A non-parametric approach to identifying anomalies in Bitcoin mining

Eduardo Augusto de Medeiros Silva [1], **Ivan da Silva Sendin**[1]

School of Computer Science– Federal University of Uberlandia(UFU)

{eduardomedeiros,sendin}@ufu.br

Abstract. *Selfish Mining is an attack on the proof-of-work-based cryptocurrency consensus mechanism, enabling attackers to gain more than their fair share of rewards. Its existence indicates that the Nakamoto consensus is not incentive compatible and could jeopardize blockchain security. Recently, a method employing the Z-Score to detect selfish mining was proposed. This paper introduces a non-parametric statistical technique to identify traces of selfish miners on the blockchain without assuming any specific statistical distribution for the analyzed data. Additionally, the applicability of this type of analysis is discussed.*

1. Introduction

In Bitcoin and most existing cryptocurrencies, miners are responsible for maintaining blockchain integrity by validating transactions and executing the proof-of-work, thereby achieving Nakamoto Consensus. This process ensures blockchain security, correctness of its content, and immutability of previous blocks. The core of blockchain security is rooted in a game-theoretic approach where the distribution of computational power among participants ensures that no single entity can defraud the system. Miners are rewarded with cryptocurrencies from Coinbase transactions and transaction fees. Given that it is a probabilistic protocol, over time, each miner should receive a number of coins proportional to their computational power, adhering to the "one CPU one vote" policy [Nakamoto 2009].

Shortly after Bitcoin's popularization, Eyal and Sirer discovered that deviating from the protocol could increase miners' profits and potentially compromise blockchain security; this approach is known as Selfish Mining (SM) [Eyal and Sirer 2014]. Subsequent studies analyzed this behavior and proposed modifications to Bitcoin to mitigate this attack [Heilman 2014]. However, it was not until [Li et al. 2020a, Li et al. 2020b] that a statistical method utilizing blockchain data alone to detect selfish behavior was presented. This method is based on the Z-Score Test and assumes a Gaussian distribution for the evidence left by miners on the blockchain.

In this paper, we propose an alternative method that does not rely on any assumptions about the statistical distribution of the data. We discuss the feasibility of this non-parametric approach and explore other techniques that can be employed to detect Selfish Mining using blockchain data exclusively.

2. Background

2.1. Distributed Consensus

Bitcoin was the first technology to satisfactorily solve the problem of distributed consensus: using proof of work (PoW), miners keep the blockchain up-to-date and consistent. After validating transactions, each miner includes their own address in the set of valid transactions, creating a special transaction called Coinbase. This initiates "a race" for a `nonce` that, when combined with the block header, produces an appropriate hash value, thereby validating the new block. This search is random, and the only way to increase the chances of generating a correct `nonce` is to enhance computational power, thereby increasing the rate of nonce generation and testing.

Once a new block is generated, the miner must broadcast this block to other miners via the peer-to-peer Bitcoin network to secure their earnings. This process then restarts with the search for the next block on the blockchain.

2.2. Selfish Mining

For a long time, it was believed that the protocol proposed by Nakamoto was incentive compatible, meaning that following the protocol was the best strategy to maximize a miner's profits. However, a documented attack on the mining system, known as Selfish Mining (SM), demonstrated that deviating from the protocol could yield greater benefits than expected for a given computational power. It is well known that any deviation from the proposed protocol can compromise blockchain security.

In simplified terms, whenever a selfish miner finds the correct nonce for a block of height n, instead of disclosing it to the P2P network (which would prompt all miners to start searching for the block at height $n + 1$), the selfish miner withholds the mined block and begins working on height $n+1$ while other miners continue working on height n. This gives the selfish miner an advantage in the search for block $n + 1$. A direct consequence of SM is that the selfish miner produces blocks "in a row" disproportionately, which can generally be identified in the Coinbase transaction where the miner records their own address to receive payment.

Since its appearance, SM has garnered significant attention from the blockchain community, leading to the discovery of other attacks on the consensus mechanism [Kwon et al. 2017, Nayak et al. 2016] and proposals to enhance blockchain robustness against such attacks [Heilman 2014, Ketsdever and Fischer 2019]. Some analyses indicate that SM is only profitable when the miner controls at least 25% or 33% of the total computational power. Miners without this level of computational power might take significant risks if they attempt SM [Negy et al. 2020, Li et al. 2020a, Li et al. 2020b].

Given the complexity of the mining process, the poetic phrase "one CPU, one vote" no longer applies to Bitcoin mining. Currently, mining is conducted by a complex network of mining pools, with computational power shifting from one pool to another. This makes it challenging to analyze these mechanisms and search for evidence of SM. Additionally, there is limited research on the use of off-chain data for this purpose.

