Short Paper: An Empirical Analysis of Blockchain Forks in Bitcoin

Till Neudecker and Hannes Hartenstein

Institute of Telematics, Karlsruhe Institute of Technology, Germany
{till.neudecker,hannes.hartenstein}@kit.edu

Abstract. Temporary blockchain forks are part of the regular consensus process in permissionless blockchains such as Bitcoin. As forks can be caused by numerous factors such as latency and miner behavior, their analysis provides insights into these factors, which are otherwise unknown. In this paper we provide an empirical analysis of the announcement and propagation of blocks that led to forks of the Bitcoin blockchain. By analyzing the time differences in the publication of competing blocks, we show that the block propagation delay between miners can be of similar order as the block propagation delay of the average Bitcoin peer. Furthermore, we show that the probability of a block to become part of the main chain increases roughly linearly in the time the block has been published before the competing block. Additionally, we show that the observed frequency of short block intervals between two consecutive blocks mined by the same miner after a fork is conspicuously large. While selfish mining can be a cause for this observation, other causes are also possible. Finally, we show that not only the time difference of the publication of competing blocks but also their propagation speeds vary greatly.

1 Introduction

Blockchain forks, which occur when two miners independently find and publish a new block referencing the same previous block, occur regularly in permissionless blockchains such as Bitcoin [7]. As subsequent blocks resolve the temporary inconsistency, forks are part of a blockchain's normal operation. While the existence of delay between miners inevitably leads to blockchain forks, deviating mining strategies such as selfish mining [3] can also lead to forks. Recent discussions on block size, the feasibility of selfish mining (*negative gamma*), and speculations on the network topology between miners are all related to factors affecting the security of permissionless blockchains [5]. As forks are affected by many of these factors, the analysis of forks that actually took place may help to improve the understanding of these factors.

Based on measurements of the Bitcoin peer-to-peer (P2P) network since 2015 we analyze the announcement and propagation of blocks that led to blockchain forks. Specifically, we compare the time differences between the first announcement of competing blocks to the average block propagation delay. Furthermore, we analyze the effect of a *headstart* of one block over competing blocks (i.e., how much earlier a block was published) on the block's probability to become part of the main chain. In order to assess whether deviating mining strategies were performed, we analyze the block intervals immediately after blockchain forks. Finally, we study the differences in the propagation of blocks of four selected forks through the Bitcoin P2P network.

2 Fundamentals & Related Work

We will now briefly sketch the relevant aspects of mining and block propagation in Bitcoin. A thorough introduction can be found in, e.g., [8]. Bitcoin blocks are generated in the process of mining by aggregating a set of previously published transactions into a block and solving a proof-of-work puzzle for that block. Each block contains the hash value of the previous block, which creates a chain of blocks. Miners are expected to work on top of the longest valid blockchain known to them, i.e., when a miner receives a new block extending the current blockchain, the miner should update the block she is working on by changing the reference to the newly received block.

A *blockchain fork* occurs if two new blocks that reference the same previous block are independently found at the same time by different miners. Because solving the proof-of-work puzzle is a random process and block propagation between miners is subject to network and processing delays, such forks occur regularly. However, forks can also be the result of selfish mining [3], a mining strategy in which a miner withholds new blocks instead of immediately publishing them in order to gain an advantage in finding the next block. Another strategy that can create blockchain forks is the *fork after withholding* attack [6].

Propagation of new blocks and transactions is performed by flooding via the Bitcoin peer-to-peer network, and by transmission via additional, possibly private networks (e.g., the *Fibre* network) [2]. Several characterizations of the Bitcoin P2P network have been published in the past [1,4]. Furthermore, there are several websites that publish statistics such as block propagation delays, i.e., the time it takes blocks to propagate through a certain share of the network.[1]

3 Measurement and Analysis Method

Since 2015 we have operated two monitor nodes that establish connections to all reachable peers of the Bitcoin P2P network. The number of connections varied between around 6,000 and 14,000 in the considered period. The monitor nodes stay mostly passive (except for establishing connections and sending and answering PING messages) and log the announcement of new transactions and blocks via *inventory messages* (*INV*) by remote peers. Therefore, our dataset contains tuples consisting of (time, hash value, IP address). From this data the

[1] E.g., `https://blockchain.info`, `https://bitnodes.earn.com`, `http://bitcoinstats.com/network/propagation`, `https://dsn.tm.kit.edu/bitcoin`

