

Exploiting Ethereum after “The Merge”: The Interplay between PoS and MEV Strategies

Davide Mancino*, Alberto Leporati, Marco Viviani and Giovanni Denaro

University of Milano-Bicocca, Department of Informatics, Systems, and Communication, Edificio U14 (ABACUS), Viale Sarca, 336 – 20126, Milan, Italy

Abstract

On September 15, 2022, Ethereum changed its consensus algorithm, moving from Proof-of-Work (PoW) to Proof-of-Stake (PoS). This event is commonly known as “The Merge”. While this change has considerably reduced the overall energy consumption, it has been observed that 40% of the first 1,000 blocks since then have been validated by a very limited amount of cryptocurrency holders. This raises serious questions about the decentralized nature of the system. To address this concern, and to mitigate the potential negative externalities of current Maximal Extractable Value (MEV) extraction strategies, the community has agreed to implement a permissionless, transparent, and fair ecosystem for MEV extraction. In this paper, after describing the new ecosystem, we perform an analysis of the blocks validated after Ethereum’s “The Merge”. In particular, we analyze the actual distribution of validators by accounting for their activity both inside and outside the MEV ecosystem, and discuss how the scenario has changed in recent months.

Keywords

Blockchain, Proof-of-Stake, Ethereum, MEV, Transparency

1. Introduction

Ethereum is a decentralized, open-source blockchain with smart contract functionality [1]. Initially based on a *Proof-of-Work* (PoW) consensus mechanism, on September 15, 2022, it has transitioned to *Proof-of-Stake* (PoS), after merging with a separate blockchain called *Beacon Chain*. This update, known as “The Merge” [2], gives the opportunity to analyze in detail the actual activity of validators in the blockchain after this consensus mechanism change. To this aim, let us briefly recall the main characteristics and the operation of the two mechanisms.

PoW generates consensus and guarantees network security by combining computational power with cryptography, in the form of *block mining* [3]. Each node that wants to participate in mining, i.e., a *miner*, must solve a computationally difficult problem, i.e., a *work*, to ensure the validity of the newly mined block. The protocol is considered fair in the sense that a miner holding a percentage p of the total computational power can create a block with probability p . On

ITASEC 2023: The Italian Conference on CyberSecurity, May 03–05, 2023, Bari, Italy


*Corresponding author.

✉ d.mancino1@campus.unimib.it (D. Mancino); alberto.leporati@unimib.it (A. Leporati); marco.viviani@unimib.it (M. Viviani); giovanni.denaroi@unimib.it (G. Denaro)

🆔 0000-0002-9559-8476 (D. Mancino); 0000-0002-8105-4371 (A. Leporati); 0000-0002-2274-9050 (M. Viviani); 0000-0002-7566-8051 (G. Denaro)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

the one hand, this makes it very expensive to attack a cryptocurrency's network implementing PoW, because an attacker should solve the same tasks as the rest participants of the PoW-secured network. On the other hand, PoW is very inefficient from the point of view of computational resource use, with also the consequence that it tends to consolidate miners down to the few people who can afford the necessary equipment for computation.

To solve these problems, PoS was introduced, which instead of performing tangible work, relies on the existence of verifiable *stake* in the ecosystem. A user must basically prove to own a particular amount of cryptocurrency tokens that are native to the blockchain. The probability p to create a block and receive the associated reward is proportional to the amount p of the tokens put in stake. The hypothesis behind the proper functioning of this mechanism is that the users with the highest stakes in the system have the greatest interest in maintaining the network secure, as they would suffer the most if the reputation and price of cryptocurrency decreased due to attacks. Furthermore, if users act maliciously, they may lose their stake as a result of their actions. One possible disadvantage of PoS concerns the fact that the network may, in fact, be subject to the dominance of the most important token holders. This is due to the so-called *compounding of wealth* phenomenon, by which the rich simply stake their wealth and get richer with stake rewards [4].

