137.64014705	1.30100683	114.75160643	2021-09-08 06:50:12
3	public	725743...0afc45	13,183,634	Unknown	124/126	2	0.00263298	123.27216185	135.21393222	2.10805685	125.38021870	2021-09-08 07:06:31
	private	2b8ecc...15cdf1	13,183,635	Spark Pool	321/340	3	0.00257468	120.49562077	135.21393222	2.10805685	122.60367762	2021-09-08 07:06:44
4	public	e21695...2e1574	13,183,679	Ethermine	280/302	13	0.00223433	104.69510748	108.95262574	1.70164453	106.39675202	2021-09-08 07:18:37
	private	4e482b...87c76f	13,183,690	Babel Pool	150 / 212	24	0.00179917	83.97323655	108.95262574	1.70164453	85.67488108	2021-09-08 07:20:12

Appendix 2 Bitcoin Transaction Acceleration Experiment

Table 5. We conduct 10 transaction acceleration experiments in Bitcoin. If we rank the miners whose included these transactions based on their daily hash-rate power as (D) and weekly hash-rate power as (W), together these mining pools corresponds to a hash-rate power of (D: 55.2%; W: 56%).

<i>txid</i>	<i>block height</i>	<i>miner</i>	<i>tx. position</i>	<i>delay (in blocks)</i>	<i>acc. cost (BTC)</i>	<i>vsize (byte)</i>	<i>fee rate sat-per-vsize</i>	<i>Mempool # of txs. vsize (MB)</i>		<i>timestamp in UTC</i>
35b18e...52dbc1	658,805	Huobi	2 nd	2	0.001254	110	2	36,644	44.63	2020-11-26 19:10
65765c...baede2	658,898	F2Pool	73 rd	1	0.001254	110	2	20,998	32.55	2020-11-27 11:06
0c2098...29fbfd	658,912	AntPool	2 nd	2	0.001254	110	1	30,126	38.01	2020-11-27 13:38
1515a7...179af3	658,971	Binance	2 nd	3	0.001254	110	1	25,922	37.89	2020-11-27 21:55
48a0a5...0ddaec	659,335	ViaBTC	3 rd	1	0.001045	110	1	15,605	9.82	2020-11-30 10:09
9a17cf...E3734c	659,341	Huobi	2 nd	2	0.001045	110	1	14,945	9.41	2020-11-30 10:28
831b24...95d421	659,351	AntPool	2 nd	1	0.001045	110	1	10,990	8.66	2020-11-30 12:22
1f59bf...47096c	659,355	F2Pool	111 th	3	0.001045	110	1	17,093	11.40	2020-11-30 12:58
6942e0...8c06c3	659,362	Huobi	2 nd	2	0.001045	110	1	30,836	19.06	2020-11-30 14:49
8e49e2...ae825f	659,481	ViaBTC	6 th	1	0.001254	110	2	30,935	22.59	2020-12-01 10:40

We ran an active Bitcoin transaction acceleration experiment where we paid 205 EUR to ViaBTC [45] to accelerated 10 transactions from 10 different snapshots of our Mempool. To select these transactions, we checked whether the Mempool was congested (i.e., having more transactions waiting for inclusion than the next block would be able to include), with its size being at least 8 MB. Then, we considered only transactions with low fee rates—less than or equal to 2 sat-per-byte—to ensure that these transactions would be highly unlikely to be included soon in a subsequent block. Next, we sorted the remaining transactions by size to limit the experiment cost as the acceleration-service costs grow proportional to the transaction size. Finally, we select the transaction with the smallest size in bytes for our active experiment.

Table 5 summarizes the transactions used in our experiment. Most of these 10 accelerated transactions were included nearly in the next block, demonstrating the acceleration efficiency. Also, these transactions were wrongly positioned in the block: They appeared, for instance, at the top of the block, i.e., higher than the non-accelerated transactions, showing that miners indeed prioritized them (see Table 2). Further, we observed that although we had only accelerated transactions via ViaBTC, other top mining pools were also involved in confirming the accelerated transactions.