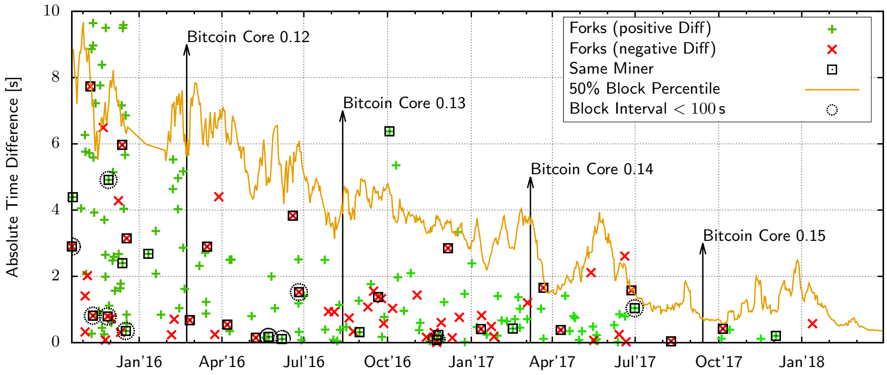

Fig. 1. Time difference between the first announcement of forking blocks. A green cross (*positive Diff*) indicates that the block that was announced first became part of the main chain, a red cross (*negative Diff*) indicates that the later announced block became part of the main chain. Boxes around blocks indicate that the subsequent block has been mined by the same miner; additional circles around blocks indicate that the subsequent block has been mined by the same miner within less than 100 seconds. Finally, the average 50 % block propagation percentile is shown.

timestamp of the first announcement of a block can be derived. Furthermore, the propagation speed of a block (i.e., how many announcements were received within a certain time) can be derived.

As our monitor nodes do not actually request blocks from remote peers, our dataset does not contain the block headers and does not indicate whether forks happened. Therefore, we combine our data with data published by *blockchain.info* that contains further information on each block hash, such as the reference to the previous block, whether the block became part of the main chain, and the miner as indicated in the coinbase transaction (set by the miner). All data used in this paper can be accessed at `https://dsn.tm.kit.edu/bitcoin/forks`.

4 Analysis of Bitcoin Blockchain Forks

As a first step, we analyze the time differences between the first announcements of the competing blocks that cause a fork. If all miners follow the protocol and immediately start working on top of any new valid block they receive, these time differences should not be larger than the block propagation delay between miners. Fig. 1 shows all forks between October 2015 and March 2018 that we have data on, and the time difference between the first announcements of both blocks of each fork: every cross indicates one fork, i.e., one blockchain height at which two blocks have been announced. A green cross indicates that the block that has been announced first became part of the main chain, a red cross indicates that the later announced block became part of the main chain.

The data confirms that the fork rate decreased substantially in the past years. Additionally, the observed announcement time differences decreased from sev-

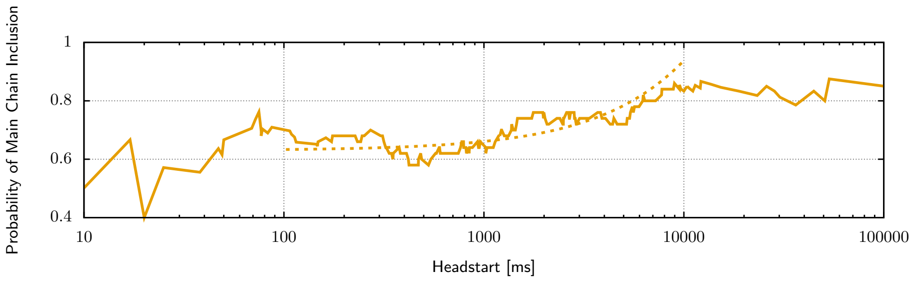

Fig. 2. Probability (moving average, binsize = 50 ms) of a forking block becoming part of the main chain depending on the headstart over the competing block.

eral seconds in late 2015 to less then two seconds since mid 2017. Fig. 1 also shows the measured average 50 % block propagation percentile, i.e., the time difference between the first announcement of a block and the time the block has been announced by 50 % of all peers. While we expect mining pools to be better connected to the Bitcoin P2P network than *the average peer*, the 50 % block propagation percentile gives an idea of the latency between peers. The decreased block propagation delay also reflects the improvements made to the block propagation mechanism of Bitcoin. The comparison of the announcement time difference to the block propagation delay shows that the announcement time difference of almost all forks is smaller than the 50 % block propagation percentile. However, some announcement time differences are still strikingly large, and a few are even larger than the 50 % block propagation percentile.