So far, the scientific literature has investigated the behavior of the Ethereum blockchain before "The Merge", addressing several issues about mining power [5] and wealth distribution [6]. However, after the major update of September 15, 2022, numerous online newspapers and magazines in the crypto community pinpointed how a very limited amount of token holders validated more than 40% percent of the first 1,000 blocks, ultimately questioning the decentralization of Ethereum after "The Merge". This raises concerns about the actual decentralized nature of the Ethereum network [7, 8, 9]. Moreover, since the PoS consensus protocol entails that the user in charge to propose the next block is known two epochs in advance with respect to the epoch in which the block will be eventually added, there are some potential risks related to the proliferation of actors in the blockchain implementing *Maximal Extractable Value* (MEV) strategies, i.e., acting on including, excluding or, rearranging transactions to obtain additional value in terms of cryptocurrency. To mitigate the potential negative impact of private MEV strategies on the Ethereum blockchain (e.g., increased centralization) a research and development organization known as *Flashbots* has recently started a set of initiatives with the main goal of enabling a permissionless, transparent, and fair ecosystem for MEV extraction [10].

In light of these issues, since (to the best of our knowledge) no in-depth scientific analysis of the decentralization and economic effects of "The Merge" and the introduction of these new MEV-related players in the Ethereum ecosystem has been published in the literature, the purpose of this article is to analyze the current situation of the network and whether the proposed solutions are actually working to not favor the few at the expense of the many. In particular, the rest of the article is organized as follows: Section 2 describes in detail Ethereum and its PoS consensus mechanism; Section 3 describes current MEV on-chain and off-chain extraction strategies; Section 4 discusses the operation of *Flashbots* in the Ethereum ecosystem; Section 5 provides an analysis related to the actual distribution of validators in Ethereum; finally, Section 6 concludes the paper and provides some perspectives on future research directions.

2. Post-Merge Ethereum: Transaction Execution and Block Validation under PoS

Ethereum is an open-source, permissionless, decentralized blockchain platform that establishes a peer-to-peer network that securely executes and verifies *transactions*. These are sent from and received by user-created Ethereum *accounts*. A sender must sign transactions and spend *Ether* (ETH), i.e., Ethereum's native cryptocurrency, as a cost of processing transactions on the network. Transactions may involve the execution of application code, called *smart contracts*. Approximately every 12 seconds, a *batch* of new transactions, known as a *block*, is processed by the network. Each block also contains a cryptographic hash identifying the series of blocks that must precede it if the block is to be considered valid. This series of blocks, from the genesis (first) block to the most recent one, is known as the *blockchain*. Transaction records are immutable, verifiable, and securely distributed across the network, giving participants full ownership and visibility into transaction data. Ethereum can be seen as a (very large) state machine, where the state is given by Ether balances and other storage values of all Ethereum accounts; transactions modify some of these values, thus altering the overall state of the machine. In this context, the *Ethereum Virtual Machine* (EVM) is the runtime environment for transaction execution. It includes a stack, memory, and persistent storage for all Ethereum accounts (including contract code), and is designed to be *deterministic*, so that given a pre-transaction state and a transaction, each network node produces the same post-transaction state, thereby enabling network consensus. The Ethereum Yellow Paper formally defines the EVM [11].

Post-Merge Ethereum consists of an *execution layer* and a *consensus layer*, both running on different client software. The *execution client* is responsible for transaction-related actions, as detailed below, while the *consensus client* implements the Proof-of-Stake consensus algorithm. There are also two types of accounts on Ethereum, i.e., *user accounts* (also known as *externally-owned accounts*, EOAs), and *contract accounts*. Both types are identified on the blockchain and in the state by an account address; they have an Ether balance and may send Ether to any account, call any public function of a contract, or create a new contract. User accounts are the only type of account that may create transactions. For a transaction to be valid, it must be signed using the sending account's *private key*, which is generated with the corresponding public key when creating the account. Figure 1 depicts the process of proposing and validating transactions (and blocks). When a transaction is submitted to an *execution client*, it verifies its legitimacy by ensuring that the balance of the sending account contains enough Ether to complete the transaction and that the digital signature is correct. If the transaction is legitimate, the execution client adds it to its local *memory pool* (*Mempool*), i.e., a list of pending transactions, and broadcasts it to the other nodes of the network. Upon hearing about the transaction, the other nodes also add it to their local Mempool.