3. Methodology

As previously mentioned, [Li et al. 2020a, Li et al. 2020b] present a method for identifying Selfish Mining (SM) based on detecting blocks produced consecutively by the same

miner. The authors' method employs the Z-Score Test to assess how far blockchain data deviates from a block distribution produced by a set of honest miners. When this threshold exceeds two standard deviations, the data in question is considered to have originated from SM with 95% confidence. This method relies heavily on the assumption of a Gaussian distribution in the sequence of blocks: whenever a standard deviation is greater than two in relation to the mean, there is high confidence that this event is not the result of honest mining.

Given that the distribution of consecutive blocks does not necessarily follow a Gaussian Distribution, we propose a non-parametric test that directly counts the sequential mining events in the blockchain and compares these values with the expected data in a scenario without SM. Our proposed method is similar to that presented by [Li et al. 2020a], but with the significant advantage of being non-parametric, thus eliminating the need to prove that the data distribution is Gaussian.

Algorithm 1 SM Non-Parametric Identification

Input B: an subsequence of Blockchain,k: number of permutations used in this test; M: miner identification
$count \leftarrow 0$
$c \leftarrow$ number of blocks mined in a row in B by M
for $i \in 1..k$ **do**
    $p \leftarrow$ random permutation of B
    $cp \leftarrow$ number of blocks mined consecutively by M in p
    if $c \geq cp$ **then**
        $count \leftarrow count + 1$
    end if
end for
if $\frac{count}{k} > 0.95$ **then**
    Miner M is suspected of using SM strategy in B
end if

The rationale for the proposed non-parametric method is based on the following facts:

- The computational power of a miner in a given period can be inferred by the proportion of blocks this miner produced in the period [Li et al. 2023, Lin et al. 2021, Campajola et al. 2022];
- Applying a permutation to blocks from a given period does not change a miner's estimated computational power;
- The number of consecutive blocks mined in a permutation of a sequence of blocks represents the expected number of consecutive blocks in a scenario without SM.

Therefore, the method, as detailed in Algorithm 1, performs a permutation test on data obtained from the blockchain (input parameter B). The test involves checking whether the occurrence of sequential mining is more frequent in B than in its permutated versions. Since the permutation maintains the miner's computational power information from the blockchain while erasing any traces of selfish mining, a higher frequency of se-

quential mining in the real data compared to the set of permutations indicates the presence of selfish mining.

4. Experiments and Results

To search for evidence of SM, we obtained Blockchain data through HTTP requests to Blockchain.info using its API. We collected approximately one year of block information, spanning from block height 684,192 to block height 737,327, totaling 56,008 blocks. This data was stored in a local instance of `MongoDB`. The scripts for data capture and storage, as well as for processing and analysis, were written in `Python`.

The most relevant characteristic of SM is the sequence of blocks mined by the same miner. Table 1 displays the results of applying Algorithm 1 to the blocks in the observed period, with $k = 2,000$ permutations for each test, this value of k was chosen empirically, taking into consideration the experiment already conducted in [Li et al. 2020a, Li et al. 2020b]. For each month, we identified miners who mined a sequence of blocks that was inconsistent with their computational power, using a p-value threshold of 0.05 for this selection, this is a standard value also used in [Li et al. 2020a, Li et al. 2020b].

Table 1. Analysis of the probability of the miner having obtained the consecutive minings considering its computational power compared to a random process. The column `ZScore` is calculated according to [Li et al. 2020a]. To identify miners, the Blockchain.info service was used.

Month	Miner	Mined Blocks	Power (%)	Blocks Mined in a row	p-value	Z Score
1	Binance Pool	19	0.4	2	0.003	–
4	Poolin	464	9.9	59	0.023	2.1
7	F2Pool	766	16.4	144	0.028	2.1
7	147Sw-z3aPq	21	0.4	3	0.000	9.2
7	1JhAQ-HGCDp	14	0.3	1	0.038	–
8	Antpool	812	17.4	158	0.046	1.7
9	Binance Pool	401	8.6	47	0.013	2.3
9	12cKi-tYjmR	13	0.3	1	0.034	–
9	125m2-qRyYu	13	0.3	1	0.031	0.2
10	35y82-bMj83	129	2.8	9	0.008	2.9
10	F2Pool	32	0.7	4	0.000	9.8
11	BTC.com	486	10.4	63	0.036	2.0
12	Braiins Pool	240	5.1	22	0.007	2.8
12	Huobi Pool	65	1.4	4	0.010	3.2
12	191sN-XDAGo	121	2.6	11	0.000	4.6
12	Poolin	322	6.9	48	0.000	5.8
12	SBI Crypto	35	0.7	2	0.027	2.4

5. Discussion

In evaluating the use of sequential mining as a detection mechanism for selfish mining (SM), several critical considerations must be addressed:

p-value issues By partitioning the data on a monthly basis to maintain the assumption of stable computational power among miners, we inadvertently create a statistical artifact. This artifact can yield numerically significant p-values that lack true statistical significance, as highlighted by [P et al. 2016];