Assuming that all miners always mine on top of the longest blockchain they received, the data indicates that the block propagation delay between miners that caused forks was not substantially lower than the block propagation delay of average Bitcoin peers. While this might be surprising, we emphasize that the observed announcement delays might be caused by single miners that temporarily suffer from a high link latency, i.e., they represent worst cases, whereas the shown 50 % block propagation percentile represents an average case.

Several questions arise from the discussion of the data shown in Fig. 1. First, while the block that is announced first is regularly included in the main chain, the effect of the *headstart* of one block over another block on the probability to become included in the main chain is unclear. Secondly, the data is not sufficient to assess whether miners deviate from the mining strategy, e.g., by selfish mining. Finally, the effect of the P2P propagation speed of forking blocks remains unclear. We will address all three questions in the remainder of this section.

4.1 Effect of Headstart on Probability of Main Chain Inclusion

In order to analyze the relationship between the headstart of a block (i.e., the time difference the block has been announced before the competing block) and the probability that this block becomes part of the main chain, we look at the block that was announced first and check whether this block became part of the main chain. Hence, each fork can be represented by a tuple (headstart,

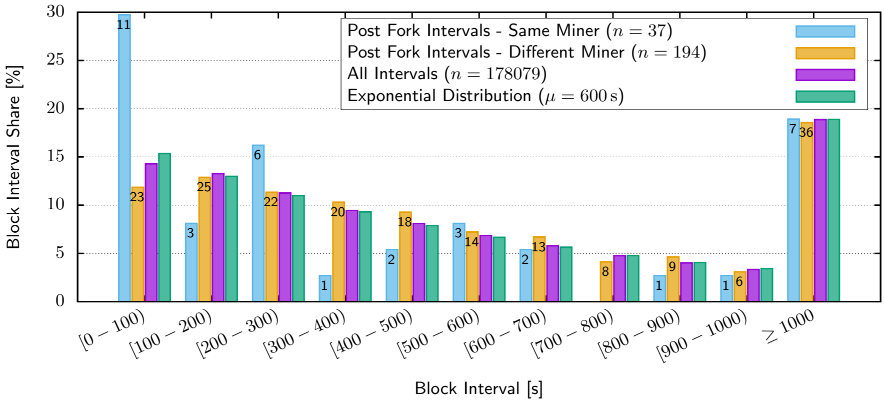

Fig. 3. Histogram of the interval to the next block after the fork. Post Fork Intervals - Same miner: only intervals where the subsequent block has been mined by the same miner as the previous block. Post Fork Intervals - Different miner: only intervals where the subsequent block has been mined by a different miner as the previous block. All Intervals: All block intervals since block 350,000. Exponential Distribution: Idealized block interval.

$i \in \{1, 0\}$). By sorting all tuples by the headstart, a moving average of the probability of main chain inclusion can be calculated. Fig. 2 shows the moving average of the probability of a block becoming part of the main chain depending on the headstart over the competing block. At the borders of the plot, the moving average window is reduced symmetrically, hence, the variance of the plot increases in these areas. Although the sample size of the data is small, a general trend can be seen, especially between 100 ms and 10 s. For this interval, Fig. 2 also shows a linear trend line ($y = 3.07 \cdot 10^{-5}x + 0.63$).

The data shows that a headstart of 100 ms results in a probability of main chain inclusion of around 70 %. After a short drop at a headstart of around 500 ms, the probability increases to more than 80 % for a headstart of 10 seconds. We emphasize that the data is dominated by the large number of forks until around mid 2017. It is likely that due to the reduced block propagation delay, today a smaller headstart leads to a much larger probability of main chain inclusion.

4.2 Deviating Mining Strategy

Consider a miner following the selfish mining strategy that withholds two blocks and receives a competing block for her first block withheld. In that case the selfish miner would publish both withheld blocks within a a short period of time in order to prevent the competing block from becoming included in the main chain, rendering the withheld blocks useless. Hence, very small block intervals after the occurrence of a fork can be caused by selfish mining. We will now analyze the block intervals after forks.