In Ethereum's implementation of Proof-of-Stake, time is divided into *slots*, each lasting 12 seconds. In each slot, a single *validator* is (pseudo-)randomly selected to be the block proposer responsible for creating a new block and sending it out to other nodes on the network. In order to participate in the selection of validators in the PoS Ethereum network, a user is required to deposit 32 Ether, and run three separate pieces of software: the previously introduced execution and consensus clients, and a *validator client*. Upon deposit, the user aiming to become a validator




Figure 1: Transaction execution and block validation in Ethereum.

is added to an *activation queue*, which limits the rate of new validators joining the network. Additionally, a *committee of validators* is randomly chosen in each slot to vote on the validity of the block being proposed [12]. So, the selected validator is responsible for constructing and broadcasting a block to be appended to the blockchain, enabling the network to agree on the sequence of blocks at the chain's head. The committee of validators receives the new block and re-executes it locally to ensure that the proposed state change is valid. Then they confirm that the block is legitimate, send an *attestation* (a vote in favor of the block), and add the block to their local database. It is considered the next block on the Ethereum blockchain, the one with the highest attestation weight, as defined by *fork choice* rules [13].

In Ethereum, each transaction has a cost, which is determined by a *base fee*, set by the protocol, and a *priority fee*, set by the user to make the transaction more appealing for being included in a block. The base fee acts as a reserve price; its value is calculated independently of the current block and is instead determined by the blocks before it – making transaction fees more predictable for users. When the block is mined this base fee is “burned”, removing it from circulation. Further, validating the transaction consumes a number of *gas units*, whose value varies according to the market's law. Overall, the *total fee* is obtained as $units\ of\ gas\ used * (base\ fee + priority\ fee)$. Tweaking the value of the *priority fee* can be done by experienced users to arrange transactions in upcoming blocks for *Maximal Extractable Value* (MEV).

3. Maximal Extractable Value

The term *Maximal Extractable Value* (MEV) is commonly used in the blockchain community to indicate the goal of maximizing the *extra value* that can derive from including ad-hoc transactions in a block, or from controlling the order of inclusion of the transactions thereby [14]. MEV resembles financial speculation in traditional finance, where the goal of some transactions is not in the actual semantics of the concrete assets being purchased, but rather in the gain due to the increase of value that an asset may undergo after the purchase, or yet in profits obtained by buying and then selling given assets at proper moments in time. A straightforward example that applies to both traditional finance and blockchain finance is *arbitrage*, in which a trader can spot that some asset (e.g., a token in the case of blockchains) has different quotations in distinct markets, and obtain profit by first buying the asset at the lowest price, and then selling it at the highest price immediately afterward.

Importantly, blockchain allows for MEV opportunities to be pursued with technical flavors

that are not there in the case of traditional finance. For instance, continuing with the example of arbitrage transactions, the blockchain trader can succeed to make the buying and selling transactions become part of the same block, thus making them appear as transactions that happened atomically at the very same time. In this way, they obtain the corresponding profit immediately as that block gets included in the blockchain, incurring no risks that prices may oscillate in between. In the blockchain community, the term MEV is used to refer by generalization to both the behavior of pursuing such type of speculative revenues and the resulting profit obtained by actuating behavior of this type. Furthermore, the word *maximal* is to be interpreted in a best-effort fashion. Thus it makes sense to refer to arbitrage actions as a type of MEV opportunity or quantify the MEV that resulted from actuating a given arbitrage action.