References

1. Antonio Pierro, G., Rocha, H., Tonelli, R., Ducasse, S.: Are the gas prices oracle reliable? a case study using the ethgasstation. In: 2020 IEEE International Workshop on Blockchain Oriented Software Engineering (IWBOSE) (2020)
2. Babel Finance: Economic Daily: Babel Finance Launches Ethereum Mining Pool (2021)

3. Basu, S., Easley, D., O'Hara, M., Siner, E.G.: Towards a Functional Fee Market for Cryptocurrencies. CoRR abs/1901.06830 (2019)
4. BTC.com: BTC.com Transaction Accelerator (2022). <https://pushtx.btc.com>
5. Buterin, V., Conner, E., Dudley, R., Slipper, M., Norden, I., Bakhta, A.: "EIP-1559: Fee market change for ETH 1.0 chain. Ethereum Improvement Proposals (2019)
6. Carlsten, M., Kalodner, H., Weinberg, S.M., Narayanan, A.: On the instability of bitcoin without the block reward. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. CCS 2016 (2016)
7. Chen, X., Papadimitriou, C., Roughgarden, T.: An axiomatic approach to block rewards. In: Proceedings of the 1st ACM Conference on Advances in Financial Technologies. AFT 2019 (2019)
8. Piatt, C., Quesnelle, J., Sheridan, C.: Eden Network (2021)
9. Daian, P., et al.: Flash boys 2.0: frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In: 2020 IEEE Symposium on Security and Privacy (SP) (2020)
10. Easley, D., O'Hara, M., Basu, S.: From mining to markets: the evolution of bitcoin transaction fees. *J. Financ. Econ.* (2019)
11. EigenPhi: EigenPhi Crypto & DeFi Analytics (2022). <https://eigenphi.io>
12. Eskandari, S., Moosavi, S., Clark, J.: SOK: transparent dishonesty: front-running attacks on blockchain. In: Bracciali, A., Clark, J., Pintore, F., Rønne, P.B., Sala, M. (eds.) *Financial Cryptography and Data Security* (2020)
13. Ethereum Foundation: Proof-of-Stake (PoS) (2022). <https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/>
14. Ethereum Foundation: The Merge (2022). <https://ethereum.org/en/upgrades/merge/>
15. Ethermine: Ethermine MEV-Relay (2022). <https://ethermine.org/mev-relay>
16. Eyal, I., Siner, E.G.: Majority is not enough: bitcoin mining is vulnerable. *Commun. ACM* (2018)
17. Fiat, A., Karlin, A., Koutsoupias, E., Papadimitriou, C.: Energy equilibria in proof-of-work mining. In: Proceedings of the 2019 ACM Conference on Economics and Computation (EC 2019) (2019)
18. Flashbots: Flashbots Docs (2022). <https://docs.flashbots.net>
19. Goren, G., Spiegelman, A.: Mind the mining. In: Proceedings of the 2019 ACM Conference on Economics and Computation. EC 2019 (2019)
20. Partz, H.: Second-largest Ethereum mining pool to suspend all operations. *Coin-telegraph* (2021)
21. Huberman, G., Leshno, J.D., Moallemi, C.: Monopoly without a monopolist: an economic analysis of the bitcoin payment system. *Rev. Econ. Stud.* **88**, 3011–3040 (2021)
22. Judmayer, A., Zamyatin, A., Stifter, N., Voyiatzis, A.G., Weippl, E.: Merged mining: curse or cure? In: *Data Privacy Management, Cryptocurrencies and Blockchain Technology* (2017)
23. Kiayias, A., Koutsoupias, E., Kyropoulou, M., Tselekounis, Y.: Blockchain mining games. In: Proceedings of the 2016 ACM Conference on Economics and Computation (EC 2016) (2016)
24. Kiffer, L., Levin, D., Mislove, A.: Stick a fork in it: analyzing the Ethereum network partition. In: Proceedings of the 16th ACM Workshop on Hot Topics in Networks (2017)
25. Lavi, R., Sattath, O., Zohar, A.: Redesigning bitcoin's fee market. In: *The World Wide Web Conference. WWW 2019* (2019)