The real hashpower A direct consequence of SM is the disproportionate production of blocks relative to hash power. Paradoxically, when identifying a miner with a proportion α of blocks as engaging in SM - because it produces more sequential blocks than expected based on its hash power - one finds that the miner's actual hash power fraction is less than α;

Miners Ecossystem The relationships and agreements among miners can be complex, involving computational power leasing and mining pools that can migrate from one master miner to another. These dynamics obscure block frequency patterns, rendering the proposed detection approach less effective;

Given these observations, alternative methods for detecting SM must be considered. Effective SM detection should account for the role of luck in obtaining proof of work. Consequently, blocks generated by SM should exhibit shorter intervals between them. The set of mined transactions should possess characteristics that facilitate the identification of SM, such as their volume, timestamps, financial signatures, and associated addresses.

6. Conclusion

In this study, we introduced a non-parametric method for identifying selfish mining (SM) using only Blockchain data. A non-parametric approach is generally more robust than methods relying on specific underlying distributions, such as the Gaussian distribution of consecutive mining by the same miner.

An analysis of approximately one year of Blockchain data suggests that up to 15 miners may have engaged in SM during this period. However, despite this evidence, the problem requires additional methodologies for more robust results, as the current approach may distort the calculated p-values. Next steps could include a choice of different mining windows, seeking greater stability in the computational power of miners. The complexities involved in detecting SM are not fully addressed by the presented method.

Furthermore, considering the security concerns associated with SM and the intricacies of mining systems, incorporating additional data sources is essential to enhance the accuracy and reliability of the results.

References

Campajola, C., Cristodaro, R., Collibus, F. M. D., Yan, T., Vallarano, N., and Tessone, C. J. (2022). The evolution of centralisation on cryptocurrency platforms. *ArXiv*, abs/2206.05081.

Eyal, I. and Sirer, E. G. (2014). Majority is not enough: Bitcoin mining is vulnerable. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8437:436–454.

Heilman, E. (2014). One weird trick to stop selfish miners: Fresh bitcoins, a solution for the honest miner (poster abstract). *Lecture Notes in Computer Science (including*

subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 8438:161–162.

Ketsdever, S. and Fischer, M. J. (2019). Incentives don't solve blockchain's problems. In *2019 IEEE 10th Annual Ubiquitous Computing, Electronics Mobile Communication Conference (UEMCON)*, pages 0873–0876.

Kwon, Y., Kim, D., Son, Y., Vasserman, E., and Kim, Y. (2017). Be selfish and avoid dilemmas: Fork after withholding (faw) attacks on bitcoin. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, page 195–209, New York, NY, USA. Association for Computing Machinery.

Li, C., Palanisamy, B., Xu, R., and Duan, L. (2023). Cross-consensus measurement of individual-level decentralization in blockchains. In *2023 IEEE 9th Intl Conference on Big Data Security on Cloud (BigDataSecurity), IEEE Intl Conference on High Performance and Smart Computing, (HPSC) and IEEE Intl Conference on Intelligent Data and Security (IDS)*, pages 45–50.

Li, S.-N., Yang, Z., and Tessone, C. J. (2020a). Mining blocks in a row: A statistical study of fairness in bitcoin mining. In *2020 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–4.

Li, S.-N., Yang, Z., and Tessone, C. J. (2020b). Proof-of-work cryptocurrency mining: a statistical approach to fairness. In *2020 IEEE/CIC International Conference on Communications in China (ICCC Workshops)*, pages 156–161.

Lin, Q., Li, C., Zhao, X., and Chen, X. (2021). Measuring decentralization in bitcoin and ethereum using multiple metrics and granularities. In *2021 IEEE 37th International Conference on Data Engineering Workshops (ICDEW)*, pages 80–87.

Nakamoto, S. (2009). Bitcoin: A peer-to-peer electronic cash system.

Nayak, K., Kumar, S., Miller, A., and Shi, E. (2016). Stubborn mining: Generalizing selfish mining and combining with an eclipse attack. *Proceedings - 2016 IEEE European Symposium on Security and Privacy, EURO S and P 2016*, pages 305–320.

Negy, K. A., Rizun, P. R., and Sirer, E. G. (2020). Selfish mining re-examined. In *Financial Cryptography and Data Security: 24th International Conference, FC 2020 , Kota Kinabalu, Malaysia, February 10–14, 2020 Revised Selected Papers*, page 61–78, Berlin, Heidelberg. Springer-Verlag.

P, R., CS, P., and M., B. (2016). Common pitfalls in statistical analysis: The perils of multiple testing. *Perspect Clin Res.*, 7:106–107.