3.1. Exploiting MEV Opportunities

In the simplest instance of the phenomenon, the *validator* in charge of proposing the next block at a given blockchain epoch is directly in the position to exploit MEV opportunities, provided that they identified any viable MEV opportunity at that time.¹ However, the activities of *searching* for MEV opportunities and *proposing* the next blocks are inherently separable, and indeed, in blockchain finance, MEV opportunities are often pursued by traders that are not necessarily validators themselves. Blockchain traders searching for MEV opportunities are commonly referred to as *searchers*.

Because of the strong technical flavors related to identifying and exploiting MEV opportunities, a searcher is most often reified in the form of a software agent. Following programmed algorithms, a searcher monitors the evolution of the blockchain status in real-time, while new transactions get included, inspects the transactions submitted in the Mempool as candidates for future inclusion, identifies if the inclusion of some specific transactions generates MEV opportunities, and readily deploys the further transactions needed to exploit those MEV opportunities.

3.2. Interaction between Searchers and Validators

When a searcher agent eventually succeeds in identifying a MEV opportunity, it must face the problem of compelling some validators to propose a block that includes the MEV transactions. This requires: (i) establishing the interaction between the searcher and a validator, (ii) convincing the validator to choose the transactions deployed by the searcher, in particular, because there can be other searchers that are trying to exploit the same MEV opportunities, and (iii) making the validator build a proper block, i.e., a block that includes all needed transactions in the proper order. To reach these goals, searchers can opt for either *on-chain* or *off-chain* solutions [16], the latter being often the preferable choices to avoid the risk that the MEVs get stolen from *generalized frontrunners* [15], as we explain below.

¹In a blockchain using Proof-of-Work as consensus protocol the block-proposing role competes to the miner that won the crypto-puzzle competition at the given epoch. This is why the acronym MEV is sometimes spelled as the *Miner Extracted Value* [15].

3.2.1. On-chain MEV

On-chain solutions ground on the mechanisms that the blockchain naturally offers for traders to accomplish the execution of transactions. Upon identifying a MEV opportunity, searchers put MEV transactions in the public Mempool, such that the validators can look up those transactions. To improve the chances of succeeding, the searchers set high gas fees for the MEV transactions, thus incentivizing validators to select those transactions for inclusion in the next blocks. Packaging the transactions in a smart contract may assist in achieving a proper sequencing of the transactions, although the actual control of the searcher on the ordering of the transactions might be admittedly insufficient for some types of MEVs. Unfortunately, the publication of the MEV transactions in the Mempool allows for others to spot those transactions before their inclusion in the blockchain. This opens the possibility that malicious traders *frontrun* those transactions, meaning that they put in the Mempool their own transactions for addressing the same MEV, but setting higher fees than the original searcher, such that validators will prefer the frontrunning transactions to the searcher's ones. This attack can be technically sophisticated by implementing *generalized frontrunners*, i.e., software agents that inspect the Mempool, simulate the transactions (in any possible sequence) in sandboxes in order to readily identify possible MEVs, and indiscriminately frontrun any set of MEV transactions.²

3.2.2. Off-chain MEV

As the risk of incurring frontrunning discourages the use of On-chain solutions to accomplish MEVs, off-chain solutions have become increasingly popular over time. In off-chain solutions, searchers establish private channels with the validators that are in charge of proposing the next blocks and use the private channels to advertise the MEV-augmented blocks, maintaining that the fees in their blocks overcome the fees that the validators could gain by selecting the transactions out of the public Mempool. In particular, the off-chain solutions work well in the context of blockchains based on PoS as a consensus protocol, where the validator in charge of proposing the next block is designated an epoch in advance with respect to the epoch in which the block will get included in the blockchain, and thus the negotiation between the validator and the searchers can take place at that epoch.

Off-chain solutions have led to the rise of off-chain platforms assisting validators and searchers to meet and negotiate on the MEV blocks. This way, searchers compete with each other in *fee auctions*, trying to win the assistance of the in-charge validators by offering them higher fees than other searchers. These platforms are commonly referred to as *private mempools*, or also *permissioned mempools*: like the public Mempool, they allow for validators to look up transactions for inclusion in the blockchain; unlike the public Mempool, they can be controlled by parties that can limit the access to subsets of validators, e.g., the ones considered well-reputed, or censor some transactions on the basis of commercial or political reasons.