26. Liu, F., Wang, X., Li, Z., Xu, J., Gao, Y.: Effective GasPrice prediction for carrying out economical Ethereum transaction. In: 2019 6th International Conference on Dependable Systems and Their Applications (DSA) (2020)
27. Mars, R., Abid, A., Cheikhrouhou, S., Kallel, S.: A machine learning approach for gas price prediction in Ethereum blockchain. In: 2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC) (2021)
28. Messias, J., Alzayat, M., Chandrasekaran, B., Gummadi, K.P.: On blockchain commit times: an analysis of how miners choose bitcoin transactions. In: KDD Workshop on Smart Data for Blockchain and Distributed Ledger. SDBD 2020 (2020)
29. Messias, J., Alzayat, M., Chandrasekaran, B., Gummadi, K.P., Loiseau, P., Mislove, A.: Selfish & opaque transaction ordering in the bitcoin blockchain: the case for chain neutrality. In: Proceedings of the ACM Internet Measurement Conference (IMC 2021) (2021)
30. Messias, J., Pahari, V., Chandrasekaran, B., Gummadi, K.P., Loiseau, P.: Data sets and scripts used to analyze the contention and prioritization transparency in both Bitcoin and Ethereum blockchains (2023). <https://github.com/johnnatan-messias/blockchain-transaction-ordering>
31. Nakamoto, S.: Bitcoin: A Peer-to-Peer Electronic Cash System (2008)
32. Noda, S., Okumura, K., Hashimoto, Y.: An economic analysis of difficulty adjustment algorithms in proof-of-work blockchain systems. In: Proceedings of the 21st ACM Conference on Economics and Computation (EC 2020) (2020)
33. Pass, R., Seeman, L., Shelat, A.: Analysis of the blockchain protocol in asynchronous networks. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10211, pp. 643–673. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_22
34. Qin, K., Zhou, L., Gervais, A.: Quantifying blockchain extractable value: how dark is the forest? In: 2022 IEEE Symposium on Security and Privacy (SP) (2022)
35. Qin, K., Zhou, L., Livshits, B., Gervais, A.: Attacking the DeFi ecosystem with flash loans for fun and profit. In: Financial Cryptography and Data Security. FC 2021 (2021)
36. Romiti, M., Judmayer, A., Zamyatin, A., Haslhofer, B.: A deep dive into bitcoin mining pools: An empirical analysis of mining shares. In: Workshop on the Economics of Information Security. WEIS 2019 (2019)
37. Roughgarden, T.: Transaction fee mechanism design for the Ethereum blockchain: an economic analysis of EIP-1559. In: Proceedings of the 2021 ACM Conference on Economics and Computation. EC '21 (2021)
38. Siddiqui, S., Vanahalli, G., Gujar, S.: Bitcoinf: achieving fairness for bitcoin in transaction fee only model. In: Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems. AAMAS 2020 (2020)
39. Sompolinsky, Y., Zohar, A.: Secure high-rate transaction processing in bitcoin. In: Böhme, R., Okamoto, T. (eds.) Financial Cryptography and Data Security. FC 2015 (2015)
40. SparkPool: Taichi Network (2021). <https://taichi.network>
41. Strehle, E., Ante, L.: Exclusive mining of blockchain transactions. In: Scientific Reports 2020-Conference Proceedings of the Scientific Track of the Blockchain Autumn School 2020 (2020)
42. Torres, C.F., Camino, R., State, R.: Frontrunner jones and the raiders of the dark forest: an empirical study of frontrunning on the Ethereum blockchain. In: 30th USENIX Security Symposium (2021)
43. Tsabary, I., Eyal, I.: The gap game. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. CCS 2018 (2018)

44. Turksonmez, K., Furtak, M., Wittie, M.P., Millman, D.L.: Two ways gas price oracles miss the mark. In: 2021 IEEE International Conference on Omni-Layer Intelligent Systems (COINS) (2021)
45. ViaBTC: Transaction Accelerator (2022). <https://www.viabtc.com/tools/txaccelerator/>
46. Weintraub, B., Torres, C.F., Nita-Rotaru, C., State, R.: A flash(bot) in the pan: measuring maximal extractable value in private pools. In: Proceedings of the ACM Internet Measurement Conference (IMC 2022) (2022)
47. Wood, G., et al.: Ethereum: a secure decentralised generalised transaction ledger. Ethereum project yellow paper (2014)
48. Zhang, R., Preneel, B.: Lay down the common metrics: evaluating proof-of-work consensus protocols' security. In: Proceedings - IEEE Symposium on Security and Privacy (2019)
49. Zhou, L., Qin, K., Torres, C.F., Le, D.V., Gervais, A.: High-frequency trading on decentralized on-chain exchanges. In: 2021 IEEE Symposium on Security and Privacy (SP) (2021)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