For all forks, we calculate the block interval between the first announcement of the block that got included in the main chain and the first announcement of the subsequent block (i.e., the block at the next height). We split all forks we have data about into two groups: Group *Same Miner* contains all 37 forks where the block that got included in the main chain and the subsequent block has both been mined by the same miner (also shown in Fig. 1 as rectangles). Group *Different Miner* contains all 194 forks where both blocks were mined by different miners. Please note that the miner attribution is done based on information embedded by the miner in the block, which can be freely set by the miner. For comparison, we also calculate all block intervals since block 350,000.

Fig. 3 shows histograms of the block interval for the groups *Same Miner* and *Different Miner* along with all block intervals and an idealized block interval distribution modeled by an exponential distribution. While the relative frequencies of all groups correspond well for larger block intervals, major differences can be observed for the smallest interval ($< 100\,\mathrm{s}$): Out of the 37 forks with the same miner, 11 forks (30 %) had a block interval of less than 100 seconds between the fork and the subsequent block. Contrary, only 23 forks of the 194 forks with different miners had a block interval of less than 100 seconds (12 %). The expected relative frequency is in the order of 14 % (measured) or 15 % (idealized).

We will now discuss possible reasons for the observed deviation. First, although a validation of our measurements with other data shows a high correspondence, we cannot completely rule out measurement errors. Secondly, the probability that 11 or more samples out of 37 samples of the idealized block interval distribution are smaller than 100 s is around 2 %. Therefore, while the observation seems unlikely, there is a substantial probability that the observation is simply the result of the random mining process and the small sample size.

Thirdly, the presence of block propagation delays makes the considered events statistically dependent. For instance, if a block interval is smaller than the block propagation delay, the subsequent block is definitely mined by the same miner, as other miners did not receive the previous block yet. However, the peculiar relative frequency shown in Fig. 3 corresponds to the conditional probability of observing a small block interval *given that* a fork occurred and both blocks were mined by the same miner. The existence of a fork is independent of the next block interval, as the mining power remains constant (although split). However, the block propagation delay gives the miner of the last block an advantage in finding the subsequent block, until other miners have received the block. Therefore, during block propagation, the overall mining power is reduced to the mining power of the miners that have already received the block. Hence, the overall block interval should actually increase (minimally) compared to the idealized block interval as modeled by an exponential distribution. Furthermore, all observed block intervals in the *Same Miner* group are at least 40 seconds, hence block propagation delay should not affect the interval, as the advantage of the miner vanishes as soon as other miners receive a block.

Finally, selfish mining could be the cause for the observed block intervals. Blocks of 9 of the 11 forks with block intervals below 100 seconds were mined by only two different mining pools, which had a share of the network hash rate of around 20 % and 10 %, respectively, at the time. Hence a single mining pool

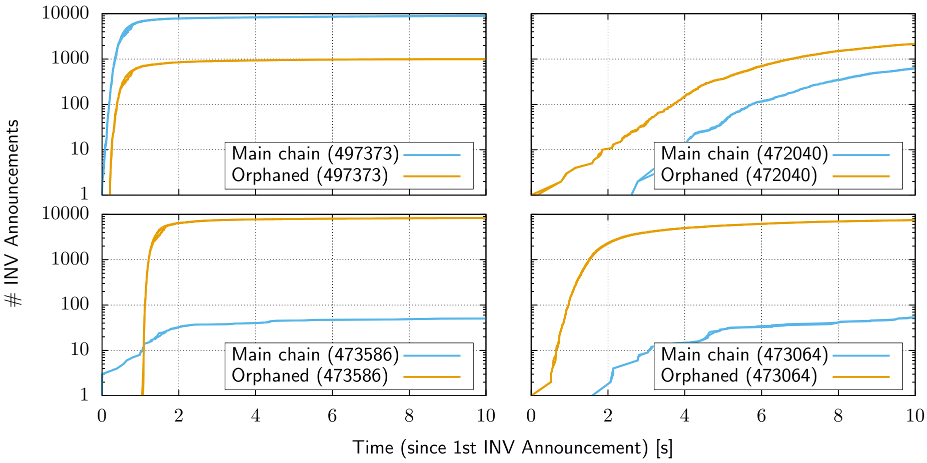

Fig. 4. Block propagation for both blocks (*Main chain*, *Orphaned*) for selected forks.

following the selfish mining strategy could have caused the observed deviations. However, the fact that all observed block intervals were at least 40 seconds raises doubts that selfish mining was actually performed, because one would expect miners to immediately publish the subsequent block. Furthermore, one would not expect a selfish miner to voluntarily include information about its identity in a mined block. Finally, the mining power shares of the pools render selfish mining only lucrative when assuming a significant network advantage γ [3].