²Anecdotally Paradigm's researchers wrote the story of their attempt to recover a relevant amount of liquidity tokens that people at their company had erroneously frozen in a smart contract [17]. To protect their liquidity-get transaction from frontrunning, they masked it as a pair of (apparently unrelated) transactions that would succeed only if executed in proper sequence. To their surprise, they were frontrunned anyway, likely because incurring a generalized frontrunning attack.

4. The Off-chain MEV Ecosystem in Ethereum

As we commented above, PoS-based blockchains enable MEV searchers to opt for off-chain solutions to the problem of establishing their interactions with validators, and indeed such off-chain solutions are most often the favorite choice of searchers in order to successfully accomplish their MEV revenues. Since the transition to PoS (“The Merge”), the Ethereum blockchain has seen a flourishing of off-chain MEV interactions, up to the emergence of a rich ecosystem of actors, roles, and platforms, which interplay among them to support the interests of each other, while attempting to contrast the possible vulnerabilities and detrimental aspects that derive from the off-chain deployment of the mechanism. Figure 2 sketches Ethereum’s off-chain MEV ecosystem, and we refer to it to explain the ecosystem’s roles, platforms, and solutions.




Figure 2: The Ethereum MEV ecosystem.

Searchers and *validators* (leftmost and rightmost sides of the figure) are the endpoints of MEV interactions. As previously introduced, searchers are trader agents, usually instantiated as bots, who identify MEV opportunities; validators are the main actors of Ethereum’s PoS protocol, who can be designated in a pseudo-random fashion as proposers of the next blocks. In particular, in Ethereum, each blockchain epoch is comprised of 32 block slots, and the execution of the PoS protocol results in designating a validator as the block proposer for each slot.

The off-chain MEV ecosystem depicted in Figure 2 aims to assist and optimize the interactions between the searchers and the block-proposer validators. In Ethereum, an ecosystem of this type is currently being promoted by the *Flashbots* organization [10], a research collective that pursues the goal of mitigating the negative impacts (e.g., increased centralization) of private MEV solutions on the Ethereum blockchain [18]. In what follows, while continuing with our exposition of the MEV ecosystem, we will refer to Flashbots to exemplify the solutions.

Builder actors address the task of block building, that is, they aggregate the MEV transactions identified by searchers as blocks ready for inclusion in the blockchain. Separating the roles of builders and searchers allows for builders to receive MEV transactions (each referred to as a *bundle*) from multiple searchers, while also monitoring the transactions in the public Mempool. With all this information available, builders can construct the blocks by optimizing the ordering of the transactions in each block, possibly putting together multiple MEV bundles along with

public transactions, in order to maximize profits.

Once a builder has created a block, it must bid for validator block spaces. Flashbots publicized the so-called *Builder* API as an open-source standard interface for builders to offer their service. As part of this standardization, they recommend a method by which builders can pay validators, which involves the builder setting (in the block header) their own address as the fee recipient of the block, while they append the block with a special last transaction that sends a payment to the validator. Thus, as soon as the block will become part of the blockchain (if it eventually indeed will), the builder will receive the fees of the transactions in the block and the validator will receive the specified payment from the builder. At the same time, the MEV bundles in the block will trigger the revenues expected from the searchers.

The *relay* and *MEV-Boost* components provide the middleware for builders and validators to connect with each other. We remark that this specific architecture, as well as the denominations of the components as relays and MEV-Boost, explicitly refer to the design of this type of middleware in the Flashbots platform [19, 20]; other designs could be easily imagined. A relay connects with one or many builders, provided that they implement the Builder API. The relays verify the validity of the blocks sent by the builders, and select the valid blocks that offer the highest bid to the validator; then they feed the selected blocks to the MEV-Boost component, which in turn is in charge to present the most profitable block to the validator. The Flashbots organization participates directly in the ecosystem with its own builder.

To participate in the ecosystem, validators install and setup the MEV-Boost client, such that they can receive the bids from the builders. In the configuration, they can specify from which relays they would like to receive block bids. MEV-Boost was activated on September 15, 2022, just 17 epochs after “The Merge”.

5. Analysis of Validator Distribution after “The Merge”

Based on what has been exposed so far about the Ethereum off-chain MEV ecosystem, we now analyze the distribution of validators in the ecosystem from September 15, 2022, to January 24, 2023.

5.1. Data Collection

Block data were downloaded from the date of Ethereum’s transition to PoS (September 15, 2022). In particular, we started from block 15,537,394, i.e., the first block of PoS Ethereum’s Paris Network Upgrade,³ to block 16,474,262 (January 24, 2023). In total, we constructed a dataset containing data from 936,869 blocks over a time period of 131 days. By employing the *Etherscan* API,⁴ we collected the following information about each considered block: *block number*, the number of the block, *timestamp*, the block time, *fee recipient*, the account address for paying block reward, and *block reward*, the reward for the block producer. For MEV-Boost data of the Flashbots relay, we used the `proposer_payload_delivered` endpoint from the Data section of the *Relay* API [21].⁵ In this way, we were able to collect all the information of the blocks

³<https://etherscan.io/block/15537394>

⁴<https://etherscan.io/apis>, using the `getBlockreward` endpoint.

⁵<https://boost-relay.flashbots.net>

in which a validator requested a block from the Flashbots relay. In particular: *block number*, *builder public key*, *proposer public key*, *proposer fee recipient*, the address chosen by the validator to receive the reward from the builder, and *value*, the amount of the reward. Finally, since Flashbots standardizes how payments are made from builders to block proposers, i.e., via a transaction contained at the end of the block proposed by the builder to an address given by the validator [22], we collected the last transaction of each block number contained in the data from the Flashbots relay. To do this, we used the Python `web3.py` library along with the *Infura* API⁶, collecting the *block number*, *sender's address*, which corresponds to the fee recipient of the block, *receiver's address*, which corresponds to the proposer fee recipient, and *value*, the amount of the reward.

5.2. Data Analysis

In this section, we analyze the distribution of validators in the Ethereum blockchain after “The Merge”. We show that the analysis may lead to different conclusions if we just look at the fee recipient information stored in the blocks, or if we extract the actual identities of validators based on the knowledge of the off-chain MEV ecosystem. Indeed we claim that the latter analysis provides a clearer picture of the phenomenon than the former one, making the distribution of validators more apparent.

5.2.1. Analysis of the first 1,000-block Fee Recipients

The first analysis we performed was on the first 1,000 blocks after “The Merge”, to verify the many websites, blogs, and online articles claiming that 40 percent of the blocks were validated by just a few entities [7, 8]. On these 1,000 blocks, we analyzed the different *fee recipient* addresses, i.e., those addresses in the validated blocks that received *priority fees* as a reward for validating the block [23]. As can be seen in Figure 3, our analysis confirmed that the first two fee recipients were actually validators present in 43% of the first 1,000 blocks. Considering the first 3 entities brings the figure up to 51.2%. The address of the recipient with the highest percentage is `0x388C818CA8B9251b393131C08a736A67ccB19297`, which corresponds to *Lido: Execution Layer Rewards Vault* [24],⁷ which we will identify as `Lido` from now on.

5.2.2. Analysis on the 936,869-block Fee Recipients

Having collected data from the 131 days following Ethereum’s transition to PoS, we analyzed the distribution of fee recipients following the introduction of actors who can carry out off-chain MEV strategies. By observing Figure 4, we can notice that the first three fee recipients are different from those of the first 1,000 blocks. `Lido` collapses from 27.9% of the first 1,000 blocks to 6.64%; the first place is now occupied by the address `0xdafea492d9c6733ae3d56b7ed1adb60692c98bc5`, which corresponds to ENS address `flashbots-builder.eth`, with a percentage of 20.9%; it is followed by the address `0x690b9a9e9aa1c9db991c7721a92d351db4fac990`, which corresponds to

⁶<https://docs.infura.io/infura/>

⁷<https://lido.fi/>




Figure 3: Fee recipients of the first 1,000 blocks after Ethereum’s “The Merge”.