Although all discussed possible causes for the observed block intervals seem unlikely, the presented data provides insights into a specific aspect of the mining process and can serve as a starting point for further research.

4.3 Peer-to-Peer Propagation Comparison

The differences in the time a block has been announced shown in Fig. 1 only show that a block has not been received by a miner within this time difference, but do not give reasons for why the block has not reached the other miner. Fig. 4 shows the P2P propagation of the blocks that caused four different selected forks.

The fork at height 497373 (top left) shows the *standard* case: The main chain block is propagated slightly before the orphaned block, which is announced only by those peers that have not already received the main chain block. In contrast, in the fork at height 472040 (top right) the orphaned block is propagated first, however, it propagates very slowly. More than one second later, the included block is propagated at a similar propagation speed. In the fork at height 473586 (bottom left) the main chain block is propagated first, but is only announced by less than 100 peers within 10 seconds. Contrary, the orphaned block is published one second later, but propagates very fast through the network. Finally, the main chain block at height 473064 (bottom right) is published more than a second after the well propagated orphaned block, but still it became part of the main chain.

The examples show that not only the first announcement of a block plays a role in which block becomes part of the main chain, but also the propagation speed of each block. However, as all combinations of headstart (positive vs. negative) and propagation speed (slow vs. fast) could be observed, P2P propagation of blocks does not seem to be the main decisive factor in which block becomes part of the main chain. A possible reason for slow propagation speeds could be extremely long validation times for these blocks. For instance, if a block contains a transaction on which other transactions that are contained in a peer's mempool depend, the peer also has to validate and order these transactions. Additionally, the examples show that the propagation of each block can differ drastically, hence a purely statistical model of block propagation can be insufficient.

5 Conclusions & Future Work

We provided an empirical analysis of the announcement and propagation of Bitcoin blocks that caused blockchain forks. The large differences in the first announcements of competing blocks indicate that the block propagation delay between miners can be of similar order as the observed 50 % block propagation percentile. The probability of a block to become part of the main chain increases linearly in the headstart (i.e., the time the block has been published before the competing block) between 100 ms and 10 s (from less than 70 % to more than 80 %). The observed frequency of block intervals between two consecutive blocks mined by the same miner to be less than 100 seconds is conspicuously large. While selfish mining can be a cause for this observation, other causes are also possible.

A better understanding of the factors influencing the propagation speed of specific blocks might be gained by an in-depth analysis of (orphaned) blocks and client implementations. Furthermore, our analysis might profit from more data, especially on recent forks. While the reduced frequency of forks is generally good for the system, it is unfortunate for empirical research.

Acknowledgments

This work was supported by the German Federal Ministry of Education and Research within the projects *KASTEL_IoE* and *KASTEL_ISE* in the Competence Center for Applied Security Technology (*KASTEL*) and by the state of Baden-Württemberg through bwHPC, bwFileStorage, and LSDF. The authors would like to thank Tim Ruffing and the anonymous reviewers for their valuable comments and suggestions.

References

1. Decker, C., Wattenhofer, R.: Information propagation in the Bitcoin network. In: Peer-to-Peer Computing (P2P), 2013 IEEE Thirteenth International Conference on. pp. 1–10. IEEE (2013)

2. Delgado-Segura, S., Pérez-Solà, C., Herrera-Joancomartí, J., Navarro-Arribas, G., Borrell, J.: Cryptocurrency networks: A new p2p paradigm. Mobile Information Systems 2018 (2018)
3. Eyal, I., Sirer, E.G.: Majority is not enough: Bitcoin mining is vulnerable. In: International Conference on Financial Cryptography and Data Security. Springer (2014)
4. Gencer, A.E., Basu, S., Eyal, I., van Renesse, R., Sirer, E.G.: Decentralization in bitcoin and ethereum networks. arXiv preprint arXiv:1801.03998 (2018)
5. Gervais, A., Karame, G.O., Wüst, K., Glykantzis, V., Ritzdorf, H., Capkun, S.: On the security and performance of proof of work blockchains. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (2016)
6. Kwon, Y., Kim, D., Son, Y., Vasserman, E., Kim, Y.: 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. ACM (2017)
7. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008)
8. Narayanan, A., Bonneau, J., Felten, E., Miller, A., Goldfeder, S.: Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction. Princeton University Press (2016), `http://bitcoinbook.cs.princeton.edu/`