ENS address `builder0x69.eth`, with 15.8%; then, in third place, we find the address `0x95222290DD7278Aa3Ddd389Cc1E1d165CC4BAfe5`, with 10.6%, for a total of 47.3%. With respect to Figure 3, it would seem that the situation has completely reversed in favor of these three builders, which are in fact fee recipient addresses of builders employing the Flashbots relay.




Figure 4: Fee recipients of the first 936,869 blocks after Ethereum’s “The Merge”.

By way of example, by further analyzing only Lido and `flashbots-builder.eth`, we can see in Figure 5 that the number of blocks per day in which Lido appears as a fee recipient begins to collapse from thousands to less than a hundred over time. In contrast, the exact opposite occurs for `flashbots-builder.eth`, as it emerges from Figure 6.




Figure 5: Number of blocks per day in which Lido appears as fee recipient.




Figure 6: Number of blocks per day in which flashbots-builder.eth appears as a fee recipient.

5.2.3. Analysis of the Actual Block Validators

Analyzing only the fee recipient addresses is not sufficient to understand who really performs the activity of validation. In fact, it is necessary to recall that Flashbots builders (and other builders based on the Flashbots relay) have a specific standard to reward actual validators. The builder is indeed the fee recipient, but it includes a transaction (the last in the block) that pays Ether tokens to the block-proposer validator [22]. Hence, we extracted the addresses of these rewarded validators within the Flashbots relay data, in order to investigate who was in charge of validation. As a result, we noticed that the top-3 validator addresses are the same as those in Figure 3. This is detailed in Figure 7, illustrating the addresses of rewarded validators instead of those of fee recipients for the top-3 builders in Figure 4. To make this concept even clearer, as an example, let us show in Figure 8 the number of blocks per day in which Lido appears as a validator rewarded by the builders; we clearly see that the number of blocks per day does not collapse after the introduction of builders, but maintains instead a constant trend. Hence, based on this preliminary analysis, it does not appear that the off-chain MEV Ethereum ecosystem contributes to modify the actual validator distribution in Ethereum, which still remains largely the preserve of the already wealthier players.




Figure 7: Top-3 rewarded validators of the first 936,869 blocks after Ethereum’s “The Merge”.




Figure 8: Number of blocks per day in which Lido appears as a rewarded validator.

6. Conclusions

In recent times, MEV has become an important topic for the Ethereum community. As a result, several initiatives have been developed to mitigate the negative impact of MEV (e.g., centralization) on the network. Flashbots’ MEV-Boost is a noteworthy proposal that aims to create a more equitable distribution of value. In this article, we conduct an analysis of the validated blocks after Ethereum’s update to PoS, known as “The Merge”, showing how the scenario has changed in recent months. In fact, the recipient of the fee no longer always indicates only the validator, but often indicates the MEV builder, and tracking rewards related to the validator requires some additional steps. From a preliminary analysis, it seems to remain confirmed what was observed in [7, 8] with respect to the first 1,000 blocks, because most share of the validation activity continues to reside in the hands of very few validators, the same who were already dominating this role. Therefore, in the future we would like to extract data from more relays belonging to the Ethereum ecosystem, trying to obtain complete data on the rewards received by the validator for the task performed.

References

- [1] Ethereum.org, Ethereum, <https://ethereum.org/>, 2023. Accessed on April, 28th, 2023.
- [2] Ethereum.org, The Merge, <https://ethereum.org/en/upgrades/merge/>, 2023. Accessed on April, 28th, 2023.
- [3] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, S. Capkun, On the security and performance of proof of work blockchains, in: Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, 2016, pp. 3–16.
- [4] G. Fanti, L. Kogan, S. Oh, K. Ruan, P. Viswanath, G. Wang, Compounding of Wealth in Proof-of-Stake Cryptocurrencies, arXiv (2018). URL: <https://arxiv.org/abs/1809.07468>. doi:10.48550/ARXIV.1809.07468.
- [5] Q. Lin, C. Li, X. Zhao, X. Chen, Measuring decentralization in bitcoin and ethereum using multiple metrics and granularities, CoRR abs/2101.10699 (2021). URL: <https://arxiv.org/abs/2101.10699>. arXiv:2101.10699.
- [6] B. Kusmierz, R. Overko, How centralized is decentralized? Comparison of wealth distribution in coins and tokens, in: 2022 IEEE International Conference on Omni-layer Intelligent Systems (COINS), IEEE, Barcelona, Spain, 2022, pp. 1–6. doi:10.1109/COINS54846.2022.9854972.
- [7] Cointelegraph, Ethereum at the center of centralization debate as SEC lays claim, <https://cointelegraph.com/news/ethereum-at-the-center-of-centralization-debate-as-sec-lays-claim>, 2022. Accessed on April, 28th, 2023.
- [8] Decrypt, Big Firms Dominate Post-Merge Ethereum Validation, <https://decrypt.co/109901/big-firms-dominate-post-merge-ethereum-validation>, 2022. Accessed on April, 28th, 2023.
- [9] V. Buterin, The Meaning of Decentralization, <https://medium.com/@VitalikButerin/the-meaning-of-decentralization-a0c92b76a274>, 2017. Accessed on April, 28th, 2023.
- [10] Flashbots, Flashbots, <https://www.flashbots.net/>, 2022. Accessed on April, 28th, 2023.
- [11] G. Wood, Ethereum: a secure decentralised generalised transaction ledger, <https://github.com/ethereum/yellowpaper>, 2014.
- [12] Ethereum.org, Proof-of-Stake (PoS), <https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/>, 2022. Accessed on April, 28th, 2023.
- [13] Ethereum.org, Fork Choice, <https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/#fork-choice>, 2023. Accessed on April, 28th, 2023.
- [14] Ethereum.org, Maximal Extractable Value (MEV), <https://ethereum.org/en/developers/docs/mev/>, 2023. Accessed on April, 28th, 2023.
- [15] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, A. Juels, Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges, CoRR abs/1904.05234 (2019). URL: <http://arxiv.org/abs/1904.05234>. arXiv:1904.05234.
- [16] I. B. et al., The Cost of Decentralization in 0x and Etherdelta, <https://hackingdistributed.com/2017/08/13/cost-of-decent/>, 2017. Accessed on April, 28th, 2023.
- [17] G. K. Dan Robinson, Ethereum is a Dark Forest, <https://www.paradigm.xyz/2020/08/ethereum-is-a-dark-forest>, 2020. Accessed on April, 28th, 2023.
- [18] A. Obadia, Flashbots—Frontrunning the MEV Crisis, <https://writings.flashbots.net/>

frontrunning-mev-crisis, 2020. Accessed on April, 28th, 2023.

- [19] Flashbots, Flashbots MEV-Boost - Introduction, <https://docs.flashbots.net/flashbots-mev-boost/introduction>, 2022. Accessed on April, 28th, 2023.
- [20] Flashbots, Relay Fundamentals, <https://docs.flashbots.net/flashbots-mev-boost/relay>, 2022. Accessed on April, 28th, 2023.
- [21] Flashbots, Flashbots Relay API, <https://flashbots.github.io/relay-specs/#/Data/getDeliveredPayloads>, 2022. Accessed on April, 28th, 2023.
- [22] Flashbots, Block Builders, <https://docs.flashbots.net/flashbots-mev-boost/block-builders>, 2022. Accessed on April, 28th, 2023.
- [23] Ethereum.org, Gas and fees, <https://ethereum.org/en/developers/docs/gas/>, 2023. Accessed on April, 28th, 2023.
- [24] Lido, Execution Layer Rewards Vault, <https://docs.lido.fi/deployed-contracts/>, 2023. Accessed on April, 28th, 2023.